



HAL
open science

Quantitative finance at the microstructure scale : algorithmic trading and regulation

Bastien Baldacci

► **To cite this version:**

Bastien Baldacci. Quantitative finance at the microstructure scale : algorithmic trading and regulation. Economics and Finance. Institut Polytechnique de Paris, 2021. English. NNT : 2021IPPAX050 . tel-03375983

HAL Id: tel-03375983

<https://theses.hal.science/tel-03375983>

Submitted on 13 Oct 2021

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



INSTITUT
POLYTECHNIQUE
DE PARIS

NNT : 2021IPPAX050

Thèse de doctorat



Quantitative finance at the microstructure scale: algorithmic trading and regulation.

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à École polytechnique

École doctorale n°574 École doctorale de mathématiques Hadamard (EDMH)
Spécialité de doctorat : Mathématiques Appliquées

Thèse présentée et soutenue à Paris, le 14/05/2021, par

BASTIEN BALDACCI

Composition du Jury :

Huyên Pham Professor, Université Paris 7 Diderot, LPSM	Rapporteur
Mike Ludkovski Professor, University of California Santa Barbara, DSAP	Rapporteur
Olivier Guéant Professor, Université Paris 1 Panthéon-Sorbonne, CES	Président
Sophie Laruelle Maître de Conférences, Université Paris-Est Créteil, LAMA	Examineur
Thibaut Mastrolia Maître de Conférences, École Polytechnique, CMAP	Examineur
Nizar Touzi Professor, École Polytechnique, CMAP	Examineur
Mathieu Rosenbaum Professor, École Polytechnique, CMAP	Directeur de thèse
Dylan Possamaï Professor, ETH Zürich, Mathematics	Co-directeur de thèse
Grégoire Loeper Quantitative analyst, BNP Paribas	Invité

Remerciements

Tout d'abord, je tiens à remercier mes deux directeurs de thèse, Mathieu Rosenbaum et Dylan Pos-samaï. Sous leur direction, j'ai pu apprendre quantité d'éléments à la fois mathématiques mais également humains. Ces années de thèse furent pour moi un véritable apprentissage de vie, à bien des égards, et je leur suis reconnaissant pour la complète autonomie qu'ils m'ont offerte ainsi que leurs encouragements. En particulier, je remercie Dylan que je connais maintenant depuis plus de six années, sans qui l'idée d'une thèse n'aurait jamais traversé mon esprit. Un grand merci pour ton introduction auprès de Mathieu, prélude à ces travaux, et surtout ta gentillesse et ta bienveillance.

Je remercie très chaleureusement Mike Ludkovski et Huyên Pham d'avoir accepté de rapporter cette thèse. Leur intérêt pour les travaux qui y sont présentés m'honore particulièrement, leurs travaux en contrôle stochastique étant parmi les principales sources d'inspiration de ce manuscrit. Un remerciement tout particulier doit être apporté à Olivier Guéant pour sa remarquable présidence de jury toute en nuances, ce qui le caractérise parfaitement, mais également pour notre collaboration durant la thèse. J'ai tout particulièrement apprécié son attachement à faire cohabiter recherche académique et industrielle. Je remercie également Thibaut Mastrolia, avec qui j'ai pris beaucoup de plaisir à travailler durant ces trois années sur ce qui constitue le premier chapitre du manuscrit, pour sa gentillesse à toute épreuve : en témoigne sa participation de dernière minute au jury, et je l'en remercie grandement pour cela. Je remercie également Grégoire Loeper, dont j'apprécie particulièrement les travaux précurseurs liant valorisation et exécution, pour avoir accepté de faire partie du jury, également dans un court délai. Enfin, je remercie Sophie Laruelle et Nizar Touzi pour leur bienveillance tout au long de la soutenance, et plus largement pendant ces trois années où j'ai pu les côtoyer. En particulier, je remercie Sophie pour ses conseils avisés, l'intérêt porté à mes travaux et notre très agréable (et pluvieux) séjour islandais.

Cette thèse est le fruit d'un nombre certain de collaborations, aussi j'aimerais remercier l'ensemble de mes co-auteurs : Philippe Bergault, Olivier Guéant, Thibaut Mastrolia, Paul Jusselin, Joffrey Derchu, Jerome Benveniste, Gordon Ritter et Mike Weber. En premier lieu, Philippe, que je connais depuis plus de 8 ans, avec qui nous avons des trajectoires assez similaires, et un certain nombre de points communs (dont certains assez inattendus). Merci pour ta camaraderie, qui se manifeste beaucoup plus après quelques litres de bière, ton humour, nos débats politiques, ton souci du détail, et ton amitié. Merci également à Joffrey pour nos collaborations ainsi que ton flegme incomparable. Merci à Jerome et Gordon pour notre collaboration extrêmement stimulante, et vos idées novatrices, toujours en prise avec les marchés. Pour ce qui est de notre collaboration académique, merci à Iulia pour notre prix "Rising Star in Quant Finance" et nos nombreux articles sans lesquels cette thèse serait bien moins intéressante. Enfin, merci à Mike, dont la collaboration chez J.P. Morgan m'a permis d'aiguiser mes raisonnements.

Je remercie également l'ensemble des personnes que j'ai pu côtoyer pendant la durée de cette thèse. À l'École Polytechnique, je remercie Kaitong Hu pour son calme et sa gentillesse à toute épreuve, Stefano De Marco pour sa sympathie et son bon goût vestimentaire, Giovanni Conforti et Benoit Dagallier pour nos deux années de cours de probabilité auprès des Bachelor, riches d'un point de vue sociologique, Pamela Saliba pour sa joie de vivre, Othmane Mounjid pour son optimisme et sa gentillesse, Mehdi Tomas pour nos débats politiques, Mehdi Talbi pour nos débats politiques (dans un sens légèrement différent), Heythem Farhat pour sa gentillesse, Antoine Fosset pour avoir sauvé l'honneur au trading game, et Marcos Carreira pour son haut patronage de l'équipe Quantreg. Un merci tout particulier à Eduardo Abi Jaber (et bien sûr Maya) pour ses margaritas et sa rigueur (un brin obsessive) mathématique que j'admire, et à Sergio Pulido pour son rire sans pareil et son côté dilettante combiné à une technicité qui force l'admiration également. Merci à Anthony Darmon pour nos aventures immobilières dignes d'un livre de science fiction, et à Albert Darmon pour son bon sens et sa sagesse. Merci à Benjamin Chen pour sa modestie, Stéphane Durand pour son abnégation et Paul Pelletier pour l'ensemble de son oeuvre.

Un remerciement tout particulier se doit d'être apporté à certains membres du personnel administratif de l'École Polytechnique. Ainsi, un immense merci à Alexandra Noiret et Nassera Naar pour leur grande disponibilité, leur gentillesse, leur écoute et leur sens de la diplomatie. Merci également à Thierry Bodineau pour sa direction du département de mathématiques appliquées.

Cette thèse fut également l'occasion de découvrir, par le fruit du hasard et un déjeuner organisé par Kaitong, un monde relativement orthogonal au trading algorithmique : celui de la réassurance. Je remercie très chaleureusement Yann Ait-Mokhtar avec qui nous travaillons depuis un certain temps sur ces sujets passionnants d'économie quantitative, mais au-delà pour son amitié, nos discussions politiques, philosophiques, nos déjeuners et dîners. Je remercie également Vincent Gros pour sa gentillesse et son talent littéraire incomparable. Un très grand merci à Jacques Tierny pour son optimisme et sa passion de l'analyse financière, et Nathalie Mercier pour son œil extrêmement attentif. Merci également à Thibaut Gress pour nos nombreux mails très variés. Je remercie également l'ensemble des personnes avec qui j'ai pu interagir au sein de la Caisse Centrale de Réassurance, et en particulier Bertrand Labilloy pour la confiance qu'il m'a accordé, Pascale Demarcq pour son aide tout du long, ainsi que David Moncoulon, Thierry Cohignac et Antoine Quantin pour leur bienveillance et leurs précieux conseils.

Ces années furent également l'occasion de travailler sur un projet européen Digital Europe en Corse. Monter ce projet avec les personnes associées est un réel plaisir qui, je l'espère, se concrétisera de la meilleure des manières. Ainsi, un grand merci à Sébastien Simoni, notamment pour sa résistance aux spritz, Dominique Siacci, Benjamin Pereney et Laetitia Sabatini pour nos magnifiques réunions Zoom/Teams afin d'avancer dans la co-construction et le développement de solutions capacitantes et éthiques. Merci également à Marie-Pascale Simoni de me supporter lors de nos soirées "travail" avec Sébastien.

Je tiens également à remercier très chaleureusement mes nouveaux collègues de l'équipe Equity Derivatives d'HSBC Paris pour leur accueil : Manuel Abellan-Lopez, Arnaud Gocsei, Arnaud Ou, André Chhoa, Jaehong Kim et Hind Zegmout. En particulier, je remercie Manuel pour sa bienveillance et ses très nombreuses idées sur les problématiques de calibration.

Il reste alors à remercier ma famille. Par peur d'absence d'exhaustivité, je me limiterai ici au premier cercle. Ainsi, merci à Louise (notamment pour ses cours du week-end qui ne sont pas totalement décorrélés de mon parcours), Jean-Claude, Marie, Dume, Cosima, Paola, Dominique, Pierre-Marie, Cosima, Julie, Antoine, Cristina (pour énormément de choses), Uma, Flam, Paule, Marie-Sé, Mo, Lina, Yassine, Dominique, Philippe, Mado, Joseph, Antoine-Marcel et bien sûr Thibaut. Un immense merci à mes parents, Paul-Marie et Marie-Françoise, pour tout ce que vous avez pu (et continuez) à m'apporter chaque jour. Enfin merci à Iuliia d'être à mes côtés et d'enchanter ma vie depuis plus de 2 ans.

Abstract

This thesis is split into three parts. In the first part, we apply the Principal-Agent theory to some problems of market microstructure. In Chapter 1, we show how to design an incentives policy in order to improve the market quality in the context of market-making activity in a lit and a dark pool managed by the same exchange. We derive a general form of incentives offered to the market-maker, indexed on the transactions and the dynamics of the price process. At the optimum, the incentives are obtained via the resolution of a Hamilton-Jacobi-Bellman (HJB) equation with deep reinforcement learning methods. We show numerically that the optimal incentives lead to an improvement of market liquidity on the dark pool. Chapter 2 is devoted to the regulation of market-making activity when several market-makers compete on an underlying listed on a single liquidity platform. A specific form of incentives indexed on the transactions of all market participants, which can be computed using classical finite differences, leads to a unique Nash equilibrium between the market-makers. We then obtain an analytic approximation of the optimal number of market-makers in terms of profits for the exchange. One finding is that infinitely many market-makers is not optimal for the exchange. In Chapter 3, we propose a form of incentives based on the choice of tick sizes on the bid and ask sides of a single asset. We first develop a high-frequency market-making model following the model with uncertainty zones. Then, we show that interactions between the market-maker and the exchange lead to optimal asymmetric tick sizes: If the market-maker is reluctant to end the trading period with a short position, the exchange sets a higher tick size on the bid side. In Chapter 4, we tackle the issue of designing a derivatives market. First, we propose a quantization method to select the options listed on the exchange, in order to match the market demand with a finite number of options. Second, we use the Principal-Agent framework developed in Chapters 1 and 2 to create an incentives mechanism for an option market-maker. The goal is to increase the global liquidity provision and to reduce the width of the spreads on illiquid options. By considering the aggregate inventory of the market-maker, the exchange can propose tractable incentives, whose efficiency is illustrated numerically. To conclude the first part of the thesis, we develop in Chapter 5 an incentives mechanism to increase the investment in green bonds. The optimal incentives proposed by a government to an investor are function of the price dynamics of the bonds and of its portfolio process, as well as the (co)-variations between these quantities. In the case of deterministic drift and volatility processes, the optimal incentives provided to the investor easily obtained using classical root-finding algorithms. We show that our methodology is robust to model specification, by considering stochastic drift of the price processes, and that it outperforms the current tax-incentives policies of the governments.

The second part of this thesis is dedicated to option market-making in high dimension. In Chapter 6, we propose a framework using stochastic control to deal with long-dated options. In this case, the risk faced by the market-maker is the Vega of its portfolio. By using a constant Greek assumption, we show that the optimal bid and ask quotes of the market-maker on an arbitrary large portfolio of options on a single underlying are obtained through the resolution of a two-dimensional HJB equation. Then, in Chapter 7, we propose an approximation of the value function enabling to deal with time-varying Greeks. We show the evolution of the bid and ask quotes with respect to the spot and stochastic volatility and that it outperforms the strategy developed in Chapter 6 when dealing with short-dated options. Finally, in Chapter 8, we develop a framework for the high-frequency dynamics of the implied volatility surface. Using multidimensional Hawkes processes, we show how this setting can reproduce easily well-known stylized facts such as the skew, smile and term structure of the surface. We provide sufficient conditions to ensure absence of statistical arbitrage of the volatility surface. At the macroscopic limit, that is for a large time horizon, we show that the dynamics of the surface is a sum of risk factors with rough volatility. We illustrate the usefulness of our methodology with two applications: backtesting of option market-making strategies and computation of the associated market impact curves.

The third and last part of this thesis is devoted to optimal trading problems in high dimension. In Chapter 9, we develop a framework for optimal trading on stocks listed on several limit order

books. For a small number of venues, the optimal volumes and limits selected by the trader on each liquidity platform are obtained by solving a system of ordinary differential equations. This can be carried through finite differences methods. When the model is less parcimonious, we use a deep reinforcement learning approach to compute the optimal controls of the trader. Moreover, we propose a Bayesian update of the market parameters decoupled from the control problem: We divide the trading horizon into slices over which the parameters remain constant. At the end of each slice, given the quantities observed by the trader (spread, imbalance, price process, filling of limit orders), the parameters are updated for the next slice. In the case of optimal trading on a portfolio of assets listed on several liquidity venues, the control problem suffers from the curse of dimensionality. Thus, in Chapter 10, we propose a methodology to solve approximately optimal trading problems without using stochastic control theory. We start by using a “high-level” value function, for example the value function corresponding to the control problem in Almgren and Chriss [10], that can be computed analytically. The optimal trajectory coming from this control problem neglects the microstructure effects such as passive filling, spread and imbalance of the venues. We then propose a framework in which a myopic agent (that is, an agent that only maximizes the expectation of its next trade) exhibits approximately an optimal behavior, in the sense that it follows the trajectory described by the high-level value function, if he uses the gradient of the value function as short-term alpha. The problem then boils down to a simple static optimization, in which we can include all the desired microstructure effects. In particular, we can adjust the short-term alphas of the myopic trader such that the long-term trading curve is tracked optimally, and we provide uniform bounds of the error. Finally, in Chapters 11 and 12, we present two new developments on the optimal execution literature. First, we show that we can obtain a closed-form solution for the Almgren-Chriss execution problem with geometric Brownian motion and quadratic penalty. This solution is obtained using some tools from operator’s theory, and is robust to the curse of dimensionality: the execution of a portfolio of correlated assets following geometric Brownian motions is solved using a system of ODEs. Second, we propose an application of the latent order book model to the problem of optimal execution of a portfolio of assets, in the context of liquidity stress testing. We find that the optimal terminal time for a linear execution of a small order is proportional to the square root of the ratio between the amount being bought or sold and the average daily volume. The accuracy of the numerical method is illustrated with real market data on a portfolio of corporate bonds.

Résumé

Cette thèse est divisée en trois parties. Dans la première partie, nous appliquons la théorie Principal-Agent à certains problèmes de microstructure du marché. Dans le chapitre 1, nous montrons comment concevoir une politique d'incitation afin d'améliorer la qualité de la liquidité de marché dans le cadre d'une activité de market-making dans un lit et un dark pool gérés par la même bourse d'échange. Nous dérivons une forme générale d'incitations offertes au market-maker, indexées sur les transactions et la dynamique du processus de prix. À l'optimum, les incitations sont obtenues via la résolution d'une équation de Hamilton-Jacobi-Bellman (HJB) par des méthodes d'apprentissage par renforcement profond. Nous montrons numériquement que les incitations optimales conduisent à une amélioration de la liquidité de marché sur le dark pool. Le chapitre 2 est consacré à la régulation de l'activité de market-making lorsque plusieurs market-makers sont en concurrence sur un actif coté sur une seule plateforme. Une forme spécifique d'incitations indexées sur les transactions de tous les acteurs du marché, pouvant être calculées à l'aide de différences finies classiques, conduit à un équilibre de Nash unique entre les market-makers. On obtient alors une approximation analytique du nombre optimal de market-makers en termes de bénéfices profits pour la bourse d'échange. En particulier, il n'est pas optimal d'avoir un nombre infini de market-makers concurrents. Dans le chapitre 3, nous proposons une forme d'incitation basée sur le choix des tailles de ticks asymétriques à l'achat et à la vente sur un actif. Nous développons d'abord un modèle de market-making haute fréquence en suivant le modèle à zones d'incertitude. Ensuite, nous montrons que les interactions entre le market-maker et la bourse d'échange conduisent à des tailles de ticks asymétriques optimales: si le market-maker est réticent à terminer la période de trading avec une position courte, la bourse fixe une taille de tick plus élevée du côté achat. Dans le chapitre 4, nous abordons la question de la conception d'un marché de produits dérivés. Dans un premier temps, nous proposons une méthode de quantification pour sélectionner les options cotées en bourse qui correspondent à la demande de marché. Dans un second temps, nous utilisons la modélisation Principal-Agent développée dans les chapitres 1 et 2 pour fournir des incitations à un market-maker d'options. L'objectif est d'augmenter la provision de liquidité et de réduire la taille des spreads sur les options illiquides. En considérant l'inventaire agrégé du market-maker, la bourse peut proposer des incitations sur un grand nombre d'options, dont l'efficacité est illustrée numériquement. Pour conclure la première partie de la thèse, nous développons dans le chapitre 5 un mécanisme d'incitations pour augmenter l'investissement dans les obligations vertes. Les incitations optimales proposées par un gouvernement à un investisseur sont fonction de la dynamique des prix des obligations et du portefeuille d'obligations, ainsi que des (co)-variations entre ces quantités. Dans le cas où les processus de drift et de volatilité sont déterministes, les incitations optimales fournies à l'investisseur sont facilement obtenues à l'aide d'algorithmes classiques de recherche de zéros. Nous montrons que notre méthodologie est robuste à la spécification du modèle, en prenant en compte le drift stochastique dans les processus de prix, et qu'elle surpasse les politiques d'incitations fiscales actuelles des gouvernements.

La deuxième partie de cette thèse est consacrée au market-making d'options. Dans le chapitre 6, nous proposons un modèle de contrôle stochastique pour traiter les options de longue maturité. Dans ce cas, le risque auquel le market-maker est exposé est celui de Vega. En faisant l'hypothèse de grecques constants, nous montrons que les politiques optimales d'achat et vente du market-maker sur un large portefeuille d'options sur un seul sous-jacent sont obtenus par la résolution d'une équation HJB bidimensionnelle. Puis, dans le chapitre 7, nous proposons une approximation de la fonction valeur permettant de traiter les grecques non-constants. Nous montrons l'évolution des politiques d'achat et de vente par rapport au niveau du spot et à la volatilité stochastique et observons que ces politiques surperforment les stratégies développées au chapitre 6 en ce qui concerne les options de courtes maturités. Enfin, dans le chapitre 8, nous développons un modèle pour la dynamique haute fréquence de la surface de volatilité implicite. En utilisant des processus Hawkes multidimensionnels, nous montrons comment ce modèle peut reproduire de nombreux faits stylisés tels que le skew, le smile et la structure par termes de la surface. Nous fournissons également des conditions suffisantes pour

garantir l'absence d'arbitrage statistique de la surface de volatilité. A la limite macroscopique, c'est-à-dire pour un horizon temporel lointain, nous montrons que la dynamique de la surface est une somme de facteurs de risques à volatilité rugueuse. Nous illustrons l'efficacité de notre méthodologie avec deux applications: le backtesting des stratégies de market-making d'options et le calcul des courbes d'impact de marché associées.

La troisième et dernière partie de cette thèse est consacrée aux problèmes de trading optimal en grande dimension. Dans le chapitre 9, nous développons un modèle pour le trading optimal d'actions listées sur plusieurs plateformes. Pour un petit nombre de plateformes, les volumes et limites optimaux sélectionnés par le trader sur chaque plateforme sont obtenus en résolvant un système d'équations différentielles classiques via des méthodes de différences finies. Lorsque le modèle est moins parcimonieux, nous utilisons une méthode d'apprentissage par renforcement profond pour calculer les contrôles optimaux du trader. De plus, nous proposons un update bayésien des paramètres de marché qui est découplé du problème de contrôle: nous divisons l'horizon de trading en tranches d'exécution sur lesquelles les paramètres restent constants. A la fin de chaque tranche, compte tenu des quantités observées par le trader (spread, imbalance, processus de prix, exécution des ordres limites), les paramètres sont mis à jour pour la tranche suivante. Dans le cas où l'on veut traiter sur un portefeuille d'actifs cotés sur plusieurs plateformes, le problème de contrôle souffre du fléau de la dimension. Ainsi, dans le chapitre 10, nous proposons une méthodologie pour résoudre des problèmes de trading de façon approximativement optimale sans utiliser la théorie du contrôle stochastique. Nous commençons par utiliser une fonction valeur macroscopique, par exemple la fonction valeur correspondant au problème de contrôle de Almgren and Chriss [10], qui peut être calculée explicitement. La trajectoire optimale de trading issue de ce problème de contrôle néglige les effets de microstructure tels que l'utilisation d'ordres limites, le spread et l'imbalance des différentes plateformes. Nous proposons alors un modèle dans lequel un agent myope (c'est-à-dire un agent qui ne maximise que l'espérance de son prochain trade) présente un comportement approximativement optimal, en ce sens qu'il suit la trajectoire décrite par la fonction valeur macroscopique, s'il utilise le gradient de cette fonction valeur comme signal de court terme. Le problème devient alors une simple optimisation statique, dans laquelle nous pouvons inclure tous les effets de microstructure souhaités. En particulier, nous pouvons ajuster les signaux de court terme du trader myope de telle sorte que la courbe de trading de long terme soit suivie de manière optimale, et nous fournissons des bornes uniformes de l'erreur d'approximation. Enfin, dans les chapitres 11 et 12, nous présentons deux nouveaux développements sur la littérature d'exécution optimale. Tout d'abord, nous montrons que nous pouvons obtenir une solution analytique au problème d'exécution d'Almgren-Chriss avec mouvement Brownien géométrique et pénalité quadratique. Cette solution est obtenue à l'aide d'outils de la théorie des opérateurs, et est robuste au fléau de la dimension : le problème d'exécution d'un portefeuille d'actifs corrélés suivant des mouvements Browniens géométriques est résolu à l'aide d'un système d'EDO. Deuxièmement, nous proposons une application du modèle de carnet d'ordres latent au problème d'exécution optimale d'un portefeuille d'actifs, dans le cadre de stress tests de liquidité. Nous constatons que le temps terminal optimal pour une exécution linéaire d'un ordre de faible taille est proportionnel à la racine carrée du quotient entre le montant acheté ou vendu et le volume quotidien moyen échangé. La précision de la méthode numérique est illustrée sur un portefeuille d'obligations d'entreprises.

Contents

Introduction	2
Part I Market-making regulation activity via Principal-Agent theory	3
I.1 Market-making and incentives design in the presence of a dark pool: a deep reinforcement learning approach	3
I.2 Optimal make-take fees in a multi-market-maker environment	8
I.3 On bid and ask side-specific tick sizes	10
I.4 Derivatives market design	12
I.5 Governmental incentives for green bonds investment	14
Part II Options market-making in high dimension	17
II.1 Option market-making algorithms in high dimension	17
II.2 High-frequency dynamics of the volatility surface	22
Part III Optimal trading in high dimension	25
III.1 Adaptive trading strategies across liquidity pools	25
III.2 Optimal trading without optimal control	28
III.3 Two new developments on optimal execution	31
Part I Regulation of market-making activity via Principal-Agent theory	34
1 Market-making and incentives design in the presence of a dark pool: a deep reinforcement learning approach	36
1.1 Introduction	36
1.2 The market model	39
1.2.1 Stochastic basis	39
1.2.2 Traded volumes, market impact and latency	40
1.3 Market-making without the intervention of the exchange	43
1.4 Market-making with the intervention of the exchange	44
1.4.1 Modified PnL of the market-maker	44
1.4.2 Objective function of the exchange	45
1.4.3 Design of an optimal make-take fees policy	45
1.4.4 Problem of the exchange	47
1.5 Numerical solution: a deep reinforcement learning approach	50
1.5.1 Description	50
1.5.2 Numerical results: influence of the make-take fees policy	54
1.5.3 Numerical results: sensitivity analysis	57
1.6 Conclusion	58
1.A Appendix	59
1.A.1 Dynamic programming principle and contract representation	59
1.A.2 Proof of Theorem 1.2	60
1.A.3 Proof of Theorem 1.3	63
2 Optimal make-take fees in a multi-market-maker environment	65
2.1 Introduction	65

2.2	The model	67
2.2.1	Framework	67
2.2.2	Market-makers' problem	70
2.2.3	The exchange optimal contracting problem	72
2.2.4	Stackelberg games in a nutshell	72
2.3	Solving the market-maker's problem	73
2.3.1	Preliminaries	73
2.3.2	Contract representation	74
2.3.3	On the shape of compensation proposed and contractible variables.	75
2.4	Solving the principal's problem	76
2.4.1	Saturation of utility constraint	76
2.4.2	The HJB equation for the reduced exchange problem	77
2.4.3	Change of variable and verification theorem	78
2.4.4	Discussion	79
2.5	Impact of the presence of several market-makers	82
2.5.1	One market-maker	82
2.5.2	Two market-makers	83
2.5.3	Five market-makers	85
2.A	Appendix	86
2.A.1	Dynamic programming principle	86
2.A.2	Proof of Theorem 2.1	88
2.A.3	Proof of Lemma 2.1	93
2.A.4	Exchange's Hamiltonian maximisation	93
2.A.5	Proof of Lemma 2.3	94
2.A.6	Proof of Theorem 2.2	94
2.A.7	First-best exchange problem	96
3	On bid and ask side-specific tick sizes	99
3.1	Introduction	99
3.2	The model with uncertainty zones	101
3.3	High frequency market-making under side-specific tick values and interaction with the exchange	102
3.3.1	The market-maker's problem	102
3.3.2	The platform's problem	106
3.4	Numerical results	106
3.4.1	Similar tick values on both sides	106
3.4.2	Side-specific tick values: additional opportunities for the market-maker	107
3.4.3	Side-specific tick values: effect of ϕ_-	108
3.5	Conclusion	110
3.A	Appendix	111
3.A.1	Proof of Proposition 3.1	111
3.A.2	Proof of Theorem 3.1	113
3.A.3	Effects of the uncertainty zones on h	119
4	How to design a derivatives market?	120
4.1	Introduction	120
4.2	Market driven selection of the listed options	121
4.2.1	How to choose the strikes in order to match market demand?	122
4.2.2	Solving the quantization problem	122
4.2.3	Application	124
4.3	Incentive policy of the exchange	127
4.3.1	The market	127

4.3.2	Market maker’s problem and contract representation	130
4.3.3	Solving the exchange’s problem	132
4.3.4	Numerical results	134
4.3.5	Conclusion	135
4.A	Appendix	136
4.A.1	Proof of the convergence of Lloyd’s algorithm	136
4.A.2	Stochastic basis	136
4.A.3	Well-posedness of the optimization problems	137
4.A.4	Dynamic programming principle	137
4.A.5	Proof of Lemma 4.1	139
4.A.6	Proof of Theorem 4.1	141
4.A.7	Proof of Theorem 4.2	142
4.A.8	Perfect and imperfect Delta-hedging	144
5	Governmental incentives for green bonds investment	147
5.1	Introduction	147
5.2	Framework	150
5.3	Solving the optimisation problem	153
5.3.1	The optimal contract	153
5.3.2	Discussion	158
5.4	Numerical results	160
5.4.1	Data, key results and remarks for the policy-maker	160
5.4.2	Reference case	162
5.4.3	Influence of the green target	166
5.4.4	Sensitivity analysis	168
5.A	Appendix	171
5.A.1	Weak formulation of the problem	171
5.A.2	Proof of Theorem 5.1	172
5.A.3	Green investments with stochastic interest rates	173
Part II	Options market-making in high dimension	177
6	Algorithmic market-making for options	178
6.1	Introduction	178
6.2	Description of the problem	180
6.2.1	The market	180
6.2.2	The optimization problem of the market-maker	181
6.2.3	Assumptions and approximations	183
6.3	An approximate solution to the problem	184
6.3.1	Change of variables: beating the curse of dimensionality	184
6.3.2	Hamilton-Jacobi-Bellman equation and optimal controls	184
6.4	Numerical results	185
6.4.1	Model parameters	185
6.4.2	Optimal quotes	186
6.4.3	Conclusion	190
6.A	Appendix	190
6.A.1	An alternative to the Δ -hedging assumption	190
6.A.2	Beyond the constant-Vega assumption	191
6.A.3	On the construction of the processes $N^{i,b}$ and $N^{i,a}$	192
7	An approximate solution for options market-making in high dimension	194

7.1	Introduction	194
7.2	Framework	196
7.2.1	The option book	196
7.2.2	The market-maker's problem on OTC markets	197
7.3	Solving the market-maker's problem with a system of non-linear PDEs	199
7.4	Numerical results	200
7.A	Appendix	206
7.A.1	The market-maker's problem for large number of underlyings	206
8	High-frequency dynamics of the volatility surface	207
8.1	Introduction	207
8.2	Microscopic modelling of the volatility surface	209
8.2.1	Framework and no-arbitrage conditions	209
8.2.2	Parametrization of the volatility surface	215
8.3	Macroscopic limit of the volatility surface	217
8.3.1	Limiting processes in the separable, one-dimensional case	218
8.3.2	Limiting process in the semi-separable, factor case	219
8.3.3	General case	220
8.4	Applications	222
8.4.1	Backtesting of option market-making strategies	222
8.4.2	Market impact curves	224
8.5	Conclusion	225
Part III	Optimal trading in high dimension	226
9	Adaptive trading strategies across liquidity pools	227
9.1	Introduction	227
9.2	Optimal trading on several liquidity pools	229
9.2.1	Framework	229
9.2.2	The Hamilton-Jacobi-Bellman quasi-variational inequality	231
9.3	Adaptive trading strategies with Bayesian update	232
9.3.1	Bayesian update of the model parameters	232
9.3.2	Algorithm description	235
9.4	Model extensions	235
9.4.1	Extension 1: incorporation of signals in the price process	236
9.4.2	Extension 2: market impact	236
9.4.3	Extension 3: hidden liquidity	237
9.5	Numerical results	237
9.5.1	Global parameters	237
9.5.2	Numerical methods	238
9.5.3	Two identical venues	242
9.5.4	Two different venues	247
9.5.5	Bayesian update	250
9.A	Appendix	255
9.A.1	Proof of Theorem 9.3	255
9.A.2	Application to OTC market-making	260
9.A.3	Bayesian update for OTC market-makers	261
10	Optimal trading without optimal control	264
10.1	Introduction	264
10.2	The long-term trading curve	266

10.3	From smooth relaxation to microstructure decision	269
10.4	A general microstructure trading model with long-term trading schedule	272
	10.4.1 Optimal trading	272
	10.4.2 Market-making	274
10.5	Numerical results	278
	10.5.1 Transaction cost model and microstructure simulation	278
	10.5.2 Results	279
10.6	Conclusion	282
11	A note on Almgren-Chriss optimal execution problem with geometric Brownian motion	283
11.1	Introduction	283
11.2	The model	284
	11.2.1 Almgren-Chriss framework in continuous time	284
	11.2.2 Reformulation in terms of cash	285
11.3	Solving explicitly the Almgren-Chriss problem with GBM	286
11.4	Numerical results	288
11.5	Extensions of the model	290
	11.5.1 Stochastic drift	290
	11.5.2 Multi-dimensional case	291
11.6	Conclusion	293
11.A	Appendix	293
	11.A.1 Proof of Lemma 11.1	293
	11.A.2 Gâteaux differentiability in (11.4)	294
	11.A.3 Proof of Lemma 11.2	294
	11.A.4 Proof of Theorem 11.2	294
12	Liquidity stress testing using optimal portfolio liquidation	296
12.1	Introduction	296
12.2	Framework	297
	12.2.1 The locally linear order book model	297
	12.2.2 Single asset liquidation	299
	12.2.3 Portfolio liquidation	300
12.3	Numerical results	300
	12.3.1 Long-short portfolio with two correlated bonds	301
	12.3.2 Long-short portfolio of 20 bonds	301
12.4	Conclusion	304

List of Figures

1	Optimal quotes of the market-maker with (right) and without (left) incentives.	7
2	Evolution of the PnL of the platform with the number of market-makers.	10
3	Value function of the market-maker(on the left) and the exchange (on the right) as functions of α^a and α^b , for $\phi_- = 0.005$	12
4	Some trajectories of the optimal portfolio process with and without contract.	17
5	Optimal mid-to-bid quotes divided by option price as a function of the portfolio Vega for $K=10$	20
6	Optimal ask quotes with respect to S and ν for options with maturuty $T = 0.6$	21
7	Evolution of the slice of volatility with sell orders on (k_3, τ) and $\phi^{(k_1, \tau)-, (k_3, \tau)-} > \phi^{(k_1, \tau)+, (k_3, \tau)-}$ (upper left), $\phi^{(k_1, \tau)-, (k_3, \tau)-} < \phi^{(k_1, \tau)+, (k_3, \tau)-}$ (upper right). Evolution of the slice of volatility with buy orders on (k_3, τ) and $\phi^{(k_1, \tau)-, (k_3, \tau)+} < \phi^{(k_1, \tau)+, (k_3, \tau)+}$ (lower left), $\phi^{(k_1, \tau)-, (k_3, \tau)+} > \phi^{(k_1, \tau)+, (k_3, \tau)+}$ (lower right).	23
8	Limit order strategy, $\psi^1 = \psi^2 = \delta$, $I^1 = -0.5, I^2 = 0.5$ using neural networks.	28
9	Volume strategy, $\psi^1 = \psi^2 = \delta$, $I^1 = -0.5, I^2 = 0.5$ using neural networks.	28
1.1	Best response of the market-maker as a function of $z^{a,l}$ and $z^{b,l}$, with $q = 50$	52
1.2	Best response of the market-maker as a function of $z^{a,l}$ and $z^{b,l}$, with $q = 150$	52
1.3	Optimal quotes of the market-maker.	56
1.4	Optimal quotes of the market-maker.	56
1.5	Optimal incentives of the exchange.	56
1.6	Optimal quotes of the market-maker without the exchange.	57
1.7	Optimal quotes of the market-maker with the exchange.	57
1.8	Optimal quotes of the market-maker without the exchange.	57
1.9	Optimal quotes of the market-maker with the exchange.	57
1.10	Optimal quotes of the market-maker without the exchange.	58
1.11	Optimal quotes of the market-maker with the exchange.	58
2.1	Total spread for 1 market-maker.	83
2.2	Ask order flow for 1 market-maker.	83
2.3	Total order flow for 1 market-maker.	83
2.4	PnL of the exchange for 1 market-maker.	83
2.5	Trading cost for 1 market-maker.	83
2.6	Total spread for $N = 2$, $[0.01, 0.001]$	84
2.7	Total spread for $N = 2$, $[0.01, 0.01]$	84
2.8	Total spread for $N = 2$, $[0.01, 0.1]$	84
2.9	Total order flow for $N = 2$, $[0.01, 0.01]$	84
2.10	Ask order flow for $N = 2$, $[0.01, 0.01]$	84
2.11	PnL of the exchange for $N = 2$, $[0.01, 0.01]$	84
2.12	Trading cost for $N = 2$, $[0.01, 0.01]$	84
2.13	Total spread for $N = 5$	85
2.14	Total order flow for $N = 5$	85

2.15	Ask order flow for $N = 5$.	85
2.16	PnL of the exchange for $N = 5$.	85
2.17	Trading cost for $N = 5$.	85
2.18	Evolution of the PnL of the platform with the number of market-makers.	85
3.1	Value function h (on the left) and v (on the right) for $\phi_- = 0$ in blue, $\phi_- = 0.0005$ in orange, $\phi_- = 0.005$ in green, as a function of $\alpha = \alpha^a = \alpha^b$.	107
3.2	Variation of the functions h and v (difference between $\phi_- = 0$ in blue, $\phi_- = 0.0005$ in orange, $\phi_- = 0.005$ in green, and $\phi_- = 0$ as a function of $\alpha = \alpha^a = \alpha^b$.	107
3.3	Value function h (on the left) and v (on the right) as functions of α^a , for $\alpha^b = 0.0124$, for $\phi_- = 0$ in blue, $\phi_- = 0.0005$ in orange, $\phi_- = 0.005$ in green.	108
3.4	Value function h (on the left) and v (on the right) as functions of α^a and α^b , for $\phi_- = 0$.	109
3.5	Value function h (on the left) and v (on the right) as functions of α^a and α^b , for $\phi_- = 0.005$.	109
3.6	Difference between the value function h (on the left) and v (on the right) as functions of α^a and α^b , between the case $\phi_- = 0.005$ and the case $\phi_- = 0$.	109
3.7	Value functions h and v for $\alpha^a = 0.0045$, as functions of α^b , for different values of ϕ .	110
3.8	Value function h of the market-maker for $q = 0$, as a function of S .	119
4.1	Empirical distribution of traded option strikes (left). Quantile plot in log-scale of traded option maturities for the whole sample set (right).	124
4.2	Empirical distribution of the strikes for different maturities.	125
4.3	Quantization of the option strikes using $p = 2$ and $\varepsilon = 10^{-8}$. Empirical distribution of traded strikes is plotted in blue or red. The dotted lines correspond to the optimal quantization of \mathbb{P}^{mkt} .	126
4.4	Quantization of the option strikes using $p = 8$ and $\varepsilon = 10^{-8}$. Empirical distribution of traded strikes is plotted in blue or red. The dotted lines correspond to the optimal quantization of \mathbb{P}^{mkt} .	127
4.5	Evolution of the spread (in ticks) at initial time with respect to inventory, $\omega = 0$.	135
4.6	Evolution of the spread (in ticks) at initial time with respect to inventory, $\omega = 0.1$.	135
4.7	Evolution of the spread (in ticks) at initial time with respect to inventory, $\omega = 0.2$.	135
5.1	Optimal investment policy (upper left), optimal incentives Z^* (upper right) and Γ^* (bottom) as a function of time.	163
5.2	Some trajectories of the optimal portfolio process with and without contract.	163
5.3	Average absolute difference of portfolio value over time, for 10000 simulations.	163
5.4	Optimal investment policy (left) and optimal incentives Z^* (right) as a function of time.	164
5.5	Optimal investment policy with stochastic rates.	164
5.6	Average relative difference (in %) of portfolio value over time for 10000 simulations.	165
5.7	Some trajectories of the optimal portfolio process with the optimal contract and with the tax-incentives policy (labeled ‘without contract’).	165
5.8	Optimal investment policy (upper left), optimal incentives Z^* (upper right) and Γ^* (bottom) as a function of time.	166
5.9	Average difference of portfolio value over time, for 10000 simulations.	166
5.10	Some trajectories of the optimal portfolio process with and without contract.	167
5.11	Average relative difference (in %) of portfolio value over time (with the optimal contract and tax-incentives policy), for 10000 simulations.	167
5.12	Some trajectories of the optimal portfolio process with the optimal contract and with the tax-incentives policy (labeled ‘without contract’).	168
5.13	Optimal investment policy (upper left), optimal incentives Z^* (upper right) and Γ^* (bottom) as a function of time.	168

5.14	Optimal investment policy (upper left), optimal incentives Z^* (upper right) and Γ^* (bottom) as a function of time.	169
5.15	Optimal investment policy (upper left), optimal incentives Z^* (upper right) and Γ^* (bottom) as a function of time.	169
5.16	Optimal investment policy (upper left), optimal incentives Z^* (upper right) and Γ^* (bottom) as a function of time.	170
5.17	Optimal investment policy (upper left), optimal incentives Z^* (upper right) and Γ^* (bottom) as a function of time.	170
6.1	Implied volatility surface associated with the above parameters.	186
6.2	Value function as a function of instantaneous variance and portfolio Vega.	186
6.3	Optimal mid-to-bid quotes as a function of time for option 1: $(K^1, T^1) = (8, 1) - \nu = 0.04$	187
6.4	Optimal mid-to-bid quotes divided by option price as a function of the portfolio Vega for $K=8$	187
6.5	Optimal mid-to-bid quotes divided by option price as a function of the portfolio Vega for $K=9$	187
6.6	Optimal mid-to-bid quotes divided by option price as a function of the portfolio Vega for $K=10$	188
6.7	Optimal mid-to-bid quotes divided by option price as a function of the portfolio Vega for $K=11$	188
6.8	Optimal mid-to-bid quotes divided by option price as a function of the portfolio Vega for $K=12$	188
6.9	Optimal (relative) bid implied volatility as a function of the portfolio Vega for $K=8$	188
6.10	Optimal (relative) bid implied volatility as a function of the portfolio Vega for $K=9$	189
6.11	Optimal (relative) bid implied volatility as a function of the portfolio Vega for $K=10$	189
6.12	Optimal (relative) bid implied volatility as a function of the portfolio Vega for $K=11$	189
6.13	Optimal (relative) bid implied volatility as a function of the portfolio Vega for $K=12$	189
7.1	Implied volatility surface associated with the market parameters.	201
7.2	Value function for different inventories in (97, 0.3) and (98, 0.3) options, inventories in other options assumed to be equal 0, for different values of ν	202
7.3	Value function for different inventories in (97, 0.3) and (100, 0.7) options, inventories in other options assumed to be equal 0, for different values of ν	203
7.4	Optimal ask quotes with respect to ν for different options maturities.	203
7.5	Optimal ask quotes with respect to S for different options maturities.	204
7.6	Average PnL per request over the trading day using constant and non-constant Greek approximations.	205
7.7	Example of Vega trajectories.	205
7.8	Cumulative distribution functions of the PnL over the trading day using both methods.	205
8.1	Evolution of the slice of volatility when $\phi^{(k_1, \tau)-, (k_2, \tau)+} > \phi^{(k_1, \tau)+, (k_2, \tau)+}$ (left), $\phi^{(k_1, \tau)-, (k_2, \tau)+} < \phi^{(k_1, \tau)+, (k_2, \tau)+}$ and there is less buy orders on option (k_2, τ) (center), $\phi^{(k_1, \tau)-, (k_2, \tau)+} < \phi^{(k_1, \tau)+, (k_2, \tau)+}$ and there is more buy orders on option (k_2, τ) (right)	211
8.2	Evolution of the slice of volatility with sell orders on (k_3, τ) and $\phi^{(k_1, \tau)-, (k_3, \tau)-} > \phi^{(k_1, \tau)+, (k_3, \tau)-}$ (upper left), $\phi^{(k_1, \tau)-, (k_3, \tau)-} < \phi^{(k_1, \tau)+, (k_3, \tau)-}$ (upper right). Evolution of the slice of volatility with buy orders on (k_3, τ) and $\phi^{(k_1, \tau)-, (k_3, \tau)+} < \phi^{(k_1, \tau)+, (k_3, \tau)+}$ (lower left), $\phi^{(k_1, \tau)-, (k_3, \tau)+} > \phi^{(k_1, \tau)+, (k_3, \tau)+}$ (lower right).	212
8.3	PnL with respect to time (in seconds) under the Poisson and Hawkes assumption.	224
8.4	Cross impact curves (in basis points) with respect to time (in seconds) in the Poisson and Hawkes cases.	225
9.1	Target value function for increasing inventory and random market states.	240

9.2	Comparison of the target value with approximation continuous in spread and imbalance.	240
9.3	Comparison of the target value with approximation discrete in spread and imbalance.	240
9.4	Neural network structure for the trader's strategy.	242
9.5	Neural network structure for the value function.	242
9.6	Value function with respect to the inventory between $t = 0$ and $t = 0.4$.	242
9.7	Evolution of the value function v between $t = 0.5$ and $t = 0.9$.	242
9.8	Evolution of the value function v between $t = 0$ and $t = 0.4$ using neural networks.	242
9.9	Evolution of the value function v between $t = 0.5$ and $t = 0.9$ using neural networks.	242
9.10	Limit strategy in the first venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$.	243
9.11	Limit strategy in the second venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$.	243
9.12	Volume sent to the first venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$.	244
9.13	Volume sent to the second venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$.	244
9.14	Limit strategy, $\psi^1 = \delta, \psi^2 = 2\delta$, $I^1 = I^2 = 0$.	244
9.15	Volume strategy, $\psi^1 = \delta, \psi^2 = 2\delta$, $I^1 = I^2 = 0$.	244
9.16	Limit order strategy, $\psi^1 = \psi^2 = \delta$, $I^1 = -0.5, I^2 = 0.5$.	245
9.17	Volume strategy, $\psi^1 = \psi^2 = \delta$, $I^1 = -0.5, I^2 = 0.5$.	245
9.18	Limit order strategy in the first venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$ using neural networks.	245
9.19	Limit order strategy in the second venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$ using neural networks.	245
9.20	Volume posted in the first venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$ using neural networks.	246
9.21	Volume posted in the second venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$ using neural networks.	246
9.22	Limit order strategy, $\psi^1 = \delta, \psi^2 = 2\delta$, $I^1 = I^2 = 0$ using neural networks.	246
9.23	Volume strategy, $\psi^1 = \delta, \psi^2 = 2\delta$, $I^1 = I^2 = 0$ using neural networks.	246
9.24	Limit order strategy, $\psi^1 = \psi^2 = \delta$, $I^1 = -0.5, I^2 = 0.5$ using neural networks.	246
9.25	Volume strategy, $\psi^1 = \psi^2 = \delta$, $I^1 = -0.5, I^2 = 0.5$ using neural networks.	246
9.26	Evolution of the value function v between $t = 0$ and $t = 0.4$.	247
9.27	Evolution of the value function v between $t = 0.5$ and $t = 0.9$.	247
9.28	Evolution of the value function v between $t = 0$ and $t = 0.4$ using neural networks.	247
9.29	Evolution of the value function v between $t = 0.5$ and $t = 0.9$ using neural networks.	247
9.30	Limit order strategy in the first venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$.	248
9.31	Limit order strategy in the second venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$.	248
9.32	Volume posted in the first venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$.	248
9.33	Volume posted in the second venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$.	248
9.34	Limit order strategy, $\psi^1 = \delta, \psi^2 = 2\delta$, $I^1 = I^2 = 0$.	248
9.35	Volume strategy, $\psi^1 = \delta, \psi^2 = 2\delta$, $I^1 = I^2 = 0$.	248
9.36	Limit order strategy, $\psi^1 = \psi^2 = \delta$, $I^1 = -0.5, I^2 = 0.5$.	249
9.37	Volume strategy, $\psi^1 = \psi^2 = \delta$, $I^1 = -0.5, I^2 = 0.5$.	249
9.38	Limit order strategy in the first venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$ using neural networks.	249

9.39	Limit order strategy in the second venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$ using neural networks.	249
9.40	Volume posted in the first venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$ using neural networks.	250
9.41	Volume posted in the second venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$ using neural networks.	250
9.42	Limit order strategy, $\psi^1 = \delta, \psi^2 = 2\delta,$ $I^1 = I^2 = 0$ using neural networks.	250
9.43	Volume strategy, $\psi^1 = \delta, \psi^2 = 2\delta,$ $I^1 = I^2 = 0$ using neural networks.	250
9.44	Limit order strategy, $\psi^1 = \psi^2 = \delta,$ $I^1 = -0.5, I^2 = 0.5$ using neural networks.	250
9.45	Volume strategy, $\psi^1 = \psi^2 = \delta,$ $I^1 = -0.5, I^2 = 0.5$ using neural networks.	250
9.46	Market simulation: spreads (upper left), imbalances (upper right), volumes (lower left) and limits (lower right) in both venues.	251
9.47	Evolution of the inventory of the trader on a slice of execution.	251
9.48	Bayesian update of the executed proportion in the first venue.	252
9.49	Bayesian update of the executed proportion in the second venue.	252
9.50	Bayesian update of the transition matrix $r^{1,\psi}$	252
9.51	Bayesian update of the transition matrix $r^{2,\psi}$	252
9.52	Bayesian update of the transition matrix $r^{1,I}$	252
9.53	Bayesian update of the transition matrix $r^{2,I}$	253
9.54	Bayesian update of the drift of the asset.	253
9.55	Bayesian update of the intensity of limit orders in the first venue.	254
9.56	Bayesian update of the intensity of limit orders in the second venue.	255
10.1	Portfolio holdings in the “always passive” model, and portfolio characteristics: gross market value (gmv), net, and β_t given by (10.29).	281
10.2	Portfolio holdings in the sophisticated model, and portfolio characteristics: gross market value (gmv _t), net _t , and β_t given by (10.29).	281
10.3	Portfolio value (10.28) over the lifetime of the execution, for both execution methods.	281
11.1	Evolution of the cash position with respect to time.	288
11.2	Trading strategy in cash with respect to time.	288
11.3	Evolution of the share’s position with respect to time.	289
11.4	Trading strategy in shares with respect to time.	289
11.5	Evolution of the stock price with respect to time.	289
11.6	Evolution of the share’s position with respect to time using (11.7).	289
11.7	Evolution of the cash position with respect to time.	289
11.8	Trading strategy in cash with respect to time.	289
11.9	Evolution of the share’s position with respect to time.	289
11.10	Trading strategy in shares with respect to time.	289
11.11	Evolution of the cash position with respect to time.	290
11.12	Trading strategy in cash with respect to time.	290
11.13	Evolution of the share’s position with respect to time.	290
11.14	Trading strategy in shares with respect to time.	290
12.1	Optimal portfolio liquidation costs with respect to correlation for different strategies.	301
12.2	Liquidation times for 2-bond case.	301
12.3	Optimal liquidation costs with portfolio optimizer.	302
12.4	Optimal liquidation times with portfolio optimizer.	302
12.5	Short deadline premium.	303

List of Tables

- 4.1 Number of options traded by maturity and month. 124
- 4.2 Number of options traded by strike and month. 126

- 12.1 Test portfolio characteristics. 302
- 12.2 Short deadline costs and times comparison. 303
- 12.3 Comparison between three type of liquidation. 304
- 12.4 Correlation matrix for the example set of bonds. 304

Part
Introduction

Part I Market-making regulation activity via Principal-Agent theory

I.1 Market-making and incentives design in the presence of a dark pool: a deep reinforcement learning approach

The first part of this thesis is dedicated to the application of Principal-Agent theory to the regulation of market-making activity. Before diving into the solutions we propose, we begin with a historical review of the Principal-Agent literature in order to put into context our modeling assumptions. The Principal-Agent problem, also known as agency dilemma, arises between two entities, the Agent and the Principal. The dilemma occurs because the former is motivated to act in his best interests, which may contradict those of the latter. The problem consists in designing a compensation system to produce a behavior of the Agent consistent with the Principal's preferences. Principal-Agent problems typically arise where the two parties have different interests and asymmetric information. The Principal, usually having less information, cannot ensure that the Agent is acting in the Principal's best interest, as soon as these actions induce a cost for the Agent. In this situation, the goal of the Principal is to design a contract, a remuneration given to the Agent, maximizing the Principal's utility while the Agent's utility is fixed to a certain level.

In the absence of moral hazard, that is when the actions of the Agent are perfectly known by the Principal, the problem is to share the project's risks with the Agent optimally. The risk-sharing problem has been widely studied, notably in Borch [51], Cadenillas, Cvitanić, and Zapatero [61], and Ou-Yang [227], where the authors study the optimal compensation of portfolio managers. More complicated situations arise in the presence of moral hazard, where the Principal cannot observe (or contract upon) the actions chosen by the Agent. Initially, these problems were studied in discrete-time or static settings with the so-called first-order approach, see Grossman and Hart [134], Holmström [157], Jewitt [176], Mirrlees [212, 214, 215], Rogerson [240], Spence and Zeckhauser [256], Zeckhauser [272] and more recently in Alvi [18], Conlon [78], Fagart and Fluet [114], Sinclair-Desgagné [250], Chade and Swinkels [73], Kirkegaard [183, 184]. In the discrete-time framework, the Principal observes the realized output of the Agent's project and pays him a fraction of this output. Based on the agreed payment schedule, the Agent chooses his effort level. The Principal's problem is to maximize the difference between the output's value and the payment offered to the Agent, given the so-called incentives compatibility and participation constraint. The Agent must have the correct incentives to go along with the Principal's idea of how much effort should be provided, and his utility must be above a pre-defined level under which he has no interest in receiving the Principal's incentives. The first-order approach replaces the incentives compatibility constraint with the condition that the Agent's expected utility should be stationary in the effort. Under the assumption that the Agent's effort is chosen from an open interval, all global maxima are stationary points. The solution to the original problem and the first-order approach problem have no reason to be the same in general. In particular, Rogerson [240] provided the first sufficient conditions, the monotone likelihood ratio condition and convexity of the distribution function condition, so that the solution provided by the first-order approach is the solution of the original moral hazard problem. More general conditions have been proposed since then, for example in Alvi [18], Conlon [78], Kirkegaard [183, 184].

The main issue of the first order analysis approach for Principal-Agent problems is that the conditions allowing its use are very restrictive. Moreover, there are examples where the solution of the first order analysis differ from the solution of the initial problem, see Mirrlees [215]. Alternatives approaches have been developed in discrete-time, see for example Araujo and Moreira [22], Kadan and Swinkels [179], Kirkegaard [183], Ke and Ryan [182], as well as numerical-based approaches, see Renner and Schmedders [236, 237]. However, the applications of discrete-time methods are very limited and not well-suited to the regulation problems we face, contrary to a continuous-time framework. One had to wait for the seminal paper by Holmström and Milgrom [159] to witness the treatment of specific moral hazard problems in such framework. They considered the problem of providing incentives over time for

an Agent with constant absolute risk aversion who controls the drift rate of a vector of accounts that is subject to frequent small random fluctuations. The optimal compensation scheme is a linear function of a vector of N accounts, which count the number of times that each of the N kinds of observable events occurs. This is the so-called “drift-control problem”, which has been generalized in Müller [220], Schättler and Sung [246, 247], Sung [259, 260], Sannikov [243, 244], DeMarzo and Sannikov [95, 96], Williams [268, 269, 270], Cvitanić, Wan, and Zhang [85, 86, 87], notably using stochastic control theory tools such as dynamic programming and martingales optimality principles. The key result of these papers is that, in the continuous-time setting where the Agent controls only the drift of the output process, it is possible to derive tractable characterizations of optimal contracts.

Due to the high information asymmetry between participants in financial markets, financial regulation is one of promising areas of applying incentives theory. As electronic markets are fragmented, meaning that the same asset can be traded in several platforms, the exchanges compete to increase trading activity on their liquidity pools, see Laruelle and Lehalle [192]. The usual regulation mechanism is the make-take fees system: a fee rebate is associated with executed limit orders while transaction costs are applied to market orders. This way, exchanges charge in an asymmetric way liquidity provision and liquidity consumption. The use of this mechanism facilitated the expansion of high-frequency traders, see Arnuk and Saluzzi [23], who have become responsible for the vast majority of liquidity provided in electronic markets. One of the main concerns is the quality of liquidity provided by these market participants: in times of stress, they tend to leave the market, see Bellia [40], Megarbane, Saliba, Lehalle, and Rosenbaum [210], Menkveld [211], Mirrlees [213]. The study of the make-take fees mechanism has been mostly empirical or with very stylized models in discrete time, see Angel, Harris, and Spatt [21], Brolley and Malinova [57], Colliard and Foucault [76], Harris [149], showing that several market parameters such as depths, volumes, or price impact do depend on the make-take fees structure. In the recent work El Euch, Mastrolia, Rosenbaum, and Touzi [109], the authors adapt the Principal-Agent framework with drift control to optimize a make-take fees schedule to increase market liquidity. In their formulation, the Principal is the exchange, the Agent is a market-maker, and the effort corresponds to the liquidity provided by the market-maker quality. The quality stands for the size of the bid-ask spread proposed by the market-maker. The contract depends on the realized flow of transactions. The exchange is looking for the best make-take fees policy to offer to a single market-maker to maximize its utility. In other words, it wants to establish an optimal contract with the market-maker to incentivize the liquidity increase. The optimal contract is obtained in quasi closed-form, depending on the market-maker’s inventory trajectory and market volatility. Numerical results show that such contracts lead to reduced spreads and lower trading costs for investors.

This make-take fees mechanism based on contract theory is designed for traditional liquidity venues, called “lit pool”, where the flow of transactions is visible by the market participants. Recent regulatory changes have induced a rise of different types of alternative trading mechanisms, notably “dark pools”, which have gained a significant market share. Nowadays, many major exchanges, such as Bats Chi-X and Turquoise, have their dark pools in addition to their major trading platforms. Furthermore, several traditional exchanges such as NYSE and Euronext offer trading platforms whose functioning is inspired mainly by dark pools, see for example Laruelle and Lehalle [192]. Though trading rules for dark pools are very diverse, they share at least two essential properties: dark pools have no visible order book for market participants and usually set prices different from those in the lit pool. The former leads to smaller or no market impact compared to traditional lit pools (due to reporting of transactions imposed by the regulation, there is usually a delayed market impact). The latter is aimed at giving better prices to clients (closer to mid-price). A remarkable phenomenon is that dark pools are prone to a latency effect: the price displayed in the lit pool’s order book can change between the time of a request in the dark pool and that of the corresponding transaction. Such price discrepancy is particularly frequent in the presence of high imbalance because the price is likely to move when liquidity is scarce on one side of the book.

The fact that trades on dark pools are executed, in the absence of latency, at a better price than in a lit pool and reported with delays makes it a suitable platform to liquidate a large position but does not incentivize market-making activity. Indeed, market-makers would be present on one side of the dark pool only to take advantage of the latency effect or liquidate a large position.

Problem 1: How to increase liquidity in dark pools?

To address this problem, we build a Principal-Agent framework in which a market-maker (the Agent) controls the volume posted on the bid and ask sides of lit and dark venue of the same exchange (the Principal). The latter proposes a contract to the former to increase the trading activity in the dark pool by providing a remuneration proportional to the volume executed. We emphasize that, from a technical point of view, the first part of this thesis is based on reinterpretations and extensions of the results of Sannikov [243], Cvitanić, Possamai, and Touzi [89].

In the lit market, we assume that there is an efficient price process S_t and that the market-maker always posts volumes on the bid and ask sides at prices $S_t + \frac{T}{2}$ and $S_t - \frac{T}{2}$, where $\frac{T}{2}$ represents the half-tick of the market. The market-maker also provides liquidity in the dark pool where the transaction price is the efficient price S_t (possibly with the latency effect). In addition to market impact and latency phenomena, we also take into account transaction costs for market orders in both venues, which can be smaller in the dark pool. Thus, in our setting, a single market-maker only needs to select the volumes to post on the bid and ask sides of both lit and dark pools.

The market-maker's inventory Q_t is the aggregated sum of the volumes filled on both sides of the lit and dark pools. The volumes posted by the market-maker are $\mathcal{L}_t := (\mathcal{L}_t^l, \mathcal{L}_t^d)$, where $\mathcal{L}_t^l = (\ell_t^{a,l}, \ell_t^{b,l})$ and $\mathcal{L}_t^d = (\ell_t^{a,d}, \ell_t^{b,d})$ with $\ell_t^{i,j}$ corresponding to the volume posted by the market-maker at time t on side $i \in \{a, b\}$ (ask or bid) of the pool $j \in \{l, d\}$ (lit or dark).

The intensities of the total number of trades of size k are

$$\lambda_t^{\mathcal{L}, i, j, k} := \lambda^{i, j}(\mathcal{L}_t) \mathbf{1}_{\{\phi(i)Q_t - > -\bar{q}, \ell_t^{i, j} = k\}}, \quad \phi(i) := \begin{cases} 1 & \text{if } i = a, \\ -1 & \text{if } i = b, \end{cases}$$

where $\bar{q} > 0$ and $\lambda^{i, j}(\mathcal{L}_t)$ is decreasing in ask (respectively bid) side imbalance of the market for ask lit and bid dark (respectively bid lit and ask dark) type of orders. In other words, a high imbalance on the ask side decreases the probability that an ask limit order is filled in the lit pool and conversely for the bid side. Moreover, when the imbalance on the ask (resp. bid) side of the lit pool is high, if a market-taker wants to buy, it is worth trying it in the dark pool because the high imbalance indicates that the ask price in the lit may not be competitive.

The mid-price of the asset at time $t \in [0, T]$ takes into account a market impact from trades:

$$S_t := \tilde{S}_t + \sum_{j \in \{l, d\}} \int_0^t \Gamma^j \ell_u^{a, j} dN_u^{a, j} - \Gamma^j \ell_u^{b, j} dN_u^{b, j},$$

where $\Gamma^l, \Gamma^d > 0$ are fixed constants representing the magnitude of market impact in the lit and dark pools, $N_t^{i, j} := \sum_{k \in \mathcal{V}^j} N_t^{i, j, k}$ where \mathcal{V}^j is the set of order size of the trader on pool j and \tilde{S}_t is an arithmetic Brownian motion with volatility σ . Trade in the lit pool occurs at the best bid and the best ask, whereas in the dark pool, orders may be executed at the mid-price, but due to the latency the mid-price can change by one half tick (or more) before the transaction is made. High imbalance can induce latency, so we model this with Bernoulli random variables $\nu_t^a \sim \text{Ber}(I^a(\mathcal{L}_t^l))$, $\nu_t^b \sim \text{Ber}(I^b(\mathcal{L}_t^l))$, which are associated to each incoming market order in the dark pool. If $\nu_t = 1$, there is no latency,

and conversely for $\nu_t = 0$. The market-maker generates the profit and loss defined as the sum of the cash earned from his executed orders and the value of his inventory:

$$PL_t^{\mathcal{L}} := \mathcal{W}_t^{\mathcal{L}} + Q_t S_t,$$

where $\mathcal{W}_t^{\mathcal{L}}$, at time $t \in [0, T]$, represents his cash process and $Q_t S_t$ is the Mark-to-Market value of his inventory. Note that market-making activity in the dark pool without latency does not generate PnL through spread collection. We consider a risk-averse market-maker with exponential utility function and risk aversion parameter $\gamma > 0$. We define his optimization problem as

$$\sup_{\mathcal{L}} \mathbb{E} \left[-\exp \left(-\gamma (PL_T^{\mathcal{L}} - PL_0^{\mathcal{L}}) \right) \right]. \quad (1)$$

By the means of viscosity solution theory, we obtain that there exists a unique viscosity solution with polynomial growth to the HJB equation associated with the problem of the market-maker (1).

However, the market-maker acts in the lit and dark pools regulated by an exchange wishing to attract transactions. The wealth of the exchange depends on the market order flow, which depend on the volumes posted by the market-maker. However, the exchange cannot control those volumes and may only provide incentives to influence the market-maker's behavior.

Regarding the contract between the both parties, the market-maker's function of interest becomes

$$V_0^{\text{MM}}(\xi) := \sup_{\mathcal{L}} \mathbb{E} \left[-\exp \left(-\gamma (PL_T^{\mathcal{L}} - PL_0^{\mathcal{L}} + \xi) \right) \right].$$

The exchange receives fixed fees for each market order occurring in the lit and dark pools respectively. His goal is essentially to maximize the total number of market orders sent during the period of interest. As the arrival intensities of market orders are controlled by the market-maker through \mathcal{L} , the contract ξ should aim at increasing these intensities. Thus, the exchange subsidizes the Agent at time T with the compensation ξ . Then we specify the set of admissible contracts potentially offered by the exchange. Firstly the problem of the market-maker should not be degenerate. Secondly, we ensure the well-posedness of the exchange's problem. Finally, to accept the contract the exchange needs to assure some reservation utility for the market-maker. The set of admissible contracts is the set of contract satisfying all three of the above conditions.

Result 1. *Any admissible contract is of the form*

$$Y_T^{Y_0, Z} := Y_0 + \int_0^T \left(\sum_{(i,j) \in \{a,b\} \times \{l,d\}} \sum_{k \in \mathcal{V}^j} Z_u^{i,j,k} dN_u^{i,j,k} \right) + Z_u^{\tilde{S}} d\tilde{S}_u + \left(\frac{1}{2} \gamma \sigma^2 (Z_u^{\tilde{S}} + Q_u)^2 - H(Z_u, Q_u) \right) du, \quad (2)$$

where $H(z, q)$ is the Hamiltonian of the agent's problem, defined in (1.9).

Despite the fact that we initially search a contract over a, possibly very large, subset of random variables, this result states that without loss of generality, we can consider a smaller set of contracts with a precise structure. The contract consists of the following elements:

- The constant Y_0 is calibrated by the exchange to ensure that the reservation utility of the market-maker.
- The term $Z^{\tilde{S}}$ is the compensation given to the market-maker with respect to the volatility risk induced by the efficient price \tilde{S} .
- Every time a trade of size k occurs on the ask or bid side of the lit or dark pool, the market-maker receives $Z^{i,j,k}$.
- The term $\frac{1}{2} \gamma \sigma^2 (Z^{\tilde{S}} + Q)^2 - H(Z, Q)$ is a continuous "coupon" given to the market-maker.

When the exchange proposes this form of contract, we are able to characterize the optimal response of the market-maker.

Result 2. *The market-maker's utility function satisfies $V_0^{MM}(Y_T^{Y_0, Z}) = -\exp(-\gamma Y_0)$, with associated optimal volumes $\mathcal{L}^*(Z_t, Q_t)$ solutions of Equation (1.9) evaluated along the incentives and the market-maker's inventory for all $t \in [0, T]$.*

Finally, the problem of the exchange is solved using stochastic control theory.

Result 3. *There exists a unique continuous viscosity solution to the HJB equation associated with the problem of the exchange. The optimal incentives derived from this optimization problem are Markovian functions of time and the market-maker's inventory.*

Contrary to El Euch, Mastrolia, Rosenbaum, and Touzi [109], the optimal controls of the market-maker and the exchange cannot be obtained analytically. In the numerical experiments, we apply a method inspired by the reinforcement learning (RL) theory, known as actor-critic method, where the exchange's controls and its value function are represented by neural networks. This method allows finding approximations of the optimal value function and optimal control in the neural network representation. Contrary to this hybrid method, a simple value iteration algorithm would not allow overcoming the problem of dimensionality related to the calculation of optimal controls. A policy iteration algorithm would still require estimating the value function needed to calculate the controls. Therefore the actor-critic algorithms seem a suitable way of tackling our problem. The essence of this method is the alternation of the learning phases of the controls and value function. In the late 1990s, several breakthroughs achieved by RL-algorithms, for example, the outperformance of human level in games like Backgammon and chess, allowed RL to become one of the very popular subjects of studies and applications. The applications in finance have not been an exception. From straightforward applications to optimal trading and execution strategies to resolving high-dimensional Hamilton-Jacobi-Bellman equations – RL-based methods proved to be useful in the field. For example, in Han, Jentzen, and E [148], the authors introduce a deep learning-based methodology that can handle general high-dimensional parabolic PDEs. This approach relies on the reformulation of PDEs via Backward Stochastic Differential Equations, where neural networks approximate the gradients of the unknown solution. Since then, many extensions have been proposed, see, for example, Bachouch, Huré, Langrené, and Pham [29], Huré, Pham, Bachouch, and Langrené [167], Huré, Pham, and Warin [168], Henry-Labordère [152], Germain, Pham, and Warin [129, 130], Han and E [147], Chan-Wai-Nam, Mikael, and Warin [74], Pham, Warin, and Germain [230], Li, Rao, and Shi [200].

Our numerical method consists of two stages. The first stage is to represent the volumes posted by the market-maker by a neural network. On the domain of all possible inventories of the market-maker and all possible incentives given by the exchange we optimize the continuous version of the Hamiltonian of the market-maker. The second stage is to solve the HJB equation of the exchange, taking into account the optimal response of the market-maker to given incentives.




Figure 1: Optimal quotes of the market-maker with (right) and without (left) incentives.

The numerical results show the difference in the behavior of the market-maker with and without an incentive policy. When the market-maker has an inventory near zero, incentives increase the volumes posted in the lit pool and decrease in the dark pool. When he has a very positive inventory, he posts a higher volume on the dark pool's ask side than in the case without incentives. In addition to this, he posts an equal volume (small but not negligible) on the ask and bid sides of the lit pool. Thus, the exchange prevents the market-maker from artificial manipulation of the imbalance on the ask side. As the imbalance is around $1/2$, the market-maker does not take advantage of the latency effect. These results are illustrated in Figure 1.

I.2 Optimal make-take fees in a multi-market-maker environment

In the design of the make-take fees policy, we assumed that there was a single market-maker. In most optimal market-making literature, there is a single liquidity provider on a financial asset. In this setting, the market-maker needs to manage his inventory risk and does not need to adjust his strategy depending on competitors. However, for many financial assets, market-making activity is provided by several market-makers (typically three to ten), see, for example, Mounjid, Rosenbaum, and Saliba [219]. Thus, for the vast majority of markets it is more realistic to consider a multi-Agent model. A single-Agent model may overestimate the order flow that a market-maker will produce during the trading period.

Moreover, several important features of financial markets are linked to competition between market-makers, such as spread formation and order book shape, see Dayri and Rosenbaum [91], Glosten and Milgrom [131], Madhavan, Richardson, and Roomans [203]. From the viewpoint of the exchange, considering several market-makers can be relevant since it has, in practice, access to information related to the identity of the Agents involved in each transaction. Thus, we want to develop a multi-agent framework where several market-makers compete on a platform regulated by an exchange.

Problem 2: How to increase liquidity when market-makers compete against each other?

The framework of competitive Agents with one Principal has been studied in one-period models, see Demski and Sappington [97], Green and Stokey [133], Harris, Kriebel, and Raviv [150], Holmström [158], Nalebuff and Stiglitz [221], and in continuous time in Koo, Shim, and Sung [186], where the authors consider the same framework as Holmström [158]. The closest work related to the issue of competing market-makers is Élie and Possamaï [110], where the authors study the impact of competitiveness among Agents in a firm on optimal incentives and contracts. In their work, the Agents control the drift of the output process and their utility is increasing in their wages and performance compared to other Agents in the firm. This model is inspired by Espinosa and Touzi [113], which studied a classical problem of portfolio optimization in financial markets, where the investors' utility also depends on how well they performed compared to other related investors. One Principal hires many Agents to manage several, possibly correlated projects, on his behalf. Furthermore, given the intrinsic notion of competition, the Agents can decide to work on their projects and try to either help or decrease the value of the projects managed by the other Agents. Besides, Agents have a Nash equilibrium type behavior, that is to say, that they all compute their best reaction functions given a strategy chosen by the other Agents, and then agree upon an equilibrium. This optimization occurs in a moral hazard setting where the Principal can only observe each project's outcomes. The authors prove that the optimal contract is a linear function of all the Agents' projects' terminal values. Therefore, to provide correct incentives in a framework where Agents compare their performances, the Principal has to reward them using their projects and those managed by other Agents. We use these results to build a general class of optimal contracts for competing market-makers.

We propose a model where several market-makers compete in a liquidity platform regulated by an exchange. The market is composed of a single asset S with arithmetic Brownian dynamics and volatility

$\sigma > 0$. Based on their view on the efficient price, market-makers offer bid and ask quotes on the underlying asset. Such prices are defined by

$$P_t^{i,b} := S_t - \delta_t^{i,b}, \quad P_t^{i,a} := S_t + \delta_t^{i,a}, \quad t \in [0, T], \quad i \in \{1, \dots, N\},$$

where control processes $\delta_t^{i,b}, \delta_t^{i,a}$ are predictable and uniformly bounded.

For the i -th Agent, a filled bid order, represented by the point process $N^{i,b}$ (whose intensity is defined in Equation (2.1)), increases his inventory by one unit, and conversely, for an ask order. The intensity of buy (resp. sell) market order arrivals is a decreasing function of:

- the spread $\underline{\delta}^b$ (resp. $\underline{\delta}^a$) imposed by the market-maker who is currently trading at the best bid (resp. best ask),
- the transaction cost $c > 0$ collected by the exchange,
- the market liquidity, namely the spread quoted by all market participants.

The inventory of the i -th market-maker is denoted by Q_t^i and its PnL process is

$$PL_t^{\delta,i} := \int_0^t \sum_{j \in \{a,b\}} \left(\delta_s^{i,j} \mathbf{1}_{\{\delta_s^{i,j} = \underline{\delta}^j\}} + \sum_{\ell=1}^K \omega_\ell \delta_s^{i,j} \mathbf{1}_{\{\delta_s^{i,j} \in K_\ell\}} \right) dN_s^j + \int_0^t Q_s^i dS_s, \quad t \in [0, T], \quad (3)$$

where the $\omega_\ell > 0$ represent the fact that market-maker is remunerated with an increasing fraction of his quote when he is close to the best spread. Each Agent maximizes the sum of his PnL and the contract offered by the exchange given actions of the other market-makers to obtain his so-called best-reaction function. We represent the utility of the market-maker by a CARA utility function $U_i(x) := -e^{-\gamma_i x}$, $x \in \mathbb{R}$, where $\gamma_i > 0$ stands for the risk-aversion of the i -th market-maker. Therefore the optimization problem of the i -th market-maker, given actions δ^{-i} of the other agents, is

$$V_{\text{MM}}^i(\xi^i, \delta^{-i}) := \sup_{\delta^i} \mathbb{E} \left[U_i \left(\xi^i + PL_T^{\delta,i} \right) \right], \quad (4)$$

where ξ^i is the contract offered by the exchange. As the N market-makers act simultaneously, we are looking for a Nash equilibrium resulting from the Agents' interactions. The equilibrium is described by a set of actions $\hat{\delta}(\xi) \in \mathcal{A}$ such that the optimum in (4) is reached for all $i \in \{1, \dots, N\}$.

The goal of the exchange is to maximize the total number of aggregated market orders $N_T^a + N_T^b$ arriving during the time interval $[0, T]$. We assume a CARA utility function with risk-aversion parameter $\eta > 0$, and define the set of admissible contracts \mathcal{C} so that the problems of the exchange and the market-makers are non-degenerated, while the utility of each market-maker is above a certain level $R_i < 0$. Moreover, there must exist at least one Nash equilibrium for a given set of contracts. The optimization problem of the exchange is then defined as

$$V_0^E := \sup_{\xi \in \mathcal{C}} \sup_{\delta \in \text{NA}(\xi)} \mathbb{E} \left[-\exp \left(-\eta \left(c(N_T^a + N_T^b) - \sum_{i=1}^N \xi^i \right) \right) \right],$$

with $\text{NA}(\xi)$ the set of Nash equilibria for a set of contract ξ . Under technical assumptions, we can show that any admissible contract can be written, in the spirit of (2), as a sum of incentives indexed on the number of bid and ask filled orders of each market-maker and the dynamics of the efficient price, see Equation (2.14). Moreover, if we assume that the compensations given to each market-maker with respect to the number of filled orders are the same, we obtain a closed form for the best response of the i -th market-maker to a given contract, as a linear function of the incentives Z . Thus, we can characterize in semi-explicit form (up to the value function) the optimal incentives proposed to the market-makers.

Result 4. *The optimal incentives of the Principal are Markovian functions of the time and the inventories of the market-makers. They can be obtained using the value function of the Principal, which boils down to the resolution of a system of $(2\bar{q} + 1)^N$ ODEs, where \bar{q} is the risk limit of the market-makers' inventory.*

The numerical results show that the incentive policy lead to a reduction of the average spread therefore increasing market liquidity. In terms of profits for the exchange, it is not optimal to add an infinite number of market-makers, as we can conclude from the PnLs illustrated in Figure 2. In particular, we




Figure 2: Evolution of the PnL of the platform with the number of market-makers.

provide in Equation (2.21) an analytical approximation of the optimal number of market-makers, which is of particular interest for an exchange wishing to choose acting market-makers and their incentives.

I.3 On bid and ask side-specific tick sizes

So far, we have developed incentives to increase market liquidity based on the market-makers' number of transactions. In this framework, each time a limit order of the market-maker is filled, he receives compensation whose value depends on his inventory process and the market parameters. The incentives force the market-maker to quote tighter spreads to benefit from the Principal's payments. However, the liquidity quality is also linked to the tick size, which is the smallest increment between two consecutive prices on a trading instrument. The exchange or the regulator usually fixes its value as a function of the asset price and the traded volume, see Huang, Lehalle, and Rosenbaum [163], Laruelle, Rosenbaum, and Savku [195]. If the tick size is too small, the price changes frequently and, conversely, for a large tick size. Its value is of particular importance for market-makers: stocks with smaller relative tick sizes attract a more significant proportion of market-makers since they can rapidly adjust their quotes to seize price priority, see Dayri and Rosenbaum [91], Megarbane, Saliba, Lehalle, and Rosenbaum [210]. The influence of the tick size on market liquidity leads us to the following problem.

Problem 3: *How to increase liquidity using the tick size?*

To answer this question, we build an Agent-based model where a market-maker sets bid and ask quotes on a single asset, on a liquidity platform regulated by an exchange. Contrary to the models described above, we work at the high-frequency scale. This specificity implies in particular that transaction prices are lying on a discrete tick grid. We stylize this fact by assuming that there exists a continuous-time efficient price process S with arithmetic Brownian dynamics and volatility σ . The market-maker sets his bid and ask "fair" prices as

$$S_t^a = S_{t-}^a + \alpha^a \mathbf{1}_{\{S_t - S_{t-}^a > (\frac{1}{2} + \eta^a)\alpha^a\}} - \alpha^a \mathbf{1}_{\{S_t - S_{t-}^a < -(\frac{1}{2} + \eta^a)\alpha^a\}},$$

$$S_t^b = S_{t-}^b + \alpha^b \mathbf{1}_{\{S_t - S_{t-}^b > (\frac{1}{2} + \eta^b)\alpha^b\}} - \alpha^b \mathbf{1}_{\{S_t - S_{t-}^b < -(\frac{1}{2} + \eta^b)\alpha^b\}}.$$

The market-maker increases (resp. decreases) his bid price if the efficient price is "sufficiently" higher (resp. lower) than his current fair bid price. The notion of "sufficiently" higher or lower is determined

by some parameters η^a, η^b , and the tick sizes on the bid and ask sides α^a, α^b . If η^a is small (resp. large), the market-maker changes more (resp. less) frequently his ask price, and similarly for the bid price with η^b . Thus the fair bid (resp. ask) is modified when the efficient price is close enough to a new tradable price on the tick grid with mesh α^b (resp. α^a).

The volume of transaction is assumed to be equal to one so that the cash process of the market-maker at terminal time T is

$$X_T = \int_0^T \left(S_t^a dN_t^a - S_t^b dN_t^b \right),$$

where N_t^i are point processes representing the number of transactions on the bid or ask side between 0 and t . Their intensity are decreasing functions of the tick sizes, and can be zero or strictly positive depending on the control processes $\ell_t^i \in \{0, 1\}$ of the market-maker: when he does not want to be present on the bid (resp. ask side) at price S^b (resp. S^a) he sets $\ell^b = 0$ (resp. $\ell^a = 0$) and conversely so that there are no incoming transactions. Similarly to the previous problems, we denote by Q_t the inventory of the market-maker at time t .

The optimization problem faced by the market-maker is

$$\sup_{\ell \in \mathcal{L}} \mathbb{E} \left[X_T + Q_T(S_T - AQ_T) - \int_0^T \left(\phi Q_s^2 + \phi_- Q_s^2 \mathbf{1}_{Q_s < 0} \right) ds \right],$$

where $\phi > 0$ represents the risk-aversion parameter of the market-maker, $\phi_- > 0$ is the additional risk aversion of the market-maker toward short position on $[t, T]$ and AQ_T^2 , with $A > 0$, is a penalty term for the terminal inventory position regardless of its sign. In this setting, the market-maker wishes to hold a terminal inventory close to zero because of the quadratic penalty AQ_T^2 . The term $\phi \int_t^T Q_s^2 ds$ penalizes long or short positions over the trading period. Using the theory of viscosity solutions we show that there exists a unique viscosity solution associated to the control problem of the market-maker.

Moreover, the optimal controls of the market-maker are such that he posts a limit order on the bid (resp. ask) side of the market if the gain obtained by a counterpart sell (resp. buy) market order produces a change in his value function higher than the value of the bid (resp. ask) price.

The market-maker acts on a platform whose goal is to maximize the number of market orders on $[0, T]$, by choosing the tick sizes (α^a, α^b) . We assume that the platform is risk-neutral and earns a fixed taker cost $c > 0$ for each market order. Therefore its optimization problem is defined as

$$\sup_{(\alpha^a, \alpha^b) \in \mathbb{R}_+^2} \mathbb{E} \left[c(N_T^a + N_T^b) \right],$$

given the optimal controls $(\ell^{*,a}, \ell^{*,b})$ of the market-maker.

The following result summarizes the numerical findings on the importance of tick size selection by the exchange.

Result 5.

- *If the platform imposes the same tick value on the bid and ask sides, it has to be sufficiently large to ensure significant PnL per trade for the market-maker and sufficiently small to attract market orders from market-takers.*
- *When allowing for side-specific tick values, the market-maker can take advantage of additional trading opportunities and increase his activity. The exchange benefits from this situation because of the higher number of trades on its platform.*

- When there is a penalty for short inventory positions of the market-maker, there is only one optimal couple of tick values where $\alpha^b > \alpha^a$.

Thus, a suitable choice of tick sizes (symmetric or not) leads to a better market liquidity quality. We show in Figure 3 that both the exchange and the market-maker benefit from an asymmetric choice of tick sizes.




Figure 3: Value function of the market-maker (on the left) and the exchange (on the right) as functions of α^a and α^b , for $\phi_- = 0.005$.

I.4 Derivatives market design

The incentives mechanism, based on the remuneration of the number of transactions or suitable tick sizes developed above, are limited to the cash markets. However, many exchanges also offer their clients the possibility to trade more complex financial instruments such as derivatives. There is very few articles addressing the design of a derivatives exchange, most of them focusing the relationship between stock and option markets, see among others Bollen [50], Conrad [79], Damodaran and Subrahmanyam [90], Detemple and Jorion [100] and Mayhew and Mihov [209]. In the following, we will focus on the vanilla options markets being the most liquid traded derivatives.

An exchange dealing in the options market faces two major problems. The first issue is the choice of the derivatives offered to clients. Clearly, the exchange cannot propose all maturities and strikes on its platform, it is hard to manage from a technical point of view, and it is even harder to guarantee the liquidity. The maturities are relatively standardized, so the main challenge is the selection of strikes. The exchange needs to offer a certain range of strikes to satisfy its clients, including some out-of-the-money options. So the second issue is the lack of liquidity on specific options since market-makers tend to quote large spreads for far-from-the-money options, but the exchange nevertheless needs to provide sufficient volume to attract clients. Moreover, in European exchanges, the status of “designated market-maker” is related to the obligation to quote a specific list of options below some fixed threshold. These two issues are addressed in the following problem.

Problem 4: How to design a derivatives market?

We answer this question in two steps. First, we propose a methodology of strike selections for the exchange based on quantization techniques. Second, we develop a make-take fees system to increase the liquidity of far-from-the-money options.

We consider that the exchange’s problem is to select n call options with a fixed maturity (or equivalently n strikes) to maximize the clients’ satisfaction. Let us assume that the client would like to buy a call

option with strike K and that the only available strikes are $K_1 < \dots < K_n$, $K \neq K_i, \forall i \in \{1, \dots, n\}$. Then he buys the call with strike K_i where

$$K_i = \arg \min_{1 \leq j \leq n} |K - K_j|.$$

We assume that the strike K is chosen according to a distribution \mathbb{P}^{mkt} which represents the historical law of market demand. The problem of the exchange is then to find the n -uplet $(K_i)_{1 \leq i \leq n}$ that solves

$$\arg \min_{K_1 \leq \dots \leq K_n} \mathbb{E}^{mkt} \left[\min_{1 \leq j \leq n} |K - K_j|^p \right]. \quad (5)$$

Among the n available strikes, the client wishes to minimize the distance with the strike K drawn from the market distribution. Thus, the exchange must choose the set of strikes that minimize this distance. We choose a power-law function which has the advantage to be symmetric and convex. As greater errors are increasingly penalized, the solution of this problem should capture the tail features of the market distribution. We face here a so-called quantization problem for which we prove the following result:

Result 6. *Under technical assumptions, for any $p \geq 2$, Problem (5) admits a unique solution of the form $K_1 < \dots < K_n$ which is a fixed point of Lloyd's algorithm, see Graf and Luschgy [132], Pagès, Pham, and Printems [228].*

Using Lloyd's algorithm, we obtain the set of strikes (K_1, \dots, K_n) matching the client's demand in the sense of Equation (5).

Given the obtained set of strikes, denoted by \mathcal{K} , and a set of maturities \mathcal{T} , we now design make-take fees via the Principal-Agent paradigm. The Agents, the market-makers providing liquidity on the listed options, receive incentives from the Principal, the exchange willing to lower the spreads on specific options. We assume that the price of the option with strike $k \in \mathcal{K}, \tau \in \mathcal{T}$ is denoted by $C^{k,\tau}$ and has the following dynamics:

$$dC_t^{k,\tau} = \Delta^{k,\tau} dS_t,$$

where $\Delta^{k,\tau}$ is the Bachelier delta of the option at the beginning of the trading period, and S is the underlying asset, following an arithmetic Brownian motion. The market-maker's best bid and best ask price at time t on the option $C^{k,\tau}$ are respectively

$$P_t^{k,\tau,b} := C_t^{k,\tau} - \delta_t^{k,\tau,b}, \quad P_t^{k,\tau,a} = C_t^{k,\tau} + \delta_t^{k,\tau,a}, \quad t \in [0, T],$$

where $\delta_t = (\delta_t^{k,\tau,i})_{k,\tau,i} \in \mathcal{A}$ are his control processes, and \mathcal{A} is the set of predictable control processes, uniformly bounded by $\delta_\infty > 0$.

For every listed option, the arrival of the bid (resp. ask) market orders of the unitary volume is modeled by a point process $N^{k,\tau,b}$ (resp. $N^{k,\tau,a}$) with intensity that are decreasing functions of the spread $\delta^{k,\tau,b}$ (resp. $\delta^{k,\tau,a}$) quoted by the market-maker, and the fee $f^{k,\tau}$ fixed by the exchange. The Mark-to-Market value of the market-maker's portfolio is defined by

$$Q_t := \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \Delta^{k,\tau} Q_t^{k,\tau},$$

where $Q_t^{k,\tau}$ is the number of options $C^{k,\tau}$ held by the market-maker at time t .

The market-maker's PnL process is

$$PL_t^\delta := W_t^\delta + Q_t S_t,$$

where \mathcal{W}_t^δ is the sum of cash earned from executed orders. In addition to his PnL, at terminal time T , the market-maker receives a remuneration ξ from the exchange for the number of transactions. Assuming a CARA utility function with risk-aversion parameter $\gamma > 0$, the optimization problem of the market-maker is

$$V_{\text{MM}}(\xi) := \sup_{\delta \in \mathcal{A}} \mathbb{E} \left[-\exp \left(-\gamma(\xi + PL_T^\delta) \right) \right].$$

In the spirit of the previous sections, we can prove under some technical conditions that any admissible contract (for which the problems of the Principal and the Agent are not degenerate) can be written as a sum of stochastic integrals indexed on the price process S and the number of executed bid and ask limit orders $N^{k,\tau,b}, N^{k,\tau,a}$ for every (k, τ) . Moreover, we obtain the best response of the market-maker in closed form.

Result 7. *Any admissible contract has a unique representation $\xi = Y_T^{Y_0, Z}$ for which*

$$V_{\text{MM}}(Y_T^{Y_0, Z}) = -\exp(-\gamma Y_0),$$

and the associated optimal bid-ask policy is a linear function of the incentives Z .

We now turn to the problem of the exchange. First, it wants to carry on a high number of trades to collect associated fees. Second, it wants to have smaller spreads, in particular, for far-from-the-money options.

In order to quantify the first objective, we introduce a weighted version of the total number of trades:

$$\mathcal{N}_t = \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} c^{k,\tau} N_t^{k,\tau,i},$$

where for any $(k, \tau) \in \mathcal{K} \times \mathcal{T}$, $c^{k,\tau} \geq 0$ represents the value attributed to a trade in the option $C^{k,\tau}$ by the exchange. To quantify the second objective, we consider the following quantity

$$\mathcal{L}_T^Z := \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \int_0^T \omega(\delta_\infty^{k,\tau,i}(Z_t) - \delta_\infty^{k,\tau}) dN_t^{k,\tau,i},$$

where $\omega \in [0, 1)$ and $\delta_\infty^{k,\tau}$ is a spread threshold the exchange would like to impose to the market-maker's quotes on option $C^{k,\tau}$. Assuming a CARA utility function with risk-aversion parameter $\eta > 0$, the optimization problem of the exchange given the optimal response of the market-maker is

$$V_0^E := \sup_{Z \in \mathcal{Z}} \mathbb{E} \left[-\exp \left(-\eta(\mathcal{N}_T - \mathcal{L}_T^Z - Y_T^{\hat{Y}_0, Z}) \right) \right], \quad (6)$$

where $\hat{Y}_0 = -\frac{\log(-R)}{\gamma}$ and \mathcal{Z} is defined in Section 4.A.3.

Result 8. *The maximization problem (6) admits a solution Z^* which is a Markovian function of time and the aggregated inventory \mathcal{Q} . This solution is obtained by solving a 2-dimensional PDE.*

Even though the exchange's problem initially has as many state variables as there are listed options, the design of incentives is reduced to the resolution of a two-dimensional PDE, so the problem of designing a make-take fees schedule for a large number of options is tractable.

I.5 Governmental incentives for green bonds investment

We developed a series of incentive mechanisms based on Principal-Agent theory to increase the quality of market liquidity in financial markets. An exchange provides these incentives to one (or several) market-makers acting in its platform. However, there are financial instruments where other Agents can provide incentives, for example, government bonds, sometimes called "govies", and, in particular,

green bonds. The latter are issued by governments to finance renewable projects and represent 9% of the total of issued green bonds total, see the report of the Financial Stability Board [49] or the reports of OECD [225, 226]. The influence of green bonds on gas emissions and environmental ratings is a well-documented topic, see Flammer [118, 117], Baker, Bergstresser, Serafeim, and Wurgler [34], Tang and Zhang [262] and de Angelis, Tankov, and Zerbib [92]. This type of financing environmental projects is a hot topic among institutional investors, particularly pension funds and asset managers, considering the possibility of including sustainable environmental investments in their assets, see for example the report of the GSIA [6]. The motivations of green investors are diverse: search of higher alpha or lower risk (see Nilsson [223], Bauer and Smeets [37], Krüger [191]) or the will for a more socially responsible image (see Hong and Kacperczyk [160]).

There are still several barriers to developing the green bond market, such as a lack of green bond definition, framework, and transparency. In Zerbib [273, 274], the author investigates the existence of a yield premium for green bonds: the results show a small negative premium, meaning that the yield of a green bond is lower than that of a comparable conventional bond. This premium is generally attributed to intangible asset creation, imperfectly captured in the models of rating agencies, see for example Porter and Van der Linde [231], Ambec and Lanoie [19], or Brooks and Oikonomou [58]. In order to develop the green bonds market, policymakers put in place several type incentives, see Morel and Bordier [218] and Della Croce, Kaminker, and Stewart [94], generally in the form of tax incentives. All the incentives currently proposed can be modelled as a function of the amount invested in green bonds. However, policy-makers cannot necessarily control or monitor directly the actions of the investor, which make the investor's actions subject to moral hazard. This leads for example to the so-called "green-washing" practice, see Della Croce, Kaminker, and Stewart [94]. We, therefore, want to provide an operational and quantitative answer to the following problem.

Problem 5: How to increase investments in green bonds?

To answer this question, we propose a Principal-Agent framework in which an investor (the Agent) manages a portfolio of green and conventional bonds. The government entity (the Principal) proposes incentives to the investor in order to increase the proportion of green bonds in his portfolio. We consider that he manages, over a time horizon $[0, T]$, $d^g \geq 1$ green bonds, $d^c \geq 1$ conventional bonds and an index of conventional bonds of geometric Brownian dynamics with deterministic drift and volatility, respectively defined by $P^g(t, T^g), P^c(t, T^c), I$. We also define

$$W_t = \begin{pmatrix} W_t^g \\ W_t^c \\ W_t^I \end{pmatrix}$$

as a $d^g + d^c + 1$ Brownian motion with correlation matrix Σ . The investment policy is defined by a vector of strictly positive and uniformly bounded control processes $\pi_t = (\pi_t^g, \pi_t^c, \pi_t^I)_{t \in [0, T]} \in \mathcal{A}$, representing the amount of money invested at time t . The corresponding dynamics of the portfolio of the investor is

$$dX_t = \pi_t^g \cdot \frac{dP^g(t, T^g)}{P^g(t, T^g)} + \pi_t^c \cdot \frac{dP^c(t, T^c)}{P^c(t, T^c)} + \pi_t^I \frac{dI_t}{I_t}.$$

Throughout the investment period $[0, T]$, the investor wants to maintain his investment in bonds at some pre-defined levels, which can be seen as his investment profile. We introduce the cost function:

$$k(\pi) = \frac{1}{2} \beta^g \cdot (\pi^g - \alpha^g)^2 + \frac{1}{2} \beta^c \cdot (\pi^c - \alpha^c)^2 + \frac{1}{2} \beta^I (\pi^I - \alpha^I)^2, \quad \pi \in \mathbb{R}^{d^g + d^c + 1},$$

where vectors $\alpha = (\alpha^g, \alpha^c, \alpha^I) \in \mathbb{R}^{d^g + d^c + 1}$ represent the targets of the investor and $(\beta^g, \beta^c, \beta^I)$ represent the cost intensity of changing the investments of the Agent.

In order to change an investment policy $\pi \in \mathcal{A}$, the government proposes a remuneration to the investor, which takes the general form of a random variable ξ . For a given contract provided by the government, the optimization problem of the investor with CARA utility function is

$$V_0^A(\xi) = \sup_{\pi \in \mathcal{A}} \mathbb{E} \left[- \exp \left(- \gamma \left(\xi - \int_0^T k(\pi_s) ds \right) \right) \right],$$

where $\gamma > 0$ is his risk aversion parameter. On the other hand, the government wishes to maximize the portfolio value of the issued bonds while increasing the amount invested in green bonds. Thus, in average, it wants to maximize

$$X_T - \sum_{i \in \{1, \dots, d^g\}} \int_0^T \kappa (G^i - \pi_t^{*,i,g})^2 dt,$$

where G^i is the investment target in the i -th green bond of the government entity, $\kappa > 0$ is the cost of moving away from the targets G^1, \dots, G^{d^g} and π^* is the best response of the investor to a given contract ξ . The government also subtracts from this quantity the contract ξ offered to the investor. Thus, his optimisation problem with CARA utility function is function is

$$V_0^P = \sup_{\xi \in \mathcal{P}} \mathbb{E} \left[- \exp \left(- \nu \left(X_T - \sum_{i \in \{1, \dots, d^g\}} \int_0^T \kappa (G^i - \pi_t^{*,i,g})^2 dt - \xi \right) \right) \right], \quad (7)$$

where $\nu > 0$ is the risk aversion parameter of the Principal and \mathcal{P} is the set of admissible contracts for the government.

As it would be unrealistic (and hardly tractable) to offer a compensation based on the whole universe of governmental bonds, we suggest a remuneration based on the green bonds, the value of the portfolio and an index of conventional bonds. This way, the contract is only indexed on the $d^g + 2$ variables $B^{\text{obs}} = (X, W^g, W^I)^\top$ and is of the form

$$\xi = \xi_0 + \int_0^T Z_t^{\text{obs}} \cdot dB_t^{\text{obs}} + \frac{1}{2} \text{Tr}((\Gamma_t^{\text{obs}} + \gamma Z_t^{\text{obs}} (Z_t^{\text{obs}})^\top) d\langle B \rangle_t) - h(t, Z_t^{\text{obs}}, \Gamma_t^{\text{obs}}, \pi^*(t, Z_t^{\text{obs}}, \Gamma_t^{\text{obs}})) dt, \quad (8)$$

where $Z^{\text{obs}}, \Gamma^{\text{obs}}$ are the incentives provided by the government, and h is the Hamiltonian of the investor for a best response π^* to a given contract.

Contrary to the Principal-Agent developed in previous chapters, the admissible contracts (8) are not only linear functions of the state variables but depend also linearly on their quadratic variation and covariation. This comes from the fact that by investing in the bonds, the investor controls directly the volatility of the portfolio process X . As we used only deterministic functions to model the risk premium, short term rate and volatility processes, the optimal incentives of the government can be obtained by maximizing a deterministic function, which leads to the following result.

Result 9. *The optimal incentives $Z^{*\text{obs}}, \Gamma^{*\text{obs}}$ are deterministic functions of time, solution of Equation (5.8).*

Numerical results on a portfolio of French governmental bonds show, among others, that the processes (π^*, Z^*, Γ^*) show a rather constant behavior through the period $[0, T]$. Thus, the optimal contract does not need a frequent re-calibration throughout the year and we show how to replicate the contract with several financial instruments: bond prices, variance, and covariance swaps on the bonds, log-contracts and vanilla options. The last result is a summary of the numerical findings.

Result 10.

- *The optimal controls show a rather constant behavior throughout the year: the government does not have to frequently recalibrate the optimal contract.*

- The government can increase the amount invested in the green bonds by the mean of G, κ , at the expense of a utility loss.
- The most important incentive with respect to the contractible variables is Z_X^* : the government always encourage a higher value of the portfolio of bonds by setting $Z_X^* > 0$. If the cost intensity κ decreases, the government reaches the investment's target by increasing the incentive Z_X^* and encouraging a higher volatility of the portfolio process.
- The incentives Z_G^*, Z_I^* have a small, but non-negligible, impact on the investor's choices compared to Z_X^* .
- At the expense of some substantial utility loss, the government can propose a contract indexed only on the contractible variables. This results in a higher incentive Z_X^* .
- The government has no interest in encouraging or penalizing a high volatility on the contractible variables apart from the portfolio process X .

We present in Figure 4 some trajectories of the portfolio process with optimal incentives. We obtain a higher value of the portfolio and a more investment in green bonds.




Figure 4: Some trajectories of the optimal portfolio process with and without contract.

This result concludes the first part, dedicated to improving the quality of liquidity on several financial markets. All the problems are studied from the viewpoint of the exchange or the governmental entity.

Part II Options market-making in high dimension

II.1 Option market-making algorithms in high dimension

In this part of the thesis we leave aside the viewpoint of the exchange to focus on the problem of the trader. Two central topics for traders are optimal market-making and optimal trading. While the latter will be treated in the third and last part of this thesis, the former is the main focus of the second part of this thesis. For more than three decades, the optimal market-making problem on cash markets has been the object of many academic studies, see, for example, Grossman and Miller [135], Ho and Stoll [156]. In Grossman and Miller [135], the authors studied a three-period model representing the interaction between market-makers and market-takers and analyzed its equilibrium state. In Ho and Stoll [156], the authors examined the behavior of a market-maker facing a stochastic demand and an inventory risk. The subject was revived in 2008 in the seminal paper Avellaneda and Stoikov [26], where the authors address the quoting and inventory management problems of market-makers using the stochastic optimal control tools. In this paper, the traded asset follows an arithmetic Brownian

motion, and the spread processes are continuous, which makes the model more suited to quote-driven markets. For order-driven markets, the more suitable model is the one developed in Guilbaud and Pham [143]: the market-maker can send limit and market-orders at the first and second best limit of an order book, and the optimal controls (the volumes and limits at which he posts) are obtained through the resolution of a HJBQVI equation.

There are numerous extensions of the work of Avellaneda and Stoikov tackling the problem of single-asset market-making. For example, in Guéant, Lehalle, and Fernandez-Tapia [142], the authors show that, in the case of exponential intensity functions, the market-making problem boils down to the resolution of a system of linear ordinary differential equations. A significant proportion of the contribution to the market-making literature comes from Cartea and Jaimungal. By considering a risk-adjusted expectation instead of Avellaneda-Stoikov's CARA objective function, they manage to enrich the initial model by introducing alpha signals, ambiguity aversion, competition with other agents, see, for example, Cartea, Donnelly, and Jaimungal [71, 72], Cartea and Jaimungal [68], Cartea, Jaimungal, and Penalva [70].

Regardless of how rich is the academic literature considering linear markets, the one studying optimal market-making on options is far less extensive. Market-making models on options, or other derivatives, are intrinsically more complicated because they must consider both the market for an underlying asset and the derivatives market. Consequently, one needs to first impose a factorial stochastic volatility model, possibly with jumps, on the underlying asset. Second, option market-makers need to manage several thousands of positions, which leads to very high-dimensional problems that cannot be solved using classical numerical schemes. Even if machine learning techniques are used, involving, for example, deep reinforcement learning methods in the spirit of Guéant and Manziuk [140], Huré, Pham, and Warin [168], E, Han, and Jentzen [103], the computation time can still be an obstacle for a large portfolio of options.

At the time of writing this thesis, the only papers addressing option market-making are El Aoud and Abergel [106], Stoikov and Sağlam [258]. In the former, the authors consider a single-option market driven by a stochastic volatility model and assume that the position is always Δ -hedged. They provide optimal bid and ask quotes for the option and focus on the risk of model misspecification. In the latter, the authors consider three different settings for a market-maker managing a single option and its underlying. The first setting is a complete market with continuous trading in the perfectly liquid underlying. The second is a complete market with an illiquid underlying, where the market-maker sets bid and ask quotes in the option and the stock. The third is an incomplete market with residual risks due to stochastic volatility and overnight jumps in the stock price. These two articles treat only the case of a single option on a single underlying. Thus, we want to solve the following problem.

Problem 6: How can we use stochastic control to design option market-making algorithms in high dimension?

We propose two models to build option market-making strategies. In the first one, we consider the case of a market-maker in charge of a book of options (on a single asset) whose prices are driven by a stochastic volatility model. We assume that trading in continuous time can be carried out in the underlying asset so that the residual risk is only that of the Vega associated with the inventory.

The dynamics of the underlying asset under a risk-neutral measure \mathbb{Q} is given by the following one-factor stochastic volatility model:

$$\begin{cases} dS_t = \sqrt{\nu_t} S_t d\widehat{W}_t^S \\ d\nu_t = a_{\mathbb{Q}}(t, \nu_t) dt + \xi \sqrt{\nu_t} d\widehat{W}_t^\nu, \end{cases}$$

where $(\widehat{W}_t^S, \widehat{W}_t^\nu)$ is a couple of correlated Brownian motions and $a_{\mathbb{Q}}$ is such that the processes are well-defined.

We consider an option market-maker in charge of providing bid and ask quotes for N options on S over the period $[0, T]$ where $T < \min_{i=1, \dots, N} T^i$. For all $i = 1, \dots, N$ we denote by $\mathcal{O}_t^i - \delta_t^{i,b}$ and $\mathcal{O}_t^i + \delta_t^{i,a}$ the bid and ask prices proposed by the market-maker for the i -th option, where which are predictable and uniformly bounded from below.

The dynamics of the inventory process $(q_t)_{t \in [0, T]} := (q_t^1, \dots, q_t^N)'$ of the market-maker is given by

$$dq_t^i := z^i (dN_t^{i,b} - dN_t^{i,a}), \quad i \in \{1, \dots, N\},$$

where z^i is the size of transactions for option i , and $(N_t^{i,b}, N_t^{i,a})$ are the transaction processes of the i -th option on the bid and ask side with intensity

$$\lambda_t^{i,b} := \Lambda^{i,b}(\delta_t^{i,b}) \mathbf{1}_{\{q_t - z^i e^i \in \mathcal{Q}\}}, \quad \lambda_t^{i,a} := \Lambda^{i,a}(\delta_t^{i,a}) \mathbf{1}_{\{q_t - z^i e^i \in \mathcal{Q}\}},$$

where $(e^i)_{i \in \{1, \dots, N\}}$ is the canonical basis of \mathbb{R}^N , \mathcal{Q} the compact set of authorized inventories, and $\Lambda^{i,b}, \Lambda^{i,a}$ are positive functions satisfying general growth and continuity conditions (this includes in particular exponential, logistic and SU Johnson intensity functions). The Delta of the portfolio is defined by

$$\Delta_t := \sum_{i \in \{1, \dots, N\}} \partial_S O^i(t, S_t, \nu_t) q_t^i \text{ for all } t \in [0, T].$$

The resulting dynamics for the cash process $(X_t)_{t \in [0, T]}$ of the market-maker is:

$$dX_t := \sum_{i \in \{1, \dots, N\}} \left(z^i \left(\delta_t^{i,b} dN_t^{i,b} + \delta_t^{i,a} dN_t^{i,a} \right) - \mathcal{O}_t^i dq_t^i \right) + S_t d\Delta_t + d\langle \Delta, S \rangle_t,$$

and the Mark-to-Market value of the market-maker's portfolio at time $t \in [0, T]$ is

$$V_t := X_t - \Delta_t S_t + \sum_{i \in \{1, \dots, N\}} q_t^i \mathcal{O}_t^i.$$

As in Cartea, Jaimungal, and Ricci [69], Cartea, Donnelly, and Jaimungal [71], Cartea, Jaimungal, and Penalva [70], we consider a risk-adjusted expectation as the objective function for the problem of the market-maker:

$$u(0, S, \nu, q) = \sup_{\delta \in \mathcal{A}} \mathbb{E}[V_T] - \frac{1}{2} \mathbb{V}[V_T],$$

where $S_0 = S, \nu_0 = \nu, q_0 = q$.

When it comes to approximating the optimal quotes of the market-maker for the N options, classic numerical methods are ineffective since the value function u has $N + 2$ variables (in addition to the time variable). In order to make the problem tractable, we approximate the Vega of each option over $[0, T]$ by its value at time $t = 0$, namely

$$\mathcal{V}_t^i = \mathcal{V}_0^i =: \mathcal{V}^i \in \mathbb{R}, \text{ for all } i \in \{1, \dots, N\}.$$

This is acceptable if T is not too large and if we deal with long-dates options, because the Vega of the portfolio will not change drastically. Moreover, we assume a global risk limit on the sum of inventories on each option weighted by their Vega.

Under the above assumptions, the $N + 2$ state variables can be replaced by a single one: the Vega of the portfolio $\mathcal{V}^\pi := \sum_{i \in \{1, \dots, N\}} q^i \mathcal{V}^i$ whose dynamics is

$$d\mathcal{V}_t^\pi = \sum_{i \in \{1, \dots, N\}} z^i \mathcal{V}^i (dN_t^{i,b} - dN_t^{i,a}).$$

Then, the value function u verifies the ansatz

$$u(t, S, \nu, q) = v\left(t, \nu, \sum_{i \in \{1, \dots, N\}} q^i \mathcal{V}^i\right),$$

where v solves the two-dimensional HJB equation (6.5). The optimal quotes can be derived easily using classic numerical schemes. In the next result, we show how the assumptions of perfect Δ -hedging and constant Vega can be relaxed.

Result 11.

1. *If the market-maker can be Δ -hedged in continuous-time, the optimization problem of the market-maker is the same except that the volatility of volatility parameter ξ is multiplied by $\sqrt{1 - \rho^2}$ to account for the reduction of risk made possible by the optimal trading strategy in the underlying asset in presence of vol-spot correlation.*
2. *We consider a perturbative approach, that is*

$$\sum_{i \in \{1, \dots, N\}} q_i^i \partial_{\sqrt{\nu}} O^i(t, S_t, \nu_t) = \mathcal{V}_t^\pi + \epsilon \mathcal{W}(t, S_t, \nu, t, q_t),$$

for a general function \mathcal{W} and

$$u(t, S, \nu, q) = v(t, \nu, \mathcal{V}^\pi) + \epsilon \varphi(t, S, \nu, q).$$

Then, φ solves a linear PDE thus it admits a Feynmac-Kac representation and can be computed using Monte-Carlo methods. Therefore we can compute the first-order approximation (in ϵ) of quotes accounting for the variation of the Vega.

We show in Figure 5 the quotes of the market-makers for a set of call options with same strike and different maturities as a function of the Vega of the portfolio




Figure 5: Optimal mid-to-bid quotes divided by option price as a function of the portfolio Vega for $K=10$.

The model can also be extended to more complex underlying dynamics (for example, the two-factor model of Bergomi [44]) to perform market-making on variance derivatives. However, the constant Vega assumption, which makes the control problem time-inconsistent, is only valid for a market-maker in charge of long-dated options where possible jumps in the underlying do not drastically influence

the global risk position. Thus this model cannot be used for market-making strategies on short-dated options where the vast majority of the liquidity is. If one adds other Greeks such as Vanna and Vomma, the model becomes hardly tractable as the HJB equation would be in dimension 5. Finally, option market-making on several underlying would be impossible for the same reasons.

To tackle the above challenges, we build an option market-making model where the ansatz is not on the state variables but directly on the value function. We present here the single-asset case, but the model treats the multi-asset case as well. Consider the new dynamics for the stock process

$$\begin{cases} dS_t = b_{\mathbb{P}}(t, S_t)dt + \sigma(t, S_t, \nu_t)dW_t^S + \int_{\mathbb{R}} Z(dt, dz), \\ d\nu_t = a_{\mathbb{P}}(t, \nu_t)dt + v_{\mathbb{P}}(t, \nu_t)dW_t^{\nu}, \end{cases}$$

where $Z(dt, dz)$ is a marked point process independent of the Brownian motions, with intensity kernel $\kappa_t(dz)$. There also exists a risk-neutral probability measure \mathbb{Q} such that

$$\begin{cases} dS_t = \sigma(t, S_t, \nu_t)d\hat{W}_t^S + \int_{\mathbb{R}} Z(dt, dz), \\ d\nu_t = a_{\mathbb{Q}}(t, \nu_t)dt + v_{\mathbb{Q}}(t, \nu_t)d\hat{W}_t^{\nu}, \end{cases}$$

These new dynamics take into account the Vega, the Vomma and the Vanna of the options \mathcal{O}_t^i for all $i \in \{1, \dots, N\}$. The HJB equation associated to the new control problem is still intractable for a large number of options. However by using a Taylor expansion on the value function as well as an ansatz quadratic with respect to the inventories, we obtain the following result.

Result 12. *We assume symmetry of the buy and sell intensities. By using the ansatz*

$$u(t, S, \nu, q) = \theta^0(t, S, \nu) + q^\top \theta^1(t, S, \nu) - q^\top \theta^2(t, S, \nu)q,$$

where $\theta^0 \in \mathbb{R}$, $\theta^1 \in \mathbb{R}^N$, $\theta^2 \in \mathcal{M}_N(\mathbb{R})$, the value function u solves a $N^2 + N + 1$ system of two-dimensional non-linear PDEs (7.4)

We show in Figure 6 the evolution of ask quotes for call options of same maturity as a function of the spot and the stochastic volatility.




Figure 6: Optimal ask quotes with respect to S and ν for options with maturity $T = 0.6$.

The great advantage of this method is that we only need to solve a system of non-linear PDE with respect to (S, ν) in order to obtain the value function and the optimal quotes of the market-maker. Contrary to the previous model, the Greeks are time-varying and depend on both the stock and the volatility processes. Moreover, we can increase the complexity of the underlying process (by using a multi-factor model for example) without increasing drastically the complexity of the system of non-linear PDEs. All these features make this model more flexible than the previous one, in particular for market-making on short-dated options. The numerical results show that, for a set of short- to mid-dated options, this algorithm produces a higher average PnL compared to the one using the constant Vega assumption.

The two proposed models enable the design of market-making strategies for a large portfolio of options. The first method is more suited to long-dated options, while the second offers greater flexibility for short-dated options. In order to test the efficiency of these methods, one would require a model for the microscopic behavior of the implied volatility surface.

II.2 High-frequency dynamics of the volatility surface

Modeling the behavior of the implied volatility surface has often been carried out at a coarse time scale, see among others Fengler, Härdle, and Mammen [116], Kamal and Gatheral [181], Skiadopoulos, Hodges, and Clewlow [252] and Cont and Da Fonseca [80]. For example, in the latter, the authors use a time series of option prices on the S&P500 and FTSE indices to study the volatility surface's deformation. The surface dynamics can be described by fluctuations of a small number of orthogonal random factors:

- The “level” factor corresponding to the global level of the whole surface of implied volatilities.
- The “calendar” factor corresponding to the skew of the volatility surface.
- The “butterfly” factor corresponding to the convexity of the volatility surface.

However, such a model does not exist at thin time scales, where the implied volatility moves as a discrete process. Models for thin and coarse time scales have only been developed for price processes, which have inspired our study. In Bacry and Muzy [30], the authors develop a Hawkes-based model for the high-frequency dynamics of an asset's price. Assuming that all jumps are of the same size, the microscopic price of the asset is the difference of the number of upward and downward jumps, that is

$$P_t = P_0 + N_t^+ - N_t^-,$$

where (N^+, N^-) is a bi-dimensional Hawkes process with intensity kernel

$$\phi(t) = \begin{pmatrix} \phi^{++}(t) & \phi^{+-}(t) \\ \phi^{-+}(t) & \phi^{--}(t) \end{pmatrix},$$

where ϕ is an endogenous source of price moves: for example, ϕ^{+-} raises the intensity of upward price jumps after a downward price jump, creating a mean-reversion effect, while ϕ^{++} creates a trending effect. This model accurately reproduces some stylized facts of the high-frequency price behavior, which can be encoded easily in the Hawkes kernel: no statistical arbitrage property, bid-ask asymmetry, long memory property of order flow, and a high degree of endogeneity of financial markets. At the coarse time scale, the price process defined above converges to a stochastic rough volatility process. The rough volatility models are particularly appreciated for their ability to capture key features of the implied volatility surface and its dynamics, see Bayer, Friz, and Gatheral [38], El Euch, Gatheral, and Rosenbaum [108], Horvath, Jacquier, and Tankov [161], Jacquier, Martini, and Muguruza [171], Alòs, León, and Vives [15]. Their properties have raised interest in building microscopic models for market dynamics which reproduce rough volatility at a macroscopic scale, see among others Jaisson and Rosenbaum [174, 175], Jusselin and Rosenbaum [177], Tomas and Rosenbaum [263]. These models focus on stock prices, mainly because it is easy to stylize stock markets' characteristics with Hawkes kernels. However, one can apply similar ideas to model the microscopic behavior of the volatility surface.

Problem 7: How to model the intraday dynamics of the implied volatility surface?

To answer this question, we propose a tick-by-tick model for the high-frequency dynamics of the volatility surface. Assuming a set \mathcal{K} of strikes and \mathcal{T} of maturities, the microscopic volatility surface is modeled as a multi-dimensional process $(\sigma_t^{(k,\tau)})_{(k,\tau) \in \mathcal{K} \times \mathcal{T}}$ where

$$\sigma_t^{(k,\tau)} = N_t^{(k,\tau)+} - N_t^{(k,\tau)-},$$

with $\mathbf{N} = (N^{(k,\tau)+}, N^{(k,\tau)-})_{(k,\tau) \in \mathcal{K} \times \mathcal{T}}$ is a multi-dimensional Hawkes process. The processes $N_t^{(k,\tau)+}$ (resp. $N_t^{(k,\tau)-}$) count the number of upward (resp. downward) moves of the implied volatility of option with strike k and maturity τ . This modelling is especially well suited to FX options markets, where options are directly quoted in terms of implied volatility. This convention is related to the notion of market impact: a net buy pressure on option (k, τ) will lead to a higher implied volatility $\sigma^{(k,\tau)}$ and conversely for a net sell pressure. The intensity of the vector \mathbf{N} is

$$\lambda_t = \mu_t + \int_0^t \phi(t-s) d\mathbf{N}_s.$$

We show in Figure 7 an example of smile modeling on a slice of volatility with three options. To ensure meaningful results, it is crucial to enforce no-arbitrage conditions of the volatility surface at the microscopic scale.

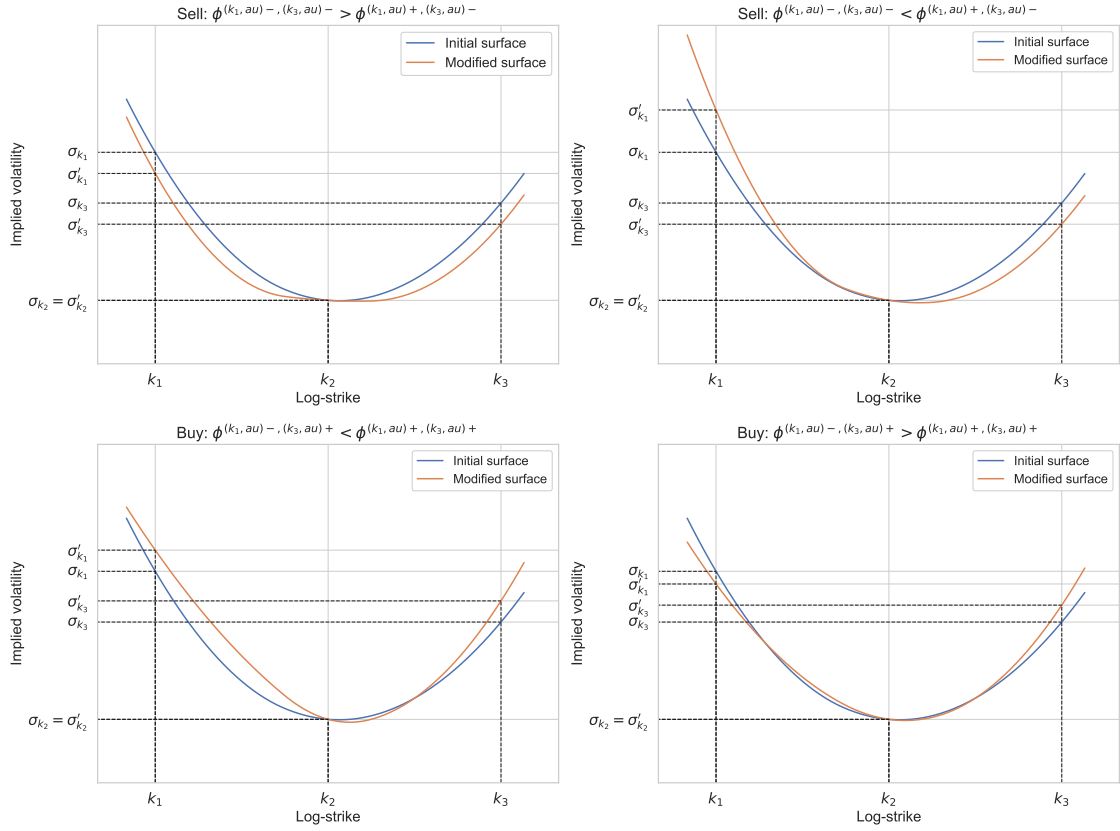


Figure 7: Evolution of the slice of volatility with sell orders on (k_3, τ) and $\phi^{(k_1, \tau)-, (k_3, \tau)-} > \phi^{(k_1, \tau)+, (k_3, \tau)-}$ (upper left), $\phi^{(k_1, \tau)-, (k_3, \tau)-} < \phi^{(k_1, \tau)+, (k_3, \tau)-}$ (upper right). Evolution of the slice of volatility with buy orders on (k_3, τ) and $\phi^{(k_1, \tau)-, (k_3, \tau)+} < \phi^{(k_1, \tau)+, (k_3, \tau)+}$ (lower left), $\phi^{(k_1, \tau)-, (k_3, \tau)+} > \phi^{(k_1, \tau)+, (k_3, \tau)+}$ (lower right).

Following Gatheral and Jacquier [125], we consider that a volatility surface is arbitrage-free if it is free of calendar-spread and butterfly arbitrage. There is absence of arbitrage opportunity in a calendar spread strategy if European call option prices are monotonous with respect to the maturity. By rewriting this condition in terms of intensities of Hawkes processes, we obtain the following result.

Result 13. *The volatility surface $\sigma_t = (\sigma_t^{(k,\tau)})_{(k,\tau) \in \mathcal{K} \times \mathcal{T}}$ is free of calendar-spread arbitrage if*

$$\begin{aligned} \phi^{(k, \tau_i)+, \cdot} &= \beta^{\phi,+} \sqrt{\frac{\tau_j}{\tau_i}} \phi^{(k, \tau_j)+, \cdot}, & \phi^{(k, \tau_i)-, \cdot} &= \beta^{\phi,+} \sqrt{\frac{\tau_j}{\tau_i}} \phi^{(k, \tau_j)-, \cdot}, \\ \mu^{(k, \tau_i)+} &= \beta^{\mu,+} \sqrt{\frac{\tau_j}{\tau_i}} \mu^{(k, \tau_j)+}, & \mu^{(k, \tau_i)-} &= \beta^{\mu,-} \sqrt{\frac{\tau_j}{\tau_i}} \mu^{(k, \tau_j)-}, \end{aligned} \quad (9)$$

where $(\beta^{\mu,+}, \beta^{\mu,-}, \beta^{\phi,+}, \beta^{\phi,-}) \in (1, +\infty)^4$.

Condition (9) leads to a maturity parametrization for the Hawkes kernel, thus reducing the number of parameters to calibrate.

There is absence of butterfly arbitrage if and only if the function $f_S(k) = \partial_{kk}C(k)$, where $C(k)$ is the market price of a call option with strike k , is a probability density. The following result is an analogous to the Roger Lee's moment formula, see Lee [197], for our discrete implied volatility surface.

Result 14. *We have $f_S(k) \rightarrow_{k \rightarrow +\infty} 0$ if and only if*

$$(\sigma^{(k,\tau)})^2 < \frac{2k}{\tau}.$$

A sufficient condition is

$$\begin{aligned} \phi^{(k_i,\tau)+,\cdot} &= \beta \sqrt{\frac{k_i}{k_j}} \phi^{(k_j,\tau)+,\cdot}, & \phi^{(k_i,\tau)-,\cdot} &= \beta \sqrt{\frac{k_i}{k_j}} \phi^{(k_j,\tau)-,\cdot}, \\ \mu^{(k_i,\tau)+} &= \beta \sqrt{\frac{k_i}{k_j}} \mu^{(k_j,\tau)+}, & \mu^{(k_i,\tau)-} &= \beta \sqrt{\frac{k_i}{k_j}} \mu^{(k_j,\tau)-}, \end{aligned} \tag{10}$$

for large $(k_i, k_j) \in \mathcal{K}^2, \tau \in \mathcal{T}$.

Ensuring the positivity of the function f_S is hard in the general case. Thus, we propose an alternative to this condition which ensures at-the-money convexity of the volatility slices. First, we consider a three-point volatility surface meaning that we choose three strikes corresponding to options $-25\Delta P, 50\Delta$ and $25\Delta C$:

- the put option $-25\Delta P$ has a strike such that its delta is equal to -0.25 ,
- the call option 50Δ has a strike such that the delta of a call option is equal to 0.5 ,
- the call option $25\Delta C$ has a strike such that its delta is equal to 0.25 .

We denote by $\sigma^{-25\Delta P,\tau}, \sigma^{50\Delta,\tau}, \sigma^{25\Delta C,\tau}$ the volatilities associated to these options. The implied volatilities of the so-called risk-reversal and butterfly (which are in fact measures of the skew and the convexity of the slice) with maturity τ are given by

$$BF_{25\Delta,\tau} = \frac{\sigma^{25\Delta C,\tau} + \sigma^{-25\Delta P,\tau}}{2} - \sigma^{50\Delta,\tau}, \quad RR_{25\Delta,\tau} = \sigma^{25\Delta C,\tau} - \sigma^{-25\Delta P,\tau}.$$

In the spirit of the previously mentioned results, we find sufficient conditions on the coefficients of the Hawkes kernel to ensure the convexity of the at-the-money strike, and therefore absence of arbitrage of the three-points volatility surface.

The simplicity of the formulas allows to give simple conditions on the kernel parameters so that the volatility slices are right or left skewed.

Result 15. *Consider the following general power-law kernel*

$$\phi^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}}(t) = \frac{\alpha^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}}}{(1+t)^{1+\gamma^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}}}},$$

with $(k, \tilde{k}) \in \{-25\Delta P, 50\Delta, 25\Delta C\}^2, (\tau, \tilde{\tau}) \in \mathcal{T}, (s, s') \in \{+, -\}^2$ and assume stationary increments of the Hawkes processes. We have

$$\bar{\lambda} = \mathbb{E}[\lambda] = (I - \tilde{\phi})^{-1} \mu,$$

where $\tilde{\phi}^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}} = \frac{\alpha^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}}}{\gamma^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}}}$. Thus, the three-point volatility surface is right (resp. left) skewed if for every $\tau \in \mathcal{T}$

$$\bar{\lambda}^{(25\Delta C,\tau)+} - \bar{\lambda}^{(25\Delta C,\tau)-} = \beta^{RR}(\bar{\lambda}^{(-25\Delta P,\tau)+} - \bar{\lambda}^{(-25\Delta P,\tau)-}), \quad \beta^{RR} > 1 \text{ (resp. } \beta^{RR} \in (0, 1)).$$

In particular, a high β^{RR} increases the right skewness of the volatility surface and conversely for β^{RR} close to zero.

We obtain similar results for a five-point volatility surface.

Finally, we can use the results on multidimensional Hawkes processes in Tomas and Rosenbaum [263], to prove the following result on the macroscopic limit of the implied volatility surface.

Result 16. Assume that, for all $(\tilde{k}, \tilde{\tau}) \in \mathcal{K} \times \mathcal{T}$, and $s, \tilde{s} \in \{+, -\}$, we have

$$\phi^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}}(t) = \sqrt{\frac{\tau}{\tilde{\tau}}} \sum_{i \in \{1, \dots, r\}} \tilde{z}_i(k) \tilde{z}_i(\tilde{k})^\top \varphi_i(t),$$

meaning that the self $(++, --)$ and cross $(+-, -+)$ exciting terms are the same.

At the limit, the macroscopic volatility surface is of the form $\sigma_t = \sum_{i \in \{1, \dots, r\}} v_i F_t^i$, where for $i \in \{1, \dots, r\}$, v_i are the eigenvectors of the option kernel and

$$\begin{aligned} F_t^i &= C^i \int_0^t \sqrt{V_s^i} dW_s^i, \\ V_t^i &= \frac{1}{\Gamma(\alpha)} \int_0^t (t-s)^{\alpha-1} (\theta^i - V_s^i) ds + \frac{1}{\Gamma(\alpha)} \int_0^t (t-s)^{\alpha-1} \lambda^i \sqrt{V_s^i} dZ_s^i, \end{aligned}$$

where (W^i, Z^i) are independent Brownian motion for $i \in \{1, \dots, r\}$.

The macroscopic limit of the volatility surface dynamics is given by a sum of risk factors with rough volatility. That does not contradict to the studies on the roughness of the volatility as the implied volatility has no reason to be rough contrary to the realized volatility. This implied volatility modelling has several applications. We present two of them in this thesis: backtest of option market-making strategies and computation of market impact curves of option market-making strategies.

This result concludes the second part of this thesis. We now move to the last part dedicated to the second main area of the literature on systematic strategies: optimal trading.

Part III Optimal trading in high dimension

III.1 Adaptive trading strategies across liquidity pools

Most of the first studies of optimal trading examined the problem of optimal scheduling. The scheduling consists in finding an optimal plan for a large buy or sell order minimizing the overall trading costs. There is a trade-off between trading fast, which leads to high execution costs and market impact and trading slow, which leads to low market impact but high uncertainty on price change. In Almgren and Chriss [9, 10], the authors propose a framework to address this issue of large order optimal splitting. In their model, the authors assume a Bachelier price dynamics with linear execution costs and discrete periods at which the trader chooses the volume to execute. Since then, the Almgren-Chriss model has been widely used in practice with or without numerous generalizations such as different price dynamics, non-linear or transient market impact, illiquidity of the traded asset, see among others Almgren [8, 7], Forsyth [119], Forsyth, Kennedy, Tse, and Windcliff [120], Gatheral, Schied, and Slynko [127], Lehalle [198], Schied and Schöneborn [248]. In all these models, the market's interactions are hidden: their

aim is only to find the right quantity to buy or sell at a given time for certain market conditions. The first main contribution in this direction, considering aggressive orders only, is Obizhaeva and Wang [224], extended in Alfonsi, Fruth, and Schied [5], Predoiu, Shaikhet, and Shreve [234]. In Bayraktar and Ludkovski [39], the authors introduce an optimal liquidation framework using passive orders. This idea has been used notably in Guéant [136], Guéant and Lehalle [139], Guéant, Lehalle, and Fernandez-Tapia [141] to solve the optimal scheduling and trading problems together. The framework used to solve optimal liquidation using passive orders is directly inspired by the market-making model of Avellaneda and Stoikov [26], and it is more suitable for liquidation on quote-driven markets. Optimal liquidation models in limit order books have been notably developed by Guilbaud and Pham [143] using stochastic control theory. The trader can choose the order's volume, aggressiveness, and the limit at which to post. One important drawback of these models is the impossibility of providing simple expressions for the trader's optimal control: they rely either on first order approximations or on numerical approximations for PDEs. Moreover, they do not address the problem of optimal trading in several liquidity platforms.

The problem of optimal splitting of orders across liquidity pools has been treated for instance in Almgren and Harts [11], Cont and Kukanov [81], Laruelle, Lehalle, and Pagès [193, 194]. In Almgren and Harts [11], the authors develop a dynamic estimate of the hidden liquidity present on several venues and use this information to make order splitting decisions (Smart Order Routing). The paper Cont and Kukanov [81] solve a general order placement problem and provide an explicit solution for the optimal split between limit and market orders in different venues. Finally, in Laruelle, Lehalle, and Pagès [193, 194], the authors build a stochastic algorithm to find the optimal splitting between liquidity pools, including dark pools. These models are mostly “static” because they do not combine the optimal scheduling of orders with a long term inventory target. Moreover, an optimal trading model's quality mainly relies on estimating the market parameters (cross-dependence between the imbalance and spread of each venue and the probability and the proportion of execution of limit orders). Some optimal liquidation models treat online updates of market parameters, see Almgren and Lorenz [12, 13], however, those are usually rather parsimonious models, where the parameters are, for example, only the drift and the volatility of the price process, updated via Bayesian updates. There is still a lack of literature treating optimal trading in a more dynamic fashion, which leads to the following problem.

Problem 8: How to design adaptive trading strategies across liquidity pools?

We consider the trader dealing in a stock listed in several venues by placing limit and market orders. To answer the question, we first formulate the trading problem via stochastic control theory – the controls are volumes of limit and market orders sent to each venue and limits chosen for passive orders. The optima are obtained from a classical HJB quasi-variational inequality (HJBQVI), which, for a parsimonious market modeling, can be easily solved by grid methods. Then we propose a Bayesian update of each market parameter, decoupled from the control problem. One of the advantages of this method is the simplicity of the formulae for each parameter's posterior estimate. In particular, we do not need to use Markov chain Monte-Carlo. This type of update of the market parameters, contrary to the continuous one, does not increase the number of state variables of the control problem drastically and does not lead to computation time handicaps. Continuous Bayesian updates would require first to compute the conditional expectation of the value function given the market parameters and then integrate it over their posterior distribution. This last integration brings multiple non-linearities in the equation, making this fully Bayesian control problem almost impossible to solve numerically.

We consider a trader acting in N liquidity platforms operating with limit order books over time interval $[0, T]$. He trades continuously in each venue by sending limit and market orders. For $n \in \{1, \dots, N\}$, the n -th venue is characterized by the bid-ask spread process $(\psi_t^n)_{t \in [0, T]}$ and the imbalance process $(I_t^n)_{t \in [0, T]}$ which are continuous-time Markov chains. The general formulation we use allows for a full coupling between the spread and imbalance of all venues. The number of, possibly partially, filled ask

orders in the venue n is modeled by a Cox process denoted by $N^n, n \in \{1, \dots, N\}$ with intensities $\lambda^n(\psi_t, I_t, p_t^n, \ell_t)$, where p_t^n represent the limit at which the trader sends a limit order of size ℓ_t^n .

Practically for $n \in \{1, \dots, N\}$, when the spread is equal to the tick size, the trader can post at the first best limit ($p^n = 0$) or the second best limit (if $p^n = 1$). When the spread is equal to two ticks or more, the trader can either create a new best limit ($p^n = -1$) or post at the best or the second best limit as previously. We also allow for partial execution of limit orders, represented by random variables $\epsilon_t^n \in [0, 1]$ following also a multi-regime categorical distribution.

Finally, we allow for the execution of market orders (denoted by a point process $(J_t^n)_{t \in [0, T]}$) in each venue of size $(m_t^n)_{t \in [0, T]} \in [0, \bar{m}]$, where $\bar{m} > 0$ and $J_t^n = J_{t-}^n + 1$. We assume that market orders are always fully executed. The cash process X_t of the trader at time $t \in [0, T]$ is obtained by summing the volume executed on each venue with limit and market orders, multiplied by their transaction price. The inventory process of the trader at time $t \in [0, T]$ is defined by

$$q_t = q_0 - \sum_{n=1}^N \left(\int_0^t \ell_u^n \epsilon_u^n dN_u^n + \int_0^t m_u^n dJ_u^n \right). \quad (11)$$

We also assume that the trader has a pre-computed trading curve q^* that he wants to follow (coming from a high-level execution problem, such as the Almgren-Chriss trading curve or VWAP strategy, for example). Then the trader's optimization problem is

$$\sup_{p, \ell, m} \mathbb{E} \left[X_T + q_T S_T - \int_0^T g(q_t - q_t^*) dt \right], \quad (12)$$

where the function g penalizes deviation from the pre-computed optimal trading curve and S is the asset price, following arithmetic Brownian motion. By using the theory of viscosity solutions, we obtain that the value function of the control problem of the trader is characterized by a unique viscosity solution to Equation (9.3).

Conditionally on the market parameters such as the transition matrix of both the spread and the imbalance processes, the drift and volatility of the underlying asset and the execution proportion probabilities, solving the HJBQVI is done using simple finite difference schemes and the optimal splitting of volumes as well as the optimal limits can be computed in advance. We propose a practical way to update the market parameters according to trader's observations in a Bayesian way. This method, which is performed separately from the optimization procedure, allows to update, at the end of a trading time slice, the trading strategy according to changing market conditions. We present as a result the description of our algorithm to obtain adaptive trading strategies.

Result 17. *Let us consider $\mathcal{V} > 0$ slices $\mathcal{T}_v = [T_v, T_{v+1}]$, $v = 0, \dots, \mathcal{V} - 1$, such that $T_0 = 0, T_{\mathcal{V}} = T$. We define for each slice $v \in \mathcal{V}$ a set of market parameters θ_v^m . At each time slice $v \in \{0, \mathcal{V} - 1\}$ starting from $v = 0$ we perform the following algorithm:*

1. *Take the best estimation of market parameters θ_v^m from the prior distribution for the current slice v .*
2. *Compute the optimal trading strategy on \mathcal{T}_v using the set of parameters θ_v^m .*
3. *Observe market events during the current slice (executions, changes of the state).*
4. *At T_{v+1} , update the parameters θ_{v+1}^m following the Bayes rules.*

In Figures 8 and 9, we illustrate the example output of the above algorithm for two venues case: limits and sizes for passive orders respectively, for some precise market conditions, such as spreads at one tick for both venues, negative imbalance in the first venue, and positive imbalance in the second venue.




Figure 8: Limit order strategy, $\psi^1 = \psi^2 = \delta$, $I^1 = -0.5, I^2 = 0.5$ using neural networks.




Figure 9: Volume strategy, $\psi^1 = \psi^2 = \delta$, $I^1 = -0.5, I^2 = 0.5$ using neural networks.

The method we propose aims at giving a useful and applicable solution for practitioners who work on cross-trading strategies. The control model is flexible enough for a quantitative firm to reproduce the market’s main stylized facts and design trading strategies based on real signals. Moreover, the Bayesian update of market parameters in the control problem enables reevaluating the optimal strategy when the market conditions may differ from the trader’s prior empirical estimation. We show that one can compute the optimal trading strategy on a slice given new market conditions in a couple of minutes using finite difference schemes or deep reinforcement learning methods (which could also be mixed).

III.2 Optimal trading without optimal control

The model we presented combines order scheduling (through the pre-computed trading curve) and order placement (through the interaction with the limit order books of venues). Let us assume that the trader observes a long term signal indicating that he must change his inventory target. He has to recompute the solution of the HJB equation to obtain the new optimal strategy. Thus, if the long-term target varies frequently, it is hard to use the proposed model because of the computation time. Another limitation is due to the curse of dimensionality: if one wants to use this model on a portfolio of assets listed on several liquidity venues, the complexity of the associated HJB equation will make the problem intractable. This intractability for high-dimensional problems is intrinsically linked to optimal control methods, so we consider the following question.

Problem 9: How to solve optimal trading problems without optimal control?

If an execution desk is concerned only about the optimal splitting of orders, then the problem boils down to a static maximization as we do not need the inventory trajectory. The true problem is interacting with the order books of the assets while taking into account a long-term trading trajectory. According to Cont and Kukanov [81], who developed a model for orders’ splitting, ‘Although simultaneous optimization of order timing, type, and routing decisions is an interesting problem, it also appears to be intractable’ [81, p. 4].

The method we present in the following breaks the problem up into two parts. The first part is standard: we build a smooth relaxation of the trading problem; the second part is an adjustment of microstructure decisions to follow the smooth relaxation optimally. The key feature of all smooth relaxations is that they hide microstructure details behind generic cost functions meant to represent the average cost of trading at a given rate. The relaxations do not guide microstructure-level decisions, effectively assuming all executions are made with market orders. In particular, if we can predict the probability of a passive fill at any given instant (e.g., based on order book imbalance), it is not clear how to use this information in the context of a smooth relaxation problem.

We first introduce some notation. We denote by $R_{i,t}(v, a)$ the (arbitrary) PnL of an order quantity v on stock i using action a over a short interval $[t, t + \delta t]$. The action a includes the trader’s choice

of whether to trade passively or aggressively. A trader's expected profit $\mathbb{E}[R_{i,t}(v, a)]$ depends on the trading cost associated with the pair (v, a) , and also on the trader's views concerning the short-horizon midpoint price return

$$r_i^{\text{mid}} = \text{mid}_i(t + \delta t) / \text{mid}_i(t) - 1.$$

With no subscript, r^{mid} denotes the d -dimensional vector of all midpoint returns for all assets. The term effective microstructure alphas denote a set of parameters given to a trader as a short or long term view on $\mathbb{E}[r^{\text{mid}}]$. Microstructure alphas could be a simple prediction of r^{mid} , or they could be deliberately skewed, for example, to match long or short term trading target.

First, we choose a long-term trading curve corresponding to the so-called smooth relaxation problem. For optimal trading problems, we choose the multi-asset Almgren-Chriss trading curve: we consider $d \geq 1$ assets, with initial positions $q_0 = (q_0^1, \dots, q_0^d)^\top$, where $q_0^i \in \mathbb{R}$ for all $i \in \{1, \dots, d\}$. The trader wants to obtain the portfolio $q_T \in \mathbb{R}^d$ at time $T > 0$. Given a control process $(v_t^i)_{t \in [0, T]}$ representing the trading rate on asset i , the inventory process of the i -th asset is given by

$$q_t^i = q_0^i - \int_0^t v_s^i ds, \quad i \in \{1, \dots, d\}.$$

For each stock, we consider Gaussian price dynamics:

$$dS_t^i = \sigma^i dW_t^i,$$

where the Brownian motions $(W_t^i)_{t \in [0, T]}$ are such that $(\sigma^1 W_t^1, \dots, \sigma^d W_t^d)_{t \in [0, T]}$ has a nonsingular covariance matrix Σ . We define the cost function $c(v) = \lambda \|v\|_2^2$. The smooth relaxation problem associated to $c(\cdot)$, Σ , q_T is defined to be:

$$V(0, q_0) = \min_{q \in C^2([0, T], \mathbb{R}^d)} \int_0^T L(q_s, \dot{q}_s) ds \text{ subject to } q_0 = q_0, q_T = q_T, \quad (13)$$

where $\kappa > 0$ is the risk-aversion constant, and the (autonomous) Lagrangian is given by:

$$L(q, v) = c(v) + \frac{1}{2} \kappa (q - q_T)^\top \Sigma (q - q_T).$$

It is known that the solution to (13) is given by

$$q_t^* = (C^\top)^{-1} \Omega \left(e^{D^{\frac{1}{2}}(T-t)} - e^{-D^{\frac{1}{2}}(T-t)} \right) \left(e^{D^{\frac{1}{2}}T} - e^{-D^{\frac{1}{2}}T} \right)^{-1} \Omega^\top C^\top q_0, \quad (14)$$

where $\Sigma = CC^\top$ is the Cholesky decomposition of C and $\Omega D \Omega^\top$ is a spectral decomposition of the positive definite matrix $\frac{\kappa}{2} C^\top \Lambda^{-1} C$. In particular, the value function V of the smooth relaxation problem can be computed analytically for an arbitrary large number of asset. The associated trading curve (14) can be computed in advance, and corresponds to the order scheduling decision. Finally, the function V satisfies the Hamilton-Jacobi differential equation:

$$\partial_t V(t, q) + H(q, \nabla V) = 0,$$

where $H(q, p) = \sup_v \{ \langle p, v \rangle - L(q, v) \}$, with the singular final condition:

$$V(T, q) = \begin{cases} 0, & \text{if } q = q_T \\ \infty, & \text{if } q \neq q_T. \end{cases}$$

Along an optimal trajectory q^* , we have

$$\nabla V|_{x=q_t^*} = L_v(q_t^*, \dot{q}_t^*) = p_t \quad \forall t \in [0, T],$$

where $p = \partial_v L(q, \dot{q})$ is the so-called generalized momenta. In our case of separable and autonomous Lagrangian L , it implies that the optimal instantaneous trade $v^* := \dot{q}^*(t)$ at each time t solves

$$\dot{q}^* = \underset{v}{\operatorname{argmax}} \{ \langle p, v \rangle - c(v) \}. \quad (15)$$

One can interpret (15) as optimization in a risk-neutral world. Indeed, if a risk-neutral agent had a vector of expected returns which happened to exactly equal the vector of generalized momenta, p , and sought only to maximize net profit, irrespective of risk, then (15) is the problem faced by this agent. This agent can be considered “myopic” because any information concerning more than one period ahead is available to the agent only indirectly, insofar as p depends on the rest of the trading path. The value function of the myopic trader at time $t \in [0, T]$ for an inventory $q \in \mathbb{R}^d$ is defined as

$$W(t, q) = \int_t^T \left(\langle p_s, v_s^* \rangle - c(v_s^*) \right) ds. \quad (16)$$

In other words, the value function of a myopic trader defined in (16) is simply the sum of his instantaneous trading gains over time.

Result 18. *Assume that L is separable. A myopic agent sending market orders with instantaneous cost function $c(\cdot)$ must choose microstructure alphas $p = \nabla V$ in order to minimize the absolute error between its value function and the long-term objective function V defining the trading scheduling. Moreover, we have bounds on the error uniform with respect to time and inventories.*

This simple result has several important consequences. If one wants to avoid the use of optimal control and still wants to follow the Almgren-Chriss trading curve, one can simply solve the static optimization problem (15) at discrete times (the times of trading), using $p = \nabla V$. Equation (15) does not give a full set of instructions for the trader with a long-term trading schedule who has to interact with a continuous limit order book market, but it can serve as a guide. Indeed, the order routing problem, treated notably in Cont and Kukanov [81], takes into account the possibility to send limit, market or cancel orders to several liquidity venues, depending on their spread and imbalance. Stochastic control appears to be inefficient for this problem, as one needs to solve a high-dimensional HJB equation. The advantage of the methodology we present is that one can avoid optimal control and solve a simple static optimization problem to determine the optimal action at each discrete trading time.

The use of the generalized momenta as effective microstructure alpha has a wide range of practical implications. First, it offers a way to bridge the gap between order placement decisions and scheduling decisions, usually decoupled in practice. Second, the microstructure formulation helps to tackle classic optimal control in limit order books. For example, a realistic optimal trading framework dealing with a portfolio of assets in several liquidity venues is in practice intractable due to the problem’s dimensionality. In the method presented above, the trader’s optimal controls (the order’s asset, volume, and level of aggressiveness) are derived through a simple static optimization problem, which can be solved for a large portfolio of assets traded in numerous venues. The convergence to the trading schedule is guaranteed by choice of the effective microstructure alpha p .

Taking the same notation as in the previous chapter, we can formulate a myopic version of the optimal trading problem across liquidity pools. Consider a trading schedule for $d = 1$ asset $q^* \in \mathbb{R}$ (the Almgren-Chriss trading schedule described in Equation (14), for example) with associated value function $V(t, q_t^*)$. The model is directly extended to the case of multiple assets listed in multiple venues in Chapter 10. The cost of a limit order of size ℓ sent at limit p in venue n is defined as $c^{n, \mathcal{L}}(\ell, p)$, and $c^{n, \mathcal{M}}(v)$ is the cost function of a market order of size v in venue n .

The myopic trader acts at discrete times and at time $t \in [0, T]$ for all $\mathbf{m} \in \mathcal{M}$, his optimization problem is given by the static optimization (Opt-Trd). For each state \mathbf{m} , the model’s output are the optimal volumes and limits $\ell^{*n}(\mathbf{m}), p^{*n}(\mathbf{m})$ for each liquidity venue. In this general framework,

order scheduling with a long-term target is easily tractable even for a large portfolio of assets and a large number of liquidity venues, as the trader has to solve a static optimization problem at each trading time. The effective microstructure alpha p^{eff} must be recomputed each time the portfolio holdings change. We also show in Chapter 10 that we can apply similar ideas to solve the problem of multi-asset market-making.

III.3 Two new developments on optimal execution

We conclude this thesis with two new developments related to optimal execution in Chapter 11. We first show that there exists an analytic solution to the Almgren-Chriss execution problem when the price process is a geometric Brownian motion and the cost function is quadratic. This offers an answer to the conjecture formulated by Gatheral and Schied [126], where the authors called this problem analytically intractable. The method we propose, which uses tools from the operators' theory, is also an alternative to the use of BSDE when the drift of the asset is a stochastic process and classical HJB equation, see [120]. We define the price process S with volatility $\sigma > 0$ and the trader's position expressed in dollars $\theta_t = Q_t S_t$, where Q_t is his inventory at time t . The trading's rate in dollar at time t is $u_t = \dot{q}_t S_t$ where \dot{q}_t is the trading's rate in shares. We also consider temporary linear market impact with intensity parameter $\lambda > 0$. The optimization problem faced by the trader is

$$\lim_{a \rightarrow +\infty} \sup_u \mathbb{E} \left[\int_0^T - \left(\frac{\lambda}{2} u_t^2 + \frac{\kappa \sigma^2}{2} (\theta_t^u)^2 \right) dt - \frac{a}{2} \theta_T^u \int_0^T u_t dt \right],$$

where $\kappa > 0$ is the risk-aversion parameter of the trader and the limit over $a > 0$ aims at representing the singular condition $\theta_T^u = 0$. The following result is the main contribution of the paper.

Result 19. *The optimal control is given analytically for all $t \in [0, T]$ by*

$$u_t^* = \theta_t^* \Gamma(t),$$

where $\Gamma(\cdot)$ is a deterministic function. Moreover, the optimal trader's position satisfies

$$\theta_t^{u^*} = \theta_t^{u^*} \exp \left(\int_0^t \left(\Gamma(s) - \frac{\sigma^2}{2} \right) ds + \sigma W_t \right).$$

An extension that we provide in the appendix of this article is that liquidation of a portfolio of correlated assets is equivalent to the resolution of a system of Riccati ODEs, which is easily tractable for a large number of assets.

The second development we propose is an application of the latent order book model (LLOB) of [264] to liquidity stress testing on OTC markets. The general idea of this model is that there exists a latent order book which at time t aggregates the total intended volume to be potentially sold (resp. bought) at price $p > 0$ or above (resp. below) $\mathcal{V}_+(t, p)$ (resp. $\mathcal{V}_-(t, p)$). These latent volumes are the volumes that would be revealed as limit or market orders if the price comes closer to p at some point. By assuming that the price process is a Brownian motion with volatility $\sigma > 0$ and that new buy and sell orders of unit volume arrive randomly with intensity $\lambda > 0$, the density of order book trades $\rho(u)$ is obtained in closed form. If we integrate this quantity from mid-price to a change of price Δp , we obtain the order book volume as a function of price change. By assuming a linear liquidation, in which the block of assets is unwound in equal parts, and a mean-variance objective function with risk-aversion parameter $\gamma > 0$, we obtain the following result.

Result 20. *The optimal liquidation time of a linear liquidation on a single asset is given by*

$$T^* = \frac{\sqrt{6}}{\gamma P_0} \sqrt{\frac{N}{ADV}},$$

where N is the face amount held by the trader, ADV is the average daily volume on the asset and P_0 is its initial price. Moreover, the cost dependence on trade size is of magnitude $N^{\frac{1}{4}}$ for small trades and $N^{\frac{1}{3}}$ for large trades.

This result is extended to the case of optimal liquidation of a portfolio of correlated assets, and is illustrated on a portfolio of corporate bonds.

List of papers

- Chapter 1. B. Baldacci, I. Manziuk, T. Mastroia, M. Rosenbaum, *Market-making and incentives design in the presence of a dark pool: a deep reinforcement learning approach*, submitted, 2019. ArXiv preprint arXiv:1912.01129.
- Chapter 2. B. Baldacci, D. Possamaï, M. Rosenbaum *Optimal make-take fees in a multi market-maker environment* to appear in SIAM Journal on Financial Mathematics, 2021.
- Chapter 3. B. Baldacci, P. Bergault, J. Derchu, M. Rosenbaum, *On bid and ask side-specific tick sizes*, submitted, 2020. ArXiv preprint arXiv:2005.14126.
- Chapter 4. B. Baldacci, P. Jusselin, M. Rosenbaum *How to design a derivatives market?*, submitted, 2019. ArXiv preprint arXiv:1909.09257.
- Chapter 5. B. Baldacci, D. Possamaï, *Governmental incentives for green bonds investment*, submitted, 2020. ArXiv preprint arXiv:2101.00648.
- Chapter 6. B. Baldacci, P. Bergault, O. Guéant, *Algorithmic market-making for options*, Quantitative Finance, 2020.
- Chapter 7. B. Baldacci, J. Derchu, I. Manziuk, *An approximate solution for options market-making in high dimension*, to appear in Risk Magazine, 2021.
- Chapter 8. B. Baldacci, *High-frequency dynamics of the volatility surface*, submitted, 2020. ArXiv preprint arXiv:2012.10875.
- Chapter 9. B. Baldacci, I. Manziuk, *Adaptive trading strategies across liquidity pools*, submitted, 2020. ArXiv preprint arXiv:2008.07807.
- Chapter 10. B. Baldacci, J. Benveniste, G. Ritter *Optimal trading without optimal control*, submitted, 2020. ArXiv preprint arXiv:2012.12945.
- Chapter 11. B. Baldacci, J. Benveniste, *A note on Almgren-Chriss optimal execution problem with geometric Brownian motion*, to appear in Market Microstructure and Liquidity, 2021.
- Chapter 12. B. Baldacci, I. Manziuk, M. Weber, *Liquidity stress testing using optimal portfolio liquidation*, submitted, 2020. ArXiv preprint arXiv:2102.02877.

Part I

Regulation of market-making activity via Principal-Agent theory

Notations: For $(v_1, v_2) \in \mathbb{R}^d$, $v_1 \cdot v_2 \in \mathbb{R}$ denote the scalar product between v_1 and v_2 whereas $v_1 \circ v_2 \in \mathbb{R}^d$ is the component-wise multiplication of the vectors. Let \mathbb{N}^* be the set of all positive integers. For any $(\ell, c) \in \mathbb{N}^* \times \mathbb{N}^*$, $\mathcal{M}_{\ell, c}(\mathbb{R})$ will denote the space of $\ell \times c$ matrices with real entries. Elements of the matrix $M \in \mathcal{M}_{\ell, c}$ are denoted $(M^{i,j})_{1 \leq i \leq \ell, 1 \leq j \leq c}$ and the transpose of M is denoted M^\top . We identify $\mathcal{M}_{\ell, 1}$ with \mathbb{R}^ℓ . When $\ell = c$, we let $\mathcal{M}_\ell(\mathbb{R}) := \mathcal{M}_{\ell, \ell}(\mathbb{R})$. For any $x \in \mathcal{M}_{\ell, c}(\mathbb{R})$, and for any $1 \leq i \leq \ell$ and $1 \leq j \leq c$, $x^{i,\cdot} \in \mathcal{M}_{1, c}(\mathbb{R})$, and $x^{\cdot, j} \in \mathbb{R}^\ell$ denote respectively the i -th row and the j -th column of M . Moreover, for any $x \in \mathcal{M}_{\ell, c}(\mathbb{R})$ and any $1 \leq j \leq c$, $x^{\cdot, -j} \in \mathcal{M}_{\ell, c-1}(\mathbb{R})$ denote the matrix x without the j -th column. For any $x \in \mathcal{M}_{\ell, c}(\mathbb{R})$ and $y \in \mathbb{R}^\ell$, we also define for $i = 1, \dots, \ell$, $y \otimes_i x \in \mathcal{M}_{\ell, c+1}(\mathbb{R})$ as the matrix whose first $i-1$ columns are equal to the first $i-1$ columns of x , such that for $j = i+1, \dots, c+1$, the j -th column is equal to the $(j-1)$ -th column of x , and whose i -th column is equal to y . For any $x \in \mathbb{R}^\ell$, we also define $\underline{x} := \min_{i \in \{1, \dots, \ell\}} x^i$. We also define 1_N the vector of \mathbb{R}^N with every component equal to one. For any $d \in \mathbb{N}^*$, \mathbb{S}_d is the space of $d \times d$ -dimensional symmetric matrices. For any $(\ell, c) \in \mathbb{N}^* \times \mathbb{N}^*$, we define I_ℓ as the identity matrix of $\mathcal{M}_\ell(\mathbb{R})$, and $\mathbf{0}_{\ell, c}$ as a matrix in $\mathcal{M}_{\ell, c}(\mathbb{R})$ with all entries equal to zero. We define the function $\text{diag} : \mathbb{R}^d \rightarrow \mathcal{M}_d(\mathbb{R})$ such that for $v \in \mathbb{R}^d$, and any $(i, j) \in \{1, \dots, d\}^2$, $\text{diag}(v)^{i,j} := v^i$ if $i = j$, and 0 otherwise. For $x \in \mathcal{M}_{\ell, c}(\mathbb{R})$, we define $\|x\|_2 := \sum_{(i,j) \in \{1, \dots, \ell\} \times \{1, \dots, c\}} (x^{i,j})^2$. The only exception to these notations is Chapter 5 where we use indices instead of exponents for the coefficients of matrices and vectors.

Chapter 1

Market-making and incentives design in the presence of a dark pool: a deep reinforcement learning approach

1.1 Introduction

Since the seminal work Avellaneda and Stoikov [26], a vast literature on optimal market-making problems has emerged. A market-maker is a liquidity provider whose role is to post orders on the bid and ask sides of the limit order book of an underlying asset. Various extensions of Avellaneda and Stoikov [26] have been considered, see for example Cartea, Jaimungal, and Ricci [69], Guéant, Lehalle, and Fernandez-Tapia [142] and the books Cartea, Jaimungal, and Penalva [70], Guéant [137] for further references. In most of these works, it is assumed that there is no make-take fees system on the market. A make-take fees policy is a contract between market participants and an exchange enabling the latter to improve liquidity on his platform. Typically these contracts subsidize liquidity provision while taxing liquidity consumption *via* fixed and/or proportional fees with respect to the amount traded. The problem of relevant make-take fees is studied quantitatively in El Euch, Mastrolia, Rosenbaum, and Touzi [109] and Chapter 2 of this thesis. In these papers, the policies are designed in the context of traditional liquidity venues, or so-called “lit pools”. On these venues, the order book is visible to market participants, and transactions are fully transparent. Market takers can in particular monitor the quotes offered by market-makers.

The initial purpose of market fragmentation, and the rise of dark pools, is to increase competition between trading platforms to the benefit of investors. An unexpected consequence is that some participants, namely the high-frequency traders, took advantage of this fragmentation structure and used their technological superiority to find new profit opportunities, see Abergel, Lehalle, and Rosenbaum [1], Laruelle and Lehalle [192] for example. Although trading rules for dark pools are very diversified, they share at least two important properties. The first one is the absence of a visible order book for market participants, which implies that investors have no information on the amount of liquidity posted by market-makers. Second, aiming at improving prices for clients compared to the lit venue, dark pools usually set prices that are different from those in the lit pool. For example, many dark pools take the mid-price of the lit pool as their transaction price. Because of these two effects, it is presumed that trades in dark pools have no or less price impact.¹ This feature enables market-makers to mitigate their inventory risk. Finally, dark pools are prone to a relative latency effect with respect to the associated lit venue. The price is monitored in the lit pool and can change between the time

¹Note however, that transactions’ reporting imposed by regulation in most markets may still induce some delayed price impact.

of a request in the dark pool and that of the corresponding transaction. The impact of latency issues on financial markets has been particularly investigated in Moallemi and Sağlam [216]. The advantages and disadvantages of trading in the dark pool for a market-maker are summarized below:

- **Advantages:** the market-maker can liquidate a large position without showing this information to other market participants (contrary to the case where he would have posted a large position on the lit pool, even at a higher limit). Moreover, this quantity will be executed at a more favorable price compared to the lit pool. Finally, the market-maker can use the relative latency effect on the dark pool to arbitrage market takers.
- **Disadvantages:** in the absence of relative latency, the market-maker does not collect the spread hence has no interest in being a liquidity provider on the dark pool. Moreover, there is uncertainty about the execution of the orders as no limit order book is visible.

As the market impact of trades on a dark pool is less important or delayed, market-makers can also use it to liquidate large positions. Therefore there is a trade-off between transacting in the dark pool at a lousy price with low impact or in the lit pool at a better price with higher impact. Dark pools are also very attractive for market takers because of the reduced market impact and the possibility to be executed at a better price than in the lit pool.

To our knowledge, most of studies treat the issue of trading in dark pools mainly from the point of view of optimal liquidation: a trader wishing to buy or sell a large number of shares of one or several stocks and needing to find an optimal order placement strategy between the liquidity pools (lit, dark or both), see for example Kratz [188], Kratz and Schöneborn [189, 190], Laruelle, Lehalle, and Pagès [193]. In Buti, Rindi, and Werner [60], the authors build a model of a limit order book and examine the consequences of adding a periodic and/or continuous dark pool from the investor's point of view. They show that large trades can be liquidated at a lower cost on the dark pool. In this chapter, we rather focus on the behavior of a market-maker, acting on both lit and dark venues of the same exchange. Also related to this work is the paper Bielağk, Horst, and Moreno-Bromberg [47], where the authors use a Principal-Agent setting to derive optimal quotes of a dealer competing with a dark pool to provide liquidity to a set of heterogeneous traders. In particular they show that, in a portfolio-liquidation setting, the presence of a dark pool reduces the dealer's spread. In the present chapter, the controls of the market-maker are the posted volumes on a lit limit order book and a dark pool.

As explained both empirically and theoretically in Foucault, Kadan, and Kandel [121], make-take fees policies are a way for trading platforms to increase their benefits charging in an asymmetric way liquidity providers and liquidity consumers. In our model, we aim at adapting this idea to an exchange managing a market-maker on a lit pool and a dark pool. In the lit market, we assume that there is an efficient price S_t and that the market-maker always posts volumes on the bid and ask sides at prices $S_t + \frac{\tau}{2}$ and $S_t - \frac{\tau}{2}$, where $\frac{\tau}{2}$ represents the half-tick of the market.² Considering only the top of the book is reasonable for many assets for which the tick size is so that most trading activity occurs at the best limit. The market-maker also provides liquidity in the dark pool where the transaction price is the efficient price S_t (possibly with the latency effect). When one trade in a lit and a dark pool, trading is impacted by three kinds of latencies: the first one corresponds to the time between the instant where the order is sent and received by the exchange; the second one is the processing time by the exchange. These two latencies occur on lit and dark pools. In our case, there is also a relative latency between the dark pool and the associated lit pool. In the following, we consider only the latter as it is the key component of arbitrage strategies between lit and dark pools. Our market-maker problem can partially be seen as the dual problem of El Euch, Mastrolia, Rosenbaum, and Touzi [109], without dark pool, where the posted volume is fixed at one unit, and the market-maker optimizes the quoted spread. In addition to market impact and latency phenomena, we also take into account transaction costs for market orders on both venues, which can be smaller in the dark pool. Thus, in our setting, a

²We have in mind here a large tick asset for which the spread equals the tick size.

single market-maker only needs to select the volumes to post on the bid and ask sides of both lit and dark pools.

Furthermore, we consider that the exchange managing the lit and dark venues, wishes to attract transactions. Inspired by the work El Euch, Mastrolia, Rosenbaum, and Touzi [109], we consider that the exchange offers a contract to the market-maker whose remuneration at a terminal time is determined according to the executed transactions on both venues. However, the exchange cannot control those volumes and may only provide incentives to influence the market-maker's behavior. These incentives take the form of a contract between the market-maker and the exchange, whose payoff depends on observed trading flows.

In our setting, the market-maker has to fix volumes in response to the incentives of the exchange. These volumes are functions of the incentives (and of the market-maker's inventory), which are the solution of a nonlinear equation. The optimization problem of the exchange is to solve

$$\sup_{\xi} \mathcal{J}(\xi, \mathcal{L}^*(\xi)),$$

where ξ denotes the compensation given to the market-maker, satisfying a reservation utility constraint for the market-maker, the quantity $\mathcal{L}^*(\xi)$ represents the volume sent by the market-maker on the lit and dark venues in response to the contract ξ and \mathcal{J} is the objective function of the exchange.³ Identifying this game with a Stackelberg leadership model, the resolution of our Principal-Agent problem consists of two stages. The first stage is to represent the volumes posted by the market-maker by a neural network. Taking into account the optimal response of the market-maker to given incentives, the exchange needs to choose the contract maximizing its utility. So the second stage is to solve the exchange problem under the best reaction of the market-maker to obtain the optimal contract. However, dimensionality and the high degree of nonlinearity of this equation make standard numerical methods hard to apply. We circumvent this difficulty by adopting a reinforcement learning method. More precisely, we use an actor-critic approach where not only the controls of the exchange, but also its value function are represented by neural networks. The essence of this method is the alternation of the learning phases of the controls and of the value function.

Technically speaking, to find an optimal contract and corresponding optimal volumes of the market-maker's orders, we need to solve a nonlinear Hamilton-Jacobi-Bellman (HJB for short) equation. Dimensionality (above four) and the resulting equations' complexity do not allow us to apply classical root-finding algorithms. Therefore we use a method based on neural networks to solve our HJB equations, see among others Han, Jentzen, and E [148], Huré, Pham, and Warin [169]. In this chapter, we present an application to optimal control Principal-Agent type problem. In particular, our work presents some similarities with Bachouch, Huré, Langrené, and Pham [29], Guéant and Manziuk [140], Huré, Pham, Bachouch, and Langrené [167]. The main difference between these papers and our work is that as we consider a Principal-Agent model, there is an additional complexity related to the Agent's problem. In Principal-Agent problems, one has to first solve the optimization problem of the Agent for any contract, then inject the solution into the optimization problem of the Principal. We tackle the exchange's problem by including the Agent's controls neural network into the optimization problem with a deep reinforcement learning approach called an actor-critic approach. With this technique, the weights of the market-maker's neural network are a part of the exchange's optimization problem. Therefore, we provide an efficient numerical approach to this kind of problems when one does not have a closed-form solution for the Agent's response.

Below is a summary of the contributions :

³Note that the best response $\mathcal{L}^*(\xi)$ is not necessarily unique given a contract ξ . In that case, in a bilevel optimization we consider that the Principal chooses the best response to maximize his objective function, see for example [159, Equation (1)].

- **Economic contributions.** We contribute to the literature on make-take fees by introducing dynamic incentives to improve the quality of liquidity on both the lit and dark pools of a same exchange. As a by-product, we obtain a way to design optimal market-making strategies in the presence of a dark pool. Our first main finding is that without the intervention of the exchange, the market-maker tends to use the lit pool to create favorable trading conditions for him on the dark pool. This is done by taking advantage of a large imbalance on one side of the lit pool so that market takers on the opposite side of the dark pool will be executed at an unfavourable price. Our second important conclusion is that, in the presence of a contract between the market-maker and the exchange, such phenomenon disappears. Our make-take fees system prevents from such price manipulation and increases the overall liquidity provision in the dark pool. Thus, our methodology to reduce market inefficiencies between the lit and dark pools of the same exchange.
- **Technical contributions.** We show that the control problem of the exchange (and that of the market-maker in the absence of contract) can be characterized by a HJB equation which admits a unique viscosity solution. These stability results ensure that we can perform numerical analysis on the control problems. Moreover, a key difference with El Euch, Mastrolia, Rosenbaum, and Touzi [109] and Chapter 2 of this thesis is the absence of a closed-form solution for the best response of the market-maker to a given contract. Therefore, the HJB equation of the exchange cannot be solved explicitly. To the best of our knowledge, we introduce the first numerical method (based on deep reinforcement learning) to solve Principal-Agent problems when the best response of the Agent cannot be obtained in closed-form. It is a major improvement in the literature, as usually the form of the Agent's problem has to be chosen so that we can derive explicitly its best response.

The chapter is organized as follows. Market dynamics are introduced in Section 1.2. In Section 1.3, we first investigate the problem of a market-maker acting on both lit and dark venues without any incentive policy from the exchange. His goal is to maximize his PnL process while managing his inventory risk. It is a stochastic control problem, where the corresponding HJB equation cannot be solved explicitly. We show existence and uniqueness of a viscosity solution for this equation. In Section 1.4, we analyze the bi-level optimization problem associated with the issue of optimal contracting between the market-maker and the exchange owning both lit and dark pools. Following recent works on make-take fees policies mentioned above, we first prove a representation theorem for the contract proposed to the market-maker. We then establish existence and uniqueness of a viscosity solution for the HJB equation corresponding to the problem of the exchange. In Section 1.5, we introduce a deep reinforcement learning method as a computational tool enabling us to address both exchange and market-maker's problems in practice. We conclude this section with numerical experiments, illustrating various behaviors of the market-maker under different market scenarios.

1.2 The market model

1.2.1 Stochastic basis

The framework considered throughout this chapter is inspired by the article Avellaneda and Stoikov [26] in which the authors investigate the problem of optimal market-making without intervention of an exchange. Let $T > 0$ be a finite horizon time and $\mathcal{V}^l, \mathcal{V}^d \subset \mathbb{N}$ the sets of possible values for volumes in the lit and dark pools, of cardinality $\#\mathcal{V}^l, \#\mathcal{V}^d$. We define $\Omega := \Omega_c \times \Omega_d^{2(\#\mathcal{V}^l + \#\mathcal{V}^d)}$ with Ω_c the set of continuous functions from $[0, T]$ into \mathbb{R} and Ω_d the set of piecewise constant càdlàg functions from $[0, T]$ into \mathbb{N} . Ω is a subspace of the Skorokhod space $\mathcal{D}([0, T], \mathbb{R}^{2(\#\mathcal{V}^l + \#\mathcal{V}^d) + 1})$ of càdlàg functions from $[0, T]$ into $\mathbb{R}^{2(\#\mathcal{V}^l + \#\mathcal{V}^d) + 1}$ and write \mathcal{F} for the trace Borel σ -algebra on Ω , where the topology is the one associated with the usual Skorokhod distance on $\mathcal{D}([0, T], \mathbb{R}^{2(\#\mathcal{V}^l + \#\mathcal{V}^d) + 1})$.

We define $(\mathcal{X}_t)_{t \in [0, T]} := (W_t, N_t^{i,j,k})_{t \in [0, T], i \in \{a, b\}, j \in \{l, d\}, k \in \mathcal{V}^j}$ as the canonical process on Ω , that is for

any $\omega := (w, n^{i,j,k}) \in \Omega$

$$W_t(\omega) := w(t), N_t^{i,j,k}(\omega) = n^{i,j,k}(t), i \in \{a, b\}, j \in \{l, d\} \text{ and } k \in \mathcal{V}^j.$$

For any $i \in \{a, b\}$, $j \in \{l, d\}$ and $k \in \mathcal{V}^j$, $N_t^{i,j,k}$ denotes the total number of trades of size k made between time 0 and time t , where a, b stand for the ask and bid side respectively and l, d for the lit and dark pools respectively. Finally the process W represents the mid-price of the traded asset.

Then we define the probability \mathbb{P}^0 on (Ω, \mathcal{F}) under which W_t and the $N_t^{i,j,k}$ are independent, W_t is a one-dimensional Brownian motion and the $N_t^{i,j,k}$, $i \in \{a, b\}, j \in \{l, d\}, k \in \mathcal{V}^j$ are Poisson processes with intensity $\epsilon > 0$ small enough.⁴ Finally, we endow the space (Ω, \mathcal{F}) with the (\mathbb{P}^0 -completed) canonical filtration $\mathbb{F} := (\mathcal{F}_t)_{t \in [0, T]}$ generated by $(\mathcal{X}_t)_{t \in [0, T]}$.

1.2.2 Traded volumes, market impact and latency

In this section, we formalize the connection between volumes posted by the market-maker and arrival intensity of market orders on the ask and bid sides of both venues. We also take into account market impact phenomenon and latency effect in the dark pool.

1.2.2.1 Admissible controls, inventory process and market takers' arrival flows

Let $2\bar{q} \in \mathbb{N}$ represent a risk limit for the market-maker, which corresponds to the maximum number of cumulated bid and ask orders the market-maker can handle. We define the volume process $(\mathcal{L}_t)_{t \in [0, T]} := (\mathcal{L}_t^l, \mathcal{L}_t^d)_{t \in [0, T]} \in (\mathcal{V}^l)^2 \times (\mathcal{V}^d)^2$, where $\mathcal{L}_t^l = (\ell_t^{a,l}, \ell_t^{b,l})_{t \in [0, T]}$ and $\mathcal{L}_t^d = (\ell_t^{a,d}, \ell_t^{b,d})_{t \in [0, T]}$ with $\ell_t^{i,j}$ corresponding to the volume posted by the market-maker at time t on side $i \in \{a, b\}$ of pool $j \in \{l, d\}$. For simplicity in our framework, the market-maker acts only on the best bid and ask limits, which is a very reasonable assumption for large tick assets. Note that we could consider more complex strategies by adding point processes taking into account the number of filled limit orders at the second limits. The set \mathcal{A} of admissible controls of the market-maker is defined as

$$\mathcal{A} := \{(\mathcal{L}_t)_{t \in [0, T]} : \text{predictable}, i \in \{a, b\}, \ell^{i,l} + \ell^{i,d} \in [0, 2\bar{q}]\}.$$

The market-maker manages his inventory Q_t , defined as the aggregated sum of the volumes filled on both sides of the lit and dark pools, namely

$$Q_t := \sum_{j \in \{l, d\}} \sum_{(k^{a,j}, k^{b,j}) \in (\mathcal{V}^j)^2} k^{b,j} N_t^{b,j,k} - k^{a,j} N_t^{a,j,k}.$$

Remark 1. *Note that for many assets, a large proportion of aggressive orders consume the whole volume posted by the market-maker on the considered side and pool, see Saliba [242]. We have in mind such asset in our modeling. Remark also that we do not take into account the effect of queue priority of limit orders, as in Huang, Rosenbaum, and Saliba [165], Moallemi and Yuan [217] for example. From a practical viewpoint, we restrict ourselves to consider an asset where there is only a few market-makers providing the vast majority of liquidity hence without race for queue position. From a technical viewpoint, modeling the queue position with point processes is an intricate task since one has to rely on stochastic control theory with delay, which is beyond the scope of this chapter.*

We define the function

$$\psi^{i,j}(\mathcal{L}_t^l) := \begin{cases} I^a(\mathcal{L}_t^l), & \text{if } (i, j) \in \{(a, l), (b, d)\}, \\ I^b(\mathcal{L}_t^l), & \text{if } (i, j) \in \{(b, l), (a, d)\}, \end{cases}$$

⁴In other words, \mathbb{P}^0 is the product measure of the Wiener measure on Ω_c and the unique measure on $\Omega_d^{2(\#\mathcal{V}^l + \#\mathcal{V}^d)}$ so that the canonical process corresponds to a multidimensional homogeneous Poisson process with arbitrary small intensity, representing a situation where no liquidity is available.

where $I^a(\mathcal{L}_t^l) := \frac{\ell_t^{a,l}}{\ell_t^{a,l} + \ell_t^{b,l}}$, $I^b(\mathcal{L}_t^l) := \frac{\ell_t^{b,l}}{\ell_t^{a,l} + \ell_t^{b,l}}$ represent the imbalances on the ask and bid sides of the lit pool respectively. To model the behavior of market takers, we define the intensities of the processes $N^{i,j,k}$ as

$$\lambda_t^{\mathcal{L},i,j,k} := \lambda^{i,j}(\mathcal{L}_t^l) \mathbf{1}_{\{\phi(i)Q_{t-} > -\bar{q}, \ell_t^{i,j}=k\}}, \quad \phi(i) := \begin{cases} 1, & \text{if } i = a, \\ -1, & \text{if } i = b, \end{cases}$$

where

$$\lambda^{i,j}(\mathcal{L}_t^l) := A^j \exp\left(-\frac{\theta^j}{\sigma} \psi^{i,j}(\mathcal{L}_t^l)\right) \mathbf{1}_{\{\mathcal{L}_t^l \neq (0,0)\}} + \epsilon \mathbf{1}_{\{\mathcal{L}_t^l = (0,0)\}},$$

where $\sigma > 0$ is the volatility of the asset's mid-price. A high imbalance on the ask side decreases the probability that an ask limit order is filled in the lit pool and conversely for the bid side. Moreover, when the imbalance on the ask (resp. bid) side of the lit pool is high, if a market taker wants to buy, it is worth trying it in the dark pool, because the high imbalance indicates that the ask price in the lit may not be competitive. The coefficients $\theta^l, \theta^d > 0$ represent the influence of the imbalance on the intensity of orders' arrivals and $A^l, A^d > 0$ are average order flow intensity parameters.

For $\mathcal{L} \in \mathcal{A}$, we introduce a new probability measure $\mathbb{P}^{\mathcal{L}}$ under which W remains a one-dimensional Brownian motion and for $i \in \{a, b\}, j \in \{l, d\}, k \in \mathcal{V}^j$ the

$$N_t^{\mathcal{L},i,j,k} := N_t^{i,j,k} - \int_0^t \lambda_u^{\mathcal{L},i,j,k} du,$$

are martingales. This probability measure is defined by the corresponding Doléans-Dade exponential

$$L_t^{\mathcal{L}} := \exp\left(\sum_{(i,j) \in \{a,b\} \times \{l,d\}} \sum_{k \in \mathcal{V}^j} \int_0^t \mathbf{1}_{\{\phi(i)Q_{u-} > -\bar{q}, \ell_u^{i,j}=k\}} \left(\log\left(\frac{\lambda^{i,j}(\mathcal{L}_u^l)}{\epsilon}\right) dN_u^{i,j,k} - (\lambda^{i,j}(\mathcal{L}_u^l) - \epsilon) du\right)\right),$$

which is a true martingale by the uniform boundedness of the $\ell^{i,j}$.⁵ We can therefore set the Girsanov change of measure with $\frac{d\mathbb{P}^{\mathcal{L}}}{d\mathbb{P}^\sigma} \Big|_{\mathcal{F}_t} = L_t^{\mathcal{L}}$ for all $t \in [0, T]$. In particular, all the probability measures $\mathbb{P}^{\mathcal{L}}$ indexed by \mathcal{L} are equivalent. We write $\mathbb{E}_t^{\mathcal{L}}$ for the conditional expectation with respect to \mathcal{F}_t under the probability measure $\mathbb{P}^{\mathcal{L}}$. We also define for $i \in \{a, b\}, j \in \{l, d\}$ the processes

$$N_t^{i,j} := \sum_{k \in \mathcal{V}^j} N_t^{i,j,k},$$

of intensities $\lambda^{i,j}(\mathcal{L}_t^l) \mathbf{1}_{\{\phi(i)Q_{t-} > -\bar{q}\}}$. These processes correspond to the total number of transactions executed on the bid or ask side of the lit or dark pools.

1.2.2.2 Efficient price and market impact

We define the efficient price of the underlying asset, observable by all market participants (in the sense that they can infer it) as

$$\tilde{S}_t := \tilde{S}_0 + \sigma W_t,$$

where $\tilde{S}_0 > 0$ is the initial price of the underlying asset and $\sigma > 0$ its volatility. On very liquid markets, the efficient price can be simply thought of as the mid-price. For more subtle views on efficient price estimation, see for example Delattre, Robert, and Rosenbaum [93], Guéant, Lehalle, and Fernandez-Tapia [142], Stoikov [257].

⁵The associated Novikov criterion is given in Sokol [253].

When a limit order on the bid side is filled, the price decreases on average and conversely for the ask side (this is the so-called market impact, see for example Bouchaud [54], Toth, Eisler, and Bouchaud [265]). Thus, we define the mid-price of the asset at time $t \in [0, T]$ by

$$S_t := \tilde{S}_t + \sum_{j \in \{l, d\}} \int_0^t \Gamma^j \ell_u^{a,j} dN_u^{a,j} - \Gamma^j \ell_u^{b,j} dN_u^{b,j}, \quad (1.1)$$

where $\Gamma^l, \Gamma^d > 0$ are fixed constants representing the magnitude of market impact in the lit and dark pools.

Remark 2. *The market impact parameters Γ^l, Γ^d are taken small enough with respect to the tick size to discard obvious arbitrage opportunities. Moreover, as the market impact in the dark pool is usually smaller or delayed compared to the lit pool, we will take $\Gamma^l \geq \Gamma^d$.*

1.2.2.3 Latency in the dark pool

We assume that in the lit pool, the best bid and best ask prices $P^{b,l}$ and $P^{a,l}$ satisfy

$$P_t^{b,l} := S_t - \frac{\mathcal{T}}{2}, \quad P_t^{a,l} := S_t + \frac{\mathcal{T}}{2}, \quad t \in [0, T],$$

where $\frac{\mathcal{T}}{2} > 0$ is the half tick of the market. In this setting, in the lit pool, the market-maker only needs to control the volumes he posts.

In the dark pool, orders may be executed at the mid-price, which is a priori beneficial for market takers. In practice, due to latency effect in the dark pool, the mid-price can change by one half tick (or more) before the transaction is made. Therefore the order may be executed at a less advantageous price for the market taker (and sometimes at an even more advantageous one but we neglect this case for the sake of simplicity). Let us introduce the corresponding prices with and without latency:

$$\begin{cases} P_t^{b,d,\text{lat}} := S_t - \frac{\mathcal{T}}{2}, & P_t^{a,d,\text{lat}} := S_t + \frac{\mathcal{T}}{2}, \\ P_t^{b,d,\text{non-lat}} := S_t, & P_t^{a,d,\text{non-lat}} := S_t. \end{cases}$$

Recall that in most dark pools, market takers are supposed to be executed at the mid-price of the lit pool. However, the higher the imbalance on the ask (resp. bid) side of the lit pool, the higher the probability that the mid-price will move down (resp. up) quickly. To model the latency effect, we introduce Bernoulli random variables $\nu_t^a \sim \text{Ber}(I^a(\mathcal{L}_t^l))$, $\nu_t^b \sim \text{Ber}(I^b(\mathcal{L}_t^l))$ which are associated to each incoming market order in the dark pool.⁶ If $\nu_t = 1$, there is no latency, and conversely for $\nu_t = 0$. So we define

$$\begin{aligned} N_t^{a,d,\text{lat}} &:= \int_0^t (1 - \nu_u^a) dN_u^{a,d}, & N_t^{a,d,\text{non-lat}} &:= \int_0^t \nu_u^a dN_u^{a,d}, \\ N_t^{b,d,\text{lat}} &:= \int_0^t (1 - \nu_u^b) dN_u^{b,d}, & N_t^{b,d,\text{non-lat}} &:= \int_0^t \nu_u^b dN_u^{b,d}. \end{aligned}$$

Note that for any $t \in [0, T]$, $N_t^{i,d,\text{lat}} + N_t^{i,d,\text{non-lat}} = N_t^{i,d}$ for $i \in \{a, b\}$. To our knowledge, our approach is the first one considering market-making in the dark pool taking into account latency effect.

Remark 3. *The canonical variables are $S, N^{a,l}, N^{b,l}, N^{a,d}, N^{b,d}$. Thus the incentives mechanism we will design will be allowed to depend on the trajectories of these quantities only. This is actually a reasonable assumption: the efficient price is a quantity any market participant can observe, whether the chosen proxy for it is the midprice, the last traded price or some volume weighted price. The point processes encode the arrivals of market orders on the lit and dark pools and therefore actual transactions, which are clearly recorded on any exchange and accessible to most participants. So the contracts will be designed on standard and easily obtained financial variables.*

⁶We take the convention $I^a(0, 0) = I^b(0, 0) = 0$.

1.3 Market-making without the intervention of the exchange

We address the problem of a market-maker acting in the lit and dark pools, without intervention of the exchange. The profit and loss (PnL for short) of the market-maker is defined as the sum of the cash earned from his executed orders and the value of his inventory. Thus it is expressed as

$$PL_t^{\mathcal{L}} := \mathcal{W}_t^{\mathcal{L}} + Q_t S_t,$$

where, at time $t \in [0, T]$,

$$\begin{aligned} \mathcal{W}_t^{\mathcal{L}} := & \int_0^T \left(S_t + \frac{\mathcal{T}}{2} \right) \ell_t^{a,l} dN_t^{a,l} - \int_0^T \left(S_t - \frac{\mathcal{T}}{2} \right) \ell_t^{b,l} dN_t^{b,l} + \int_0^T \left(S_t + \frac{\mathcal{T}}{2} \right) \ell_t^{a,d} dN_t^{a,d,\text{lat}} \\ & + \int_0^T S_t \ell_t^{a,d} dN_t^{a,d,\text{non-lat}} - \int_0^T \left(S_t - \frac{\mathcal{T}}{2} \right) \ell_t^{b,d} dN_t^{b,d,\text{lat}} - \int_0^T S_t \ell_t^{b,d} dN_t^{b,d,\text{non-lat}}, \end{aligned}$$

represents his cash process and $Q_t S_t$ is the mark-to-market value of his inventory.⁷ Note that market-making activity in the dark pool without latency does not generate PnL through spread collection.

Remark 4. *Note that, for sake of simplicity, we assume that the spread is constant and equal to one tick. The case of a stochastic spread or that of a spread controlled by the market-maker could also be treated by our approach, at the cost of heavier mathematical derivations and more involved numerical methods.*

We consider a risk averse market-maker with exponential utility function and risk aversion parameter $\gamma > 0$. We define his optimization problem as

$$V_0^{\text{MM}} = \sup_{\mathcal{L} \in \mathcal{A}} J_0^{\text{MM}}(\mathcal{L}), \quad (1.2)$$

with for all $t \in [0, T]$,

$$J_t^{\text{MM}}(\mathcal{L}) = \mathbb{E}_t^{\mathcal{L}} \left[- \exp \left(- \gamma (PL_T^{\mathcal{L}} - PL_t^{\mathcal{L}}) \right) \right].$$

Inspired by El Euch, Mastrolia, Rosenbaum, and Touzi [109], we prove a dynamic programming principle for the control problem (1.2), see Section 1.A.1, from which we derive the corresponding HJB equation. We define $\mathcal{O} = [0, 2\bar{q}]^4$. Similarly to Guéant, Lehalle, and Fernandez-Tapia [142], we use a change of variable (see Equation (1.17) for the form of the ansatz) to reduce the initial problem to the following HJB equation:

$$\begin{aligned} 0 = & \partial_t v(t, q) + v(t, q) \frac{1}{2} \sigma^2 \gamma^2 q^2 \\ & + \sup_{\mathcal{L} \in \mathcal{O}} \left\{ \sum_{(k^l, k^d) \in \mathcal{V}^l \times \mathcal{V}^d} \left(\sum_{i \in \{a, b\}} \lambda_t^{\mathcal{L}, i, l, k^l} \left(\exp \left(- \gamma \ell^{i, l} \left(\frac{\mathcal{T}}{2} + \Gamma^l(\phi(i)q - \ell^{i, l}) \right) \right) v(t, q - \phi(i)k^l) - v(t, q) \right) \right. \right. \\ & \left. \left. + \sum_{i \in \{a, b\}} \sum_{\kappa \in K} \lambda_t^{\mathcal{L}, i, d, k^d} \phi^d(i, \kappa) \left(\exp \left(- \gamma \ell^{i, d} \left(\frac{\mathcal{T}}{2} \phi^{\text{lat}}(\kappa) + \Gamma^d(\phi(i)q - \ell^{i, d}) \right) \right) v(t, q - \phi(i)k^d) - v(t, q) \right) \right) \right\}, \end{aligned} \quad (1.3)$$

with $K := \{\text{lat}, \text{non-lat}\}$,

$$\phi^{\text{lat}}(\kappa) := \begin{cases} 1, & \text{if } \kappa = \text{lat}, \\ 0, & \text{if } \kappa = \text{non-lat}, \end{cases} \quad \phi^d(i, \kappa) := \begin{cases} I^b(\mathcal{L}^l), & \text{if } (i, \kappa) \in \{(a, \text{lat}), (b, \text{non-lat})\}, \\ I^a(\mathcal{L}^l), & \text{if } (i, \kappa) \in \{(a, \text{non-lat}), (b, \text{lat})\}, \end{cases}$$

and terminal condition $v(T, \cdot) = -1$. We have the following theorem.

⁷Note that for all $t \in [0, T]$, $\int_0^t \ell_u^{i,j} dN_u^{i,j} = \sum_{k \in \mathcal{V}^j} k N_t^{i,j,k}$.

Theorem 1.1. *There exists a unique viscosity solution with polynomial growth to the HJB equation (1.3). It satisfies*

$$V_0^{\text{MM}} = v(0, Q_0).$$

The supremum in (1.3) characterizes the optimal controls $\mathcal{L}^* \in \mathcal{A}$.

The proof follows the same arguments as Theorem 1.3 in Section 1.A.3.

We see that the supremum over \mathcal{L} is not separable with respect to each control process as in El Euch, Mastrolia, Rosenbaum, and Touzi [109], Guéant, Lehalle, and Fernandez-Tapia [142]. To the best of our knowledge there is no explicit expression for the optimal controls of the market-maker. Nevertheless, as shown in Section 1.5.2, we can solve PDE (1.3) numerically. More precisely, we make use of deep reinforcement learning techniques to approximate the optimal volumes posted of the market-maker.

1.4 Market-making with the intervention of the exchange

Let us now consider the case where a make-take fees system is in place and influences the amount of liquidity provided by the market-maker on both lit and dark venues.

1.4.1 Modified PnL of the market-maker

Following the Principal-Agent approach of El Euch, Mastrolia, Rosenbaum, and Touzi [109], we now assume that the exchange gives to the market-maker a compensation ξ defined as an \mathcal{F}_T -measurable random variable, which is added to his PnL process at terminal time T . This contract, designed by the exchange, aims at creating incentives so that the market-maker attracts more transactions.

Therefore, the total payoff of the market-maker at time T is now given by $\mathcal{W}_T^\mathcal{L} + Q_T S_T + \xi$. The problem of the market-maker then becomes

$$V_0^{\text{MM}}(\xi) := \sup_{\mathcal{L} \in \mathcal{A}} J_0^{\text{MM}}(\mathcal{L}, \xi), \quad (1.4)$$

with

$$J_t^{\text{MM}}(\mathcal{L}, \xi) := \mathbb{E}_t^\mathcal{L} \left[- \exp \left(- \gamma (PL_T^\mathcal{L} - PL_t^\mathcal{L} + \xi) \right) \right].$$

Remark 5. *We consider a risk-averse market-maker with CARA utility function. Note that the case of a risk-neutral market-maker is not really meaningful here as it would neglect the inventory management of the market-maker, which is the key component of the problem.*

To ensure that this functional is non-degenerate, we impose the following technical condition on ξ (see the next section for the definition of an admissible contract):

$$\sup_{\mathcal{L} \in \mathcal{A}} \mathbb{E}^\mathcal{L} \left[\exp(-\gamma' \xi) \right] < +\infty, \text{ for some } \gamma' > \gamma, \quad (1.5)$$

so that the optimization problem of the market-maker is well-posed.

For a fixed compensation ξ , the optimal response \mathcal{L}^* associated with the market-maker's problem (1.4) is defined as

$$J_0^{\text{MM}}(\mathcal{L}^*, \xi) = V_0^{\text{MM}}(\xi), \text{ for } \mathcal{L}^* \in \mathcal{A}. \quad (\text{OC})$$

We now consider the problem of the exchange wishing to attract liquidity on its platforms.

1.4.2 Objective function of the exchange

We assume that the exchange receives fixed fees $c^l, c^d > 0$ for each market order occurring in the lit and dark pools respectively. As in El Euch, Mastrolia, Rosenbaum, and Touzi [109], since we are working on a short time interval, we take c^l, c^d independent of the price of the asset.

The goal of the exchange is essentially to maximize the total number of market orders sent during the period of interest. As the arrival intensities of market orders are controlled by the market-maker through \mathcal{L} , the contract ξ should aim at increasing these intensities. Thus, the exchange subsidizes the Agent at time T with the compensation ξ so that its PnL is given by

$$\sum_{(i,j) \in \{a,b\} \times \{l,d\}} c^j \int_0^T \ell_t^{i,j} dN_t^{i,j} - \xi.$$

We now need to specify the set of admissible contracts potentially offered by the exchange. We assume that the exchange has exponential utility function with risk aversion parameter $\eta > 0$. The natural well-posedness condition for the problem of the exchange is, for any \mathcal{L}^* satisfying condition (OC),

$$\mathbb{E}^{\mathcal{L}^*} \left[\exp(\eta' \xi) \right] < +\infty, \text{ for some } \eta' > \eta. \quad (1.6)$$

Since the $N^{i,j}$ are point processes with bounded intensities, this condition, together with Hölder inequality, ensure that the problem of the exchange is well-defined. We also assume that the market-maker only accepts contracts ξ such that $V_0^{\text{MM}}(\xi)$ is above some threshold value $R < 0$, that is ξ must satisfy

$$V_0^{\text{MM}}(\xi) \geq R. \quad (\text{R})$$

This threshold, called reservation utility of the Agent, is the critical utility value under which the market-maker has no interest in the contract. This quantity has to be taken into account carefully by the exchange when proposing a contract to the market-maker.

Remark 6. *We consider that the exchange does not have a reservation utility. For example, we can assume that R is so that $\xi = 0$ is an admissible contract, which would lead to positive PnL for the exchange.*

We can therefore define the space of admissible contracts \mathcal{C} by

$$\mathcal{C} := \left\{ \xi : \mathcal{F}_T\text{-measurable verifying (1.5), (1.6) and (R)} \right\}.$$

Thus the contracting problem the exchange has to solve is

$$V_0^E := \sup_{\xi \in \mathcal{C}} \mathbb{E}^{\mathcal{L}^*} \left[- \exp \left(- \eta \left(\sum_{(i,j) \in \{a,b\} \times \{l,d\}} c^j \int_0^T \ell_t^{i,j} dN_t^{i,j} - \xi \right) \right) \right]. \quad (1.7)$$

In the next section, we characterize the form of an admissible contract $\xi \in \mathcal{C}$.⁸

1.4.3 Design of an optimal make-take fees policy

1.4.3.1 A class of contracts built on transactions

Inspired by El Euch, Mastrolia, Rosenbaum, and Touzi [109], we prove in this section that without loss of generality, we can consider a specific form of contracts. From an economical viewpoint, we expect that the contract is indexed on the canonical variables \tilde{S} , $N^{a,l}$, $N^{b,l}$, $N^{a,d}$, $N^{b,d}$, see Remark 3. The contract consists of the following elements:

⁸Note that for fixed ξ , the control \mathcal{L}^* is not necessarily unique. However, numerical results seem to indicate its uniqueness. Otherwise we could also consider a supremum over all \mathcal{L}^* satisfying (OC), as it is usually done in Principal-Agent theory (see for instance [89, Section 2.4]).

- A constant Y_0 calibrated by the exchange to ensure that the reservation utility constraint (R) of the market-maker is satisfied.⁹
- A term $Z^{\tilde{S}}$ denoting the compensation given to the market-maker with respect to the volatility risk induced by the efficient price \tilde{S} .
- Compensation $Z^{i,j,k}$ with respect to the number of trades of size k on the ask or bid side of the lit or the dark pool.
- A continuous coupon $\frac{1}{2}\gamma\sigma^2(Z^{\tilde{S}} + Q)^2 - H(Z, Q)$ composed by a compensation to balance the risk aversion of the market-maker as a risk sharing term minus his Hamiltonian function. This penalization term is very classical in contract theory and corresponds to a gain resulting from the optimization of the market-maker. We consider that the exchange receives all the benefit of this optimization and reallocate it in the contract proposed to the market-maker.

Remark 7. *In our setting, the volumes of limit orders do not belong to the canonical process and so the Principal does not contract on the volumes displayed by the market-maker. It is very reasonable as in practice, a large part of volumes sent by market-makers are not executed or rapidly canceled. Therefore it is clearly preferable to build contracts based on actual transactions. Moreover, note that \tilde{S} , and not the mid-price S , appears in the contract (1.1). This is not an issue since S can be decomposed into elements of the canonical process. Note that we could also write the contracts in terms of the bid and ask prices $P_t^{b,l}, P_t^{a,l}$ on the lit pool instead of \tilde{S} in view of (1.1).*

Technically speaking, a contract ξ is defined by some chosen constant by the Principal $Y_0 \in \mathbb{R}$ and a predictable process $Z = (Z^{\tilde{S}}, Z^{i,j,k})_{i \in \{a,b\}, j \in \{l,d\}, k \in \mathcal{V}^j}$, which can be written as

$$\begin{aligned} \xi = Y_T^{Y_0, Z} := Y_0 + \int_0^T & \left(\sum_{(i,j) \in \{a,b\} \times \{l,d\}} \sum_{k \in \mathcal{V}^j} Z_u^{i,j,k} dN_u^{i,j,k} \right) + Z_u^{\tilde{S}} d\tilde{S}_u \\ & + \left(\frac{1}{2} \gamma \sigma^2 (Z_u^{\tilde{S}} + Q_u)^2 - H(Z_u, Q_u) \right) du, \end{aligned} \quad (1.8)$$

where

$$H(z, q) := \sup_{\mathcal{L} \in \mathcal{A}} h(\mathcal{L}, z, q), \quad (1.9)$$

and $h : \mathbb{R}_+^4 \times \mathbb{R}^{2(\#\mathcal{V}^l + \#\mathcal{V}^d)} \times \mathbb{Z} \rightarrow \mathbb{R}$ is the Hamiltonian of the Agent's problem.¹⁰ To ensure admissibility of the contract, the process $(Z_t)_{t \in [0, T]}$ has to satisfy the following technical conditions:

$$\sup_{\mathcal{L} \in \mathcal{A}} \mathbb{E}^{\mathcal{L}} \left[\sup_{t \in [0, T]} \exp \left(-\gamma' Y_t^{0, Z} \right) \right] < +\infty, \text{ for some } \gamma' > \gamma, \quad (1.10)$$

and

$$\int_0^T |Z_t^{\tilde{S}}|^2 + |H(Z_t, Q_t)| dt < +\infty. \quad (1.11)$$

Given this integrability condition, the process $(Y_t^{0, Z})_{t \in [0, T]}$ is well-defined. Formally stated, the definition of the space Ξ of contracts of the form (1.8) is

$$\Xi = \left\{ Y_T^{Y_0, Z} : \mathcal{F}_T\text{-measurable}, Y_0 \in \mathbb{R}, Z \in \mathcal{Z}, \text{ s.t. (R) holds} \right\},$$

where

$$\mathcal{Z} := \left\{ (Z_t^{\tilde{S}}, Z_t^{i,j,k}) : t \in [0, T], i \in \{a, b\}, j \in \{l, d\}, k \in \mathcal{V}^j, (1.6), (1.10), (1.11) \text{ are satisfied} \right\}. \quad (1.12)$$

⁹From Theorem 1.2, $\hat{Y}_0 = -\gamma^{-1} \log(-R)$ ensures that the reservation utility constraint of the market-maker is satisfied.

¹⁰Its form is defined in (1.13). This Hamiltonian term appears naturally when applying the dynamic programming principle for the market-maker's problem.

1.4.3.2 Solving the market-maker's problem

For $(\mathcal{L}, z, q) \in (\mathcal{V}^l)^2 \times (\mathcal{V}^d)^2 \times \mathbb{R}^{2(\#\mathcal{V}^l + \#\mathcal{V}^d)} \times \mathbb{Z}$ we define the Hamiltonian of the market-maker, which appears in the contract $Y_T^{Y_0, Z}$ via the continuous coupon $H(Z, Q)$, by

$$h(\mathcal{L}, z, q) := \sum_{(k^l, k^d) \in \mathcal{V}^l \times \mathcal{V}^d} \left(\sum_{i \in \{a, b\}} \gamma^{-1} \left(\left(1 - \exp \left(-\gamma \left(z^{i, l, k^l} + \ell^{i, l} \left(\frac{\mathcal{T}}{2} + \phi(i) \Gamma^l q \right) - \Gamma^l (\ell^{i, l})^2 \right) \right) \right) \lambda_t^{\mathcal{L}, i, l, k^l} \right. \\ \left. + \sum_{\kappa \in K} \left(1 - \exp \left(-\gamma \left(z^{i, d, k^d} + \ell^{i, d} \left(\frac{\mathcal{T}}{2} \phi^{\text{lat}}(\kappa) + \phi(i) \Gamma^d q \right) - \Gamma^d (\ell^{i, d})^2 \right) \right) \lambda_t^{\mathcal{L}, i, d, k^d} \phi^d(i, \kappa) \right) \right). \quad (1.13)$$

The next theorem states that the two sets \mathcal{C} and Ξ are, in fact, equal. Moreover, the contract representation (1.8) enables us to provide a solution to the market-maker's problem (1.4). The proof is given in Section 1.A.2.

Theorem 1.2. *Any admissible contract can be written under the form (1.8), that is $\mathcal{C} = \Xi$. Moreover, for any $Y_T^{Y_0, Z} \in \Xi$ we have*

$$V_0^{\text{MM}}(Y_T^{Y_0, Z}) = -\exp(-\gamma Y_0),$$

and Condition (OC) with $\xi = Y_T^{Y_0, Z}$ is equivalent to the fact that \mathcal{L} satisfies $h(\mathcal{L}_t, Z_t, Q_t) = H(Z_t, Q_t)$ for any $t \in [0, T]$.

This theorem provides a tractable form of contracts for the design of a suitable make-take fees policy. Given the knowledge of the market-maker's response to a given contract, we reformulate the problem of the exchange and prove the existence and uniqueness of the associated value function.

1.4.4 Problem of the exchange

1.4.4.1 Reformulation of the problem

Following Theorem 1.2, the contracting problem (1.7) is reduced to

$$V_0^E := \sup_{\substack{(Y_0, Z, \mathcal{L}) \in \mathbb{R} \times \mathcal{Z} \times \mathcal{A}, \\ \mathcal{L} \text{ satisfies (1.9)}}} \mathbb{E}^{\mathcal{L}} \left[-\exp \left(-\eta \left(\sum_{(i, j) \in \{a, b\} \times \{l, d\}} c^j \int_0^T \ell_t^{i, j} dN_t^{i, j} - Y_T^{Y_0, Z} \right) \right) \right]. \quad (1.14)$$

For a given contract $Y^{Y_0, Z}$, due to the form of (1.13), the market-maker's optimal response does not depend on Y_0 . With the exchange's objective function being decreasing in Y_0 , the maximization with respect to Y_0 is achieved at the level $\hat{Y}_0 = -\gamma^{-1} \log(-R)$. Therefore Problem (1.14) can be reduced to

$$v_0^E := \sup_{\substack{(Z, \mathcal{L}) \in \mathcal{Z} \times \mathcal{A}, \\ \mathcal{L} \text{ satisfies (1.9)}}} \mathcal{J}(Z, \mathcal{L}), \quad (1.15)$$

where

$$\mathcal{J}(Z, \mathcal{L}) = \mathbb{E}^{\mathcal{L}} \left[-\exp \left(-\eta \left(\sum_{(i, j) \in \{a, b\} \times \{l, d\}} c^j \int_0^T \ell_t^{i, j} dN_t^{i, j} - Y_T^{0, Z} \right) \right) \right].$$

1.4.4.2 A bi-level optimization problem

We define $\mathcal{D} := [0, T] \times \mathbb{R} \times \mathbb{N}^{2(\#\mathcal{V}^l + \#\mathcal{V}^d)} \times \mathbb{N}^{2(\#\mathcal{V}^l + \#\mathcal{V}^d)} \times \mathbb{R}$ and for any vector p , $i \in \{1, \dots, \#p\}$, $p^{-i} = (p^1, \dots, p^{i-1}, p^{i+1}, \dots, p^{\#p}) \in \mathbb{N}^{\#p-1}$. By Equation (1.7) and the corresponding footnote, there might be more than one optimal response \mathcal{L}^* of the market-maker. We show here how to solve the

Principal's problem for a specific optimal response \mathcal{L}^* .¹¹ Using a dynamic programming principle similar to the one in Lemma 1.1, we write the value function of the exchange's problem, $v^E : \mathfrak{D} \rightarrow \mathbb{R}$, as

$$v^E(t, \tilde{S}_t, \bar{N}_t, N_t, Y_t) := \sup_{Z \in \tilde{\mathcal{Z}}} \mathbb{E}_t^{\mathcal{L}^*} \left[- \exp \left(- \eta \left(\sum_{(i,j) \in \{a,b\} \times \{l,d\}} \sum_{k \in \mathcal{V}^j} c^j (\bar{N}_T^{i,j,k} - \bar{N}_t^{i,j,k}) - Y_T^{0,Z} \right) \right) \right],$$

with

$$(\bar{N}_t, N_t) := (kN_t^{i,j,k}, N_t^{i,j,k})_{i \in \{a,b\}, j \in \{l,d\}, k \in \mathcal{V}^j},$$

and $\mathcal{L}^* = (\ell^{*b,l}(z, q), \ell^{*a,l}(z, q), \ell^{*b,d}(z, q), \ell^{*a,d}(z, q))$ the optimal response of the market-maker, in the sense of (1.9), displayed at time t for a given inventory q and given incentives z of the exchange.

Recall that $Q_t = \sum_{j \in \{l,d\}} \sum_{k \in \mathcal{V}^j} (\bar{N}_t^{b,j,k} - \bar{N}_t^{a,j,k})$. Usual arguments enables us to show that v^E is a viscosity solution of the HJB equation defined on \mathfrak{D} by

$$\begin{aligned} 0 = & \partial_t v^E + \frac{1}{2} \sigma^2 \partial_{\tilde{S}\tilde{S}} v^E + \sup_{z^{\tilde{S}} \in \mathbb{R}} \frac{\gamma \sigma^2}{2} (z^{\tilde{S}} + q)^2 \partial_y v^E + \frac{\sigma^2}{2} (z^{\tilde{S}})^2 \partial_{yy} v^E + \sigma^2 z^{\tilde{S}} \partial_{\tilde{S}y} v^E \\ & + \sup_{z \in \mathbb{R}^{2(\#\mathcal{V}^l + \#\mathcal{V}^d)}} \sum_{(i,j) \in \{a,b\} \times \{l,d\}} \sum_{k \in \mathcal{V}^j} \lambda_t^{\mathcal{L}^*, i,j,k} \left(\Delta_{i,j,k}(z) v^E - \partial_y v^E \mathcal{E}(z^{i,j,k}, \ell^{*i,j}(z, q)) \right), \end{aligned} \quad (1.16)$$

where, for $i \in \{a, b\}, j \in \{l, d\}$,

$$\Delta_{i,j,k}(z) v^E(t, \tilde{s}, \bar{n}, n, y) := v^E(t, \tilde{s}, \bar{n}^{i,j,k} + k, \bar{n}^{-(i,j,k)}, n^{i,j,k} + 1, n^{-(i,j,k)}, y + z^{i,j,k}) - v^E(t, \tilde{s}, \bar{n}, n, y),$$

$$\mathcal{E}(z^{i,l,k}, \ell^{*i,l}(z, q)) := \frac{1}{\gamma} \left(1 - \exp \left(- \gamma \left(z^{i,l,k} + \ell^{*i,l}(z, q) \left(\frac{\mathcal{T}}{2} + \phi(i) \Gamma^l q \right) - \Gamma^l (\ell^{*i,l}(z, q))^2 \right) \right) \right),$$

$$\mathcal{E}(z^{i,d,k}, \ell^{*i,d}(z, q)) := \frac{1}{\gamma} \sum_{\kappa \in K} \left(1 - \exp \left(- \gamma \left(z^{i,d,k} + \ell^{*i,d}(z, q) \left(\frac{\mathcal{T}}{2} \phi^{\text{lat}}(\kappa) + \phi(i) \Gamma^d q \right) - \Gamma^d (\ell^{*i,d}(z, q))^2 \right) \right) \right) \phi^d(i, \kappa),$$

and terminal condition

$$v^E(T, \tilde{s}, \bar{n}, n, y) = - \exp \left(- \eta \left(\sum_{(i,j) \in \{a,b\} \times \{l,d\}} \sum_{k \in \mathcal{V}^j} c^j \bar{n}^{i,j,k} - y \right) \right).$$

Remark that the best response of the market-maker, for which we do not have explicit expression, appears in the value function of the exchange. Inspired by El Euch, Mastrolia, Rosenbaum, and Touzi [109], Guéant, Lehalle, and Fernandez-Tapia [142], we use the following ansatz for Equation (1.16):

$$v^E(t, s, \bar{n}, n, y) = v(t, q) \exp \left(- \eta \left(\sum_{(i,j) \in \{a,b\} \times \{l,d\}} \sum_{k \in \mathcal{V}^j} c^j \bar{n}^{i,j,k} - y \right) \right), \quad (1.17)$$

where v is a solution of the following HJB equation

$$\begin{cases} 0 = \partial_t v(t, q) + \mathcal{H}(q, \mathcal{L}^*, v(t, \cdot)), & q \in \{-\bar{q}, \dots, \bar{q}\}, t \in [0, T], \\ v(T, q) = -1, \end{cases} \quad (1.18)$$

with

$$\mathcal{H}(q, \mathcal{L}_t^*, v(t, \cdot)) := \sup_{z \in \mathbb{R}^{2(\#\mathcal{V}^l + \#\mathcal{V}^d) + 1}} \mathcal{U}(z, q, \mathcal{L}^*(z, q), v(t, \cdot)), \quad (1.19)$$

¹¹If there are several optimal responses \mathcal{L}^* , the exchange should solve the HJB equation (1.16) for every \mathcal{L}^* and, according to Principal-Agent theory, choose the optimal response that maximizes its own utility.

and

$$\begin{aligned} \mathcal{U}(z, q, \mathcal{L}^*(z, q), v(t, \cdot)) &:= v(t, q) \left(\frac{\eta}{2} \sigma^2 \gamma (z^{\tilde{S}} + q)^2 + \frac{\eta^2 \sigma^2}{2} (z^{\tilde{S}})^2 \right) \\ &+ \sum_{(i,j) \in \{a,b\} \times \{l,d\}} \sum_{k \in \mathcal{V}^j} \lambda_t^{\mathcal{L}^*, i, j, k} \left(\exp(\eta(z^{i,j,k} - kc^j)) v(t, q - \phi(i)k) - v(t, q) (1 + \eta \mathcal{E}(z^{i,j,k}, \ell^{*i,j}(z, q))) \right). \end{aligned}$$

This ansatz leads to dimensionality reduction from five to two parameters. Using [53, Corollary 1.4.2], there exists a unique continuous viscosity solution associated to (1.18).

Remark 8. Note that the supremum over $z^{\tilde{S}}$ is explicit and given by $z^{\tilde{S}} = -\frac{\gamma}{\gamma+\eta}q$ as in El Euch, Mastrolia, Rosenbaum, and Touzi [109].

1.4.4.3 Solving the exchange's problem

We use the ansatz for v and reduce the bi-level optimization problem (1.15) to the following system:

$$\begin{cases} 0 = \partial_t v(t, q) + \mathcal{H}(q, \mathcal{L}_t^*, v(t, \cdot)), & \text{with final condition } v(T, q) = -1, \\ h(\mathcal{L}^*, z, q) = H(z, q), & q \in \{-\bar{q}, \dots, \bar{q}\}. \end{cases} \quad (1.20)$$

We have the following theorem.

Theorem 1.3. *There exists a unique continuous viscosity solution to HJB equation (1.18). It satisfies*

$$v_0^E = v(0, Q_0) = v^E(0, \tilde{S}_0, \bar{N}_0, N_0, Y_0).$$

Moreover, the optimal incentives of the Principal Z^* are solutions of the supremum in (1.18).

The proof can be found in Section 1.A.3.

Although the convergence of the deep reinforcement learning technique used in Section 1.5 cannot be fully guaranteed, Theorem 1.3 allows us to tackle the bi-level problem (1.20) with numerical methods to obtain the optimizers

$$\left(Z^*(t, Q_{t-}), \mathcal{L}^*(Z^*(t, Q_{t-}), Q_{t-}) \right)_{t \in [0, T]}, \quad (1.21)$$

of the bi-level problem (1.20). Moreover, the optimal contract is given by

$$\xi^* = \hat{Y}_0 + \int_0^T \left(\sum_{(i,j) \in \{a,b\} \times \{l,d\}} \sum_{k \in \mathcal{V}^j} Z_u^{*i,j,k} dN_u^{i,j,k} \right) + Z_u^{*\tilde{S}} d\tilde{S}_u + \left(\frac{1}{2} \gamma \sigma^2 (Z_u^{*\tilde{S}} + Q_u)^2 - H(Z_u^*, Q_u) \right) du.$$

The second problem in (1.20) is a classical optimization problem. Having found numerically $\mathcal{L}^*(z, q)$, we solve the Hamilton-Jacobi-Bellman (1.18) using neural networks.

Remark 9. *Theorem 1.3 characterizes only the value function of the exchange and not the optimal incentives defined in (1.21), which are computed through deep reinforcement learning techniques. In particular, there is no guarantee of admissibility of the incentives $(Z^*(t, Q_{t-}))_{t \in [0, T]}$ solving (1.19). However, we observe numerically (see Figure 1.5) that these incentive parameters are essentially linear (despite nonlinear nature of neural networks) in the inventory Q at any fixed time t . This result is indeed quite usual in the optimal market-making literature where asymptotic development of the function v is used¹², so v should be regular enough (see [142, Section 4] or [26, Section 3.2]). The linearity of the incentives Z^* , together with the uniform boundedness of the inventory, implies them to be in the*

¹²This corresponds to a first order Taylor expansion of v with respect to time, see for example Chapter 4, Equation (4.18).

set of admissible contracts \mathcal{Z} defined by (1.12). Finally, note that our framework does not take into account the case of cross-listed stocks where one financial instrument is traded on several platforms. In that case, the exchange does not know exactly the inventory of the market-maker, who can try to take advantage of the existence of a contract on one platform while trading on an alternative venue. Nevertheless, recall that for many assets, a large part of liquidity is captured by one venue.

1.5 Numerical solution: a deep reinforcement learning approach

We now turn to the description of our numerical method to solve (1.20), the optimization procedure consists of two stages. At the first stage, we optimize the controls of the market-maker for all possible values of the incentives given by the exchange. At the second stage, we use an actor-critic approach, to obtain both the optimal controls and the value function of the exchange. We conclude this section with numerical experiments showing the impact of incentives as well as that of market conditions on the volumes posted by the market-maker on both lit and dark venues regulated by the exchange. Throughout these experiments, we assume the following:

Assumption 1.1. For all $i \in \{a, b\}, j \in \{l, d\}$ and $k \in \mathcal{V}^j$, $Z^{i,j,k} = Z^{i,j} \in \mathbb{R}$.

This means that the Principal provides incentives only with respect to the number of transactions on each side of each pool independently of the volumes. In that case, HJB equation (1.18) remains valid. Recall that the optimal incentives depend on time and market-maker's inventory, therefore, implicitly they depend on the transacted volume.

There is obvious bid-ask symmetry in our model with respect to the inventory of the market-maker, as it can be seen in Hamiltonian (1.13). Thus for our numerical experiments we impose symmetry of the incentives with respect to q . As a consequence, we have symmetry of volumes posted by the market-maker with respect to q , given incentives satisfying the bid-ask symmetry property.

1.5.1 Description

1.5.1.1 Market-maker's problem

The first step to tackle our Principal-Agent problem is to find optimal volumes $\mathcal{L}^* = (\ell^{*a,l}, \ell^{*b,l}, \ell^{*a,d}, \ell^{*b,d})$, by solving for any couple (z, q) , the maximization problem of the market-maker (1.9). Hence, for fixed compensation terms z , the market-maker has to design his optimal volumes as function of z and q . Numerically, given the input z, q of our problem, we aim at finding the optimal volumes by maximizing his Hamiltonian function. To do so, we introduce a continuous version of the Hamiltonian (1.13) with respect to \mathcal{L} , that is we maximize the following functional:

$$\begin{aligned} \mathcal{L} \mapsto h^c(\mathcal{L}, z, q) := & \sum_{i \in \{a,b\}} \gamma^{-1} \left(\left(1 - \exp \left(-\gamma \left(z^{i,l} + \ell^{i,l} \left(\frac{\mathcal{T}}{2} + \phi(i) \Gamma^l q \right) - \Gamma^l (\ell^{i,l})^2 \right) \right) \right) \lambda^{i,l}(\mathcal{L}_t^l) \right. \\ & \left. + \sum_{\kappa \in K} \left(1 - \exp \left(-\gamma \left(z^{i,d} + \ell^{i,d} \left(\frac{\mathcal{T}}{2} \phi^{\text{lat}}(\kappa) + \phi(i) \Gamma^d q \right) - \Gamma^d (\ell^{i,d})^2 \right) \right) \lambda^{i,d}(\mathcal{L}_t^l) \phi^d(i, \kappa) \right) \right). \end{aligned} \quad (1.22)$$

For fixed incentives, we have that $\mathcal{L}^{*a,l}(q) = \mathcal{L}^{*b,l}(-q)$. Because of the intricate form of the function h^c , we cannot have an explicit solution to the first order condition $\nabla_{\mathcal{L}} h^c = 0$, which is four-dimensional. Moreover, we do not have an *a priori* knowledge on the functional form of optimizers $\mathcal{L}^* : \mathbb{R}^4 \times [-\bar{q}, \bar{q}] \rightarrow \mathbb{R}_+^4$, so we cannot apply canonical root-finding methods. Therefore to address this problem, we approximate the best response of the market-maker by a neural network.¹³

¹³To tackle the potential non-uniqueness of the market-maker's best response function, we start the optimization procedure at different initial points.

Although we do not use a purely grid-based method, we need to define a domain for arguments q and z . In our model inventory q of the market-maker is bounded and evolves between risk limits $-\bar{q}$ and \bar{q} . We also define a bound \bar{z} for the incentives $z \in \mathbb{R}^4$, so that $z \in [-\bar{z}, \bar{z}]^4$. This is reasonable as the optimal incentives defined in Equation (1.21) are Markovian functions of time and market-maker's inventory which are themselves bounded by T and \bar{q} respectively. This is in fact justified also by the paper El Euch, Mastrolia, Rosenbaum, and Touzi [109] in which optimal incentives are proved to be bounded.

We approximate the best response function \mathcal{L}^* by a neural network $l[\omega^l]$, where ω^l are the weights of the neural network.¹⁴ The neural network $l[\omega^l]$ takes as inputs Principal's incentives and the market-maker's current inventory $(z^{a,l}, z^{b,l}, z^{a,d}, z^{b,d}, q)$, which are normalized by \bar{z} and \bar{q} respectively. The network is composed of 2 hidden layers with 10 nodes in each of them and with ELU activation functions. ELU activation function is of the form

$$\text{ELU}(x) = \begin{cases} \alpha(e^x - 1), & \text{for } x \leq 0 \\ x, & \text{for } x > 0, \end{cases}$$

where α is a non-negative parameter, usually taken equal to 1.

The final layer of the network contains four outputs, and the activation function is sigmoid (for the outputs to be between 0 and 1). The output of $l[\omega^l]$ is then renormalized via multiplication by \bar{q} to obtain volumes between 0 and \bar{q} .

To obtain optimal volumes of the market-maker, we minimize the opposite of the Hamiltonian function defined by Equation (1.22). We generate $K > 0$ random samples of z and q , and conduct several epochs of batch learning with the following weights update:

$$\omega^l \leftarrow \omega^l + \mu^l \frac{1}{K} \sum_{k=1}^K \nabla_{\omega^l} l[\omega^l](z_k, q_k) \left(\nabla_l h^c(l[\omega^l](z_k, q_k), z_k, q_k) - \rho((q_k + l[\omega^l]^{b,l} + l[\omega^l]^{b,d} - \bar{q})_+ + (q_k - l[\omega^l]^{a,l} - l[\omega^l]^{a,d} + \bar{q})_-) \right), \quad (1.23)$$

where μ^l is the learning rate. The term scaled by ρ corresponds to a penalty employed to force quotes to stay in \mathcal{A} , so that $l[\omega^l]^{i,l} + l[\omega^l]^{i,d} \in [0, 2\bar{q}]$, $i \in \{a, b\}$. In our computations we use $\rho = 0.1$. In our tests we used a uniform grid for inventories q at each training epoch (no randomization), while for incentives z we applied a uniform sampling for each of its components. This kind of sampling allows us to focus on the impact of different incentives on each inventory state.

Algorithm 1 Market-Maker's optimal response ℓ^* calculation

Input: A market-maker's model (1.22), maximal risk limits \bar{q} .

1: **Initialization**

2: Initialize the network by default (Glorot kernel initializer, bias set to 0).

3: **Updates**

4: **repeat**

5: Sample z and q uniformly.

6: Update weights according to (1.23).

7: **until** convergence

Output: Neural network for the market-maker's response to incentives depending on his inventory.

Let us denote by $l^*[\omega^l]$ the approximated optimal response function of the market-maker \mathcal{L}^* (the result of the above optimization procedure). In Figure 1.1, we see an example of the best response

¹⁴Here we slightly abuse notation denoting by $l[\omega^l]$ the response of the market-maker obtained via neural network parametrized by weights ω^l .

$l^*[\omega^l]$ as a function of $z^{a,l} = -z^{b,l}$, when the market-maker's inventory $q = 50$ and other incentives $z^{a,d} = z^{b,d} = 0.05$ (close to zero). Remark that the choice of incentives is arbitrary only and aimed at reflecting the main properties of $l^*[\omega^l]$.

The observed behavior has quite natural interpretation. The incentive $z^{a,l}$ is a remuneration of the market-maker when his limit order is executed on the ask side of the lit pool. When this incentive increases, the market-maker ensures to have a small imbalance on the ask side of the lit pool so that he can earn $z^{a,l}$. Because of his positive inventory, the volume posted on the ask side of the dark pool is higher than on the bid side of the dark pool: the market-maker wants to liquidate his long position. Similarly when the incentive $z^{b,l}$ increases, the market-maker wants to benefit from it when transacting on the bid side of the lit pool. This explains the small imbalance on the bid side of the lit pool for positive $z^{b,l}$. Mathematically speaking, the function h^c is increasing in $z^{a,l}$. Thus for a high $z^{a,l}$, the value of the term $\mathcal{E}(z^{a,l}, l^*[\omega^l]^{a,l})$ in the Hamiltonian is high. To benefit from the remuneration $z^{a,l}$, the intensity $\lambda^{a,l}$ must be high, which implies a small imbalance on the ask side, hence I^a should be small. Similarly for $z^{b,l}$.




Figure 1.1: Best response of the market-maker as a function of $z^{a,l}$ and $z^{b,l}$, with $q = 50$.

For $q = 150$, $z^{b,l} = -z^{a,l}$, and other incentives $z^{a,d} = z^{b,d} = 0.05$ (close to zero), we display the volumes in Figure 1.2:

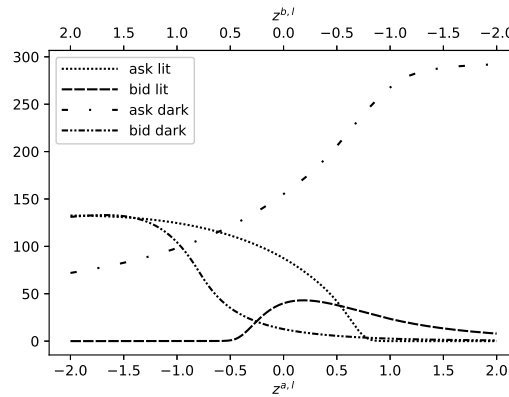


Figure 1.2: Best response of the market-maker as a function of $z^{a,l}$ and $z^{b,l}$, with $q = 150$.

As the market-maker has a higher inventory, his quotes on the bid side of both pools decrease because of the inventory risk. Moreover, his quotes on the ask side of both pools increase to liquidate his long position. For high incentives $z^{b,l}$, a small volume on the bid side of the lit pool leads to a low imbalance on the bid side, hence a high probability of execution for passive ask orders in the dark pool, where the market-maker tries to liquidate his position. Note that for high $z^{a,l}$, the imbalance is approximately equal to one half, because the market-maker does not want to suffer from the latency effect (to be executed at the mid-price in the dark pool).

We now move to the problem of the Principal.

1.5.1.2 An actor-critic approach to solve HJB equation (1.18)

A numerical approximation of the optimal incentives z^* can be obtained by

1. solving (numerically) the static maximization problem (1.22), which provides the approximation of the optimal response \mathcal{L}^* of the market-maker,
2. plugging this approximation in the continuous (with respect to \mathcal{L}^*) version of Hamilton-Jacobi-Bellman equation (1.18), that is to say:

$$\begin{cases} 0 = \partial_t v(t, q) + \mathcal{H}^c(q, \mathcal{L}^*, v(t, \cdot)), & q \in [-\bar{q}, \bar{q}], t \in [0, T), \\ v(T, q) = -1, \end{cases}$$

with

$$\mathcal{H}^c(q, \mathcal{L}^*, v(t, \cdot)) := \sup_{z \in \mathbb{R}^5} \mathcal{U}^c(z, q, \mathcal{L}^*(z, q), v(t, \cdot)),$$

and abusing the notation with \mathcal{L}^* denoting $\mathcal{L}^*(z, q)$

$$\begin{aligned} \mathcal{U}^c(z, q, \mathcal{L}^*, v(t, \cdot)) &:= v(t, q) \left(\frac{\eta}{2} \sigma^2 \gamma (z^{\tilde{S}} + q)^2 + \frac{\eta^2 \sigma^2}{2} (z^{\tilde{S}})^2 \right) \\ &+ \sum_{(i,j) \in \{a,b\} \times \{l,d\}} \lambda^{i,j}(\mathcal{L}^{*l}) \left(\exp(\eta(z^{i,j} - \ell^{*i,j})) v(t, q - \phi(i) \ell^{*i,j}) - v(t, q) (1 + \eta \mathcal{E}(z^{i,j}, \ell^{*i,j})) \right). \end{aligned}$$

We obtain explicitly $z^{\tilde{S}} = -\frac{\gamma}{\gamma + \eta} q$, so we are only interested in finding optimal $(z^{a,l}, z^{b,l}, z^{a,d}, z^{b,d})$. The classical method to solve the above problem is to obtain an approximation of the value function *via* a finite difference scheme on a grid. Since the size of the grid increases exponentially with the number of dimensions, using this approach is not possible for a high dimension. Therefore, to address our five-dimensional optimization problem, we resort to neural networks.

We use an algorithm known in reinforcement learning literature as the actor-critic method. The core of this approach is the representation of the value function and optimal controls with deep neural networks. The learning procedure itself consists of two stages: value function update (also called critic update) and controls update (actor update).

We first split our problem into sub-problems corresponding to different time steps. We consider a time step Δt . For each time step Δt , we represent the value function and the incentives with neural networks. Going backward in time, we approximate optimal controls taking into account the value function of an already resolved sub-problem. Then we approximate a value function corresponding to the just found controls. Note that $v(T, \cdot) = -1$ is given.

Let us fix $t \in [0, T - \Delta t]$. The current set of incentives approximated by the feedforward neural network $z_t[\omega^{z_t}](\cdot)$ (also called an actor's network) is parametrized by weights ω^{z_t} . The corresponding value function at time t is represented by $v_t[\omega^{v_t}](\cdot)$ (also called a critic's network), which is a feedforward neural network, parameterized by weights ω^{v_t} . Actor's network is composed of 2 hidden layers with 20 nodes in each of these layers with ELU activation functions. The final layer of the network contains four outputs, and the activation is tanh (this allows the output to stay between -1 and 1), which is therefore renormalized by \bar{z} . Critic's network is composed of 2 hidden layers with 20 nodes in each of these layers with ELU activation functions. The final layer of the network contains one output, and the activation is affine.

At each time step we have the value function's neural network approximating the value function for the sub-problem on $[t + \Delta t, T]$, that allows us to update the controls for the current period $[t, t + \Delta t]$. Controls' network update consists of two different procedures. The first one is an exploitation phase where the weights are updated according to the best direction suggested by the gradient of the function $\mathcal{U}^c(z_t[\omega^{z_t}](q_k), q_k, l^*[\omega^l], v_t[\omega^{v_t}] (\cdot))$:

$$\omega^{z_t} \leftarrow \omega^{z_t} + \mu^z \frac{1}{K} \sum_{k=1}^K \nabla_{\omega^{z_t}} z_t[\omega^{z_t}](q_k) \nabla_{z_t} \mathcal{U}^c(z_t[\omega^{z_t}](q_k), q_k, l^*[\omega^l], v_{t+\Delta t}[\omega^{v_t+\Delta t}] (\cdot)), \quad (1.24)$$

where μ^z is a learning rate. This type of updates is usually called a policy gradient.

Another type of updates we use in the learning procedure is an exploration phase. During this phase, we use the current values given by the neural network of controls and introduce noise to these values, to explore the controls slightly different from those proposed by the neural network. Noise is normally distributed around 0 with standard deviation chosen beforehand (in the following examples, we use standard normal distribution). This phase could help us to quit local minima, in case the algorithm is trapped in one. The following updates characterize this phase:

$$\begin{aligned} \omega^{z_t} \leftarrow \omega^{z_t} + \hat{\mu}^z \frac{1}{K} \sum_{k=1}^K \varepsilon_k \nabla_{\omega^{z_t}} z_t[\omega^{z_t}](q_k) & \left(\mathcal{U}^c(z_t[\omega^{z_t}](q_k) + \varepsilon, q_k, l^*[\omega^l], v_{t+\Delta t}[\omega^{v_t+\Delta t}] (\cdot)) \right. \\ & \left. - \mathcal{U}^c(z_t[\omega^{z_t}](q_k), q_k, l^*[\omega^l], v_{t+\Delta t}[\omega^{v_t+\Delta t}] (\cdot)) \right), \end{aligned} \quad (1.25)$$

where ε is a vector of length K representing introduced perturbations and $\hat{\mu}^z$ is a learning rate.

Once we found the controls for the time period $[t, t + \Delta t]$ we can approximate the value function for $[t, T]$. We use the first-order approximation of the value function at time t giving

$$v(t, \cdot) \approx v(t + \Delta t, \cdot) - \partial_t v(t + \Delta t, \cdot) \Delta t.$$

This formula leads to the following update of the value function network's weights ω^{v_t} :

$$\begin{aligned} \omega^{v_t} \leftarrow \omega^{v_t} + \mu^v \frac{1}{K} \sum_{k=1}^K \nabla_{\omega^{v_t}} v_t[\omega^{v_t}](q_k) & \left(v_{t+\Delta t}[\omega^{v_t+\Delta t}](q_k) \right. \\ & \left. + \mathcal{U}^c(z_t[\omega^{z_t}](q_k), q_k, l^*[\omega^l], v_{t+\Delta t}[\omega^{v_t+\Delta t}] (\cdot)) - v_t[\omega^{v_t}](q_k) \right), \end{aligned} \quad (1.26)$$

where μ^v is a learning rate, $\mathcal{U}^c(z_t[\omega^{z_t}](q_k), q_k, l^*[\omega^l], v_t[\omega^{v_t}] (\cdot))$ corresponds to the function under the supremum of the Hamiltonian (1.19) calculated using the current controls $z_t[\omega^{z_t}]$. The quantities q_k , $k \in \{1, \dots, K\}$ are the elements of the training set, more precisely K uniformly distributed elements from the interval $[-\bar{q}, \bar{q}]$. It should be noted that the training set q_k , $k \in \{1, \dots, K\}$ could also be just a grid. We use $\mathcal{U}(z_t[\omega^{z_t}](q_k), q_k, l^*[\omega^l], v_{t+\Delta t}[\omega^{v_t+\Delta t}] (\cdot))$ as an approximation of $\partial_t v(t + \Delta t, \cdot)$ to apply the first order approximation of the value function described above.

The whole procedure can be summarized in the following algorithm:

1.5.2 Numerical results: influence of the make-take fees policy

In the following we consider $\Delta t = 1$. Since time has little impact on the quotes chosen by the market-maker (see El Euch, Mastrolia, Rosenbaum, and Touzi [109], Guéant, Lehalle, and Fernandez-Tapia [142]), we present the results only for time $T - 1$ where T corresponds to one business day, the extension to earlier time steps is straightforward. As mentioned before, the optimization problems considered are symmetric with respect to the inventory variable q . We present in the following table a summary

Algorithm 2 Principal's incentives z^* calculation**Input:** A neural network with market-maker's response to incentives ℓ^* , maximal incentives limits \bar{z}

```

1: for  $i = \text{NB TIME STEPS} - 1 \dots 0$  do
2:   Initialization (the first part of if condition can be used for every time step)
3:   if  $i == \text{NB TIME STEPS} - 1$  then
4:     Pretrain incentives to zero.
5:     Initialize value function's network by default (Glorot kernel initializer, bias set to 0).
6:   else
7:     Take weights from the incentives and the value function network of the previous iteration.
8:   end if
9:   Incentives updates
10:  repeat
11:    Generate uniform sample of  $q$ .
12:    Obtain market-maker's response to current incentives via application of  $\ell^*$  obtained above.
13:    Update weights according to (1.24).
14:    (optional) Update weights according to (1.25)
15:  until convergence
16:  Value function updates
17:  repeat
18:    Update weights according to (1.26)
19:  until convergence
20: end for

```

Output: $\frac{T}{\Delta t}$ neural networks of the incentives given by the Principal, together with $\frac{T}{\Delta t}$ networks for corresponding value functions.

of the market parameters chosen for each case we study. In all the study, we set in all the study $\gamma = 0.01, \eta = 0.02$. The results obtained are qualitatively insensitive to the parameters we used. As it is impossible to obtain closed-form solutions for the problem we study, this section provides a numerical analysis of the behavior of the market-maker with and without the intervention of the exchange, followed by a sensitivity analysis of the volatility and impact parameters.

	σ	Γ^l	Γ^d	θ^l	θ^d	c^l	c^d	A^l	A^d
Reference case	0.1	10^{-4}	5×10^{-5}	0.15	0.15	0.05	0.01	5×10^3	3×10^3
High volatility case	0.4	10^{-4}	5×10^{-5}	0.15	0.15	0.05	0.01	5×10^3	3×10^3
Low impact case	0.1	10^{-4}	2×10^{-5}	0.15	0.15	0.05	0.01	5×10^3	3×10^3
High impact case	0.1	2.5×10^{-4}	2.5×10^{-4}	0.15	0.15	0.05	0.01	5×10^3	3×10^3

1.5.2.1 Reference model without the exchange

We study the behavior of the market-maker in the case $\xi = 0$. This constitutes a benchmark to analyze the influence of our make-take fees policy in the next section. In Figure 1.3, we present the optimal quotes of the market-maker.

One can see that the market-maker splits his orders equitably between the lit and dark pools when his inventory is near zero. However, when he has a very positive (resp. negative) inventory, he has a large imbalance on the ask (resp. bid) side of the lit pool, to liquidate his position in the dark pool. Such behavior shows that the market-maker uses the dark pool as a way to liquidate a large position by adjusting the imbalance in the lit pool. Indeed, when he posts a high volume on the ask side of the lit pool, he encourages ask orders in the dark pool. Thus, as he prioritizes the execution of a large ask order, he accepts to be executed at the mid-price in the dark pool. When $q = 300$, he does not post a sell order of size 300 in the dark pool, because of the quadratic variation between the mid-price and its

inventory process (which can be seen as a quadratic penalty in the market-maker's PnL process with respect to the volumes displayed). Because of the latency generated on the ask side of the lit pool, the market takers sending market orders on the bid side of the dark pool are likely to be executed at an unfavorable price. This is why the market-maker posts a non-zero volume on the bid side of the dark pool.




Figure 1.3: Optimal quotes of the market-maker.

Remark also that for small inventories, the market-maker posts volumes on both ask and bid sides of the dark pool because he may accept to increase his inventory risk by being executed at a more favorable price in the dark pool due to the latency effect (the volumes displayed in the lit pool lead to 50 percents chance to face this effect at least on one of the sides of the dark pool). Note that the parameters $A^l > A^d$ describe the fact that there are, on average, much more orders in the lit pool than in the dark pool.¹⁵ In the following sections, we present several numerical experiments involving the incentive policy of the exchange.

1.5.2.2 Reference model with the exchange

We now show the influence of the exchange on the behavior of the market-maker. In Figures 1.4 and 1.5, we present the optimal quotes of the market-maker and the optimal incentives provided by the exchange.




Figure 1.4: Optimal quotes of the market-maker.




Figure 1.5: Optimal incentives of the exchange.

The presence of incentives has significant effects on the market-maker's behavior. When the market-maker has an inventory near zero, incentives lead to an increase of the volumes posted in the lit pool and a decrease of that in the dark pool compared to Figure 1.3. Thus the exchange improves the liquidity in the lit venue.

Moreover, the strategy of the market-maker for very positive or negative inventory is modified. When he has a very positive inventory, he posts a higher volume on the ask side of the dark pool than in the

¹⁵This assumption is consistent with the MIFID II regulation rolled out on January 3, 2018, which imposes a cap on volumes traded in the dark pools, see for example [84].

case without exchange. In addition to this, he posts an equal volume (small but not negligible) on the ask and bid sides of the lit pool. So we see that the exchange prevents the market-maker from artificial manipulation of the market, consisting in creation of high imbalance on the ask side. As the imbalance is around $1/2$, the market-maker does not take advantage of the latency effect. In Figure 1.5, we see that, even if our problem is much more intricate than those of El Euch, Mastroli, Rosenbaum, and Touzi [109] and Chapter 4, the shape of the Principal's incentives are essentially linear functions of the market-maker's inventory.

1.5.3 Numerical results: sensitivity analysis

1.5.3.1 High volatility regime

We now investigate the impact of higher volatility on the posted volumes with and without the exchange.




Figure 1.6: Optimal quotes of the market-maker without the exchange.




Figure 1.7: Optimal quotes of the market-maker with the exchange.

In Figure 1.6 we see that, compared to Figure 1.3, higher volatility does not change the strategy of the market-maker without the exchange significantly. We observe that the contract has more limited influence in the case of high volatility as the market-maker follows the same strategy as without exchange. In particular, he does not keep his imbalance equal to $1/2$ when he has a very positive or negative inventory. This is because higher volatility leads to an increase in market activity, and the market-maker is more willing to send higher volumes on the side of interest of both pools.

1.5.3.2 Low market impact in the dark pool

We now show the volumes displayed by the market-maker with and without the exchange, when the impact parameters of the dark pool make it even more appealing. In Figures 1.8 and 1.9, we see the influence of a higher market impact and transaction costs in the lit pool. Either with or without the intervention of the exchange and for small inventories, the market-maker posts higher volumes in the dark pool than in the lit pool. We recover similar behavior for the displayed volumes as in the reference case with and without the exchange in Figures 1.3 and 1.4.

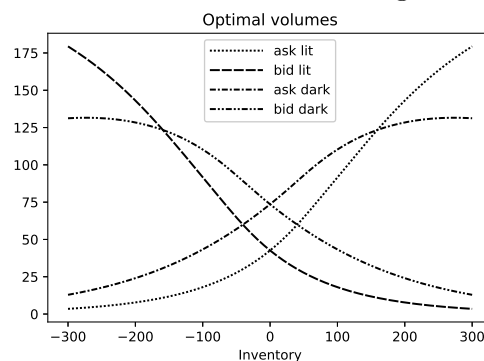


Figure 1.8: Optimal quotes of the market-maker without the exchange.




Figure 1.9: Optimal quotes of the market-maker with the exchange.

1.5.3.3 High market impact on both venues

In this last section, we show how the volumes are split between the lit and dark pools when the market impact in the lit and dark pools are equal.

In Figures 1.10 and 1.11, we see that a higher market impact reduces the volume posted on both lit and dark pools. We also recover a behavior similar to the reference case without the exchange. For the case with the exchange, for very positive (resp. negative) inventory, the market-maker has an ask (resp. bid) imbalance slightly above $1/2$, meaning that market takers on the bid (resp. ask) of the dark pool are more likely to be executed at a price unfavorable for them due to the latency effect.




Figure 1.10: Optimal quotes of the market-maker without the exchange.




Figure 1.11: Optimal quotes of the market-maker with the exchange.

1.6 Conclusion

In this chapter, we introduce an optimal make-take fees system enabling an exchange to improve liquidity provided by market-makers on both lit and dark pools. As a by-product, we obtain a way to design optimal market-making strategies in the presence of a dark pool. Mathematically, we use Principal-Agent theory and prove the existence and uniqueness of viscosity solutions associated to the control problems of the exchange and market-makers. The obtained results are easily computable in practice using recent machine learning techniques. More precisely, we use nested neural networks via a deep reinforcement learning approach called actor-critic method. The remuneration offered to market-makers is, without loss of generality, indexed on the number of transactions and the price process, which are observable quantities in practice. Our results show that our make-take fees system prevents market-makers from price manipulation on the dark pool using the imbalance of the lit pool, and increase the overall liquidity provision on the dark pool.

A natural next step would be to consider the case of one lit pool and one dark pool monitored by two different exchanges with possibly common market-makers on the two platforms aiming to attract liquidity. Some preliminary results have been obtained in El Euch, Mastrolia, Rosenbaum, and Touzi [109] for identical exchanges in competition, without considering any dark trading. Such extension to multi Principal-Agent problems is much more intricate as explained in Hu, Ren, and Yang [162], Mastrolia and Ren [206] since it would require economical stability as for example the existence of equilibria in the incentives proposed by the competitive exchanges.

1.A Appendix

1.A.1 Dynamic programming principle and contract representation

For any \mathbb{F} -predictable stopping time $\tau \in [t, T]$ and $\mathcal{L} \in \mathcal{A}_\tau$, we define:

$$J_T(\tau, \mathcal{L}) = \mathbb{E}_\tau^\mathcal{L} \left[-\mathcal{D}_{\tau, T}(\mathcal{L}) \exp(-\gamma\xi) \right], \quad V_\tau = \sup_{\mathcal{L} \in \mathcal{A}_\tau} J_T(\tau, \mathcal{L}),$$

where \mathcal{A}_τ denotes the restriction of \mathcal{A} to controls on $[\tau, T]$ and

$$\mathcal{D}_{\tau, T}(\mathcal{L}) := \exp \left(-\gamma \left(\int_\tau^T \sum_{i \in \{a, b\}} \left(\frac{\mathcal{T}}{2} \ell_t^{i, l} dN_t^{i, l} + \sum_{\kappa \in K} \phi^{lat}(\kappa) \frac{\mathcal{T}}{2} \ell_t^{i, d} dN_t^{i, d, \kappa} \right) + Q_t dS_t + d[Q \cdot, S \cdot]_t \right) \right),$$

where

$$d[Q \cdot, S \cdot]_t = - \sum_{(i, j) \in \{a, b\} \times \{l, d\}} \sum_{k \in \mathcal{V}^j} \Gamma^j k^2 dN_t^{i, j, k}.$$

We now set the dynamic programming principle associated to the control problem (1.4).

Lemma 1.1. *Let $t \in [0, T]$ and τ be an \mathbb{F} -predictable stopping time with values in $[t, T]$. Then*

$$V_t = \text{ess sup}_{\mathcal{L} \in \mathcal{A}} \mathbb{E}_t^\mathcal{L} \left[-\mathcal{D}_{t, \tau}(\mathcal{L}) V_\tau \right].$$

The proof can be found in [109, Lemma A.4]. The following technical lemma shows some integrability conditions on $\mathcal{D}_{t, T}$ and V_t , which will be useful later.

Lemma 1.2. *For all $\mathcal{L} \in \mathcal{A}$ and for a specific $\epsilon > 0$, we have*

$$\mathbb{E}^\mathcal{L} \left[\sup_{t \in [0, T]} |V_t|^{1+\epsilon} \right] < +\infty, \quad \mathbb{E}^\mathcal{L} \left[\sup_{(s, t) \in [0, T]^2} (D_{s, t}(\mathcal{L}))^{1+\epsilon} \right] < +\infty.$$

Proof. We have for $t \in [0, T]$,

$$\begin{aligned} -V_t &= \text{ess inf}_{\mathcal{L} \in \mathcal{A}} \mathbb{E}_t^\mathcal{L} \left[\exp \left(-\gamma \left(\int_t^T \sum_{i \in \{a, b\}} \left(\frac{\mathcal{T}}{2} \ell_u^{i, l} dN_u^{i, l} + \sum_{\kappa \in K} \phi^{lat}(\kappa) \frac{\mathcal{T}}{2} \ell_u^{i, d} dN_u^{i, d, \kappa} \right) + Q_u dS_u \right. \right. \right. \\ &\quad \left. \left. \left. + d[Q \cdot, S \cdot]_u \right) \right) \exp(-\gamma\xi) \right] \\ &\leq \mathbb{E}_t^\mathcal{L} \left[e^{\gamma((\bar{k}^2 + \bar{q}\bar{k}) \sum_{i \in \{a, b\}} \sum_{j \in \{l, d\}} \Gamma^j N_T^{i, j} - \int_t^T Q_u d\tilde{S}_u)} \exp(-\gamma\xi) \right], \end{aligned}$$

where $\bar{k} = \max_{j \in \{l, d\}} \max_{k \in \mathcal{V}^j} k$, and we used the fact that $N_T^{i, j} - N_t^{i, j} \leq N_T^{i, j}$, $Q_t \in [-\bar{q}, \bar{q}]$ and the control processes are uniformly bounded by \bar{k} . Moreover, for all $L > 0$,

$$\mathbb{E}_t^\mathcal{L} \left[e^{-L \int_t^T Q_u d\tilde{S}_u} \right] \leq e^{\frac{L^2 \bar{q}^2 \sigma^2 T}{2}}.$$

Thus, using Holder's inequality, we have

$$\begin{aligned} -V_t &\leq \mathbb{E}_t^\mathcal{L} \left[e^{\epsilon \gamma ((\bar{k}^2 + \bar{q}\bar{k}) \sum_{i \in \{a, b\}} \sum_{j \in \{l, d\}} \Gamma^j N_T^{i, j} - \xi)} \right]^{\frac{1}{\epsilon}} \mathbb{E}_t^\mathcal{L} \left[e^{-\gamma(1+\epsilon) \int_t^T Q_u d\tilde{S}_u} \right]^{\frac{1}{1+\epsilon}} \\ &\leq \mathbb{E}_t^\mathcal{L} \left[e^{\epsilon \gamma ((\bar{k}^2 + \bar{q}\bar{k}) \sum_{i \in \{a, b\}} \sum_{j \in \{l, d\}} \Gamma^j N_T^{i, j} - \xi)} \right]^{\frac{1}{\epsilon}} e^{\frac{(1+\epsilon)\gamma^2 \bar{q}^2 \sigma^2 T}{2}}. \end{aligned}$$

Then, we have

$$\mathbb{E}^{\mathcal{L}} \left[\sup_{t \in [0, T]} (-V_t)^{1+\epsilon} \right] \leq e^{\frac{(1+\epsilon)^2 \gamma^2 \bar{q}^2 \sigma^2 T}{2}} \mathbb{E}^{\mathcal{L}} \left[\sup_{t \in [0, T]} \mathbb{E}_t^{\mathcal{L}} \left[e^{\epsilon \gamma ((\bar{k}^2 + \bar{q}\bar{k}) \sum_{i \in \{a, b\}} \sum_{j \in \{l, d\}} \Gamma^j N_T^{i, j} - \xi)} \right]^{\frac{1+\epsilon}{\epsilon}} \right].$$

The term inside the conditional expectation is integrable¹⁶ and independent from $t \in [0, T]$ thus by Doob's inequality, we have

$$\mathbb{E}^{\mathcal{L}} \left[\sup_{t \in [0, T]} (-V_t)^{1+\epsilon} \right] \leq C e^{\frac{(1+\epsilon)^2 \gamma^2 \bar{q}^2 \sigma^2 T}{2}} \mathbb{E}^{\mathcal{L}} \left[e^{\gamma' ((\bar{k}^2 + \bar{q}\bar{k}) \sum_{i \in \{a, b\}} \sum_{j \in \{l, d\}} \Gamma^j N_T^{i, j} - \xi)} \right],$$

where $C > 0$, $\gamma' = \gamma(1 + \epsilon)$. Thanks to Hölder's inequality, together with the boundedness of the intensities of the point processes $N^{i, j}$ and Condition (1.5), the right-hand side is bounded from above by a term independent of $t \in [0, T]$. The conclusion follows.

Using the same arguments, we have

$$\mathbb{E}^{\mathcal{L}} \left[\sup_{(s, t) \in [0, T]^2} (\mathcal{D}_{s, t}(\mathcal{L}))^{1+\epsilon} \right] \leq C' \mathbb{E}^{\mathcal{L}} \left[e^{\gamma' ((\bar{k}^2 + \bar{q}\bar{k}) \sum_{i \in \{a, b\}} \sum_{j \in \{l, d\}} \Gamma^j N_T^{i, j} + \frac{(1+\epsilon)\gamma\bar{q}^2\sigma^2 T}{2})} \right] < +\infty,$$

where $C' > 0$, using boundedness of the intensities of the point processes. The conclusion follows using Hölder's inequality. \square

1.A.2 Proof of Theorem 1.2

To prove that $\mathcal{C} = \Xi$, we proceed in six steps. Our approach is largely inspired by El Euch, Mastrolia, Rosenbaum, and Touzi [109]. However, for the sake of completeness, we provide here the details.

Step 1: For $\mathcal{L} \in \mathcal{A}$ it follows from the dynamic programming principle of Lemma 1.1 and Lemma 1.2 that the process

$$U_t^{\mathcal{L}} = V_t \mathcal{D}_{0, t}(\mathcal{L})$$

is of class (D) and defines a $\mathbb{P}^{\mathcal{L}}$ -supermartingale for any $\mathcal{L} \in \mathcal{A}$. By standard analysis, we may then consider it in its càdlàg version (by taking right limits along rationals). By the Doob-Meyer decomposition, we can write $U_t^{\mathcal{L}} = M_t^{\mathcal{L}} - A_t^{\mathcal{L}}$ where $M^{\mathcal{L}}$ is a $\mathbb{P}^{\mathcal{L}}$ -martingale and $A_t^{\mathcal{L}} = A_t^{\mathcal{L}, c} + A_t^{\mathcal{L}, d}$ is an integrable non-decreasing predictable process such that $A_0^{\mathcal{L}, c} = A_0^{\mathcal{L}, d} = 0$ with pathwise continuous component $A^{\mathcal{L}, c}$ and with $A^{\mathcal{L}, d}$ a piecewise constant predictable process.

From the martingale representation theorem under $\mathbb{P}^{\mathcal{L}}$, see Appendix A.1 in El Euch, Mastrolia, Rosenbaum, and Touzi [109], there exists $\tilde{Z}^{\mathcal{L}} = (\tilde{Z}^{\mathcal{L}, S}, \tilde{Z}^{\mathcal{L}, i, j, k})_{i \in \{a, b\}, j \in \{l, d\}, k \in \mathcal{V}^j}$ predictable, such that

$$M_t^{\mathcal{L}} = V_0 + \int_0^t \tilde{Z}_r^{\mathcal{L}, S} d\tilde{S}_r + \sum_{(i, j) \in \{a, b\} \times \{l, d\}} \sum_{k \in \mathcal{V}^j} \int_0^t \tilde{Z}_r^{\mathcal{L}, i, j} dN_r^{\mathcal{L}, i, j}.$$

Step 2: We now show that V is a negative process. Thanks to the uniform boundedness of $\mathcal{L} \in \mathcal{A}$ and $I^a, I^b \in [0, 1]$ we get that

$$\frac{L_T^{\mathcal{L}}}{L_t^{\mathcal{L}}} \geq \alpha_{t, T} = \exp \left(- \sum_{j \in \{l, d\}} \frac{\theta^j}{\sigma} (N_T^{a, j} - N_t^{a, j} + N_T^{b, j} - N_t^{b, j}) - 2(A^j - \epsilon)(T - t) \right).$$

¹⁶Take $\epsilon > 1$ together with Condition (1.5) for example.

Therefore using the definition of $\mathcal{D}_{t,T}(\mathcal{L})$, we obtain

$$V_t \leq \mathbb{E}_t^0 \left[-\alpha_{t,T} \exp \left(-\gamma \left(3(\mathcal{T} - \Gamma^l - \Gamma^d) \bar{q}^2 \left(\sum_{(i,j) \in \{a,b\} \times \{l,d\}} N_T^{i,j} - N_t^{i,j} \right) + \int_t^T Q_u d\tilde{S}_u \right) \right) \exp(-\gamma\xi) \right] < 0.$$

Step 3: Let Y be the process defined for any $t \in [0, T]$ by $V_t = -\exp(-\gamma Y_t)$. As $A^{\mathcal{L},d}$ is a predictable point process and the jumps of $N^{i,j,k}$, $i \in \{a, b\}$, $j \in \{l, d\}$, $k \in \mathcal{V}^j$ are totally inaccessible stopping times under \mathbb{P}^0 , we have $\langle N^{i,j,k}, A^{\mathcal{L},d} \rangle_t = 0$ a.s. We obtain

$$Y_T = \xi \text{ and } dY_t = \left(\sum_{(i,j) \in \{a,b\} \times \{l,d\}} \sum_{k \in \mathcal{V}^j} Z_t^{i,j,k} dN_t^{i,j,k} \right) + Z_t^{\tilde{S}} d\tilde{S}_t - dI_t - d\tilde{A}_t^d.$$

Itô's formula yields to

$$\begin{aligned} Z_t^{a,l,k} &:= -\frac{1}{\gamma} \log \left(1 + \frac{\tilde{Z}_t^{\mathcal{L},a,l,k}}{U_t^{\mathcal{L}}} \right) - \ell_t^{a,l} \left(\frac{\mathcal{T}}{2} + \Gamma^l Q_{t^-} \right) + \Gamma^l k^2, \\ Z_t^{b,l,k} &:= -\frac{1}{\gamma} \log \left(1 + \frac{\tilde{Z}_t^{\mathcal{L},b,l,k}}{U_t^{\mathcal{L}}} \right) - \ell_t^{b,l} \left(\frac{\mathcal{T}}{2} - \Gamma^l Q_{t^-} \right) + \Gamma^l k^2, \\ Z_t^{a,d,k} &:= -\frac{1}{\gamma} \log \left(1 + \frac{\tilde{Z}_t^{\mathcal{L},a,d,k}}{U_t^{\mathcal{L}}} \right) - \ell_t^{a,d} \left(\frac{\mathcal{T}}{2} \mathbf{1}_{\nu_t^a=0} + \Gamma^d Q_{t^-} \right) + \Gamma^d k^2, \\ Z_t^{b,d,k} &:= -\frac{1}{\gamma} \log \left(1 + \frac{\tilde{Z}_t^{\mathcal{L},b,d,k}}{U_t^{\mathcal{L}}} \right) - \ell_t^{b,d} \left(\frac{\mathcal{T}}{2} \mathbf{1}_{\nu_t^b=0} - \Gamma^d Q_{t^-} \right) + \Gamma^d k^2, \\ Z_t^{\tilde{S}} &:= -\frac{\tilde{Z}_t^{\mathcal{L},S}}{\gamma U_t^{\mathcal{L}}} - Q_{t^-}, \\ I_t &:= \int_0^t \left(\bar{h}(\mathcal{L}_r, Z_r, Q_r) dr - \frac{1}{\gamma U_r^{\mathcal{L}}} dA_r^{\mathcal{L},c} \right), \\ \bar{h}(\mathcal{L}, Z_t, Q_t) &:= h(\mathcal{L}, Z_t, Q_t) - \frac{1}{2} \gamma \sigma^2 (Z_t^{\tilde{S}})^2, \\ \tilde{A}_t^d &:= \frac{1}{\gamma} \sum_{s \leq t} \log \left(1 - \frac{\Delta A_s^{\mathcal{L},d}}{U_s^{\mathcal{L}}} \right). \end{aligned}$$

In particular, the last relation between \tilde{A}^d and $A^{\mathcal{L},d}$ shows that $\Delta a_t \geq 0$ is independent of $\mathcal{L} \in \mathcal{A}$, with $a_t = -\frac{A_t^{\mathcal{L},d}}{U_t^{\mathcal{L}}}$ and abusing notations slightly, $\Delta a_t = -\frac{\Delta A_t^{\mathcal{L},d}}{U_t^{\mathcal{L}}}$. In order to complete the proof, we argue in the subsequent steps that $Z \in \mathcal{Z}$ and that, for $t \in [0, T]$, $A_t^{\mathcal{L},d} = -\sum_{s \leq t} U_s^{\mathcal{L}} \Delta a_s = 0$ so that $\tilde{A}_t^d = 0$ and $I_t = \int_0^t \bar{H}(Z_r, Q_r) dr$ where

$$\bar{H}(Z_t, Q_t) = H(Z_t, Q_t) - \frac{1}{2} \gamma \sigma^2 (Z_t^{\tilde{S}})^2.$$

Step 4: Since $V_T = -1$, we get that

$$\begin{aligned} 0 &= \sup_{\mathcal{L} \in \mathcal{A}} \mathbb{E}^{\mathcal{L}} [U_T^{\mathcal{L}}] - V_0 \\ &= \sup_{\mathcal{L} \in \mathcal{A}} \mathbb{E}^{\mathcal{L}} \left[U_T^{\mathcal{L}} - M_T^{\mathcal{L}} \right] \\ &= \gamma \sup_{\mathcal{L} \in \mathcal{A}} \mathbb{E}^0 \left[L_T^{\mathcal{L}} \int_0^T U_{r^-}^{\mathcal{L}} \left(dI_r - \bar{h}(\mathcal{L}, Z_r, Q_r) dr + \frac{da_r}{\gamma} \right) \right]. \end{aligned}$$

Moreover, the controls being uniformly bounded, we have

$$U_t^{\mathcal{L}} \leq -\beta_t := V_t \exp \left(-\gamma \left(3(\mathcal{T} - \Gamma^l - \Gamma^d) \bar{q}^2 \left(\sum_{(i,j) \in \{a,b\} \times \{l,d\}} N_t^{i,j} \right) + \int_0^t Q_r d\tilde{S}_r \right) \right) < 0.$$

Then, using $A^{\mathcal{L},d} \geq 0, U^{\mathcal{L}} \leq 0$ and $dI_t - \bar{h}(\mathcal{L}, Z_t, Q_t) dt \geq 0$, obtain

$$\begin{aligned} 0 &\leq \sup_{\mathcal{L} \in \mathcal{A}} \mathbb{E}^0 \left[\alpha_{0,T} \int_0^T -\beta_{r-} \left(dI_r - \bar{h}(\mathcal{L}, Z_r, Q_r) dr + \frac{da_r}{\gamma} \right) \right] \\ &= -\mathbb{E}^0 \left[\alpha_{0,T} \int_0^T \beta_{r-} \left(dI_r - \bar{H}(Z_r, Q_r) dr + \frac{da_r}{\gamma} \right) \right]. \end{aligned}$$

The quantities $\alpha_{0,T} \int_0^T \beta_{r-} (dI_r - \bar{H}(Z_r, Q_r)) dr$ and $\alpha_{0,T} \int_0^T \beta_{r-} \frac{da_r}{\gamma}$ being non-negative random variables, the result follows.

Moreover, if \mathcal{L} is such that for any $(z, q) \in \mathbb{R}^{2(\#\mathcal{V}^l + \#\mathcal{V}^d)} \times \mathbb{N}$ we have $h(\mathcal{L}, z, q) = H(z, q)$, then

$$\int_0^T U_{r-}^{\mathcal{L}} (dI_r - \bar{h}(\mathcal{L}, Z_r, Q_r)) dr = 0.$$

Therefore, $\sup_{\mathcal{L} \in \mathcal{A}} \mathbb{E}^{\mathcal{L}}[U_T^{\mathcal{L}}] = V_0$ which implies that (OC) is satisfied. Conversely, if (OC) is satisfied, the supremum is directly attained. This provides the inclusion $\mathcal{C} \supset \Xi$.

Step 5: We now prove that $Z \in \mathcal{Z}$. Using Lemma 1.2, we have that

$$\sup_{\mathcal{L} \in \mathcal{A}} \mathbb{E}^{\mathcal{L}} \left[\sup_{t \in [0, T]} |U_t^{\mathcal{L}}|^{p'+1} \right] < +\infty,$$

for some $p' > 0$. The desired conclusion follows from the fact that

$$\exp(-\gamma Y_t) = U_t^{\mathcal{L}} \exp \left(\gamma \left(\int_0^t \sum_{i \in \{a,b\}} \left(\frac{\mathcal{T}}{2} \ell_u^{i,l} dN_u^{i,l} + \sum_{\kappa \in K} \phi^{lat}(\kappa) \ell_u^{i,d} dN_u^{i,d,\kappa} \right) + Q_u dS_u + d[Q \cdot, S \cdot]_u \right) \right).$$

Consequently, $\mathcal{C} \subset \Xi$ and using Step 4 we finally get $\mathcal{C} = \Xi$.

Step 6: We prove here uniqueness of the representation. Let $(Y_0, Z), (Y'_0, Z') \in \mathbb{R} \times \mathcal{Z}$ be such that $\xi = Y_T^{Y_0, Z} = Y_T^{Y'_0, Z'}$. By following the lines of the verification argument in second part of the proof of the theorem, we obtain the equality $Y_t^{Y_0, Z} = Y_t^{Y'_0, Z'}$ using the fact that the value of the continuation utility of the market-maker satisfies

$$-\exp(-\gamma Y_t^{Y_0, Z}) = -\exp(-\gamma Y_t^{Y'_0, Z'}) = \operatorname{ess\,sup}_{\mathcal{L} \in \mathcal{A}} \mathbb{E}_t^{\mathcal{L}} \left[-\exp(-\gamma(PL_T^{\mathcal{L}} - PL_t^{\mathcal{L}} + \xi)) \right].$$

This in turn implies that $Z_t^{i,j,k} dN_t^{i,j,k} = Z_t'^{i,j,k} dN_t^{i,j,k}$ and $Z_t^{\tilde{S}} \sigma^2 dt = Z_t'^S \sigma^2 dt = d\langle Y, S \rangle_t, t \in [0, T]$. Thus $(Y_0, Z) = (Y'_0, Z')$.

We now prove the second part of Theorem 1.2. Let $\xi = Y_T^{Y_0, Z}$ with $(Y_0, Z) \in \mathbb{R} \times \mathcal{Z}$. We first show that for an arbitrary set of controls $\mathcal{L} \in \mathcal{A}$ we have $J_0^{\text{MM}}(\mathcal{L}, \xi) \leq -\exp(-\gamma Y_0)$ where we recall that $J_0^{\text{MM}}(\mathcal{L}, \xi)$ is such that $V_0^{\text{MM}}(\xi) = \sup_{\mathcal{L} \in \mathcal{A}} J_0^{\text{MM}}(\mathcal{L}, \xi)$. Then we will see that this inequality is in fact an equality when the corresponding Hamiltonian $h(\mathcal{L}, z, q)$ is maximized. Let us write

$$\begin{aligned} \bar{Y}_t &:= Y_t^{Y_0, Z} + \int_0^t \frac{\mathcal{T}}{2} \ell_u^{a,l} dN_u^{a,l} + \int_0^t \frac{\mathcal{T}}{2} \ell_u^{b,l} dN_u^{b,l} + \int_0^t Q_u dS_u + d[Q \cdot, S \cdot]_u \\ &\quad + \int_0^t \frac{\mathcal{T}}{2} \ell_u^{a,d} dN_u^{a,d,lat} + \int_0^t \frac{\mathcal{T}}{2} \ell_u^{b,d} dN_u^{b,d,lat}, \end{aligned}$$

with $t \in [0, T]$. An application of Itô's formula leads to

$$\begin{aligned} d(\exp(-\gamma \bar{Y}_t)) &= \gamma \exp(-\gamma \bar{Y}_{t-}) \left(-(Q_t + Z_t^{\bar{S}}) d\tilde{S}_t + (H(Z_t, Q_t) - h(\mathcal{L}, Z_t, Q_t)) dt \right. \\ &\quad - \sum_{(k^l, k^d) \in \mathcal{V}^l \times \mathcal{V}^d} \sum_{i \in \{a, b\}} \gamma^{-1} \left(\left(1 - \exp \left(-\gamma \left(Z_t^{i, l, k^l} + \ell_t^{i, l} \left(\frac{\mathcal{T}}{2} + \Gamma^l(\phi(i) Q_{t-} - \ell_t^{i, l}) \right) \right) \right) \right) \right) dN_t^{\mathcal{L}, i, l, k^l} \\ &\quad \left. - \sum_{\kappa \in K} \left(1 - \exp \left(-\gamma \left(Z_t^{i, d, k^d} + \ell_t^{i, d} \left(\frac{\mathcal{T}}{2} \phi^{\text{lat}}(\kappa) + \Gamma^d(\phi(i) Q_{t-} - \ell_t^{i, d}) \right) \right) \right) \right) \phi_t^d(i, \kappa) dN_t^{\mathcal{L}, i, d, k^d} \right). \end{aligned}$$

Therefore $\exp(-\gamma \bar{Y}_\cdot)$ is a $\mathbb{P}^{\mathcal{L}}$ -local submartingale. Thanks to Lemma 1.2, $\exp(-\gamma \bar{Y}_\cdot)$ is of class (D) hence a true submartingale. Doob-Meyer decomposition gives us that

$$\begin{aligned} \int_0^\cdot \gamma \exp(-\gamma \bar{Y}_{t-}) &\left(-(Q_t + Z_t^{\bar{S}}) d\tilde{S}_t \right. \\ &\quad - \sum_{(k^l, k^d) \in \mathcal{V}^l \times \mathcal{V}^d} \sum_{i \in \{a, b\}} \gamma^{-1} \left(\left(1 - \exp \left(-\gamma \left(Z_t^{i, l, k^l} + \ell_t^{i, l} \left(\frac{\mathcal{T}}{2} + \Gamma^l(\phi(i) Q_{t-} - \ell_t^{i, l}) \right) \right) \right) \right) \right) dN_t^{\mathcal{L}, i, l, k^l} \\ &\quad \left. - \sum_{\kappa \in K} \left(1 - \exp \left(-\gamma \left(Z_t^{i, d, k^d} + \ell_t^{i, d} \left(\frac{\mathcal{T}}{2} \phi^{\text{lat}}(\kappa) + \Gamma^d(\phi(i) Q_{t-} - \ell_t^{i, d}) \right) \right) \right) \right) \phi_t^d(i, \kappa) dN_t^{\mathcal{L}, i, d, k^d} \right) \end{aligned}$$

is a true martingale. Thus

$$\begin{aligned} J_0^{\text{MM}}(\mathcal{L}, \xi) &= \mathbb{E}^{\mathcal{L}} \left[-\exp(-\gamma \bar{Y}_T) \right] \\ &= -\exp(-\gamma Y_0) - \mathbb{E}^{\mathcal{L}} \left[\int_0^T \gamma \exp(-\gamma \bar{Y}_{t-}) (H(Z_t, Q_t) - h(\mathcal{L}, Z_t, Q_t)) dt \right] \\ &\leq -\exp(-\gamma Y_0). \end{aligned}$$

In addition to this, the previous inequality becomes an equality if and only if \mathcal{L} is chosen as the maximizer of the Hamiltonian h . In that case, $J_{\text{MM}}(\mathcal{L}, \xi) = -\exp(-\gamma Y_0)$. Finally we have that $V_0^{\text{MM}}(\xi) = -\exp(-\gamma Y_0)$ with optimal response $(\mathcal{L}_t^*)_{t \in [0, T]}$ defined by (OC).

1.A.3 Proof of Theorem 1.3

The aim of this proof is to provide formal justification for the numerical experiments. Namely, we want to show that the solution of the reduced equation (1.18) coincides with the value function of the exchange. By classic arguments, see for example Alvarez and Tourin [17], Hamadène and Morlais [146], the PDE (1.18) admits a unique viscosity solution, continuous in time, denoted by v .

Let $(t_0, \tilde{s}_0, \bar{n}_0, n_0, y_0) \in \mathfrak{D}$ where $\mathfrak{D} = [0, T] \times \mathbb{R} \times \mathbb{N}^{2(\#\mathcal{V}^l + \#\mathcal{V}^d)} \times \mathbb{R}$. We consider a test function $\Phi : \mathfrak{D} \rightarrow \mathbb{R}$ continuously differentiable in time, twice continuously differentiable with respect to s and y such that

$$\begin{aligned} 0 &= u(t_0, \tilde{s}_0, \bar{n}_0, n_0, y_0) - \Phi(t_0, \tilde{s}_0, \bar{n}_0, n_0, y_0) \\ &= \max_{(t, \tilde{s}, \bar{n}, n, y) \in \mathfrak{D}} \exp \left(-\eta \left(\sum_{(i, j) \in \{a, b\} \times \{l, d\}} \sum_{k \in \mathcal{V}^j} c^j \bar{n}^{i, j, k} - y \right) \right) \\ &\quad \left(v(t, q) - \Phi(t, \tilde{s}, \bar{n}, n, y) \exp \left(\eta \left(\sum_{(i, j) \in \{a, b\} \times \{l, d\}} \sum_{k \in \mathcal{V}^j} c^j \bar{n}^{i, j, k} - y \right) \right) \right). \end{aligned}$$

Therefore for all $(t, \tilde{s}, \bar{n}, n, y) \in \mathfrak{D}$

$$0 \geq v(t, q) - \Phi(t, \tilde{s}, \bar{n}, n, y) \exp\left(\eta\left(\sum_{(i,j) \in \{a,b\} \times \{l,d\}} \sum_{k \in \mathcal{V}^j} c^j \bar{n}^{i,j,k} - y\right)\right),$$

with equality at $(t_0, \tilde{s}_0, \bar{n}_0, n_0, y_0)$. Thus

$$0 = v(t_0, q_0) - \Phi(t_0, \tilde{s}_0, \bar{n}_0, n_0, y_0) = \max_{(t, \bar{n}) \in \mathfrak{D}} \left(v(t, q) - \Psi(t, \bar{n})\right),$$

where

$$\Psi(t, \bar{n}) := \Phi(t, \tilde{s}_0, \bar{n}, n_0, y_0) \exp\left(\eta\left(\sum_{(i,j) \in \{a,b\} \times \{l,d\}} \sum_{k \in \mathcal{V}^j} c^j \bar{n}^{i,j,k} - y_0\right)\right).$$

As v is the unique viscosity solution of (1.18), it is in particular a subsolution. Thus, for any $z \in \mathbb{R}^{2(\#\mathcal{V}^l + \#\mathcal{V}^d) + 1}$, Ψ satisfies

$$0 \geq \partial_t \Psi(t_0, \bar{n}_0) + \mathcal{U}(z, q_0, \mathcal{L}^*(z, q_0), \Psi(t_0, \cdot)),$$

with $q_0 := \sum_{j \in \{l,d\}} \sum_{k \in \mathcal{V}^j} (\bar{n}_0^{b,j,k} - \bar{n}_0^{a,j,k})$, \mathcal{U} is defined by (1.19) and \mathcal{L}^* is defined in Theorem 1.2. After computations, we deduce that

$$\begin{aligned} 0 \geq & \partial_t \Psi(t_0, \bar{n}_0) + \Psi(t_0, \bar{n}_0) \left(\frac{\eta}{2} \sigma^2 \gamma (z^{\tilde{S}} + q_0)^2 + \frac{\eta^2 \sigma^2}{2} (z^{\tilde{S}})^2 \right) + \sum_{(i,j) \in \{a,b\} \times \{l,d\}} \sum_{k \in \mathcal{V}^j} \lambda^{\mathcal{L}^*, i, j, k} \\ & \times \left(\exp\left(\eta(z^{i,j,k} - kc^j)\right) \Phi\left(t_0, \tilde{s}_0, \bar{n}_0^{i,j,k} + k, \bar{n}_0^{-(i,j,k)}, n_0, y_0\right) \exp\left(\eta\left(\sum_{(i,j) \in \{a,b\} \times \{l,d\}} \sum_{k \in \mathcal{V}^j} c^j \bar{n}_0^{i,j,k} - y_0\right)\right) \right. \\ & \left. - \Phi(t_0, \tilde{s}_0, \bar{n}_0, n_0, y_0) \left(1 + \eta \mathcal{E}(z^{i,j,k}, \ell^{*i,j}(z, q_0))\right) \right). \end{aligned}$$

Dividing on both sides of the equation by $\exp\left(\eta\left(\sum_{(i,j) \in \{a,b\} \times \{l,d\}} \sum_{k \in \mathcal{V}^j} c^j \bar{n}_0^{i,j,k} - y_0\right)\right) > 0$, we obtain

$$\begin{aligned} 0 \geq & \partial_t \Phi(t_0, \tilde{s}_0, \bar{n}_0, n_0, y_0) + \Phi_0 \left(\frac{\eta}{2} \sigma^2 \gamma (z^{\tilde{S}} + q_0)^2 + \frac{\eta^2 \sigma^2}{2} (z^{\tilde{S}})^2 \right) + \sum_{(i,j) \in \{a,b\} \times \{l,d\}} \sum_{k \in \mathcal{V}^j} \lambda^{\mathcal{L}^*, i, j, k} \\ & \times \left(\exp\left(\eta(z^{i,j,k} - kc^j)\right) \Phi\left(t_0, \tilde{s}_0, \bar{n}_0^{i,j,k} + k, \bar{n}_0^{-(i,j)}, n_0, y_0\right) - \Phi_0 \left(1 + \eta \mathcal{E}(z^{i,j,k}, \ell^{*i,j}(z, q_0))\right) \right), \end{aligned}$$

where $\Phi_0 := \Phi(t_0, \tilde{s}_0, \bar{n}_0, n_0, y_0)$. Therefore, u is a viscosity subsolution of (1.16). A similar argument holds to prove that u is also a viscosity supersolution of (1.16). Consequently, u is a viscosity solution of (1.16). The uniqueness of u follows from an application of [201, Theorem II.3], together with the continuity of v in time for a fixed q . Thus, we deduce that $v_0^E = u(0, \tilde{S}_0, \bar{N}_0, N_0, Y_0) = v(0, Q_0)$.

Chapter 2

Optimal make-take fees in a multi-market-maker environment

2.1 Introduction

Optimal market-making has been a topic of interest in mathematical finance since the seminal work Avellaneda and Stoikov [26]. market-makers are liquidity providers, who post limit orders on the bid and ask sides of the order book of an underlying asset, available on an exchange. They buy and sell simultaneously, earning the spread between their quotes and the mid-price, and have to dynamically manage their inventory, thus skew their quotes depending on their position. First simple market-making problems are addressed in Avellaneda and Stoikov [26], Ho and Stoll [156], where the authors use a stochastic control approach. Later in Guéant, Lehalle, and Fernandez-Tapia [142], an explicit solution is provided by imposing an inventory threshold for the market-maker. A vast literature has emerged from these articles, and various extensions have been studied, see for example Cartea and Jaimungal [67], Guéant [137]. All these models deal with the case of a market-making activity with no maker taker fees policy from the exchange.

Due to the fragmentation of financial markets, exchanges (Nasdaq, Euronext) are in competition and therefore need to find innovative ways to attract liquidity on their platforms. One of these ways is the use of a maker taker fees system: the exchange typically associates a fee rebate to executed limit orders, while charging a transaction cost for market orders. This enables it to subsidize liquidity provision and tax liquidity consumption. The problem of a relevant make take fees policy is therefore key for the quality of market liquidity and for the revenue of the corresponding exchange platform. In El Euch, Mastrolia, Rosenbaum, and Touzi [109], an optimal make take fees policy is derived, depending on the transaction flow generated by the market-maker.

Either with or without the intervention of an exchange, no optimal market-making framework addresses the issue of several market-makers competing with each other. Such consideration is of crucial importance for several reasons. The case of a single market-maker means that he has no competitor, hence only needs to manage his inventory risk. This assumption is not very realistic as there are usually several liquidity providers, see for example Mounjid, Rosenbaum, and Saliba [219]. Such models lead to an overestimation of the order flow that a market-maker will process during the trading period. Many components of financial markets, spread and order book shape for example, are linked to competition between market-makers, see Dayri and Rosenbaum [91], Glosten and Milgrom [131], Huang, Rosenbaum, and Saliba [165], Madhavan, Richardson, and Roomans [203]. Considering several market-makers is of particular interest for the exchange as it has access to information related to the identity of the agents involved in each transaction. For those reasons, in the spirit of El Euch, Mastrolia, Rosenbaum, and Touzi [109], we extend the framework of make take fees problem to the case

of an exchange (or of a regulator) wishing to attract liquidity on a market, with several market-makers trading on a single underlying asset.

From a modelling point of view, this chapter follows the same inspiration as the literature mentioned above. Here, we aim at studying the contracting problem of an exchange and several market-makers who trade on a single underlying asset. We place ourselves in a principal-agent framework similar to Holmström and Milgrom [159]. The principal wants to build a contract for the agents, enabling him to define an optimal make take fees policy, in order to maximize his revenue. market-takers send bid and ask orders, of constant volume equal to one, whereas market-makers control their quotes on the asset (no notion of volume of limit orders is considered here). We base our contract on the transaction flow, as well as the asset price. Note that, as in El Euch, Mastrolia, Rosenbaum, and Touzi [109], the spread of the agents is not a contractible variable. It is indeed preferable to only consider variables involved in actual transactions. Furthermore, the very definition of individual spreads is ambiguous, as in practice market-makers operates in a limit order book where they post several orders with different volumes. The market orders are executed by the market-maker with the best quotes. However, other market participants receive a compensation depending on the distance between their quotes and the best bid or ask. By doing so, we aim at modelling queue position affecting the order book, see Glosten and Milgrom [131], Huang, Rosenbaum, and Saliba [165], Moallemi and Yuan [217]. As we assume constant volume, this represents the fact that an order may be splitted between agents depending on their position in the order book. Furthermore, this compensation reflects the importance for market-makers to get tight enough quotes for commercial reasons, see Section 2.3 for details. Consequently, the PnL process of each market-maker depend on the quotes of the others. Finally, in the single market-maker framework of El Euch, Mastrolia, Rosenbaum, and Touzi [109], intensity of arrival orders are a decreasing function of the agent's spread. In our case, the intensity increases with total liquidity on the market. To do so we represent the whole liquidity of the order book using a weighted sum of the spreads depending on their distance with respect to an efficient price.

As in El Euch, Mastrolia, Rosenbaum, and Touzi [109], our problem is addressed by solving a Stackelberg game between the exchange and the market-makers:

- (i) Following the approach in Élie and Possamaï [110], each market-maker computes his best-reaction function given spreads from other agents, which provides a Nash equilibrium.
- (ii) Given this equilibrium, the exchange computes his optimal response (the optimal contract) by solving the appropriate Hamilton-Jacobi-Bellman equation.
- (iii) This response is re-injected in the agents' optimal quotes, which give the optimal answer of both parties.

One can retrieve optimal quotes in a semi explicit form through a partial differential equation. Moreover, the optimal contract¹ is expressed as a sum of stochastic integrals with respect to market order and efficient price processes. We emphasize that such contracting scheme can be readily implemented and is easily interpretable, see Section 2.3.3.

An important finding is that an increase of the number of market-makers does not necessarily decrease the average spread. It means that there is, for a given set of market parameters, an "optimal" number of market-makers in the sense of PnL maximisation for the platform. It therefore provides, for an exchange, a framework to decide how many market participants they wish to attract in order to increase their profit. We also provide a simple formula to choose the "taker cost", in the same spirit as in El Euch, Mastrolia, Rosenbaum, and Touzi [109]. This is discussed through Section 2.4.4. We emphasize the fact that, to our knowledge, we provide the first model using a framework à la Avellaneda and Stoïkov dealing with the multi market-maker problem. One of our main contribution

¹As would be done in most cases in practice, we impose that every market-maker receives the same contract. This means that it can not depend on individual risk aversion parameters.

from this chapter lies in the possibility to analyse the impact of adding market-makers in a market on quantities of interest such as trading cost, total order flow, and PnL of the platform. We also see that adding a market participant with a higher risk-aversion parameter will decrease the average spread and conversely. Moreover, decreasing the taker cost when the number of market-makers increase leads to a higher PnL for the platform up to a certain point (see Section 2.4.4).

We organize the chapter as follows: in Section 2.2, we introduce some preliminaries on stochastic calculus for the main objects of the model. Then, we present the way we model the multi market-maker case, and the key differences with El Euch, Mastrolia, Rosenbaum, and Touzi [109]. We also describe the market-makers and the exchange's optimisation problem. In Section 2.3, we give the best reaction functions of each market-makers for a given contract and define the form of admissible contracts. In Section 2.4, we solve explicitly the exchange's problem and provide the form of the incentives given to each market-maker. Finally, in Section 2.5, we discuss the impact of the presence of several market-makers on market liquidity and PnL of the platform.

Throughout the chapter, we fix a constant $\delta_\infty > 0$, which is assumed to be sufficiently large (how large exactly will be made specific later on).

2.2 The model

2.2.1 Framework

2.2.1.1 Canonical process

As in El Euch, Mastrolia, Rosenbaum, and Touzi [109], the framework considered throughout this chapter is inspired by the seminal works Avellaneda and Stoikov [26] and Guéant, Lehalle, and Fernandez-Tapia [142] where there is no exchange acting on the market. Let $N \geq 1$ be an integer representing the number of market-makers in the market. We also define a positive constant δ_∞ , which is assumed to be large enough, a statement that will be made precise later on. We consider a final horizon time $T > 0$, and the space $\Omega =: \Omega_c \times \Omega_d^{2N}$, with Ω_c the set of continuous functions from $[0, T]$ into \mathbb{R} , and Ω_d the set of piecewise constant càdlàg functions from $[0, T]$ into \mathbb{N} . We consider Ω as a subspace of the Skorokhod space $\mathcal{D}([0, T], \mathbb{R}^{2N+1})$ of càdlàg functions from $[0, T]$ into \mathbb{R}^{2N+1} , and let \mathcal{F} be the trace Borel σ -algebra on Ω , where the topology is the one associated to the usual Skorokhod distance on $\mathcal{D}([0, T], \mathbb{R}^{2N+1})$. We let $(\mathcal{X}_t)_{t \in [0, T]} := (W_t, N_t^{1,a}, \dots, N_t^{N,a}, N_t^{1,b}, \dots, N_t^{N,b})_{t \in [0, T]}$ be the canonical process on Ω , that is to say

$$W_t(\omega) := w(t), \quad N_t^{i,a}(\omega) := n^{i,a}(t), \quad N_t^{i,b}(\omega) := n^{i,b}(t), \quad \text{for all } i \in \{1, \dots, N\},$$

with $\omega := (w, n^{1,a}, \dots, n^{N,a}, n^{1,b}, \dots, n^{N,b}) \in \Omega$. The following aggregated counting processes will also be useful

$$N^j := \sum_{i=1}^N N^{i,j}, \quad j \in \{a, b\}.$$

Finally, define for $i \in \{1, \dots, N\}, j \in \{a, b\}$ the maps $\lambda^{i,j} : \mathbb{R}^N \times \mathbb{Z}^N \rightarrow \mathbb{R}$ and $\lambda : \mathbb{R}^N \times \mathbb{Z}^N \rightarrow \mathbb{R}$ by

$$\begin{aligned} \lambda^{i,j}(x, q) &:= A \exp \left(-\frac{k}{\sigma} \left(c + \varpi \sum_{i=1}^N x^i \mathbf{1}_{\{x^i = \underline{x}\}} + \sum_{i=1}^N \sum_{\ell=1}^K H_\ell x^i \mathbf{1}_{\{x^i \in K_\ell\}} \mathbf{1}_{\{x^i \neq \underline{x}\}} \right) \right) \frac{\mathbf{1}_{\{x^i = \underline{x}, q^i > -\phi(j)\bar{q}\}}}{\sum_{\ell=1}^N \mathbf{1}_{\{x^\ell = \underline{x}, q^\ell > -\phi(j)\bar{q}\}}}, \\ \lambda^j(x, q) &:= \sum_{i=1}^N \lambda^{i,j}(x, q), \end{aligned} \tag{2.1}$$

where A, k, σ , and c are fixed positive constants,² K is a fixed positive integer, $\bar{q} \in \mathbb{N}$, ϖ and

$(H_\ell)_{\ell=1,\dots,K}$ are real-valued and will be fixed later, and $\{K_\ell : \ell = 1, \dots, K\}$ is an open covering of the interval $[0, \delta_\infty]$. Moreover

$$\phi(j) := \begin{cases} 1, & \text{if } j = a, \\ -1, & \text{if } j = b. \end{cases}$$

Remark 10. *The maps $\lambda^{i,j}$ are here to define the intensity of the point processes $N^{i,j}$, whereas λ plays this role for the aggregated point processes. It is a generalisation of the exponential intensity used in El Euch, Mastrolia, Rosenbaum, and Touzi [109], Guéant, Lehalle, and Fernandez-Tapia [142], in the sense that*

- all spreads are taken into account, weighted with respect to their value;
- when $N = 1$ and $\varpi = 1$, we recover the intensity for the single market-maker case.

2.2.1.2 Admissible controls, inventory process and canonical probability measure

We define the probability \mathbb{P}^0 on (Ω, \mathcal{F}) such that under \mathbb{P}^0 , W , $N^{i,a}$, and $N^{i,b}$ are independent for all $i = 1, \dots, N$, W is a one-dimensional Brownian motion, $N^{i,a}$ and $N^{i,b}$ are Poisson processes with intensity $\lambda^{i,a}(0, 0)$, and $\lambda^{i,b}(0, 0)$ respectively.³ In other words, \mathbb{P}^0 is simply the product measure of the Wiener measure on Ω_c and the unique measure on Ω_d^{2N} that makes the canonical process there into an homogeneous Poisson process with the prescribed intensity. We therefore endow the space (Ω, \mathcal{F}) with the (\mathbb{P}^0 -augmented) canonical filtration $\mathbb{F} := (\mathcal{F}_t)_{t \in [0, T]} := (\mathcal{F}_t^c \otimes (\mathcal{F}_t^d)^{\otimes d})_{t \in [0, T]}$ generated by $(\mathcal{X}_t)_{t \in [0, T]}$. It is well-known that the filtration \mathbb{F} satisfies the usual conditions and the Blumenthal 0–1 law. All notions of measurability for processes, unless otherwise stated, should be understood as being associated to \mathbb{F} .

The market-maker has a view on the efficient price (which should be understood as the mid-price) of the asset given by $(S_t)_{t \in [0, T]}$, defined as

$$S_t := S_0 + \sigma W_t, \quad t \in [0, T], \tag{2.2}$$

where $S_0 > 0$ is the known initial value of the price, and $\sigma > 0$ is its volatility.

Remark 11. *The use of an arithmetic Brownian motion for the efficient price process is motivated by its simplicity. The price (2.2) can reach negative values with non-negligible probability only on a sufficiently large time horizon T . For practical purposes, we choose $T < 1$ day so that Equation (2.2) approximates accurately an efficient price driven by a geometric Brownian, which stays positive almost surely.*

Next, we define the properties of the controlled processes of the agents. Based on their view on the efficient price (2.2), market-makers offer bid and ask quotes on the underlying asset. Such prices are defined by

$$P_t^{i,b} := S_t - \delta_t^{i,b}, \quad P_t^{i,a} := S_t + \delta_t^{i,a}, \quad t \in [0, T], \quad i \in \{1, \dots, N\},$$

where the superscript b (resp. a) accounts for bid (respectively ask). The set of admissible controls for the market-makers is therefore defined as

$$\mathcal{A} := \left\{ (\delta_t)_{t \in [0, T]} = (\delta_t^{i,a}, \delta_t^{i,b})_{t \in [0, T], i=1, \dots, N} : \mathbb{R}^{2N}\text{-valued and predictable processes bounded by } \delta_\infty \right\}. \tag{2.3}$$

The predictability of the spreads reflects the fact that each agent chooses in advance his quotes. The boundedness assumption here is technical and simply helps us to define the associated probability measures. As the optimal contract we will derive later on leads naturally to bounded spreads, this is actually without loss of generality, provided the bounds are chosen large enough⁴.

²See Section 2.2.1.4 for economical and mathematical interpretation of (2.1).

³As a direct consequence, N^a and N^b are Poisson processes with intensity $\lambda(0)$

⁴See Lemma 2.7 for the prescribed value of δ_∞ .

A market-maker manages both the spreads and his inventory process. For the i -th agent, a filled bid order, represented by $N^{i,b}$, increase its inventory by one unit, and conversely for an ask order. It leads to the following definition of an inventory process of a market-maker

$$Q_t^i := N_t^{i,b} - N_t^{i,a}, \quad t \in [0, T], \quad i \in \{1, \dots, N\}.$$

Remark 12. *Given the form of the intensities (2.1), the i -th agent will see his inventory changing only if he quotes a spread such that $\delta^{i,b} = \underline{\delta}^b$, or $\delta^{i,a} = \underline{\delta}^a$. Such quotations are called the best bid, and best ask spread respectively.*

The term \bar{q} defined in (2.1) acts as a critical absolute inventory, which is the same for each agent. Assume the i -th market-maker cross this threshold on the bid side, then $\lambda^{i,b}(\delta, q) = 0$ and he can only receive ask orders to decrease his inventory below \bar{q} .

2.2.1.3 Change of probability measure

Given our technical assumptions, we introduce for any $\delta \in \mathcal{A}$ a new probability measure \mathbb{P}^δ on (Ω, \mathcal{F}) under which S follows (2.2) and

$$\tilde{N}_t^{\delta,i,a} := N_t^{i,a} - \int_0^t \lambda^{i,a}(\delta_r^a, Q_t) dr, \quad \tilde{N}_t^{\delta,i,b} := N_t^{i,b} - \int_0^t \lambda^{i,b}(\delta_r^b, Q_t) dr, \quad t \in [0, T], \quad i \in \{1, \dots, N\}, \quad (2.4)$$

are martingales. This probability measure is defined by the corresponding Doléans-Dade exponential

$$L_t^\delta := \exp \left(\sum_{i=1}^N \sum_{j \in \{a,b\}} \int_0^t \log \left(\frac{\lambda^{i,j}(\delta_r^j, Q_r)}{A} \right) dN_r^{i,j} - \int_0^t (\lambda^{i,j}(\delta_r^j, Q_r) - A) dr \right), \quad (2.5)$$

where $Q := (Q^1, \dots, Q^N)^\top$. By direct application of Itô's formula, and the uniform boundedness of δ^a , and δ^b , this local martingale satisfies the Novikov-type criterion given in Sokol [253], and thus is a martingale.

Remark 13. *By definition, the compensated aggregated point processes*

$$\tilde{N}_t^{\delta,a} := N_t^a - \int_0^t \lambda(\delta_r^a) dr, \quad \tilde{N}_t^{\delta,b} := N_t^b - \int_0^t \lambda(\delta_r^b) dr, \quad t \in [0, T],$$

are also martingales under \mathbb{P}^δ .

We can therefore define the Girsanov change of measure with $\frac{d\mathbb{P}^\delta}{d\mathbb{P}^0} = L_T^\delta$ (see for instance [170, Theorem III.3.1]). In particular, all the probability measures \mathbb{P}^δ indexed by $\delta \in \mathcal{A}$ are equivalent. The notation a.s., for almost surely, can be used without ambiguity. Throughout the chapter, we write \mathbb{E}_t^δ for the conditional expectation with respect to \mathcal{F}_t under the probability measure \mathbb{P}^δ .

Hence, the arrival of ask (resp. bid) market orders for the i -th market-maker is represented by the point process $(N_t^{i,a})_{t \in [0, T]}$ (resp. $(N_t^{i,b})_{t \in [0, T]}$) of intensity $(\lambda^{i,a}(\delta_t^a))_{t \in [0, T]}$ (resp. $(\lambda^{i,b}(\delta_t^a))_{t \in [0, T]}$) and the total arrival of ask (resp. bid) market orders is represented by the point process $(N_t^a)_{t \in [0, T]}$ (resp. $(N_t^b)_{t \in [0, T]}$) of intensity $(\lambda(\delta_t^a))_{t \in [0, T]}$ (resp. $(\lambda(\delta_t^b))_{t \in [0, T]}$).

2.2.1.4 Interpretations

First, we comment on the shapes of the intensities in (2.1). The intensity of buy (resp. sell) market order arrivals depends on the extra cost of each trade paid by the market-taker compared to the efficient price. This extra cost is the sum of the spread $\underline{\delta}^b$ (resp. $\underline{\delta}^a$) imposed by the market-maker who is currently trading at the best bid (resp. best ask), and the transaction cost $c > 0$ collected by the exchange. Moreover, following classical financial economics results, the average number of trades

per unit of time is a decreasing function of the ratio between the spread and the volatility (see Dayri and Rosenbaum [91], Madhavan, Richardson, and Roomans [203], and Wyart, Bouchaud, Kockelkoren, Potters, and Vettorazzo [271]).

The intensity of order arrivals depends also on the market liquidity, namely the spread quoted by all market participants. Hence, the spread of the i -th market-maker is weighted by a constant H_ℓ , $\ell \in \{1, \dots, K\}$. Such constant will be chosen later on. For the moment, note that it is a decreasing function of ℓ . Hence, a small spread corresponds to a high weight and conversely. Recall that in our model, we make the approximation that we can only have orders of size 1. Hence, an increase of the intensity represents the fact that we can send bigger orders (many orders of size 1 corresponds to one large order).

Such weights depends on the open covering $\{K_\ell : \ell \in \{1, \dots, K\}\}$, introduced in Subsection 2.2.1.1. Several forms can be chosen (thinner intervals around the first Tick for instance), and throughout this chapter we use the following definition

$$K_\ell := ((\ell - 1)\text{Tick}, \min(\ell\text{Tick}, \delta_\infty)), \ell \in \{1, \dots, K\}. \quad (2.6)$$

Note that we can choose $K = 1$, which leads to a unique zone $K_1 = [0, \delta_\infty]$. In that case, the penalisation on the intensity of arrival orders is the same for all $\delta \neq \min_{i=1, \dots, N} \delta^i$. A larger number of intervals K is used in order to have a penalisation increasing with respect to the value of the spread of the market-makers. We will see that at the optimum, the covering has no impact on the PnL of both principal and agents, due to the form of both PnL process and intensity.

Remark 14. Equation (2.6) indicates that the open covering is not a dynamic function of the vector of spreads. Hence, we add the indicator function $\mathbf{1}_{\{x_i \neq \underline{x}\}}$ to ensure that the coordinate \underline{x} is not associated to a weight $\varpi + H_\ell$, for $\ell \in \{1, \dots, K\}$, but only to ϖ . In addition to this, fixing the open covering makes the model more tractable from a numerical point of view. Otherwise, at each time step, the corresponding areas should be computed again.

2.2.2 Market-makers' problem

2.2.2.1 PnL process of the agents

First, we define the PnL process of the i -th market-maker when the market-makers play $\delta \in \mathcal{A}$ as

$$PL_t^{\delta, i} := \int_0^t \sum_{j \in \{a, b\}} \left(\delta_s^{i, j} \mathbf{1}_{\{\delta_s^{i, j} = \delta^j\}} + \sum_{\ell=1}^K \omega_\ell \delta_s^{i, j} \mathbf{1}_{\{\delta_s^{i, j} \in K_\ell\}} \right) dN_s^j + \int_0^t Q_s^i dS_s, \quad t \in [0, T]. \quad (2.7)$$

The first integral corresponds to the cash flow process, whereas the other represents the inventory risk process of the i -th agent, and $\omega_\ell \in (0, 1)$, $\ell \in \{1, \dots, K\}$ are weights that decrease toward zero as ℓ increases. The market-maker is remunerated with an increasing fraction of his quote when he is near from the best spread. Such remuneration decreases as the spread quoted moves away from the best spread. This is an incentive for the agents to quote a lower spread in general, and represents the fact that an order may be splitted between agents depending on their position in the order book.

Such form of incentive is particularly well suited to certain markets where there are hundreds of market-makers, and a selection of market participants has to be done by the platform. Such selection is required by some clients, who ask for a specific number of market-maker to ensure competition on the exchange. Moreover, good position on external rankings attracts new clients on the platform. To do so, a selection criteria is the price quality provided by market-makers.

Moreover, we assume no partial execution for the agents. Hence, if two market-makers are playing the best spread, they are both fully executed, in the sense that they receive the whole trade.

Remark 15. Whereas the market-maker is remunerated with a portion of his spread when trading in a certain area compared to the best bid-ask, it is not reflected in his inventory process. Indeed, the terms $\omega_\ell \delta^{i,j}$ accounts only for the cash process: the market-maker's inventory does not move if he does not quote at the best bid-ask.

2.2.2.2 Best reaction functions

Equation (2.7) represents the PnL of an agent in the absence of contract from the exchange. Following the principal-agent approach, the exchange proposes a remuneration ξ^i defined by an \mathcal{F}_T -measurable random variable to each market-maker, in addition to their PnL process. These aim at creating an incentive to attract liquidity on the platform by reducing the market-makers' spreads. We will prove a certain representation theorem for the set of admissible contracts.

Given their contracts, the agents are facing a stochastic differential game, since the relative rankings of their spreads directly impact whether their orders are executed or not. Since we will later be looking for Nash equilibria for this game, we consider that the optimisation problem of the i -th market-maker is a function of the spread vector δ^{-i} quoted by the $N - 1$ other agents. Hence, each agent will maximize his PnL given actions of the other market-makers to obtain his so-called best-reaction function. By denoting $\gamma_i > 0$ the risk-aversion of the i -th market-maker, and using a CARA utility function $U_i(x) := -e^{-\gamma_i x}$, $x \in \mathbb{R}$, we are left with the following maximisation problem

$$V_{\text{MM}}^i(\xi^i, \delta^{-i}) := \sup_{\delta^i \in \mathcal{A}^i(\delta^{-i})} \mathbb{E}^{\mathbb{P}^{\delta \otimes_i \delta^{-i}}} \left[U_i \left(\xi^i + \sum_{j \in \{a,b\}} \int_0^T \delta_t^{i,j} \left(\mathbf{1}_{\{\delta_t^{i,j} = \delta_t^j \otimes_i \delta_t^{j,-i}\}} + \sum_{\ell=1}^K \int_0^T \omega_\ell \mathbf{1}_{\{\delta_t^{i,j} \in K_\ell\}} \right) dN_t^j + \int_0^T Q_t^i dS_t \right) \right],$$

where $\mathcal{A}^i(\delta^{-i}) := \{\delta : \delta \otimes_i \delta^{-i} \in \mathcal{A}\}$.

To ensure that this quantity is not degenerate, for $i \in \{1, \dots, N\}$, we assume that for all $\delta \in \mathcal{A}$

$$\mathbb{E}^\delta [\exp(-\gamma' \xi^i)] < +\infty, \text{ for some } \gamma' > \max(\gamma_1, \dots, \gamma_N). \quad (2.8)$$

We call \mathbb{R}^N -valued \mathcal{F}_T -measurable random variables ξ satisfying (2.8) contracts. The integrability condition ensures that the market-maker's problem is well-defined (that is the sup remains finite). As the N market-makers play simultaneously, we are looking for a Nash equilibrium resulting of the interactions between agents. We now provide the appropriate definition of such an equilibrium

2.2.2.3 Nash equilibrium

A Nash equilibrium is a set of admissible controls such that each market-maker has no interest in deviating from its current position, given a contract offered by the principal. It is formalized with the following definition.

Definition 2.1. For a given contract ξ , a Nash equilibrium for the N agents is a set of actions $\hat{\delta}(\xi) \in \mathcal{A}$ such that for all $i \in \{1, \dots, N\}$

$$V_{\text{MM}}^i(\xi^i, \hat{\delta}^{-i}(\xi)) = \mathbb{E}^{\mathbb{P}^{\hat{\delta}(\xi)}} \left[U_i \left(\xi^i + \sum_{j \in \{a,b\}} \int_0^T \hat{\delta}_t^{i,j}(\xi) \left(\mathbf{1}_{\{\hat{\delta}_t^{i,j}(\xi) = \hat{\delta}_t^j(\xi)\}} + \sum_{\ell=1}^K \int_0^T \omega_\ell \mathbf{1}_{\{\hat{\delta}_t^{i,j}(\xi) \in K_\ell\}} \right) dN_t^j + \int_0^T Q_t^i dS_t \right) \right]. \quad (2.9)$$

We introduce for any contract ξ the set $\text{NA}(\xi)$ of all associated Nash equilibria. This set is of particular importance as it will be imposed to be non-empty, to ensure the existence of, at least, one Nash equilibrium. We now turn to the exchange's contracting problem.

2.2.3 The exchange optimal contracting problem

In our framework, the exchange is compensated by a fixed amount $c > 0$ for each market order that occurs in the market. As in El Euch, Mastrolia, Rosenbaum, and Touzi [109], since we are anyway working on a short time interval, we take c independent of the price of the asset. The goal of the exchange is to maximize the total number of aggregated market orders $N_T^a + N_T^b$ arriving during the time interval $[0, T]$. As the arrival intensities are only controlled by the market-makers, the contract vector ξ aim at increasing these intensities, which are decreasing functions of the spreads. Hence, the exchange will pay to each market-maker this contract at time T and the form of his PnL, using a CARA utility function, is given by

$$-\exp\left(-\eta(c(N_T^a + N_T^b) - \xi \cdot \mathbf{1}_N)\right),$$

where $\eta > 0$ denote the risk aversion parameter of the principal.

We now provide a suitable definition of the set of admissible contracts offered by the exchange. First, we need to ensure that the problem of the exchange does not degenerate. Hence, we assume that, for all $\delta \in \mathcal{A}$

$$\mathbb{E}^\delta[\exp(\eta' N \xi^i)] < +\infty, \text{ for some } \eta' > \eta. \quad (2.10)$$

Since N^a and N^b are point processes with bounded intensities, this condition, together with Hölder's inequality, ensure that the problem of the exchange is well-defined.⁵ We also assume that the market-makers only accept contracts ξ^i such that their maximal utility $V_{\text{MM}}^i(\xi^i, \hat{\delta}^{-i})$, taken at a Nash equilibrium $\hat{\delta} \in \text{NA}(\xi)$, is above a threshold value $R_i < 0$. This value is known as the reservation utility of the i -th agent, and leads to the following definition.

Definition 2.2. *The set of admissible contracts \mathcal{C} is defined as the set of \mathbb{R}^N -valued, \mathcal{F}_T -measurable random variables $\xi := (\xi^1, \dots, \xi^N)^\top$, such that for all $i \in \{1, \dots, N\}$, (2.8) and (2.10) hold, and the participation constraints of all agents are satisfied for at least one Nash equilibrium in $\text{NA}(\xi)$ (which is then automatically non-empty).*

In the set of admissible contracts, the participation constraints of the agents are satisfied for at least one Nash equilibrium generated by ξ . As we anticipate that the participation constraints will be binding for any optimal contract ξ , the agents are indifferent between the possible different Nash equilibria generated by ξ . This means that we can use the same convention as the one used in the classical principal-agent literature, where the principal has enough bargaining power to impose to the agents which equilibrium he wants them to use. His optimisation problem is thus written as

$$V_0^E := \sup_{\xi \in \mathcal{C}} \sup_{\hat{\delta} \in \text{NA}(\xi)} \mathbb{E}^{\hat{\delta}(\xi)} \left[-\exp\left(-\eta(c(N_T^a + N_T^b) - \xi \cdot \mathbf{1}_N)\right) \right]. \quad (2.11)$$

Now that we have properly defined the two problems of the Stackelberg game, we can move towards the resolution of the market-maker's problem. Before solving this two-steps problem, we first sketch the approach we undertake.

2.2.4 Stackelberg games in a nutshell

Each market-maker has an optimisation problem which depends on the control processes of the $N - 1$ others agents, and on the contract given by the principal. Hence, solving the i -th agent's problem is done by searching the best reaction functions of each market-makers, given a set of actions δ^{-i} of the other agents. Hence, the spreads quoted by every agent will both depend on the incentives given by the principal, and the spreads of opponents.

⁵We will see in the verification Theorem 2.2 that such a condition is required for a uniform integrability type argument

As stated before, the market-makers fix their quotes simultaneously, so they must agree on an equilibrium between their reaction functions. To solve this problem, we will make use of an equivalent definition of a Nash equilibrium, given in Élie and Possamaï [110] and recalled in the next section. In particular, there is a direct link between the existence of a Nash equilibrium and a solution to a multidimensional system of BSDEs.

The key point is that, for a specific choice of weights H_ℓ in (2.1), the existence and uniqueness of a Nash equilibrium is direct, because of two important facts. First, the use of indicator functions $\mathbf{1}_{\{\delta^{i,j}=\delta^j\}}$, and $\mathbf{1}_{\{\delta^{i,j}\neq\delta^j\}}$ acts as a decoupling effect on the agents' Hamiltonian. Second, such effect can be achieved only in the case of a restriction to a specific form of contracts. This restriction will be explained and commented in Section 2.3. For the moment, note that it enables to compute explicitly a unique Nash equilibrium for the market-maker's problem.

Given that admissible contracts generate at least one Nash equilibrium, the principal solve his optimisation problem by choosing the incentives given to each market-maker, as a result of its associated HJB equation. This provides explicitly the optimal quotes of the agents, and solve the two steps Stackelberg game.

2.3 Solving the market-maker's problem

We start by solving the problem of the i -th market-maker facing an arbitrary admissible contract proposed by the exchange. This section is mainly devoted to Theorem 2.1. First, we introduce a certain form of contracts proposed by the principal to the i -th agent. This \mathcal{F}_T -measurable random variable takes the form of the terminal condition to a specific BSDE, although we do not use this theory to solve the problem. We then prove that this is the only form of contracts that can be proposed to the market-makers. Then, given this specific form, we derive the optimal response of each agents, other actions being fixed.

2.3.1 Preliminaries

For notational simplicity, let us define $\mathcal{R} := \mathbb{R}^N \times \mathbb{R}^N \times \mathbb{R}$, $\mathcal{B}_\infty := [-\delta_\infty, \delta_\infty]$.

Definition 2.3. Fix some $i \in \{1, \dots, N\}$. For any $(d^i, d^{-i}, z^i, q) \in \mathcal{B}_\infty^2 \times \mathcal{B}_\infty^{2(N-1)} \times \mathcal{R} \times \mathbb{Z}^N$, where we have $z^i := ((z^{i,j,a})_{j=1,\dots,N}, (z^{i,j,b})_{j=1,\dots,N}, z^{S,i})$, and $d^i := (d^{i,a}, d^{i,b})$, the Hamiltonian of the i -th agent is defined by

$$H^i(d^{-i}, z^i, q) := \sup_{d^i \in \mathcal{B}_\infty^2} h^i(d^i, d^{-i}, z^i, q), \quad (2.12)$$

where

$$h^i(d^i, d^{-i}, z^i, q) := \sum_{\ell=1}^N \sum_{j \in \{a,b\}} \gamma_i^{-1} \left(1 - \exp \left(-\gamma_i \left(z^{i,\ell,j} + d^{i,j} \mathbf{1}_{\{d^{i,j}=\underline{d}^j \otimes_i d^{j,-i}\}} + \sum_{k=1}^K \omega_k d^{i,j} \mathbf{1}_{\{d^{i,j} \in K_k\}} \right) \right) \right) \lambda^{\ell,j}(d^j, q).$$

As in every stochastic control problem, such quantity is of particular importance. It is naturally derived from an application of Itô's formula to $e^{-\gamma_i Y_t}$, for $i = 1, \dots, N$ where Y is defined by (2.14). Then, maximising this quantity will give the optimal spreads quoted by the i -th market-maker, given the spreads quoted by the other agents. The form of the maps $(h^i)_{i=1,\dots,N}$ requires that a maximizer in (2.12) is a function of d^{-i} , z^i , and q . This suggests a proper definition of a fixed point of the Hamiltonian vector $(H^i(d^{-i}, z^i, q))_{i=1,\dots,N}$.

Definition 2.4. For every $(z, q) \in \mathcal{R}^N \times \mathbb{Z}^N$ a fixed point of the Hamiltonian is defined by a matrix $\delta^*(z, q) \in \mathcal{M}_{N,2}(\mathbb{R})$ such that for any $1 \leq i \leq N$

$$\delta^{*i}(z, q) \in \operatorname{argmax}_{\delta^i \in \mathcal{B}_\infty^2} h^i(\delta^i, \delta^{*-i}, z^i, q). \quad (2.13)$$

For every $(z, q) \in \mathcal{R}^N \times \mathbb{Z}^N$, we denote by $\mathcal{O}(z, q)$ the set of all fixed points.

We need the following standing technical assumption.

Assumption 2.1. *There exists at least one Borel-measurable map $\delta^* : \mathcal{R}^N \times \mathbb{Z}^N \rightarrow \mathcal{M}_{N,2}(\mathbb{R})$ such that for every $(z, q) \in \mathcal{R}^N \times \mathbb{Z}^N$, $\delta^*(z, q) \in \mathcal{O}(z, q)$. The corresponding set of maps is denoted by \mathcal{O} .*

Remark 16. *We will see that in the specific case where $z^{i,j,a} = z^{i,a}$, and $z^{i,j,b} = z^{i,b}$ for all $(i, j) \in \{1, \dots, N\}^2$, there exists a unique fixed point in $\mathcal{O}(z, q)$ for any $(z, q) \in \mathcal{R}^N \times \mathbb{Z}^N$. This specification is used in Corollary 2.1, where the Nash equilibrium is provided explicitly.*

We now define a family of processes which represents the form of contract given to the agents.

Definition 2.5. *Given $y_0 \in \mathbb{R}^N$, and \mathcal{R} -valued predictable process $Z^i := (Z^{i,j,a}, Z^{i,j,b}, Z^{S,i})_{j=1, \dots, N}$, for $i \in \{1, \dots, N\}$, we introduce the family of \mathbb{R}^N -valued processes $(Y^{y_0, Z, \hat{\delta}})_{\hat{\delta} \in \mathcal{O}}$ indexed by fixed point maps $\hat{\delta} \in \mathcal{O}$, whose i -th coordinate is given by, for $i \in \{1, \dots, N\}$ and $t \in [0, T]$*

$$Y_t^{i, y_0, Z, \hat{\delta}} := y_0^i + \sum_{j=1}^N \int_0^t Z_r^{i,j,a} dN_r^{j,a} + Z_r^{i,j,b} dN_r^{j,b} + Z_r^{S,i} dS_r + \left(\frac{\gamma_i \sigma^2}{2} (Z_r^{S,i} + Q_r^i)^2 - H^i(\hat{\delta}^{-i}(Z_r, Q_r), Z_r^i, Q_r) \right) dr. \quad (2.14)$$

We say that $Z := (Z^i)_{i=1, \dots, N}$ belongs to the set \mathcal{Z} , if $Y_T^{y_0, Z, \hat{\delta}}$ satisfies (2.8), (2.10) and for all $\delta \in \mathcal{A}$,

$$\mathbb{E}^\delta \left[\sup_{t \in [0, T]} \exp \left(- \gamma_i' Y_t^{i, y_0, Z, \hat{\delta}} \right) \right] < +\infty.$$

This condition ensures that the market-maker's problem is not degenerated given this specific form of contract. Moreover, given the integrability conditions on the coefficients, the processes $(Y^{y_0, Z, \hat{\delta}})_{\hat{\delta} \in \mathcal{O}}$ are well defined and $(e^{-\gamma_i Y^i})$ is a uniformly integrable process under \mathbb{P}^δ , for every $\delta \in \mathcal{A}$, and $i \in \{1, \dots, N\}$.⁶ To link an admissible vector contract $\xi \in \mathcal{C}$ to the processes defined in (2.14), we define the following set.

Definition 2.6. *We define Ξ as the set of random variables $Y_T^{y_0, Z, \hat{\delta}}$ where $(y_0, Z, \hat{\delta})$ ranges in $\mathbb{R}^N \times \mathcal{Z} \times \mathcal{O}$, and such that $e^{-\gamma_i y_0^i} \geq R_i$ for any $i \in \{1, \dots, N\}$.⁷*

Since by definition all bounded predictable processes are contained in \mathcal{Z} , it is clearly nonempty. To prove equality of these sets, we are reduced to the problem of representing any contract ξ^i as $Y_T^{i, y_0, Z, \hat{\delta}}$ for some $(y_0, Z) \in \mathbb{R}^N \times \mathcal{Z}$ and some $\hat{\delta} \in \mathcal{O}$. Following the approach of Sannikov [243], we derive a dynamic programming principle for the utility function of the market-maker, and then prove the equality of the sets by identification of the coefficients.

2.3.2 Contract representation

The following theorem provides solution to the market-maker's problem, and a complete characterisation of the set of admissible contracts.

Theorem 2.1. *Any contract vector $\xi = Y_T^{y_0, Z, \hat{\delta}}$ with $(y_0, Z, \hat{\delta}) \in \mathbb{R}^N \times \mathcal{Z} \times \mathcal{O}$ leads to a unique Nash equilibrium for the agents, given by $(\hat{\delta}(Z_t, Q_t))_{t \in [0, T]}$. Conversely, any admissible contract $\xi \in \mathcal{C}$ is of the form $\xi = Y_T^{y_0, Z, \hat{\delta}}$ for some $(y_0, Z) \in \mathbb{R}^N \times \mathcal{Z}$ and a certain $\hat{\delta} \in \mathcal{O}$.*

In the next corollary, we restrict ourselves to a subset of admissible contracts under which each agent earns at least his reservation utility, and where we can derive explicitly their best-response. For such purpose, we introduce the following set

$$\Xi' := \left\{ Y_T^{y_0, Z, \hat{\delta}} : (\hat{\delta}, y_0, Z) \in \mathcal{O} \times \mathbb{R}^N \times \mathcal{Z}, \text{ s.t. for all } (i, j, k) \in \{1, \dots, N\}^2 \times \{a, b\}, e^{-\gamma_i y_0^i} \geq R_i, Z^{i,j,k} =: Z^k \right\}.$$

The main interest of the subset Ξ' is the following result.

⁶Such condition is used to provide explicitly the best response of the agents in Corollary 2.1.

⁷Theorem 2.1 proves that such contract generates at least one equilibrium.

Lemma 2.1. *Assume that, for $(z, q) \in \mathcal{R}^N \times \mathbb{Z}^N$, we have $z^{i,\ell,j} = z^j$ for all $(i, \ell) \in \{1, \dots, N\}^2$ and $j \in \{a, b\}$. We define*

$$\Gamma^{i,j}(z) := -z^j + \frac{1}{\gamma_i} \log \left(1 + \frac{\sigma \gamma_i}{k \varpi} \right), \quad z \in \mathcal{R}^N.$$

We also introduce the function $\Delta : \mathcal{R}^N \times \mathbb{Z}^N \rightarrow \mathcal{M}_{N,2}(\mathbb{R})$ defined by, for $i \in \{1, \dots, N\}$, $j \in \{a, b\}$, $(z, q) \in \mathcal{R}^N \times \mathbb{Z}^N$

$$\Delta^{i,j}(z, q) := \begin{cases} (-\delta_\infty) \vee \Gamma^{i,j}(z) \wedge \delta_\infty, & \text{if } \Gamma^{i,j}(z) < \Gamma^{\ell,j}(z), \quad -\bar{q} < q < \bar{q}, \text{ for all } \ell \neq i, \\ (-\delta_\infty) \vee \frac{1}{\omega_\ell} \Gamma^{i,j}(z) \wedge \delta_\infty, & \text{if } \frac{1}{\omega_\ell} \Gamma^{i,j}(z) \in K_\ell, \quad -\bar{q} < q < \bar{q}, \text{ for } \ell \in \{1, \dots, K\}, \\ 0, & \text{otherwise.} \end{cases} \quad (2.15)$$

Then, $\mathcal{O}(z, q)$ is reduced to the singleton $\{\Delta(z, q)\}$.

The proof is reported in the appendix, and follows from standard computations on the Hamiltonian (2.3). In particular, it leads to existence and uniqueness of the maximizer of (2.3). We can now conclude with the announced corollary.

Corollary 2.1. *For any admissible contract $Y_T^{y_0, Z, \hat{\delta}} \in \Xi'$ offered by the principal, there exist a unique Nash equilibrium, given by $(\Delta(Z_t, Q_t))_{t \in [0, T]}$, where the map Δ is defined in (2.15).*

This result states that, at the optimum, the utility function of each market-maker corresponds to its reservation utility, that is to say the quantity such that the N agents accept their contract. Moreover, it enables us to characterize explicitly a unique Nash equilibrium for the market-maker's problem. We end the section with some comments on the shape of admissible contracts (2.5).

2.3.3 On the shape of compensation proposed and contractible variables.

In this section, we would like to highlight some interpretation on the classes Ξ and Ξ' of "smooth" contracts $\xi = (\xi^1, \dots, \xi^N)^\top$ controlled by $(y_0, Z, \hat{\delta}) \in \mathbb{R}^N \times \mathcal{Z} \times \mathcal{O}$ and having the form

$$\xi^i = y_0^i + \sum_{j=1}^N \int_0^T Z_r^{i,j,a} dN_r^{j,a} + Z_r^{i,j,b} dN_r^{j,b} + Z_r^{S,i} dS_r + \left(\frac{1}{2} \gamma_i \sigma^2 (Z_r^{S,i} + Q_r^i)^2 - H^i(\hat{\delta}^{-i}(Z_r, Q_r), Z_r^i, Q_r) \right) dr.$$

Note that the contracts are indexed on the number of transactions and on the efficient price of the asset. The spreads of the market-makers are observed by the platform but are not contractible variables, since a contract depending on them would be unrealistic in practice, see El Euch, Mastroliia, Rosenbaum, and Touzi [109]. Mathematically speaking, allowing the contract to depend on the spreads would correspond to a first best problem. For sake of completeness, we compute the solution of the first best problem in Appendix 2.A.7 and show that it differs from the one considered here.

- The compensation y_0^i is calibrated by the exchange to ensure the reservation utility constraint with level R_i of the i -th market-maker, we refer to Section 2.4 for more details on it.
- The term $\int_0^T Z_r^{S,i} dS_r$ is the compensation given to the market-maker with respect to the efficient price.
- The terms $\int_0^T Z_r^{i,j,a} dN_r^{j,a}$, and $\int_0^T Z_r^{i,j,b} dN_r^{j,b}$ are the compensation of the i -th market-maker with respect to the number of trades made on the ask side or bid side by the j -th market-maker.
- The term $\int_0^T H^i(\hat{\delta}^{-i}(Z_r, Q_r), Z_r^i, Q_r) dr$ is the certain gain of the i -th agent induced by his maximisation problem. The principal anticipates that the agent will earn money coming from his maximisation strategy. Hence, he deducts such corresponding amount to the salary of the agent. This justifies why this term appears with a minus sign in the compensation ξ^i .

- The term $\int_0^T \frac{1}{2} \gamma_i \sigma^2 (Z_r^{S,i} + Q_r^i)^2 dr$ is the compensation⁸ to balance the risk aversion of the agent with respect to the efficient price and his inventory.

From a representation viewpoint, the subset Ξ' means that we index the contract of the i -th market-maker only on the aggregated order processes N^a and N^b , and the efficient price S . Moreover, we restrict ourselves to a subset of the admissible contracts where the incentives with respect to the bid and ask arrival orders are equal for every agent. Hence, we do not discriminate a priori one market-maker compared to another. However, the discrimination is done in the market risk part, namely $\int_0^T Z_t^{S,i} dS_t$, for $i \in \{1, \dots, N\}$. This assumption is in force until the end of the chapter (except in the appendix). Practically, the incentives for the i -th agent are only functions of his own inventory process, the aggregated order flow, and the efficient price. It appears reasonable from a practical point of view, as it means that the platform does not need to monitor cross incentives $Z^{i,j,a}$ or $Z^{i,j,b}$ for $j \neq i$, which is hard to do in practice. In addition to this, the exchange give the same incentives to the agents on the part driven by the market orders sent by market-takers, but can discriminate with respect to the risk aversion parameters on the part driven by the efficient price around which market-makers adjust their quotes.

From the technical point of view, this simplification enables to obtain an explicit formula for the fixed points of the Hamiltonian, which is not the case in the general framework. We will also see in the next section that this restriction drastically simplify computations to derive explicitly the optimal incentives that the principal provides to each market-makers.

2.4 Solving the principal's problem

Denote for all $i \in \{1, \dots, N\}$, $\hat{y}_0^i := -\frac{1}{\gamma_i} \log(-R_i)$. By Theorem 2.1 and Corollary 2.1, the exchange problem (2.11), when restricted to contracts in Ξ' , reduces to the control problem

$$\tilde{V}_0^E := \sup_{y_0 \geq \hat{y}_0} \sup_{Z \in \mathcal{Z}} \mathbb{E}^{\Delta(Z,Q)} \left[-\exp \left(-\eta (c(N_T^a + N_T^b) - Y_T^{y_0, Z, \Delta(Z,Q)} \cdot 1_N) \right) \right]. \quad (2.16)$$

Corollary 2.1 provides the best responses of the agents as a function of the control process $Z \in \mathcal{Z}$ of the principal.⁹ Given such response, the exchange solve (2.11), with $\xi \in \Xi'$, in two steps

- Due to the form of utility function, the optimisation with respect to y_0 ensures the reservation utility constraint of the agents is satisfied.
- Optimisation with respect to $Z \in \mathcal{Z}$ is done by solving a classical Hamilton-Jacobi-Bellman equation associated to the reformulated control problem.

The section ends with a verification argument to ensure that the value function coincides with (2.11), and some comments on switching policy between market-makers and physical interpretation of the results.

2.4.1 Saturation of utility constraint

Note that the market-makers' optimal response in Corollary 2.1 does not depend on y_0 . The exponential linear framework for the PnL of the principal enables to state directly that this objective function is clearly decreasing in all coordinates of y_0 , implying that the maximisation under the participation constraint is achieved at \hat{y}_0

$$\tilde{V}_0^E = e^{\eta \hat{y}_0 \cdot 1_N} \sup_{Z \in \mathcal{Z}} \mathbb{E}^{\Delta(Z,Q)} \left[-\exp \left(-\eta (c(N_T^a + N_T^b) - Y_T^{y_0, Z, \Delta(Z,Q)} \cdot 1_N) \right) \right].$$

⁸It corresponds to the quadratic variation of the sum of the incentive indexed on S and the inventory process of the i -th market-maker, integrated against S and weighted by its risk aversion.

⁹In this section δ is the unique Nash equilibrium of the agent's problem coming from Corollary 2.1.

Hence, we are left with a maximisation problem with respect to $Z \in \mathcal{Z}$, which is nothing else than a standard stochastic control problem with the state variables Q , N^a , N^b and $Y^{y_0, Z, \Delta(Z, Q)}$.

2.4.2 The HJB equation for the reduced exchange problem

We study in this section the HJB equation corresponding to the stochastic control problem

$$v_0^E := \sup_{Z \in \mathcal{Z}} \mathbb{E}^{\Delta(Z, Q)} \left[- \exp \left(- \eta (c(N_T^a + N_T^b) - Y_T^{y_0, Z, \Delta(Z, Q)} \cdot \mathbf{1}_N) \right) \right]. \quad (2.17)$$

For the sake of simplicity, we define for any map $v : [0, T] \times \mathbb{Z}^N \rightarrow (-\infty, 0)$, any $x \in \mathbb{R}$, any $i \in \{1, \dots, N\}$, and any $(t, q) \in [0, T] \times \mathbb{Z}^N$

$$v(t, q \oplus_i x) := v(t, q^1, \dots, q^{i-1}, q^i + x, q^{i+1}, \dots, q^N), \quad v(t, q \ominus_i x) := v(t, q^1, \dots, q^{i-1}, q^i - x, q^{i+1}, \dots, q^N).$$

We also define the maps $\mathcal{V}^+(t, q) := (v(t, q \oplus_i 1))_{i \in \{1, \dots, N\}}$, and $\mathcal{V}^-(t, q) := (v(t, q \ominus_i 1))_{i \in \{1, \dots, N\}}$, for $(t, q) \in [0, T] \times \mathbb{Z}^N$, as well as the set $\mathcal{Q} := \{-\bar{q}, \dots, \bar{q}\}$. The HJB equation associated to (2.17) is

$$\begin{cases} \partial_t v(t, q) + \mathcal{H}(q, \mathcal{V}^+(t, q), \mathcal{V}^-(t, q), v(t, q)) = 0, & (t, q) \in [0, T] \times \mathcal{Q}^N, \\ v(T, q) = -1, & q \in \mathcal{Q}^N, \end{cases} \quad (2.18)$$

where $\mathcal{H}(q, p, m, v) := \mathcal{H}^S(q, v) + \mathcal{H}^b(q, p, v) + \mathcal{H}^a(q, m, v)$, with, for any $(q, p, \ell) \in \mathcal{Q}^N \times \mathbb{R}^N \times \{a, b\}$

$$\begin{aligned} \mathcal{H}^S(q, v) &= \sup_{z^S \in \mathbb{R}^N} v \left(\sum_{i=1}^N \frac{\eta}{2} \sigma^2 \gamma_i (z^{S,i} + q^i)^2 + \frac{\eta^2 \sigma^2}{2} \|z^S\|^2 \right), \\ \mathcal{H}^\ell(q, p, v) &= \sup_{z^\ell \in \mathbb{R}} \sum_{i=1}^N \lambda^{i, \ell} (\Delta^{i, \ell}(z, q), q) \left(e^{\eta(Nz^\ell - c)} p^i - v \mathcal{L}^\ell(\Delta(z, q)) \right), \end{aligned}$$

where

$$\mathcal{L}^\ell(\Delta(z, q)) := 1 + \eta \sum_{i=1}^N \gamma_i^{-1} \left(1 - \exp \left(-\gamma_i \left(z + \Delta^{i, \ell}(z, q) \mathbf{1}_{\{\Delta^{i, \ell}(z, q) = \Delta^i(z, q)\}} + \sum_{j=1}^K \omega_j \Delta^{i, \ell}(z, q) \mathbf{1}_{\{\Delta^{i, \ell}(z, q) \in K_j\}} \right) \right) \right).$$

We now provide the optimal incentives corresponding to the solution of (2.18).

Lemma 2.2. *Assume δ_∞ is large enough so that the condition of Lemma 2.7 is verified. The optimizers in the supremum appearing in PDE (2.18) are given, for any $(t, q) \in [0, T] \times \mathcal{Q}^N$, by*

$$\begin{aligned} z^{*,a}(t, q) &:= \frac{1}{N} \left(c + \frac{1}{\eta} \log \left(\frac{v(t, q)}{\sum_{i \in \mathcal{G}} v(t, q \ominus_i 1)} \right) + \frac{1}{\eta} \log \left(\frac{k\varpi}{k\varpi + \sigma\eta} \text{Card}(\mathcal{G}) \left(1 + \eta\sigma \sum_{i=1}^N \frac{1}{k\varpi + \sigma\gamma_i} \right) \right) \right), \\ z^{*,b}(t, q) &:= \frac{1}{N} \left(c + \frac{1}{\eta} \log \left(\frac{v(t, q)}{\sum_{i \in \mathcal{G}} v(t, q \oplus_i 1)} \right) + \frac{1}{\eta} \log \left(\frac{k\varpi}{k\varpi + \sigma\eta} \text{Card}(\mathcal{G}) \left(1 + \eta\sigma \sum_{i=1}^N \frac{1}{k\varpi + \sigma\gamma_i} \right) \right) \right), \\ z^{*,S,i}(q) &:= - \sum_{j=1}^N \mu_{i,j} \gamma_j q^j, \quad \forall i \in \{1, \dots, N\}, \end{aligned}$$

where for all $(i, j) \in \{1, \dots, N\}^2$

$$\mu_{i,j} := -\eta\kappa \prod_{k \in \{1, \dots, N\} \setminus \{i, j\}} \gamma_k, \quad \text{if } i \neq j, \quad \mu_{i,i} := \kappa \left(\prod_{j \in \{1, \dots, N\} \setminus \{i\}} \gamma_j + \eta \sum_{j \in \{1, \dots, N\} \setminus \{i\}} \prod_{k \in \{1, \dots, N\} \setminus \{i, j\}} \gamma_k \right),$$

with

$$\kappa^{-1} := \prod_{i=1}^N \gamma_i + \eta \sum_{j=1}^N \prod_{k \in \{1, \dots, N\} \setminus \{j\}} \gamma_k, \quad \text{and } \mathcal{G} := \left\{ i \in \{1, \dots, N\} : \gamma_i = \max_{j \in \{1, \dots, N\}} \gamma_j \right\}.$$

The optimizers $z^{*,a}$ and $z^{*,b}$ are only functions of time and the inventory of the agents. Moreover, they are very similar and share common properties with the optimal incentives $z^{*,a}$, and $z^{*,b}$ for the single market-maker case in El Euch, Mastrolia, Rosenbaum, and Touzi [109]: for example, for small inventories, they are decreasing function of the risk aversion parameters γ^i . However, the dependence on the number of market-makers and their risk aversion is represented by the term

$$\frac{1}{\eta} \log \left(\frac{k\varpi}{k\varpi + \sigma\eta} \text{Card}(\mathcal{G}) \left(1 + \eta\sigma \sum_{i=1}^N \frac{1}{k\varpi + \sigma\gamma_i} \right) \right).$$

It is an increasing function of N and ϖ , which implies that when we increase the number of market-makers and ϖ , this term decreases the average spread. Notice also that the optimal $z^{*,S}$ depends on a weighted combination of all risk aversions and inventory processes of the agents. This is discussed more in Section 2.5, where we also present our numerical results.

2.4.3 Change of variable and verification theorem

Substituting the optima given by Lemma 2.2, PDE (2.18) boils down to

$$\begin{cases} \partial_t v(t, q) + v(t, q) C^S(q) - v(t, q) C \sum_{j \in \{a, b\}} \left(\frac{v(t, q)}{\sum_{i \in \mathcal{G}} v(t, q \ominus_i \phi(j)) \mathbf{1}_{\{\phi(j)q^i > -\bar{q}\}}} \right)^{\frac{k\varpi}{\sigma\eta}} = 0, \\ v(T, q) = -1, \end{cases} \quad (2.19)$$

for $(t, q) \in [0, T] \times \mathcal{Q}^N$, with

$$\begin{aligned} C &:= A \exp \left(-\frac{k}{\sigma} \left(c(1 - \varpi) - \frac{\varpi}{\eta} \log \left(\frac{k\varpi}{k\varpi + \sigma\eta} \text{Card}(\mathcal{G}) \left(1 + \eta\sigma \sum_{i=1}^N \frac{1}{k\varpi + \sigma\gamma_i} \right) \right) \right. \right. \\ &\quad \left. \left. + \varpi \sum_{i=1}^N \gamma_i^{-1} \log \left(1 + \frac{\sigma\gamma_i}{k\varpi} \right) \right) \right) \frac{\sigma\eta}{k\varpi + \sigma\eta} \left(1 + \eta\sigma \sum_{i=1}^N \frac{1}{k\varpi + \sigma\gamma_i} \right), \\ C^S(q) &:= \sum_{i=1}^N \frac{\eta}{2} \sigma^2 \gamma_i \left(q^i - \sum_{j=1}^N \mu_{i,j} \gamma_j q^j \right)^2 + \frac{\eta^2 \sigma^2}{2} \left(\sum_{i=1}^N \sum_{j=1}^N \mu_{i,j} \gamma_j q^j \right)^2. \end{aligned}$$

Lemma 2.3. *There exists a unique bounded solution to (2.19), which is also negative.*

The solution of (2.19) will be linked to the value function (2.17) using a verification argument in the next section. Note that, if the agents have different risk aversion parameters, one market-maker is both best bid and best ask. Indeed, when a market-maker is simultaneously the single best bid and best ask at some time $t \in [0, T]$, the HJB equation reduces to the following linear PDE

$$0 = \partial_t u - u(t, q) \tilde{C}^S(q) + \tilde{C} u(t, q \oplus_i 1) \mathbf{1}_{\{q^i < \bar{q}\}} + \tilde{C} u(t, q \ominus_i 1) \mathbf{1}_{\{q^i > -\bar{q}\}},$$

where $u := (-v)^{-\frac{k\varpi}{\sigma\eta}}$, $\tilde{C}^S(q) := (k\varpi)/(\sigma\eta)C^S(q)$, and $\tilde{C} := (k\varpi)/(\sigma\eta)C$. As the other inventories are fixed when the i -th market-maker is quoting, we obtain a tridiagonal matrix similar to the one in El Euch, Mastrolia, Rosenbaum, and Touzi [109], indexed by $q^i \in \mathcal{Q}$. We emphasize that such form is valid only at the fixed time t . We conclude with the following verification theorem, which leads to the description of a unique optimal contract to be proposed by the exchange to each market-maker.

Theorem 2.2. *Assume that $\delta_\infty \geq \Delta_\infty$, as defined in Lemma 2.7, and let v be the unique solution to (2.19) given by Lemma 2.3. Then, for any $i \in \{1, \dots, N\}$, the optimal contract for the i -th agent is given by*

$$\xi^{*,i} := \hat{y}_0^i + \int_0^T Z_r^{*,a} dN_r^a + Z_r^{*,b} dN_r^b + Z_r^{*,S,i} dS_r + \left(\frac{1}{2} \sigma^2 \gamma_i (Z_r^{*,S,i} + Q_r^i)^2 - H^i(\Delta(Z_r^*, Q_r), Z_r^*, Q_r) \right) dr,$$

where for any $r \in [0, T]$, $Z_r^{*,S} := z^{*,S}(r, Q_{r-})$, $Z_r^{*,a} := z^{*,a}(r, Q_{r-})$, $Z_r^{*,b} = z^{*,b}(r, Q_{r-})$, and we note $Z_r^* := (Z_r^{*,a}, Z_r^{*,b}, Z_r^{*,S})$. Moreover, the optimal equilibrium is given by $(\Delta(Z_r^*, Q_r))_{r \in [0, T]}$, see Corollary 2.1.

2.4.4 Discussion

2.4.4.1 Switching policy

We want to determine which market-maker is the best one at the beginning of the trading period. For any $(i, j) \in \{1, \dots, N\}^2$ such that $i \neq j$, the i -th market-maker has the best ask quotation at time $t \in [0, T]$ if and only if

$$-Z_t^{*,a} + \frac{1}{\gamma_i} \log \left(1 + \frac{\sigma \gamma_i}{k \varpi} \right) < -Z_t^{*,a} + \frac{1}{\gamma_j} \log \left(1 + \frac{\sigma \gamma_j}{k \varpi} \right).$$

Since the term $\frac{1}{\gamma_i} \log \left(1 + \frac{\sigma \gamma_i}{k \varpi} \right)$ is a decreasing function of γ_i , we conclude that the i -th market-maker trades first if and only if $\gamma_i = \max_{j \in \{1, \dots, N\}} \gamma_j$, and we have

$$\left(\frac{v(t, Q_t)}{\sum_{j \in \mathcal{G}} v(t, Q_t \ominus_j 1)} \right)^{\frac{k \varpi}{\sigma \eta}} = \left(\frac{v(t, Q_t)}{v(t, Q_t \ominus_i 1)} \right)^{\frac{k \varpi}{\sigma \eta}}.$$

We now define when there is a switching between two agents on the ask side. The $N - 1$ other market makers (recall that $j \neq i$) will place their quotes among the open covering of $[0, \delta_\infty]$. Assume that $\Delta^{i,a}(Z_t^*, Q_t) \in K_u$, for some $u \in \{1, \dots, K\}$. Then

$$\Delta^{i,a}(Z_t^*, Q_t) > \Delta^{j,a}(Z_t^*, Q_t) \iff -Z_t^{*,a} + \frac{1}{\gamma_i} \log \left(1 + \frac{\sigma \gamma_i}{k \varpi} \right) > \frac{1}{\omega_u} \left(-Z_t^{*,a} + \frac{1}{\gamma_j} \log \left(1 + \frac{\sigma \gamma_j}{k \varpi} \right) \right), \quad (2.20)$$

which can be rewritten as

$$\begin{aligned} \log \left(\frac{u(t, Q_{t-})}{u(t, Q_{t-} \ominus_i 1)} \right) &> \frac{k \varpi}{\sigma} \left(\log \left(\frac{k \varpi}{k \varpi + \sigma \eta} \left(1 + \eta \sigma \sum_{i=1}^N \frac{1}{k \varpi + \sigma \gamma_i} \right) \right) + c \right) \\ &+ \frac{k N \varpi \omega_u}{(\omega_u - 1) \sigma \gamma_i} \log \left(1 + \frac{\sigma \gamma_i}{k \varpi} \right) + \frac{k N \varpi}{\sigma \gamma_j (\omega_u - 1)} \log \left(1 + \frac{\sigma \gamma_j}{k \varpi} \right). \end{aligned}$$

The right-hand side of the inequality is an increasing function of ϖ and a decreasing function of ω_u and the volatility σ . These results are completely symmetric for the bid side. Following, El Euch, Mastrolia, Rosenbaum, and Touzi [109] the previous equations shows that there is a switching between market-makers on the ask side when the i -th market-maker holds a sufficiently negative inventory. This is because he is willing to attract bid order to mean revert his inventory towards zero. Hence, he proposes a lower spread on the bid side, and a higher spread on the ask side to discourage ask orders. Symmetric conclusion holds for the bid side.

2.4.4.2 On the number of market-makers

As the value function v depends implicitly on the number of market-makers through the terms $C^S(q)$ and C , we cannot directly maximize it with respect to N . However, using additional assumptions and working in an asymptotic setting, we show in this section that the optimal number of market-makers is finite. Numerical computations of N will then be given in Section 2.5. We use N as a subscript to highlight the dependence of the functions with respect to the number of market-makers.

Let us define

$$v_N(0, Q_0) = c(N_T^a + N_T^b) - Y_T^{y_0, Z^*, \Delta(Z^*, Q)} \cdot 1_N.$$

As the market-makers' inventory mean revert toward zero, it is reasonable to study the behaviour of $v_N(t, Q_t)$ when $Q_t = 0$. In that case, $\log\left(\frac{v(t,q)}{\sum_{i=1}^N v(t,q\oplus i)}\right) = \log\left(\frac{v(t,q)}{\sum_{i=1}^N v(t,q\oplus i)}\right) \approx 0$ and

$$Z_t^{*,j} \underset{Q_t \rightarrow 0}{=} \frac{1}{N} \left(c + \frac{1}{\eta} \log \left(\frac{k\varpi}{k\varpi + \sigma\eta} N \left(1 + \eta\sigma \sum_{i=1}^N \frac{1}{k\varpi + \sigma\gamma_i} \right) \right) \right), \quad j \in \{a, b\}.$$

By taking expectations, we obtain

$$\mathbb{E}^{\Delta(Z^*,0)} [v_N(0, Q_0)] = \mathbb{E}^{\Delta(Z^*,0)} \left[\int_0^T \left(-B_N \left(\lambda_N^a(\Delta^a(Z_t^*, 0), 0) + \lambda_N^b(\Delta^b(Z_t^*, 0), 0) \right) + \mathcal{W}_N(Z_t^*) \right) dt \right].$$

where

$$B_N = \frac{1}{\eta} \log \left(\frac{k\varpi}{k\varpi + \sigma\eta} N \left(1 + \eta\sigma \sum_{i=1}^N \frac{1}{k\varpi + \sigma\gamma_i} \right) \right),$$

$$\mathcal{W}_N(Z_t^*) = \sum_{i=1}^N \left(H_N^i \left(\Delta(Z_t^*, 0) - \frac{\hat{y}_0^i}{T} \right) \right).$$

For the sake of simplicity, assume that all the risk aversion parameters γ_i are of the same magnitude (i.e $\gamma_i = \gamma$ for all i) so that

$$\sum_{i=1}^N H_N^i(\Delta(Z_t^*, 0), Z_t^*, 0) = \sum_{j=a,b} N \frac{\sigma}{k\varpi + \sigma\gamma} \lambda_N^j(\Delta^j(Z_t^*, 0), 0),$$

$$B_N = \frac{1}{\eta} \log \left(\frac{k\varpi}{k\varpi + \sigma\eta} N \left(1 + \frac{N\eta\sigma}{k\varpi + \sigma\gamma} \right) \right),$$

$$\lambda_N^j(\Delta^j(Z_t^*, 0), 0) = A \exp \left(-\frac{k}{\sigma} \left(c(1 - \varpi) - \varpi \left(-\frac{N}{\gamma} \log \left(1 + \frac{\sigma\gamma}{k\varpi} \right) \right) \right) \right. \\ \left. + \frac{1}{\eta} \log \left(\frac{k\varpi}{k\varpi + \sigma\eta} N \left(1 + \frac{N\eta\sigma}{k\varpi + \sigma\gamma} \right) \right) \right).$$

We see that $\lambda_N^j(\Delta^j(Z_t^*, 0), 0) \rightarrow_{N \rightarrow +\infty} 0$. This implies that a too high number of market-makers with comparable risk aversion will, on average, decrease the liquidity available on the market and therefore decrease the profits of the platform.

Finally, define the (same) reservation utility of the market-makers as $\hat{y}_0^i = \frac{k\varpi}{\sigma} \log(w_N(0, Q_0))$, where $w_N(0, Q_0)$ is the value function of the market-maker when $\xi = 0$ and every agent has the same risk aversion parameter. We obtain

$$\mathcal{W}_N(Z_t^*) = N \left(\frac{\sigma}{k\varpi + \sigma\gamma} \sum_{j=a,b} \lambda_N^j(\Delta^j(Z_t^*, 0), 0) - \frac{k\varpi}{\sigma\eta} \log(w_N(0, Q_0)) \right).$$

As the optimal market-making solutions have a stationary behavior when T is sufficiently large, we approximate the value function w_N using a Taylor expansion with respect to T

$$w_N(t, Q_t) \approx 1 + 2\hat{C}_N(T - t),$$

where $\hat{C}_N \approx 2 \frac{\sigma}{k\varpi + \sigma\gamma} A \exp \left(-\frac{k}{\sigma} \left(c + \frac{\varpi N}{\gamma} \log \left(1 + \frac{\sigma\gamma}{k\varpi} \right) \right) \right)$. Therefore, we obtain

$$\mathcal{W}_N(Z^*) \approx 2 \frac{A\sigma}{k\varpi + \sigma\gamma} N \exp \left(-\frac{k}{\sigma} \left(c + \frac{\varpi N}{\gamma} \log \left(1 + \frac{\sigma\gamma}{k\varpi} \right) \right) \right) \\ \times \left(\exp \left(\frac{k\varpi}{\sigma} \left(c + \frac{1}{\eta} \log \left(\frac{k\varpi}{k\varpi + \sigma\eta} N \left(1 + \frac{N\eta\sigma}{k\varpi + \sigma\gamma} \right) \right) \right) \right) - \frac{k\varpi}{\sigma\eta} \right).$$

We finally set

$$\mathcal{S}(N) = \mathcal{W}_N(Z^*) - 2\lambda_N^a(\Delta^a(Z^*, 0), 0)B_N. \quad (2.21)$$

This function is differentiable with respect to N , and we obtain $\mathcal{S}(0) > 0$. By decreasing of the exponential, $\lim_{N \rightarrow +\infty} \mathcal{S}(N) = 0$. Moreover, thanks to simple but tedious computations, we have $\lim_{N \rightarrow 0^+} \frac{\partial \mathcal{S}(N)}{\partial N} > 0$. These conditions guarantee the existence of, at least, one global maximum of the function $\mathcal{S}(N)$.

We observe that, for the same set of parameters than in Section 2.5, the function $\mathcal{S}(N)$ attains its supremum for $N \approx 3$, which corresponds to the optimal number of market-makers found numerically with the resolution of the HJB equation.

2.4.4.3 On the choice of the weights

Numerical experiments show, when $N \rightarrow +\infty$, a decrease of the intensity of the market orders, and a slight increase of the average bid-ask spread, other parameters being fixed. However, an increase of ϖ decrease the average best bid-ask spread, as well as the PnL of the platform, increase the total order flow and decrease the trading cost. Recall that there is a trade-off between an increase of the order flow, and the amount of incentive given to the market participants. Moreover, increasing the competition between market-makers leads to an increase of their reservation utility, which is costly for the principal. Recall that we designed the aggregated intensity to be a decreasing function of a weighted sum of the spreads quoted by the agents. In practice, the intensity of arrival orders mainly depends on the best quote $\underline{\delta}$, that is to say for $j \in \{a, b\}$, and $t \in [0, T]$

$$\lambda^j(\delta_t^{i,j}, Q_t) = A \exp \left(-\frac{k}{\sigma} \left(c + \varpi \sum_{i=1}^N \delta_t^{i,j} \mathbf{1}_{\{\delta_t^{i,j} = \underline{\delta}^j\}} + \sum_{i=1}^N \sum_{\ell=1}^K H_\ell \delta_t^{i,j} \mathbf{1}_{\{\delta_t^{i,j} \in K_\ell\}} \right) \right) \approx A \exp \left(-\frac{k}{\sigma} (c + \underline{\delta}_t^j) \right).$$

Assume that $\mathcal{G} = \{i\}$, $\varpi = \frac{1}{N}$. The optimal quotes in Theorem 2.2 become, for $j \in \{a, b\}$

$$\underline{\Delta}^j(Z_t^*, Q_t) = \frac{\sigma}{k} \log \left(\frac{u(t, Q_{t-})}{u(t, Q_{t-} - \Theta_i \phi(j))} \right) + \frac{1}{\gamma_i} \log \left(1 + \frac{\sigma \gamma_i N}{k} \right) - \frac{1}{N} \left(c + \frac{1}{\eta} \log \left(\frac{k}{k + \sigma \eta N} \left(1 + \eta \sigma \sum_{i=1}^N \frac{N}{k + \sigma \gamma_i N} \right) \right) \right).$$

Hence, when the number of market-maker increases, the last term corresponding to the incentive given by the principal vanishes to zero and we are left with, for $j \in \{a, b\}$

$$\underline{\Delta}^j(Z_t^*, Q_t) \approx \frac{\sigma}{k} \log \left(\frac{u(t, Q_{t-})}{u(t, Q_{t-} - \Theta_i \phi(j))} \right) + \frac{1}{\gamma_i} \log \left(1 + \frac{\sigma \gamma_i N}{k} \right).$$

It therefore converges toward the form of spread given when there is no contract, but with a different value function.

2.4.4.4 On the form of the incentives

The quantities $z^{*,j}$, $j \in \{a, b\}$, defined in Lemma 2.2, are decreasing functions of the number of market-makers. Hence, the principal is limited in the amount of incentives he can provide to the agents. This can be viewed as a cake whose size increases slower than the number of people who eat it. Hence, each market-maker receives less incentive to decrease their spread in our case of a uniform incentive and an increasing number of market-makers.

About the risk aversion of the additional market-makers, adding a player with a small risk aversion increases the quantity $z^{*,j}$, $j \in \{a, b\}$. This means that adding a less risk-averse player increases the capacity of the principal to offer incentives to reduce the average spread and conversely.

We have found processes $Z^{*,a}, Z^{*,b}, Z^{*,S}$ fixed by the principal in order to build optimal contracts for every market-makers. The assumption that the exchange chooses a priori the same incentives on the arrival orders for each market-maker is quite natural, since in practice the principal may not know the risk aversions of each market-maker. When a market-maker is simultaneously best bid and best ask, we recover the result from [109, Proposition 4.1], that the terms $-\log\left(\frac{u(t, Q_{t-})}{u(t, Q_{t-} \ominus_i 1)}\right)$ and $-\log\left(\frac{u(t, Q_{t-})}{u(t, Q_{t-} \oplus_i 1)}\right)$ are roughly proportional to, respectively, Q_{t-}^i and $-Q_{t-}^i$. The interpretation is the same: the exchange provides incentives to the market-makers to keep their inventory not too large.

An interesting difference comes from the integrals $\int_0^T Z_r^{*,S,i} dS_r$. As in El Euch, Mastrolia, Rosenbaum, and Touzi [109], it is still understood as a risk sharing term. However, each of the $Z^{*,S,i}$ is a weighted function of both γ_i and the other risk aversions γ_j , $j \neq i$. Indeed, when the risk aversion of the i -th market-maker increases, $Z^{*,S,i}$ decreases. When the risk aversions of the $N - 1$ other market-makers increase, $Z^{*,S,i}$ increases and conversely.

2.4.4.5 On the taker cost policy

When the i -th market-maker is simultaneously best bid and best ask, the exchange can fix a relevant value of the taker cost c as in El Euch, Mastrolia, Rosenbaum, and Touzi [109]. From numerical computations

$$\frac{u(t, q)^2}{u(t, q \oplus_i 1)u(t, q \ominus_i 1)} \approx 1, \text{ for all } (t, q) \in [0, T] \times \mathbb{Z}^N.$$

Hence the exchange may fix in practice the transaction cost c so that the average best spread is close to one tick by setting

$$c \approx -\frac{1}{2N} \text{Tick} - \frac{1}{\eta N} \log\left(\frac{k\varpi}{k\varpi + \sigma\eta} \left(1 + \eta\sigma \sum_{i=1}^N \frac{1}{k\varpi + \sigma\gamma_i}\right)\right) + \frac{1}{\gamma_i N} \log\left(1 + \frac{\sigma\gamma_i}{k\varpi}\right).$$

When $\sigma\eta/k\varpi$, and $\sigma\gamma_i/k\varpi$ are small enough for all $i \in \{1, \dots, N\}$, this equation reduces to

$$c \approx \frac{1}{N} \left(\frac{\sigma}{k\varpi} - \frac{1}{2} \text{Tick}\right).$$

We therefore find a similar formula to the one in the case $N = 1$, and notice that it is a decreasing function of the number of market-makers, with $\varpi = \frac{1}{N}$. As σ and k can be estimated in practice using market data, this is a particularly useful rule of thumb to determine the taker cost c . However, when one market-maker is the best bid and another one is the best ask, the approximation $\frac{u(t, q)^2}{u(t, q \oplus_i 1)u(t, q \ominus_i 1)} \approx 1$ is no longer valid. Hence, the exchange has the choice either to stay with the previous rule of thumb, or to monitor a time-dependent taker cost given by

$$c(t, q) \approx -\frac{1}{2N} \text{Tick} - \frac{1}{\eta N} \left(\log\left(\frac{u(t, q)^2}{u(t, q \ominus_i 1)u(t, q \oplus_j 1)}\right) + \log\left(\frac{k\varpi}{k\varpi + \sigma\eta} \left(1 + \eta\sigma \sum_{i=1}^N \frac{1}{k\varpi + \sigma\gamma_i}\right)\right)\right) + \frac{1}{\gamma_i N} \log\left(1 + \frac{\sigma\gamma_i}{k\varpi}\right).$$

where the i -th agent is the best ask, and the j -th is the best bid.

2.5 Impact of the presence of several market-makers

In this section, we compare our results with the ones given in El Euch, Mastrolia, Rosenbaum, and Touzi [109].

2.5.1 One market-maker

As a sanity check, we want to recover the results of El Euch, Mastrolia, Rosenbaum, and Touzi [109]. We take the same numerical values for the parameters, namely $T = 600s$ for an asset with volatility $\sigma = 0.3 \text{ Tick} \cdot s^{-1/2}$ (unless specified differently). Market orders arrive according to the intensities




Figure 2.1: Total spread for 1 market-maker.




Figure 2.2: Ask order flow for 1 market-maker.




Figure 2.3: Total order flow for 1 market-maker.




Figure 2.4: PnL of the exchange for 1 market-maker.




Figure 2.5: Trading cost for 1 market-maker.

described in Section 2.2, with $A = 1.5s^{-1}$ and $k = 0.3s^{-1/2}$. We have $\bar{q} = 50$, $\gamma = 0.01$, $\eta = 1$, $c = 0.5\text{Tick}$ and $\varpi = 1$. We directly present the results of our model

We see in Figures 2.1 to 2.5 that we recover the results obtained in El Euch, Mastroli, Rosenbaum, and Touzi [109]. We now turn to the case $N \geq 2$.

2.5.2 Two market-makers

We first begin with the average spread in the case $N = 2$, with $\varpi = \frac{1}{2}$. The brackets in the title of the figures denote the set of risk aversion of the agents.

We can see in Figures 2.6, 2.7 and 2.8 an increase of the total spread compared to the case $N = 1$. As explained in Section 2.4.4, this is due to the fact that the quantities z^{*j} , $j = a, b$ are decreasing function of N . Hence the incentive given to each market-maker is less important than in the case $N = 1$. In addition to this, adding a market-maker with a higher risk aversion decrease the total spread and conversely.

Such spread induces a decrease of total order flow, see Figure 2.9, compared to the case $N = 1$. For sake of simplicity we only present the results for two market-makers with same risk aversion.




Figure 2.6: Total spread for $N = 2$, $[0.01, 0.001]$.




Figure 2.7: Total spread for $N = 2$, $[0.01, 0.01]$.




Figure 2.8: Total spread for $N = 2$, $[0.01, 0.1]$.




Figure 2.9: Total order flow for $N = 2$, $[0.01, 0.01]$.




Figure 2.10: Ask order flow for $N = 2$, $[0.01, 0.01]$.

Similar results occurs for different risk aversion parameters, except that the decrease of order flow is less important with a second market-maker with a higher risk aversion parameter and conversely. This also has an impact on the trading cost and the PnL of the platform, as it can be seen in Figures 2.11, 2.12.




Figure 2.11: PnL of the exchange for $N = 2$, $[0.01, 0.01]$.




Figure 2.12: Trading cost for $N = 2$, $[0.01, 0.01]$.

We can see in Figure 2.12 an increase of the trading cost due to a mixed effect of the decrease of

order flow, and an increase of the total spread, see Figures 2.6,2.7 and 2.9. However, we see an increase in the PnL of the exchange, mainly due to the fact that the reservation utility for every agent $\hat{y}_0^i := \frac{k\omega}{\sigma} \log(u(0, Q_0))$ $i \in \{1, \dots, N\}$ is less important than in the case $N = 1$.

2.5.3 Five market-makers

This case aims at illustrating what happens when we increase again the number of market-makers. For sake of simplicity we only illustrate the case of market-makers having the same risk aversion parameter equal to 0.01.




Figure 2.13: Total spread for $N = 5$.




Figure 2.14: Total order flow for $N = 5$.




Figure 2.15: Ask order flow for $N = 5$.




Figure 2.16: PnL of the exchange for $N = 5$.




Figure 2.17: Trading cost for $N = 5$.




Figure 2.18: Evolution of the PnL of the platform with the number of market-makers.

As expected, we obtain in Figure 2.13 a higher total spread, which implies a decrease of the order flow, see Figure 2.14. However, in Figure 2.16, the PnL of the platform has decreased compared to the case $N = 2$. This means that it is not optimal for the platform to attract an infinite number of market-makers. We conclude these numerical experiments with Figure 2.18 showing how the PnL of the platform evolve with the number of market-makers. We emphasize here that what is important is not the risk aversion of the market-makers added to the market: this has an impact on the PnL of the platform but not on the trend of the graph. Hence, we add market-makers with the same risk aversion

equals to 0.01.

In Figure 2.18, there are two different plots. The orange one is with $\varpi = \frac{1}{N}$ and c the taker cost being fixed. The red one is with $\varpi = \frac{1}{N}$ and $c = \frac{1}{N}(\frac{\sigma}{k} - \frac{1}{2}\text{Tick})$ as stated in the previous section. We can see that without an optimal taker cost policy, the optimal number of market-makers for the platform is at $N = 2$, other parameters being fixed. However, with an optimal policy, the platform is encouraged to add another market-maker to increase its PnL. It is also worth noting that in both cases, the platform can add up to 4 market-makers and still have a PnL higher than in the case $N = 1$.

2.A Appendix

2.A.1 Dynamic programming principle

For any $i \in \{1, \dots, N\}$, any \mathbb{F} -predictable stopping times τ taking values in $[0, T]$, any admissible contract vector $\xi \in \mathcal{C}$, any $2(N-1)$ -dimensional \mathbb{F} -predictable process δ^{-i} , bounded by δ_∞ , and for all $\delta \in \mathcal{A}^i(\delta^{-i})$, we define

$$J^i(\xi^i, \tau, \delta, \delta^{-i}) := \mathbb{E}_\tau^{\delta \otimes_i \delta^{-i}} \left[- \exp \left(- \gamma_i \int_\tau^T \delta_u^a \mathbf{1}_{\{\delta_u^a = \delta_u^a \otimes_i \delta_u^{a,-i}\}} dN_u^a + \delta_u^b \mathbf{1}_{\{\delta_u^b = \delta_u^b \otimes_i \delta_u^{b,-i}\}} dN_u^b + Q_u^i dS_u \right. \right. \\ \left. \left. + \sum_{\ell=1}^K \omega_\ell \left(\delta_u^a \mathbf{1}_{\{\delta_u^a \in K_\ell\}} dN_u^a + \delta_u^b \mathbf{1}_{\{\delta_u^b \in K_\ell\}} dN_u^b \right) \exp \left(- \gamma_i \xi^i \right) \right].$$

We also define the family $\mathcal{J}_\tau^i := (J^i(\xi^i, \tau, \delta, \delta^{-i}))_{\delta \in \mathcal{A}^i(\delta^{-i})}$. The continuation utility of the i -th market-maker is defined by

$$V_\tau^i(\xi^i, \delta^{-i}) = \operatorname{ess\,sup}_{\delta \in \mathcal{A}^i(\delta^{-i})} J^i(\xi^i, \tau, \delta, \delta^{-i}). \quad (2.22)$$

Lemma 2.4. *Let τ be a \mathbb{F} -predictable stopping time with values in $[t, T]$. Then, there exists a non-decreasing sequence $(\delta^n)_{n \in \mathbb{N}}$ in $\mathcal{A}^i(\delta^{-i})$ such that $V_\tau^i(\xi^i, \delta^{-i}) = \lim_{n \rightarrow +\infty} J^i(\xi^i, \tau, \delta^n, \delta^{-i})$.*

Proof. For $(\delta, \delta') \in \mathcal{A}^i(\delta^{-i}) \times \mathcal{A}^i(\delta^{-i})$, we define

$$\bar{\delta} := \delta \mathbf{1}_{\{J_\tau^i(\xi^i, \tau, \delta, \delta^{-i}) \geq J_\tau^i(\xi^i, \tau, \delta', \delta^{-i})\}} + \delta' \mathbf{1}_{\{J_\tau^i(\xi^i, \tau, \delta, \delta^{-i}) \leq J_\tau^i(\xi^i, \tau, \delta', \delta^{-i})\}}.$$

We have $\bar{\delta} \in \mathcal{A}^i(\delta^{-i})$ and by definition of $\bar{\delta}$, $J^i(\xi^i, \tau, \bar{\delta}, \delta^{-i}) \geq \max(J^i(\xi^i, \tau, \delta, \delta^{-i}), J^i(\xi^i, \tau, \delta', \delta^{-i}))$. Hence, \mathcal{J}_τ^i is upward directed, and the required result follows from [222, Proposition VI.I.1 p121]. \square

Lemma 2.5. *Let $t \in [0, T]$ and τ be an \mathbb{F} -predictable stopping time with values in $[t, T]$. Then*

$$V_t^i(\xi^i, \delta^{-i}) = \operatorname{ess\,sup}_{\delta^i \in \mathcal{A}^i(\delta^{-i})} \mathbb{E}_t^{\delta^i \otimes_i \delta^{-i}} \left[- \exp \left(- \gamma_i \int_t^\tau \delta_u^{a,i} \mathbf{1}_{\{\delta_u^{a,i} = \delta_u^{a,i} \otimes_i \delta_u^{a,-i}\}} dN_u^a + \delta_u^{i,b} \mathbf{1}_{\{\delta_u^{i,b} = \delta_u^{i,b} \otimes_i \delta_u^{b,-i}\}} dN_u^b \right. \right. \\ \left. \left. + Q_u^i dS_u + \sum_{\ell=1}^K \omega_\ell \left(\delta_u^{i,a} \mathbf{1}_{\{\delta_u^{i,a} \in K_\ell\}} dN_u^a + \delta_u^{i,b} \mathbf{1}_{\{\delta_u^{i,b} \in K_\ell\}} dN_u^b \right) \right) V_\tau^i(\xi^i, \delta^{-i}) \right].$$

Proof. Let $t \in [0, T]$ and fix an \mathbb{F} -predictable stopping time τ with values in $[t, T]$. To simplify the notations, we define for all $t \in [0, T]$ and $\delta \in \mathcal{A}$

$$\mathcal{D}_{t,T}^i(\delta) := e^{-\gamma_i \int_t^\tau \delta_u^{a,i} \mathbf{1}_{\{\delta_u^{a,i} = \delta_u^{a,i}\}} dN_u^a + \delta_u^{i,b} \mathbf{1}_{\{\delta_u^{i,b} = \delta_u^{i,b}\}} dN_u^b + Q_u^i dS_u + \sum_{\ell=1}^K \omega_\ell (\delta_u^{i,a} \mathbf{1}_{\{\delta_u^{i,a} \in K_\ell\}} dN_u^a + \delta_u^{i,b} \mathbf{1}_{\{\delta_u^{i,b} \in K_\ell\}} dN_u^b)}.$$

First, by the tower property, we have that

$$V_t^i(\xi^i, \delta^{-i}) = \operatorname{ess\,sup}_{\delta^i \in \mathcal{A}^i(\delta^{-i})} \mathbb{E}_t^{\delta^i \otimes_i \delta^{-i}} \left[-\mathcal{D}_{t,\tau}^i(\delta) \mathbb{E}_\tau^{\delta^i \otimes_i \delta^{-i}} \left[\mathcal{D}_{\tau,T}^i(\delta) \exp(-\gamma_i \xi^i) \right] \right].$$

For all $\delta \in \mathcal{A}$, the quotient $\frac{L_T^\delta}{L_t^\delta}$ does not depend on the value of δ before time τ . This is by definition of the integrals. Then

$$\begin{aligned} \mathbb{E}_\tau^{\delta^i \otimes_i \delta^{-i}} \left[\mathcal{D}_{\tau,T}^i(\delta) \exp(-\gamma_i \xi^i) \right] &= \mathbb{E}_\tau^0 \left[-\frac{L_T^{\delta^i \otimes_i \delta^{-i}}}{L_\tau^{\delta^i \otimes_i \delta^{-i}}} \mathcal{D}_{\tau,T}^i(\delta) \exp(-\gamma_i \xi^i) \right] \\ &\leq \operatorname{ess\,sup}_{\delta^i \in \mathcal{A}^i(\delta^{-i})} \mathbb{E}_\tau^{\delta^i \otimes_i \delta^{-i}} \left[-\mathcal{D}_{\tau,T}^i(\delta) \exp(-\gamma_i \xi^i) \right] \\ &= V_\tau^i(\xi^i, \delta^{-i}). \end{aligned}$$

Hence, we obtain that

$$V_t^i(\xi^i, \delta^{-i}) \leq \operatorname{ess\,sup}_{\delta^i \in \mathcal{A}^i(\delta^{-i})} \mathbb{E}_t^{\delta^i \otimes_i \delta^{-i}} \left[-V_\tau^i(\xi^i, \delta^{-i}) \mathcal{D}_{t,\tau}^i(\delta) \right].$$

We next prove the reverse inequality. Let $\delta^i \in \mathcal{A}^i(\delta^{-i})$ and $\delta'^i \in \mathcal{A}^i(\delta^{-i})$. We define

$$(\delta^i \otimes_\tau \delta'^i)_u := \delta_u^i \mathbf{1}_{\{0 \leq u \leq \tau\}} + \delta'_u{}^i \mathbf{1}_{\{\tau < u \leq T\}}. \quad (2.23)$$

Then, $\delta^i \otimes_\tau \delta'^i$ being predictable as a sum of two predictable processes, $\delta^i \otimes_\tau \delta'^i \in \mathcal{A}^i(\delta^{-i})$ and

$$\begin{aligned} V_t^i(\xi^i, \delta^{-i}) &\geq \mathbb{E}_t^{(\delta^i \otimes_\tau \delta'^i) \otimes_i \delta^{-i}} \left[-\mathcal{D}_{\tau,T}^i(\delta') \mathcal{D}_{t,\tau}^i(\delta) \exp(-\gamma_i \xi^i) \right] \\ &= \mathbb{E}_t^{(\delta^i \otimes_\tau \delta'^i) \otimes_i \delta^{-i}} \left[\mathbb{E}_\tau^{(\delta^i \otimes_\tau \delta'^i) \otimes_i \delta^{-i}} \left[-\mathcal{D}_{\tau,T}^i(\delta') \exp(-\gamma_i \xi^i) \right] \mathcal{D}_{t,\tau}^i(\delta) \right]. \end{aligned}$$

Using Bayes formula, and noting that $\frac{L_T^{(\delta^i \otimes_\tau \delta'^i) \otimes_i \delta^{-i}}}{L_\tau^{(\delta^i \otimes_\tau \delta'^i) \otimes_i \delta^{-i}}} = \frac{L_T^{\delta'^i \otimes_i \delta^{-i}}}{L_\tau^{\delta'^i \otimes_i \delta^{-i}}}$, we have

$$\mathbb{E}_\tau^{(\delta^i \otimes_\tau \delta'^i) \otimes_i \delta^{-i}} \left[-\mathcal{D}_{\tau,T}^i(\delta') \exp(-\gamma_i \xi^i) \right] = \mathbb{E}_\tau^0 \left[-\frac{L_T^{\delta'^i \otimes_i \delta^{-i}}}{L_\tau^{\delta'^i \otimes_i \delta^{-i}}} \mathcal{D}_{\tau,T}^i(\delta') \exp(-\gamma_i \xi^i) \right] = J_T^i(\xi^i, \tau, \delta'^i, \delta^{-i}).$$

We therefore have

$$V_t^i(\xi^i, \delta^{-i}) \geq \mathbb{E}_t^{(\delta^i \otimes_\tau \delta'^i) \otimes_i \delta^{-i}} \left[\mathcal{D}_{t,\tau}^i(\delta) J_T^i(\xi^i, \tau, \delta'^i, \delta^{-i}) \right].$$

We can therefore use Bayes's formula and the fact that $\frac{L_T^{(\delta^i \otimes_\tau \delta'^i) \otimes_i \delta^{-i}}}{L_t^{(\delta^i \otimes_\tau \delta'^i) \otimes_i \delta^{-i}}} = \frac{L_T^{\delta^i \otimes_i \delta^{-i}}}{L_t^{\delta^i \otimes_i \delta^{-i}}}$ to finally obtain

$$\begin{aligned} V_t^i(\xi^i, \delta^{-i}) &\geq \mathbb{E}_t^0 \left[\mathbb{E}_\tau^0 \left[\frac{L_T^{(\delta^i \otimes_\tau \delta'^i) \otimes_i \delta^{-i}}}{L_\tau^{(\delta^i \otimes_\tau \delta'^i) \otimes_i \delta^{-i}}} \frac{L_\tau^{(\delta^i \otimes_\tau \delta'^i) \otimes_i \delta^{-i}}}{L_t^{(\delta^i \otimes_\tau \delta'^i) \otimes_i \delta^{-i}}} \mathcal{D}_{t,\tau}^i(\delta) J_T^i(\xi^i, \tau, \delta'^i, \delta^{-i}) \right] \right] \\ &= \mathbb{E}_t^{\delta^i \otimes_i \delta^{-i}} \left[\mathcal{D}_{t,\tau}^i(\delta) J_T^i(\xi^i, \tau, \delta'^i, \delta^{-i}) \right]. \end{aligned}$$

Since the previous inequality holds for all $\delta'^i \in \mathcal{A}^i(\delta^{-i})$ we deduce from the monotone convergence theorem together with Lemma 2.4 that there exists a sequence $(\delta'^n)_{n \in \mathbb{N}}$ of controls in $\mathcal{A}^i(\delta^{-i})$ such that

$$V_t^i(\xi^i, \delta^{-i}) \geq \lim_{n \rightarrow +\infty} \mathbb{E}_t^{\delta^i \otimes_i \delta^{-i}} \left[\mathcal{D}_{t,\tau}^i(\delta) J_T^i(\xi^i, \tau, \delta'^n, \delta^{-i}) \right] = \mathbb{E}_t^{\delta^i \otimes_i \delta^{-i}} \left[\mathcal{D}_{t,\tau}^i(\delta) V_\tau^i(\xi^i, \delta^{-i}) \right],$$

thus concluding the proof. \square

2.A.2 Proof of Theorem 2.1

We begin with a lemma concerning the integrability of the continuation utility of the i -th agent defined in (2.22).

Lemma 2.6. *For all $\delta \in \mathcal{A}$ and all $i \in \{1, \dots, N\}$, the process $V^i(\xi^i, \delta^{-i})$ is negative and for a specific $\epsilon > 0$, we have*

$$\mathbb{E}^\delta \left[\sup_{t \in [0, T]} |V_t^i(\xi^i, \delta^{-i})|^{1+\epsilon} \right] < +\infty, \quad \mathbb{E}^\delta \left[\sup_{(s, t) \in [0, T]^2} (D_{s, t}^i(\delta))^{1+\epsilon} \right] < +\infty.$$

Proof. Let $\epsilon > 0$, and $\delta \in \mathcal{A}$. Thanks to the uniform boundedness of $\delta^i \in \mathcal{A}^i(\delta^{-i})$, we have that

$$\frac{L_T^{\delta^i \otimes_i \delta^{-i}}}{L_t^{\delta^i \otimes_i \delta^{-i}}} \geq \alpha_{t, T} := e^{-\frac{k}{\sigma}(c + \delta_\infty(1+H))(N_T^a - N_t^a + N_T^b - N_t^b) - 2Ae^{-\frac{kc}{\sigma}}(e^{\frac{k}{\sigma}(\delta_\infty(1+H))} + 1)(T-t)} \geq \alpha_{0, T}, \quad (2.24)$$

with $H := \max_{\ell=1, \dots, K} H_\ell$. We have

$$\begin{aligned} -V_t^i(\xi^i, \delta^{-i}) &= \operatorname{ess\,inf}_{\delta \in \mathcal{A}^i(\delta^{-i})} \mathbb{E}_t^{\delta^i \otimes_i \delta^{-i}} \left[\exp \left(-\gamma_i \int_t^T \delta_u^a \mathbf{1}_{\{\delta_u^a = \delta_u^a \otimes_i \delta_u^{a,-i}\}} dN_u^a + \delta_u^b \mathbf{1}_{\{\delta_u^b = \delta_u^b \otimes_i \delta_u^{b,-i}\}} dN_u^b + Q_u^i dS_u \right. \right. \\ &\quad \left. \left. + \sum_{\ell=1}^K \omega_\ell \left(\delta_u^a \mathbf{1}_{\{\delta_u^a \in K_\ell\}} dN_u^a + \delta_u^b \mathbf{1}_{\{\delta_u^b \in K_\ell\}} dN_u^b \right) \right) \exp(-\gamma_i \xi^i) \right] \\ &\leq \mathbb{E}_t^{\delta^i \otimes_i \delta^{-i}} \left[e^{\gamma_i (\delta_\infty(1+K\omega_1)(N_T^a + N_T^b) - \int_t^T Q_u^i dS_u)} \exp(-\gamma_i \xi^i) \right], \end{aligned}$$

with $\delta^i \otimes_i \delta^{-i} \in \mathcal{A}$. We used the fact that $N_T^j - N_t^j \leq N_T^j$ for $j \in \{a, b\}$ and for all $i \in \{1, \dots, N\}, j \in \{a, b\}, t \in [0, T]$,

$$\exp \left(-\gamma_i \int_t^T \delta_u^j \mathbf{1}_{\{\delta_u^j = \delta_u^j \otimes_i \delta_u^{j,-i}\}} dN_u^j \right) \leq \exp(\gamma_i \delta_\infty N_T^j).$$

Moreover, as Q^i is uniformly bounded by \bar{q} , we have for all $L > 0$

$$\mathbb{E}_t^{\delta^i \otimes_i \delta^{-i}} \left[e^{-L \int_t^T Q_u^i dS_u} \right] \leq e^{\frac{L^2 \bar{q}^2 \sigma^2 T}{2}}.$$

Thus, using Holder's inequality we have

$$\begin{aligned} -V_t^i(\xi^i, \delta^{-i}) &\leq \mathbb{E}_t^{\delta^i \otimes_i \delta^{-i}} \left[e^{\epsilon \gamma_i (\delta_\infty(1+K\omega_1)(N_T^a + N_T^b) - \xi^i)} \right]^{\frac{1}{\epsilon}} \mathbb{E}_t^{\delta^i \otimes_i \delta^{-i}} \left[\exp \left(-(1+\epsilon) \gamma_i \int_t^T Q_u^i dS_u \right) \right]^{\frac{1}{1+\epsilon}} \\ &\leq \mathbb{E}_t^{\delta^i \otimes_i \delta^{-i}} \left[e^{\epsilon \gamma_i (\delta_\infty(1+K\omega_1)(N_T^a + N_T^b) - \xi^i)} \right]^{\frac{1}{\epsilon}} e^{\frac{(1+\epsilon) \gamma_i^2 \bar{q}^2 \sigma^2 T}{2}}. \end{aligned}$$

Then, we have

$$\mathbb{E}^{\delta^i \otimes_i \delta^{-i}} \left[\sup_{t \in [0, T]} (-V_t^i(\xi^i, \delta^{-i}))^{1+\epsilon} \right] \leq e^{\frac{(1+\epsilon) \gamma_i^2 \bar{q}^2 \sigma^2 T}{2}} \mathbb{E}^{\delta^i \otimes_i \delta^{-i}} \left[\sup_{t \in [0, T]} \mathbb{E}_t^{\delta^i \otimes_i \delta^{-i}} \left[e^{\epsilon \gamma_i (\delta_\infty(1+K\omega_1)(N_T^a + N_T^b) - \xi^i)} \right]^{\frac{1+\epsilon}{\epsilon}} \right],$$

The term inside the conditional expectation is integrable¹⁰ and independent from $t \in [0, T]$ thus by Doob's inequality, we have

$$\mathbb{E}^{\delta^i \otimes_i \delta^{-i}} \left[\sup_{t \in [0, T]} (-V_t^i(\xi^i, \delta^{-i}))^{1+\epsilon} \right] \leq C e^{\frac{(1+\epsilon) \gamma_i^2 \bar{q}^2 \sigma^2 T}{2}} \mathbb{E}^{\delta^i \otimes_i \delta^{-i}} \left[e^{\gamma_i' (\delta_\infty(1+K\omega_1)(N_T^a + N_T^b) - \xi^i)} \right],$$

¹⁰Take $\epsilon > 1$ together with Condition (2.8) for example.

where $C > 0$ and $\gamma'_i = \gamma_i(1 + \epsilon)$. Thanks to Hölder's inequality, together with the boundedness of the intensities of the point processes $N^{i,j}$, for $i \in \{1, \dots, N\}$, and $j \in \{a, b\}$, and Condition (2.8), the right-hand side is bounded from above by a term independent of $t \in [0, T]$. The conclusion follows.

Using the same arguments, we have

$$\mathbb{E}^\delta \left[\sup_{(s,t) \in [0,T]^2} (\mathcal{D}_{s,t}^i(\delta))^{1+\epsilon} \right] \leq C' \mathbb{E}^\delta \left[e^{\gamma'_i (\delta_\infty(1+K\omega_1)(N_T^a + N_T^b) + \bar{q}^2 \gamma_i \frac{\sigma^2 T}{2})} \right] < +\infty$$

where $C' > 0$, using boundedness of the intensities of the point processes for $i \in \{1, \dots, N\}$. The conclusion follows using Hölder's inequality. \square

We introduce for all $i \in \{1, \dots, N\}$, and all $\delta \in \mathcal{A}^i(\delta^{-i})$ the process

$$U_t^{\delta \otimes_i \delta^{-i}} := V_t^i(\xi^i, \delta^{-i}) \mathcal{D}_{0,t}^i(\delta^i \otimes_i \delta^{-i}), \quad t \in [0, T].$$

which thanks to Lemma 2.6 is of class (D).

Step 1: Let $\xi \in \mathcal{C}$ be an admissible contract. By definition, there is a Nash equilibrium $\hat{\delta}(\xi) \in \mathcal{A}$. By use of the dynamic programming principle of Lemma 2.5, for all $\delta^i \in \mathcal{A}^i(\hat{\delta}^{-i}(\xi))$, the process $U^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)}$ defines a $\mathbb{P}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)}$ -supermartingale. We now check that the process $U^{\hat{\delta}(\xi)}$ is a uniformly integrable $\mathbb{P}^{\hat{\delta}(\xi)}$ -martingale.

By Definition 2.1, the control $\hat{\delta}^i(\xi)$ is optimal for the i -th market-maker in the sense that

$$V_{\text{MM}}^i(\xi^i, \hat{\delta}^{-i}(\xi)) = \mathbb{E}^{\hat{\delta}(\xi)} \left[U_i \left(\xi^i + \sum_{j \in \{a,b\}} \int_0^T \hat{\delta}_t^{i,j}(\xi) \left(\mathbf{1}_{\{\hat{\delta}_t^{i,j}(\xi) = \hat{\delta}_t^j(\xi)\}} + \sum_{\ell=1}^K \int_0^T \omega_\ell \mathbf{1}_{\{\hat{\delta}_t^{i,j}(\xi) \in K_\ell\}} \right) dN_t^j + \int_0^T Q_t^i dS_t \right) \right].$$

Hence, an application of the supermartingale property leads, for any \mathbb{F} -predictable stopping time τ taking values in $[0, T]$, to

$$\begin{aligned} V_{\text{MM}}^i(\xi^i, \hat{\delta}^{-i}(\xi)) &\geq \mathbb{E}^{\hat{\delta}(\xi)} \left[\mathcal{D}_{0,\tau}^i(\hat{\delta}(\xi)) V_\tau^i(\xi^i, \hat{\delta}^{-i}(\xi)) \right] \\ &\geq \mathbb{E}^{\hat{\delta}(\xi)} \left[U_i \left(\xi^i + \sum_{j \in \{a,b\}} \int_0^\tau \hat{\delta}_t^{i,j}(\xi) \left(\mathbf{1}_{\{\hat{\delta}_t^{i,j}(\xi) = \hat{\delta}_t^j(\xi)\}} + \sum_{\ell=1}^K \int_0^\tau \omega_\ell \mathbf{1}_{\{\hat{\delta}_t^{i,j}(\xi) \in K_\ell\}} \right) dN_t^j + \int_0^\tau Q_t^i dS_t \right) \right] \\ &= V_{\text{MM}}^i(\xi^i, \hat{\delta}^{-i}(\xi)). \end{aligned}$$

All these inequalities are thus equalities, which proves, since the filtration is right-continuous, that $(U_t^{\hat{\delta}(\xi)})_{t \in [0, T]}$ is a $\mathbb{P}^{\hat{\delta}(\xi)}$ -martingale, and thus for any $t \in [0, T]$

$$U_t^{\hat{\delta}(\xi)} = \mathbb{E}_t^{\hat{\delta}(\xi)} \left[U_i \left(\xi^i + \sum_{j \in \{a,b\}} \int_0^T \hat{\delta}_t^{i,j}(\xi) \left(\mathbf{1}_{\{\hat{\delta}_t^{i,j}(\xi) = \hat{\delta}_t^j(\xi)\}} + \sum_{\ell=1}^K \int_0^T \omega_\ell \mathbf{1}_{\{\hat{\delta}_t^{i,j}(\xi) \in K_\ell\}} \right) dN_t^j + \int_0^T Q_t^i dS_t \right) \right].$$

Using Lemma 2.6, we conclude that $U^{\hat{\delta}(\xi)}$ is a uniformly integrable $\mathbb{P}^{\hat{\delta}(\xi)}$ -martingale. Since the filtration \mathbb{F} is right-continuous, we deduce that $U^{\hat{\delta}(\xi)}$ has a càdlàg $\mathbb{P}^{\hat{\delta}(\xi)}$ -modification. Since all probability measures here are equivalent, we can assume that $U^{\hat{\delta}(\xi)}$ actually has càdlàg paths. As all the probability measures indexed by $\delta \in \mathcal{A}$ are equivalent, we deduce that $(U_t^\delta)_{t \in [0, T]}$ admits a càdlàg modification, for all $\delta \in \mathcal{A}$.

Given the above, for any $\delta^i \in \mathcal{A}^i(\hat{\delta}^{-i}(\xi))$, we can apply Doob-Meyer's decomposition to the $\mathbb{P}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)}$ supermartingale of class (D) $U^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)}$ to obtain

$$U_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} = M_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} - A_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi), c} - A_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi), d}, \quad t \in [0, T],$$

where $M^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)}$ is a uniformly integrable $\mathbb{P}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)}$ -martingale and

$$A_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} = A_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),c} + A_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),d}, \quad t \in [0, T],$$

is an integrable non-decreasing predictable process such that $A_0^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),c} = A_0^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),d} = 0$, with pathwise continuous component $A^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),c}$ and a piecewise constant predictable process $A^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),d}$.

Moreover, using the martingale representation theorem under $\mathbb{P}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)}$, see [109, Section A.1], there exist predictable processes

$$\tilde{Z}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} = (\tilde{Z}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),S}, \tilde{Z}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),i,j,a}, \tilde{Z}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),i,j,b}), \quad j \in \{1, \dots, N\},$$

such that

$$M_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} = V_0^i + \int_0^t \tilde{Z}_r^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),S} dS_r + \sum_{j=1}^N \tilde{Z}_r^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),i,j,a} d\tilde{N}_r^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),j,a} + \tilde{Z}_r^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),i,j,b} d\tilde{N}_r^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),j,b},$$

where the processes $\tilde{N}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),j,a}$, $\tilde{N}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),j,b}$, are defined by (2.4).

Let $Y^i(\xi^i, \hat{\delta}^{-i}(\xi))$ be defined by $V^i(\xi^i, \hat{\delta}^{-i}(\xi)) = -e^{-\gamma_i Y^i(\xi^i, \hat{\delta}^{-i}(\xi))}$. Since $A^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),d}$ is a predictable point process and the jump times of $(N^{i,a}, N^{i,b})$ are totally inaccessible stopping times under \mathbb{P}^0 , we have $\langle N^{i,a}, A^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),d} \rangle = \langle N^{i,b}, A^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),d} \rangle = 0$, a.s. Using Itô's formula, we obtain that

$$Y_T^i(\xi^i, \hat{\delta}^{-i}(\xi)) = \xi^i, \quad dY_t^i(\xi^i, \hat{\delta}^{-i}(\xi)) = \sum_{j=1}^N Z_t^{i,j,a} dN_t^{j,a} + Z_t^{i,j,b} dN_t^{j,b} + Z_t^{S,i} dS_t - dI_t^i - d\tilde{A}_t^{i,d}, \quad (2.25)$$

where by direct identification of the coefficients

$$Z_t^{i,j,a} := -\frac{1}{\gamma_i} \log \left(1 + \frac{\tilde{Z}_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),i,j,a}}{U_{t^-}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)}} \right) - \delta_t^{i,a} \mathbf{1}_{\{\delta_t^{i,a} = \delta_t^a \otimes_i \hat{\delta}_t^{a,-i}(\xi)\}} - \sum_{\ell=1}^K \omega_\ell \delta_t^{i,a} \mathbf{1}_{\{\delta_t^{i,a} \in K_\ell\}},$$

$$Z_t^{i,j,b} := -\frac{1}{\gamma_i} \log \left(1 + \frac{\tilde{Z}_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),i,j,b}}{U_{t^-}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)}} \right) - \delta_t^{i,b} \mathbf{1}_{\{\delta_t^{i,b} = \delta_t^b \otimes_i \hat{\delta}_t^{b,-i}(\xi)\}} - \sum_{\ell=1}^K \omega_\ell \delta_t^{i,b} \mathbf{1}_{\{\delta_t^{i,b} \in K_\ell\}},$$

$$Z_t^{S,i} := -\frac{\tilde{Z}_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),S}}{\gamma_i U_{t^-}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)}} - Q_t^i, \quad I_t^i := \int_0^t \left(\bar{h}^i(\delta_r^i, \hat{\delta}_r^{-i}(\xi), Z_r^i, Q_r) dr - \frac{1}{\gamma_i U_r^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)}} dA_r^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),c} \right),$$

$$\bar{h}^i(\delta_t^i, \hat{\delta}_t^{-i}(\xi), Z_t^i, Q_t) := h(\delta_t^i, \hat{\delta}_t^{-i}(\xi), Z_t^i, Q_t) - \frac{1}{2} \gamma_i \sigma^2 (Z_t^{S,i})^2, \quad \tilde{A}_t^{i,d} := \frac{1}{\gamma_i} \sum_{s \leq t} \log \left(1 - \frac{\Delta A_s^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),d}}{U_{s^-}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)}} \right).$$

In particular, the last relation between $\tilde{A}^{i,d}$ and $A^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),d}$ shows that the process

$$\Delta a_t^i := -\Delta A_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),d} / U_{t^-}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} \geq 0$$

is independent of $\delta^i \in \mathcal{A}^i(\hat{\delta}^{-i}(\xi))$.

We now prove that, $A^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi),d} = -\sum_{0 < s \leq \cdot} U_{s^-}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} \Delta a_s^i = 0$ so that

$$\tilde{A}^{i,d} = 0, \quad I_t = \int_0^{\cdot} \bar{H}^i(\hat{\delta}^{-i}(\xi), Z_r^i, Q_r) dr,$$

where

$$\bar{H}^i(\hat{\delta}^{-i}(\xi), Z_t^i, Q_t) = H^i(\hat{\delta}^{-i}(\xi), Z_t^i, Q_t) - \frac{1}{2}\gamma_i\sigma^2(Z_t^{S,i})^2.$$

As $V_T^i(\xi^i, \hat{\delta}^{-i}(\xi)) = -1$, note that

$$\begin{aligned} 0 &= \sup_{\delta^i \in \mathcal{A}^i(\hat{\delta}^{-i}(\xi))} \mathbb{E}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} \left[U_T^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} - V_0^i(\xi^i, \delta^{-i}) \right] \\ &= \sup_{\delta^i \in \mathcal{A}^i(\hat{\delta}^{-i}(\xi))} \mathbb{E}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} \left[U_T^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} - M_T^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} \right] \\ &= \gamma_i \sup_{\delta^i \in \mathcal{A}^i(\hat{\delta}^{-i}(\xi))} \mathbb{E}^0 \left[L_T^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} \int_0^T U_{r^-}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} \left(dI_r^i - \bar{h}^i(\delta_r^i, \hat{\delta}_r^{-i}(\xi), Z_r^i, Q_r) dr + \frac{da_r^i}{\gamma_i} \right) \right]. \end{aligned}$$

Moreover, since the controls are uniformly bounded, we have by Lemma 2.6

$$U_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} \leq -\beta_t^i = V_t^i(\xi^i, \hat{\delta}^{-i}(\xi)) e^{-2\delta_\infty(N_T^a - N_0^a + N_T^b - N_0^b) - \gamma_i \int_0^t Q_r^i dS_r} < 0.$$

Since $A^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi), d} \geq 0$, $U^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} \leq 0$, and $dI_t^i - \bar{h}^i(\delta_t^i, \hat{\delta}_t^{-i}(\xi), Z_t^i, Q_t) dt \geq 0$, we obtain

$$\begin{aligned} 0 &\leq \sup_{\delta^i \in \mathcal{A}^i(\hat{\delta}^{-i}(\xi))} \mathbb{E}^0 \left[\alpha_{0,T} \int_0^T -\beta_{r^-}^i \left(dI_r^i - \bar{h}^i(\delta_r^i, \hat{\delta}_r^{-i}(\xi), Z_r^i, Q_r) dr + \frac{da_r^i}{\gamma_i} \right) \right] \\ &= -\mathbb{E}^0 \left[\alpha_{0,T} \int_0^T \beta_{r^-}^i \left(dI_r^i - \bar{H}^i(\hat{\delta}_r^{-i}(\xi), Z_r^i, Q_r) dr + \frac{da_r^i}{\gamma_i} \right) \right]. \end{aligned}$$

The quantities $\alpha_{0,T} \int_0^T \beta_{r^-}^i (dI_r^i - \bar{H}^i(\hat{\delta}_r^{-i}(\xi), Z_r^i, Q_r)) dr$ and $\alpha_{0,T} \int_0^T \beta_{r^-}^i \frac{da_r^i}{\gamma_i}$ being non-negative random variables, this implies the announced result.

Given the dynamic under $\mathbb{P}^{\hat{\delta}(\xi)}$ of the process $U_t^{\hat{\delta}(\xi)}$, Itô's formula leads to

$$\begin{aligned} dU_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} &= \sum_{j=1}^N \tilde{Z}_t^{\hat{\delta}(\xi), i, j, a} d\tilde{N}_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi), j, a} + \tilde{Z}_t^{\hat{\delta}(\xi), i, j, b} d\tilde{N}_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi), j, b} + \tilde{Z}_t^{\hat{\delta}(\xi), i, S} dS_t \\ &\quad + U_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} \left(h(\hat{\delta}_t(\xi), \hat{\delta}_t^{-i}(\xi), Z_t^{\hat{\delta}(\xi), i}, Q_t) - h(\delta_t^i, \hat{\delta}_t^{-i}(\xi), Z_t^{\hat{\delta}(\xi), i}, Q_t) \right) dt \end{aligned}$$

and the $\mathbb{P}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)}$ -supermartingale property implies that, almost surely for all $t \in [0, T]$

$$h(\hat{\delta}_t^i(\xi), \hat{\delta}_t^{-i}(\xi), Z_t^{\hat{\delta}(\xi), i}, Q_t) - h(\delta_t^i, \hat{\delta}_t^{-i}(\xi), Z_t^{\hat{\delta}(\xi), i}, Q_t) \geq 0.$$

Hence

$$\hat{\delta}_t^i(\xi) \in \operatorname{argmax}_{\delta \in \mathcal{B}_\infty^2} h(\delta, \hat{\delta}_t^{-i}(\xi), \tilde{Z}_t^{\hat{\delta}(\xi), i}, Q_t).$$

Finally, we check that $Z \in \mathcal{Z}$. Using Lemma 2.6, we have that

$$\sup_{\delta^i \in \mathcal{A}^i(\hat{\delta}^{-i}(\xi))} \mathbb{E}^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} \left[\sup_{t \in [0, T]} |U_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)}|^{p'+1} \right] < +\infty,$$

for some $p' > 0$. The desired conclusion comes from the fact that

$$e^{-\gamma_i Y_t^i} = U_t^{\delta^i \otimes_i \hat{\delta}^{-i}(\xi)} \mathcal{D}_{0,t}(\delta^i \otimes_i \hat{\delta}^{-i}(\xi)).$$

Remark 17. Note that we described here a solution to the following system of N BSDEs given by, for all $i \in \{1, \dots, N\}$

$$Y_t^{i,y_0,Z,\hat{\delta}} := \xi^i - \sum_{j=1}^N \int_t^T Z_r^{i,j,a} dN_r^{j,a} + Z_r^{i,j,b} dN_r^{j,b} + Z_r^{S,i} dS_r + \left(\frac{1}{2} \gamma_i \sigma^2 (Z_r^{S,i} + Q_r^i)^2 - H^i(\hat{\delta}^{-i}(\xi), Z_r^i, Q_r^i) \right) dr.$$

Step 2: Conversely, let us be given a contract vector $\xi = Y_T^{y_0,Z,\hat{\delta}} \in \Xi$, with $(Y_0, Z) \in \mathbb{R}^N \times \mathcal{Z}$ and $\hat{\delta} \in \mathcal{O}$. For $i \in \{1, \dots, N\}$, we note

$$V_t^i(Y_T^{i,y_0,Z,\hat{\delta}}, \hat{\delta}^{-i}(Y_T^{y_0,Z,\hat{\delta}})) := -e^{-\gamma_i Y_t^{i,y_0,Z,\hat{\delta}}}.$$

Given an arbitrary bid-ask policy $\delta^i \in \mathcal{A}^i(\hat{\delta}^{-i})$ of the i -th agent, an application of Itô's formula leads to

$$\begin{aligned} dU_t^{\delta^i \otimes_i \hat{\delta}^{-i}} &= -\gamma_i U_t^{\delta^i \otimes_i \hat{\delta}^{-i}} \left((Q_t^i + Z_t^{S,i}) dS_t - (H^i(\hat{\delta}_t^{-i}, Z_t^i, Q_t) - h^i(\delta_t^i, \hat{\delta}_t^{-i}, Z_t^i, Q_t)) dt \right. \\ &\quad \left. + \gamma_i^{-1} \left(1 - \exp \left(-\gamma_i (Z_t^{i,a} + \delta_t^{i,a} \mathbf{1}_{\{\delta_t^{i,a} = \underline{\delta}_t^a \otimes_i \hat{\delta}_t^{a,-i}\}} + \sum_{\ell=1}^K \omega_\ell \delta_t^{i,a} \mathbf{1}_{\{\delta_t^{i,a} \in K_\ell\}}) \right) \right) d\tilde{N}_t^{\delta^i \otimes_i \hat{\delta}^{-i}, a} \right. \\ &\quad \left. + \gamma_i^{-1} \left(1 - \exp \left(-\gamma_i (Z_t^{i,b} + \delta_t^{i,b} \mathbf{1}_{\{\delta_t^{i,b} = \underline{\delta}_t^b \otimes_i \hat{\delta}_t^{b,-i}\}} + \sum_{\ell=1}^K \omega_\ell \delta_t^{i,b} \mathbf{1}_{\{\delta_t^{i,b} \in K_\ell\}}) \right) \right) d\tilde{N}_t^{\delta^i \otimes_i \hat{\delta}^{-i}, b} \right). \end{aligned}$$

Hence, $(U_t^{\delta^i \otimes_i \hat{\delta}^{-i}})_{t \in [0, T]}$ is a $\mathbb{P}^{\delta^i \otimes_i \hat{\delta}^{-i}}$ -local supermartingale. Thanks to Lemma 2.6, $U_t^{\delta^i \otimes_i \hat{\delta}^{-i}}$ is of class (D) and is a true supermartingale. Hence, we obtain that

$$\begin{aligned} & - \int_0^{\cdot} \gamma_i U_t^{\delta^i \otimes_i \hat{\delta}^{-i}} \left((Q_t^i + Z_t^{S,i}) dS_t \right. \\ & \quad \left. + \gamma_i^{-1} \left(1 - \exp \left(-\gamma_i (Z_t^{i,a} + \delta_t^{i,a} \mathbf{1}_{\{\delta_t^{i,a} = \underline{\delta}_t^a \otimes_i \hat{\delta}_t^{a,-i}\}} + \sum_{\ell=1}^K \omega_\ell \delta_t^{i,a} \mathbf{1}_{\{\delta_t^{i,a} \in K_\ell\}}) \right) \right) d\tilde{N}_t^{\delta^i \otimes_i \hat{\delta}^{-i}, a} \right. \\ & \quad \left. + \gamma_i^{-1} \left(1 - \exp \left(-\gamma_i (Z_t^{i,b} + \delta_t^{i,b} \mathbf{1}_{\{\delta_t^{i,b} = \underline{\delta}_t^b \otimes_i \hat{\delta}_t^{b,-i}\}} + \sum_{\ell=1}^K \int_0^T \omega_\ell \delta_t^{i,b} \mathbf{1}_{\{\delta_t^{i,b} \in K_\ell\}}) \right) \right) d\tilde{N}_t^{\delta^i \otimes_i \hat{\delta}^{-i}, b} \right), \end{aligned}$$

is a true martingale. Therefore

$$\begin{aligned} J_{\text{MM}}^i(\xi^i, \delta^i, \hat{\delta}^{-i}) &= \mathbb{E}^{\delta^i \otimes_i \hat{\delta}^{-i}} \left[U_T^{\delta^i \otimes_i \hat{\delta}^{-i}} \right] \\ &= -e^{-\gamma_i y_0^i} + \mathbb{E}^{\delta^i \otimes_i \hat{\delta}^{-i}} \left[\int_0^T \gamma_i U_t^{\delta^i \otimes_i \hat{\delta}^{-i}} (H^i(\hat{\delta}_t^{-i}, Z_t^i, Q_t) - h^i(\delta_t^i, \hat{\delta}_t^{-i}, Z_t^i, Q_t)) dt \right] \\ &\leq -e^{-\gamma_i y_0^i}. \end{aligned}$$

In addition to this, the previous inequality becomes an equality if and only if δ^i is chosen as the maximizer of the Hamiltonian h^i . By definition, it means that $U^{\hat{\delta}}$ is a $\mathbb{P}^{\hat{\delta}}$ -martingale and that $\hat{\delta}^i$ is the optimal control for the i -th agent, in the sense of (2.1). As this property holds for any $i \in \{1, \dots, N\}$, it means that $\hat{\delta}$ is a Nash equilibrium.

Finally as we showed that the contracts in Ξ generates at least one Nash equilibrium we have the inclusion $\mathcal{C} \supset \Xi$. Hence, the equality $\Xi = \mathcal{C}$ is proved.

2.A.3 Proof of Lemma 2.1

For $(z, q) \in \mathcal{R}^N \times \mathbb{Z}^N$, we set $z^{i,\ell,j} = z^{i,j}$ for all $(i, \ell) \in \{1, \dots, N\}^2$ and $j \in \{a, b\}$. Hence, the Hamiltonian of the i -th agent reduces to

$$h^i(d^i, d^{-i}, z^i, q) := \sum_{j \in \{a, b\}} \gamma_i^{-1} \left(1 - \exp \left(-\gamma_i \left(z^{i,j} + d^{i,j} \mathbf{1}_{\{d^{i,j} = d^j \otimes_i d^{j,-i}\}} + \sum_{k=1}^K \omega_k d^{i,j} \mathbf{1}_{\{d^{i,j} \in K_k\}} \right) \right) \right) \lambda^j(d^j, q).$$

For $i \in \{1, \dots, N\}$, an optimisation of $h^i(d^i, d^{-i}, z^i, q)$ with respect to d^i leads to a unique¹¹ maximum defined as $d^{*i,j}(z, q) = \Delta^{i,j}(z, q)$ for $i \in \{1, \dots, N\}$, and $j \in \{a, b\}$. This maximizer completely characterizes the behaviour of the i -th agent compared to the position of the $N - 1$ others.

Moreover, no matter if the i -th agent plays the best spread or not, compared to the response of the other agents, his optimal response will lead to the following value

$$h^i(\Delta^{i,\cdot}(z, q), \Delta^{-i,\cdot}(z, q), z^i, q) = \frac{\sigma}{1 + \frac{\sigma \gamma_i}{k \varpi}} \lambda^j(\Delta^{i,j}(z, q), q), \quad i \in \{1, \dots, N\}.$$

Hence, when the N agents play Δ , they have no interest in switching their bid-ask policy. Thus, it characterizes a unique fixed point of the Hamiltonian.

2.A.4 Exchange's Hamiltonian maximisation

The following technical result follows from direct but tedious computations. It provides condition on δ_∞ under which the maximizers defined in (2.2) exist.

Lemma 2.7. *Let $q \in \mathcal{Q}^N$, $c \in \mathbb{R}$, $(\eta, k, \sigma) \in (0, +\infty)^3$, $\gamma_i > 0$ for all $i \in \{1, \dots, N\}$, and $v_0, \dots, v_N < 0$. Then, for $z \in \mathcal{R}^N$ and $i \in \{1, \dots, N\}$, $j \in \{a, b\}$, we define*

$$\Phi_q^{i,j}(z) := \lambda^{i,j}(\Delta^{i,j}(z, q), q) \left(e^{\eta(Nz^j - c)} v_i - v_0 \mathcal{L}^j(\Delta(z, q)) \right), \quad \Phi_q^j(z) := \sum_{i=1}^N \Phi_q^{i,j}(z),$$

with $\Delta(z, q)$ defined as in Lemma 2.1, and $\delta_\infty > 0$. Assume that

$$\delta_\infty \geq C_\infty + \frac{N}{\eta} \left| \log \left(\frac{v_0}{\sum_{i=1}^N v_i} \right) \right|,$$

with $C_\infty := N|c| + \sum_{i=1}^N \left(\left(\frac{1}{\eta} + \frac{1}{\gamma_i} \right) \log \left(1 + \frac{\sigma \gamma_i}{k \varpi} \right) \right) - \frac{N}{\eta} \log \left(\frac{k \varpi}{k \varpi + \sigma \eta} \left(1 + \eta \sigma \sum_{i=1}^N \frac{1}{k \varpi + \sigma \gamma_i} \right) \right)$. Then, the functions Φ_q^j , $j \in \{a, b\}$, admit a maximum z^* given by

$$z^* := \frac{1}{N} \left(c + \frac{1}{\eta} \log \left(\frac{v_0}{\sum_{i \in \mathcal{G}} v_i} \right) + \frac{1}{\eta} \log \left(\frac{k \varpi}{k \varpi + \sigma \eta} \text{Card}(\mathcal{G}) \left(1 + \eta \sigma \sum_{i=1}^N \frac{1}{k \varpi + \sigma \gamma_i} \right) \right) \right).$$

Moreover

$$\Phi_q^j(z^*) = -C v_0 \exp \left(\frac{k \varpi}{\sigma \eta} \log \left(\frac{v_0}{\sum_{i \in \mathcal{G}} v_i} \right) \right),$$

where

$$C := A \exp \left(-\frac{k}{\sigma} (c(1 - \varpi) - \frac{\varpi}{\eta} \log \left(\frac{k \varpi}{k \varpi + \sigma \eta} \text{Card}(\mathcal{G}) \left(1 + \eta \sigma \sum_{i=1}^N \frac{1}{k \varpi + \sigma \gamma_i} \right) \right) \right) \\ + \varpi \sum_{i=1}^N \gamma_i^{-1} \log \left(1 + \frac{\sigma \gamma_i}{k \varpi} \right) \right) \times \frac{\sigma \eta}{k \varpi + \sigma \eta} \left(1 + \eta \sigma \sum_{i=1}^N \frac{1}{k \varpi + \sigma \gamma_i} \right).$$

¹¹Uniqueness follows from strict concavity of the vector h^i with respect to $d \in \mathcal{B}_\infty^{2N}$

2.A.5 Proof of Lemma 2.3

As the state variables $q^i, i \in \{1, \dots, N\}$ live in a discrete compact set, PDE (2.19) is in fact a system of $(2\bar{q} + 1)^N$ ordinary differential equations. Hence, a use of Cauchy-Lipschitz theorem will provide existence and unicity. We define

$$\mathcal{S} := \{x \in \mathbb{R} : l_b < x < u_b < 0, (l_b, u_b) \in \mathbb{R}_-^{*2}\}.$$

Fix some subsets I , and J of $\{1, \dots, N\}$, as well as vectors $(x_\oplus, x_\ominus, x) \in \mathbb{R}^{\text{Card}(I)} \times \mathbb{R}^{\text{Card}(J)} \times \mathcal{S}$. Then, we introduce for $q \in \mathcal{Q}^N$, the map

$$T_q(x, x_\oplus, x_\ominus) = -xC^S(q) + xC\left(\left(\frac{x}{\sum_{j \in J} x_\ominus^j}\right)^{\frac{k\varpi}{\sigma\eta}} + \left(\frac{x}{\sum_{i \in I} x_\oplus^i}\right)^{\frac{k\varpi}{\sigma\eta}}\right),$$

where $C^S(q)$, and C come from (2.19). We now show that this application is Lipschitz. Direct computations show that for any $(i, j) \in I \times J$

$$\begin{aligned} \partial_x T_q &= -C^S(q) + \left(1 + \frac{k\varpi}{\sigma\eta}\right) C\left(\left(\frac{x}{\sum_{j \in J} x_\ominus^j}\right)^{\frac{k\varpi}{\sigma\eta}} + \left(\frac{x}{\sum_{i \in I} x_\oplus^i}\right)^{\frac{k\varpi}{\sigma\eta}}\right), \\ \partial_{x_\oplus^i} T_q &= -\frac{k\varpi}{\sigma\eta} C\left(\frac{x}{\sum_{i \in I} x_\oplus^i}\right)^{1+\frac{k\varpi}{\sigma\eta}}, \quad \partial_{x_\ominus^j} T_q = -\frac{k\varpi}{\sigma\eta} C\left(\frac{x}{\sum_{j \in J} x_\ominus^j}\right)^{1+\frac{k\varpi}{\sigma\eta}}. \end{aligned}$$

By the fact that $(x, x_\oplus, x_\ominus) \in \mathcal{S} \times \mathcal{S}^{\#I} \times \mathcal{S}^{\#J}$, the gradient of T_q is uniformly bounded (in the $\|\cdot\|_\infty$ sense).

2.A.6 Proof of Theorem 2.2

We begin this section with a technical lemma.

Lemma 2.8. *Let $Z \in \mathcal{Z}$, and define $\xi := Y_T^{0,Z,\Delta(Z,Q)}$. We define*

$$K_t^Z := \exp\left(-\eta(c(N_t^a + N_t^b) - Y_t^{0,Z,\Delta} \cdot 1_N)\right), \quad t \in [0, T].$$

There exists $C > 0$, and $\epsilon > 0$ such that

$$\mathbb{E}^{\Delta(Z,Q)} \left[\sup_{t \in [0, T]} |K_t^Z|^{1+\epsilon} \right] \leq C.$$

Proof. We define for all $i \in \{1, \dots, N\}$ the processes

$$\begin{aligned} \bar{Y}_t^{i,Z,\Delta} &:= Y_t^{i,0,Z,\Delta} + \int_0^t \delta_u^{a,i} \mathbf{1}_{\{\delta_u^{i,a} = \underline{\delta}_u^a\}} dN_u^a + \delta_u^{i,b} \mathbf{1}_{\{\delta_u^{i,b} = \underline{\delta}_u^b\}} dN_u^b \\ &\quad + Q_u^i dS_u + \sum_{\ell=1}^K \omega_\ell (\delta_u^{i,a} \mathbf{1}_{\{\delta_u^{i,a} \in K_\ell\}} dN_u^a + \delta_u^{i,b} \mathbf{1}_{\{\delta_u^{i,b} \in K_\ell\}} dN_u^b), \end{aligned}$$

and we rewrite

$$\begin{aligned} K_t^Z &= \exp\left(\eta\left(\sum_{i=1}^N \bar{Y}_t^{i,Z,\Delta}\right)\right) \exp\left(-\eta\left(c(N_t^a + N_t^b) + \sum_{i=1}^N \left(\int_0^t \delta_u^{a,i} \mathbf{1}_{\{\delta_u^{i,a} = \underline{\delta}_u^a\}} dN_u^a + \delta_u^{i,b} \mathbf{1}_{\{\delta_u^{i,b} = \underline{\delta}_u^b\}} dN_u^b \right. \right. \right. \\ &\quad \left. \left. \left. + Q_u^i dS_u + \sum_{\ell=1}^K \omega_\ell (\delta_u^{i,a} \mathbf{1}_{\{\delta_u^{i,a} \in K_\ell\}} dN_u^a + \delta_u^{i,b} \mathbf{1}_{\{\delta_u^{i,b} \in K_\ell\}} dN_u^b)\right)\right)\right). \end{aligned}$$

Using Step 2 of the proof of Theorem 2.1, we know that for all $i \in \{1, \dots, N\}$, $e^{-\gamma_i \bar{Y}_t^{i,Z,\Delta}}$ is a \mathbb{P}^{δ} -martingale. Using Jensen's inequality with the convex function (on \mathbb{R}_+^*) $\phi_i(x) = x^{-\frac{\eta}{\gamma_i}}$, and condition (2.8) with $\xi^i = \bar{Y}_T^{i,Z,\Delta}$, we have

$$\mathbb{E}_t^{\delta} \left[e^{\eta \bar{Y}_T^{i,Z,\Delta}} \right] = \mathbb{E}_t^{\delta} \left[\phi_i \left(e^{-\gamma_i \bar{Y}_T^{i,Z,\Delta}} \right) \right] \geq \phi_i \left(e^{-\gamma_i \bar{Y}_t^{i,Z,\Delta}} \right) = e^{\eta \bar{Y}_t^{i,Z,\Delta}}.$$

Similar computations show that $\left(e^{\eta \bar{Y}_t^{i,Z,\Delta}} \right)_{t \in [0,T]}$ is a positive \mathbb{P}^{δ} -submartingale. By using Jensen's inequality, we have

$$\mathbb{E}^{\Delta(Z,Q)} \left[\sup_{t \in [0,T]} \exp \left(\eta' \left(\sum_{i=1}^N \bar{Y}_t^{i,Z,\Delta} \right) \right) \right] \leq \frac{1}{N} \sum_{i=1}^N \mathbb{E}^{\Delta(Z,Q)} \left[\sup_{t \in [0,T]} \exp \left(N \eta' \bar{Y}_t^{i,Z,\Delta} \right) \right],$$

where $\eta' = \eta(1 + \epsilon)$. Using Doob's inequality, there exists positive constants k_1^i independent from $t \in [0, T]$ such that for all $i \in \{1, \dots, N\}$,

$$\mathbb{E}^{\Delta(Z,Q)} \left[\sup_{t \in [0,T]} \exp \left(N \eta' \bar{Y}_t^{i,Z,\Delta} \right) \right] \leq k_1^i \mathbb{E}^{\Delta(Z,Q)} \left[\exp \left(N \eta' \bar{Y}_T^{i,Z,\Delta} \right) \right].$$

Using Holder's inequality, and noting that

$$\begin{aligned} & \exp \left(N \eta' \left(\int_0^t \delta_u^{a,i} \mathbf{1}_{\{\delta_u^{i,a} = \underline{\delta}_u^a\}} dN_u^a + \delta_u^{i,b} \mathbf{1}_{\{\delta_u^{i,b} = \underline{\delta}_u^b\}} dN_u^b + Q_u^i dS_u \right. \right. \\ & \left. \left. + \sum_{\ell=1}^K \omega_{\ell} (\delta_u^{i,a} \mathbf{1}_{\{\delta_u^{i,a} \in K_{\ell}\}} dN_u^a + \delta_u^{i,b} \mathbf{1}_{\{\delta_u^{i,b} \in K_{\ell}\}} dN_u^b) \right) \right), \end{aligned}$$

has moments of all orders by boundedness of the intensities of N^a, N^b , there exists $\eta'' > \eta'$ and a constant $k_2^i > 0$ such that

$$k_1^i \mathbb{E}^{\Delta(Z,Q)} \left[\exp \left(N \eta' \bar{Y}_T^{i,Z,\Delta} \right) \right] \leq k_2^i \mathbb{E}^{\Delta(Z,Q)} \left[\exp \left(N \eta'' \bar{Y}_T^{i,Z,\Delta} \right) \right] < +\infty,$$

where we used Condition (2.10) with $\xi^i = \bar{Y}_T^{i,Z,\Delta}$. The conclusion follows using again the boundedness of the intensities of point processes in the definition of K^Z . \square

To prove Theorem 2.2, we verify that the function v introduced in (2.18) coincides at $(0, Q_0)$ with the value function of the reduced exchange problem with maximum achieved at the optimum $z^*(t, Q_t)$ in (2.2).

The function v is negative bounded. Moreover, since $\delta_{\infty} \geq \Delta_{\infty}$, it follows that v is a solution of (2.18). A direct application of Itô's formula coupled with substitution of (2.18) leads to

$$\begin{aligned} \frac{d(v(t, Q_t) K_t^Z)}{K_t^Z} &= (h_t^Z - \mathcal{H}_t) dt + \eta \sum_{i=1}^N \left(v(t, Q_t) Z_t^{S,i} dS_t \right. \\ & \left. + \sum_{j \in \{a,b\}} \left(v(t, Q_{t-}^i + \phi(j)) e^{\eta(NZ^j - c)} - v(t, Q_{t-}) \right) d\tilde{N}_t^{\Delta,i,j} \right), \end{aligned} \quad (2.26)$$

where

$$\mathcal{H}_t := \mathcal{H}(Q_t, \mathcal{V}^+(t, q), \mathcal{V}^-(t, q), v(t, Q_t)) = \sup_{Z \in \mathcal{Z}} h_t^Z.$$

By the fact that v is bounded and $(K_t^Z)_{t \in [0, T]}$ is of class (D) , the process $(v(t, Q_t)K_t^Z)_{t \in [0, T]}$ is a $\mathbb{P}^{\Delta(Z, Q)}$ -supermartingale of class (D) and the local martingale term in (2.26) is a true martingale. Hence

$$v(0, Q_0) = \mathbb{E}^{\Delta(Z, Q)} \left[v(T, Q_T)K_T^Z + \int_0^T K_t^Z (\mathcal{H}_t - h_t) dt \right] \geq \mathbb{E}^{\Delta(Z, Q)} [v(T, Q_T)K_T^Z] = \mathbb{E}^{\Delta(Z, Q)} [-K_T^Z],$$

by the boundary condition $v(T, \cdot) = -1$. By arbitrariness of $Z \in \mathcal{Z}$, this provides the inequality

$$v(0, Q_0) \geq \sup_{Z \in \mathcal{Z}} \mathbb{E}^{\Delta(Z, Q)} [-K_T^Z] = v_0^E.$$

On the other hand, consider the maximizer $Z^*(t, Q_{t-})$ in (2.2). As $(Z^*)_{t \in [0, T]}$ is a bounded process, integrability conditions (2.8) and (2.10) are satisfied. Hence, $Z^* \in \mathcal{Z}$. By definition,

$$h^{Z^*} - \mathcal{H} = 0,$$

thus leading to

$$v(0, Q_0) = \mathbb{E}^{\Delta(Z^*, Q)} [-K_T^{Z^*}].$$

Hence, $v(0, Q_0) = v_0^E$, with optimal control Z^* .

2.A.7 First-best exchange problem

In this section, we consider the case of the first best problem. In this particular setting, the principal can control both the spreads quoted by the agents and the contracts given to them. Hence, the exchange manages all the control processes. The goal of this section is to show that the first best problem differs from the second best that we solved throughout this chapter. We first introduce the Lagrange multipliers $\lambda := (\lambda_i)_{i \in \{1, \dots, N\}}$ associated to the participation constraints of the agents. For any finite dimensional vector space E , with given norm $\|\cdot\|_E$, we also introduce the so-called Morse-Transue space on a given probability space $(\Omega, \mathcal{F}, \mathbb{P})$, defined by

$$M^\phi(E) := \left\{ \xi : \Omega \longrightarrow E \text{ measurable, } \mathbb{E}[\phi(a\xi)] < +\infty, \text{ for any } a \geq 0 \right\},$$

where $\phi : E \longrightarrow \mathbb{R}$ is the Young function, namely $\phi(x) = \exp(\|x\|_E) - 1$. Then, if $M^\phi(E)$ is endowed with the norm $\|\xi\|_\phi := \inf\{k > 0, \mathbb{E}[\phi(\xi/k)] \leq 1\}$ it is a (non-reflexive) Banach space.

The principal's problem can be reformulated as

$$V_0^{\text{FB}} := \inf_{\lambda > 0} \sup_{(\xi, \delta) \in \mathcal{C} \times \mathcal{A}} \mathbb{E}^\delta \left[-e^{-\eta(c(N_T^a + N_T^b) - \xi \cdot 1_N)} - \sum_{i=1}^N \lambda_i e^{-\gamma_i(\xi^i + X_T^i + Q_T^i S_T)} - \lambda_i R_i \right]. \quad (2.27)$$

If λ and δ are fixed, we start with the maximisation with respect to ξ . We introduce the following map $\Lambda^\delta : M^\phi(\mathbb{R}^N) \rightarrow \mathbb{R}$ defined as

$$\Lambda^\delta(\xi) := \mathbb{E}^\delta \left[-e^{-\eta(c(N_T^a + N_T^b) - \xi \cdot 1_N)} - \sum_{i=1}^N \lambda_i e^{-\gamma_i(\xi^i + X_T^i + Q_T^i S_T)} - \lambda_i R_i \right]$$

The Lagrange multipliers λ_i being strictly positive, using the boundedness of the control process $\delta \in \mathcal{A}$, the map Λ^δ is continuous, strictly concave, and Gâteaux differentiable with, for $h \in M^\phi(\mathbb{R}^N)$

$$D\Lambda^\delta(\xi)[h] = \mathbb{E}^\delta \left[-\eta h \cdot 1_N e^{-\eta(c(N_T^a + N_T^b) - \xi \cdot 1_N)} + \sum_{i=1}^N \gamma_i \lambda_i h^i e^{-\gamma_i(\xi^i + X_T^i + Q_T^i S_T)} \right]$$

For any $\delta \in \mathcal{A}$, we define

$$\begin{aligned}\xi^{*,i}(\delta) &:= \frac{\eta}{\gamma_i} (c(N_T^a + N_T^b) - \xi^* \cdot \mathbf{1}_N) + \frac{1}{\gamma_i} \log \left(\frac{\lambda_i \gamma_i}{\eta} \right) - (X_T^i + Q_T^i S_T), \\ \xi^*(\delta) \cdot \mathbf{1}_N &= \frac{1}{1 + \eta \Gamma} \left(\eta c \Gamma (N_T^a + N_T^b) + \sum_{i=1}^N \frac{1}{\gamma_i} \log \left(\frac{\lambda_i \gamma_i}{\eta} \right) - (X_T^i + Q_T^i S_T) \right)\end{aligned}$$

with $\Gamma := \sum_{i=1}^N \frac{1}{\gamma_i}$. For any $h \in M^\phi(\mathbb{R}^N)$, first order condition gives $D\Lambda^\delta(\xi)[h] = 0$. Computations show that $(\xi^{*,i})_{i=1,\dots,N}$ achieve the maximum of $\Lambda^\delta(\xi)$, hence is optimal for (2.27). Then, substituting these expressions in the main problem gives

$$V_0^{\text{FB}} = (1 + \eta \Gamma) \inf_{\lambda > 0} \prod_{i=1}^N \left(\frac{\lambda_i \gamma_i}{\eta} \right)^{\eta \gamma_i^{-1} (1 + \eta \Gamma)^{-1}} \tilde{V}_0 - \sum_{i=1}^N \lambda_i R_i, \quad (2.28)$$

where

$$\tilde{V}_0 := \sup_{\delta \in \mathcal{A}} \mathbb{E}^\delta \left[- \exp \left(- \frac{\eta}{1 + \eta \Gamma} \left(\sum_{i=1}^N (X_T^i + Q_T^i S_T) + c(N_T^a + N_T^b) \right) \right) \right]. \quad (2.29)$$

This is a stochastic control problem, see [109, Section A.7] for details, whose HJB equation is given by

$$\begin{cases} \partial_t v(t, q) + v(t, q) \frac{1}{2} \sigma^2 \tilde{\Gamma}^2 \|q\|^2 + \mathcal{H}^{\text{FB}}(q, \mathcal{V}^+(t, q), \mathcal{V}^-(t, q), v(t, q)) = 0, & (t, q) \in [0, T) \times \mathcal{Q}^N, \\ v(T, q) = -1, & q \in \mathcal{Q}^N, \end{cases} \quad (2.30)$$

where $\tilde{\Gamma} := \frac{\eta}{1 + \eta \Gamma}$, and

$$\mathcal{H}^{\text{FB}}(q, \mathcal{V}^+(t, q), \mathcal{V}^-(t, q), v(t, q)) := \mathcal{H}^{\text{FB},b}(q, \mathcal{V}^+(t, q), v(t, q)) + \mathcal{H}^{\text{FB},a}(q, \mathcal{V}^-(t, q), v(t, q)),$$

with, for any $(p, v, j) \in \mathbb{R}^N \times \mathbb{R} \times \{a, b\}$

$$\mathcal{H}^{\text{FB},j}(q, p, v) := \sup_{\delta^j \in \mathcal{A}} \sum_{i=1}^N \lambda^{i,j}(\delta^j, q) \left(\exp \left(- \tilde{\Gamma} \left(\sum_{i=1}^N \delta^{i,j} \mathbf{1}_{\{\delta^{i,j} = \delta^j\}} + \sum_{\ell=1}^K \omega_\ell \delta^{i,j} \mathbf{1}_{\{\delta^{i,j} \in K_\ell\}} \right) \right) p^i - v \right).$$

We are in a framework similar to El Euch, Mastrolia, Rosenbaum, and Touzi [109], Guéant, Lehalle, and Fernandez-Tapia [142]. First order condition gives for $j \in \{a, b\}$

$$\sum_{i=1}^N \delta^{*,j,i} \mathbf{1}_{\{\delta^{*,j,i} = \delta^{*,j,j}\}} + \sum_{\ell=1}^K \omega_\ell \delta^{*,j,i} \mathbf{1}_{\{\delta^{*,j,i} \in K_\ell\}} = \mathcal{P}^j(t, q) := \frac{1}{\tilde{\Gamma}} \left(\log \left(1 + \frac{\tilde{\Gamma} \sigma}{k \underline{\omega}} \right) + \log \left(\frac{\sum_{i \in \mathcal{G}} v(t, q \ominus_i \phi(j))}{v(t, q)} \right) \right).$$

Such conditions are satisfied with the following optimal bid-ask policy, for $j \in \{a, b\}$

$$\delta^{*,j,i}(t, q) := \begin{cases} (-\delta_\infty) \vee \frac{1}{\omega_\ell} \mathcal{P}^j(t, q) \wedge \delta_\infty, & \text{if } \frac{1}{\omega_\ell} \mathcal{P}^j(t, q) \in K_\ell, \text{ for } \ell \in \{1, \dots, K\}, \\ (-\delta_\infty) \vee \mathcal{P}^j(t, q) \wedge \delta_\infty, & \text{otherwise.} \end{cases}$$

Finally, computations show that the Hessian associated to the supremum in $\delta^{*,j}$ is symmetric definite negative, hence, δ^* is a local maximum.

Theorem 2.3. *There exists a unique negative bounded solution to the PDE*

$$\begin{cases} \partial_t v(t, q) + v(t, q) \left(\frac{\sigma^2}{2} \tilde{\Gamma}^2 \|q\|^2 - \tilde{C}^{\text{FB}} \left(\left(\frac{v(t, q)}{\sum_{i \in \mathcal{G}} v(t, q \oplus_i 1)} \right)^{\frac{k \underline{\omega}}{\sigma \tilde{\Gamma}}} + \left(\frac{v(t, q)}{\sum_{i \in \mathcal{G}} v(t, q \ominus_i 1)} \right)^{\frac{k \underline{\omega}}{\sigma \tilde{\Gamma}}} \right) \right) = 0, \\ v(T, q) = -1, \end{cases}$$

with $(t, q) \in [0, T) \times \mathcal{Q}^N$, $\tilde{C}^{\text{FB}} := A \exp \left(- \frac{k \underline{\omega}}{\sigma \tilde{\Gamma}} \log \left(1 + \frac{\tilde{\Gamma} \sigma}{k \underline{\omega}} \right) \right) \frac{\tilde{\Gamma} \sigma}{k \underline{\omega} + \tilde{\Gamma} \sigma}$. Moreover, this solution coincides with the value function of the exchange for the problem (2.27).

The proof is omitted as it relies on the same basis arguments as Lemma 2.3 and Theorem 2.2.

We now prove that the PDE satisfied by the value function of the First Best problem is different from the one verified in the second best case (2.19). Taking the special case $\gamma_i := \gamma$, i.e the case of market-makers with same risk aversion, for $i \in \{1, \dots, N\}$, the PDE boils down to

$$\begin{cases} \partial_t v(t, q) + v(t, q)C^{\text{FB}}(q) - v(t, q)\tilde{C}^{\text{FB}} \sum_{j \in \{a, b\}} \left(\frac{v(t, q)}{\sum_{i=1}^N \mathbf{1}_{\{\phi(j)q^i > -\bar{q}\}} v(t, q^i - \phi(j))} \right)^{\frac{k\varpi}{\sigma\eta}} = 0, \\ v(T, q) = -1, \end{cases}$$

where $(t, q) \in [0, T) \times \mathcal{Q}^N$ and

$$C^{\text{FB}}(q) := \frac{1}{2} \sigma^2 \tilde{\Gamma}^2 \|q\|^2.$$

By noting that for all $q \in \mathcal{Q}^N$, $C^{\text{FB}}(q) \neq C^S(q)$ and $C \neq \tilde{C}^{\text{FB}}$, we see that the value function of the exchange in the first best case does not coincide with the value function in the second best model.

Chapter 3

On bid and ask side-specific tick sizes

3.1 Introduction

The tick size is the smallest increment between two consecutive prices on a trading instrument. It is fixed by the exchange or regulator and typically depends on both the price of the asset and the traded volume, see Huang, Lehalle, and Rosenbaum [164], Laruelle, Rosenbaum, and Savku [195]. It is a crucial parameter of market microstructure and its value is often subject of debates: a too small tick size leads to very frequent price changes whereas a too large tick size prevents the price from moving freely according to the investor's views. In this chapter, we focus on so-called large tick assets, that is assets for which the spread is most of the time equal to one tick. Such assets represent a large number of financial products, especially in Europe since MIFID II regulation, see [195].

The tick size has a major influence on the ecosystem of financial markets, in particular on the activity of high frequency traders. Being usually considered as market-makers, these agents are the main liquidity providers for most heavily traded financial assets. This means that they propose prices at which they are ready to buy (bid price) and sell (ask price) units of financial products. In Frino, Mollica, and Zhang [122], the authors investigate the behavior of high frequency traders with respect to the relative tick size, which is defined as the ratio between the tick size and the price level. One of their findings is that everything else equal, stocks with a lower relative tick size attract a greater proportion of high frequency traders, see also Dayri and Rosenbaum [91], Megarbane, Saliba, Lehalle, and Rosenbaum [210]. This is because they can rapidly marginally adjust their quotes to seize price priority. In the case of a large tick asset, speed is still an important feature as market participants have to compete for queue priority in the order book, see Huang, Rosenbaum, and Saliba [165], Moallemi and Yuan [217].

Market-makers (typically high frequency traders) face a complex optimization problem: making money out of the bid-ask spread (the difference between the bid and ask prices) while mitigating the inventory risk associated to price changes. This problem is usually addressed via stochastic control theory tools, see for example Avellaneda and Stoikov [26], Cartea, Jaimungal, and Penalva [70], Cartea, Jaimungal, and Ricci [69], Guéant [137], Guéant, Lehalle, and Fernandez-Tapia [142]. In classical market-making models, the so-called efficient price, which represents the market consensus on the value of the asset at a given time, around which the market-maker posts his quotes, is a continuous semi-martingale. The quotes of the market-maker are continuous in terms of price values and not necessarily multiple of the tick size. However, in actual financial markets, transaction prices are obviously lying on the discrete tick grid. This discreteness of prices is a key feature which cannot be neglected at the high frequency scale since it plays a fundamental role in the design of market-making strategies in practice. To get a more realistic market-making model, one therefore needs to build a relevant continuous-time price dynamic with discrete state space to take into account this very important microstructural property of the asset.

To this end, we borrow the framework of the model with uncertainty zones introduced in Robert and Rosenbaum [238, 239]. In this model, transaction prices are discrete and the current transaction price is modified only when the underlying continuous efficient price process crosses some predetermined zones. In our approach, we also consider that there exists an efficient price that market participants have in mind when making their trading decisions. Based on this efficient price, market participants build “fair” bid and “fair” ask prices. These two prices are lying on the tick grid and represent the views of market participants on reasonable and tradable values for buying and selling, regardless of any inventory constraint. In our setting, depending on his views and his inventory constraint, the market-maker chooses whether or not to quote a constant volume at these fair bid and ask prices. This is a stylized viewpoint as in practice the market-maker will probably quote a larger spread rather than not quoting at all. It is particularly reasonable for large tick assets for which the spread is almost always equal to one tick. The market-takers increase (resp. decrease) their current “fair” bid price if the efficient price becomes “sufficiently” higher (resp. lower) than their current fair bid price and similarly for the ask side. The mechanism to determine whether the efficient price is sufficiently higher (resp. lower) than the current price is that of the model with uncertainty zones, described in Section 3.2.

Usual market-making models include a symmetric running penalty for the inventory process, often defined as $\phi \int_0^T Q_t^2 dt$ where Q_t is the inventory of the market-maker at time $t \in [0, T]$, $\phi > 0$ is a risk aversion parameter and T is the end of the trading period. It is well-known, see for example Adrian, Capponi, Fleming, Vogt, and Zhang [2], that for regulatory and operational reasons, market participants and especially market-makers are reluctant to have a short inventory at the end of the trading day. This is mainly due to constraints imposed by the exchange/regulator and to the overnight repo rate that they have to pay. This asymmetry between long and short terminal inventory of the market-maker gives the intuition of the potential relevance of some kind of asymmetry in the market design between buy and sell orders.

If some kind of asymmetry is implemented at the microstructure level, it can have important consequences on the profit of exchanges, as it notably depends on the number of processed orders. Typical ways to optimize the number of orders on platforms are the choice of relevant tick sizes and suitable fee schedules (which subsidize liquidity provision and tax liquidity consumption). In Foucault, Kadan, and Kandel [121], the authors highlight the importance of differentiating maker and taker fees in order to increase the trading rate. In the more recent studies El Euch, Mastrolia, Rosenbaum, and Touzi [109] and Chapter 2 of this thesis, optimal make-take fees schedules are designed based on contract theory. In this work, the asymmetry we consider is not between liquidity consumers and liquidity providers but between buyers and sellers.

The goal of this chapter is to show the possible benefits for an exchange in terms of liquidity provision of side-specific tick sizes. To this end, we build an agent-based model where a high frequency market-maker acts on a large tick asset. The exchange is mitigating the activity on its platform by choosing suitable tick sizes on the bid and ask sides. This means we have a different tick grid for buy and sell orders. For given the tick sizes chosen by the exchange, we formulate the stochastic control problem faced by the market-maker who needs to maximize his Profit and Loss (PnL for short) while controlling his inventory risk, taking into account asymmetry between short and long inventory. We show existence and uniqueness of a viscosity solution to the Hamilton-Jacobi-Bellman (HJB for short) equation associated to this problem. Then, we derive a quasi-closed form for the optimal controls of the market-maker (up to the value function). In particular, the role of the tick size in the decision of whether or not to quote is explicit: essentially, a large tick size implies a large profit per trade for the market-maker but less market orders coming from market-takers, and conversely.

Next, we solve the optimization problem of the exchange which can select optimal tick sizes knowing the associated trading response of the market-maker. In our model, the exchange earns a fixed fee when

a transaction occurs. Therefore, its remuneration is related to the quality of the liquidity provided by the market-maker on its platform. Numerical results show that side-specific tick sizes are more suitable than symmetric ones both for the market-maker and the exchange. The former is able to trigger more alternations in the sign of market orders, which is beneficial both for spread pocketing and inventory management (in contrast with the case where sequences of buy orders are followed by sequences of sell orders). The latter increases the number of transactions on its platform. We also show that a tick size asymmetry can offset short inventory constraints, therefore increasing the gains of both the market-maker and the exchange.

The chapter is organized as follows. In Section 3.2, we give a reminder on the model with uncertainty zones and explain how we revisit it for market-making purposes. The market-maker and exchange's problems are described in Section 3.3. We also state here our results about existence and uniqueness of a viscosity solution associated to the control problem of the market-maker and derive its optimal controls. Finally, Section 3.4 is devoted to numerical results and their interpretations. Proofs are relegated to an appendix.

3.2 The model with uncertainty zones

In this section, we provide a reminder on the model with uncertainty zones introduced in Robert and Rosenbaum [238, 239], and we adapt it to the framework of a market-making problem with side-specific tick values. It is commonly admitted that low frequency financial price data behave like a continuous Brownian semi-martingale. However this is clearly not the case for high frequency data. The model with uncertainty zones reproduces sparingly and accurately the behavior of ultra high frequency transaction data of a large tick asset. It is based on a continuous-time semi-martingale efficient price and a one dimensional parameter $\eta \in [0, \frac{1}{2}]$. The key idea of the model is that when a transaction occurs at some value on the tick grid, the efficient price is close enough to this value at the transaction time. This proximity is measured through the parameter η .

We define the efficient price $(S_t)_{t \in [0, T]}$ on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$ where T is the trading horizon. The logarithm of the efficient price $(Y_t)_{t \in [0, T]}$ is an \mathcal{F}_t -adapted continuous Brownian semi-martingale of the form

$$Y_t = \log(S_t) = \log(S_0) + \int_0^t a_s ds + \int_0^t \sigma_s^- dW_s,$$

where W is an \mathcal{F} -Brownian motion, and $(\sigma_t)_{t \in [0, T]}$ is an \mathcal{F} -adapted process with càdlàg paths and $(a_t)_{t \in [0, T]}$ is \mathcal{F} -progressively measurable. Transaction prices lie on two fixed tick grids, defined by $\{k\alpha^a, k\alpha^b\}$ where α^a (resp. α^b) is the tick size on the ask (resp. bid) side and $k \in \mathbb{Z}$. For $0 \leq \eta^i \leq \frac{1}{2}$ and $i \in \{a, b\}$, we define the zone $U_k^i = [0, \infty) \times (d_k^i, u_k^i)$ with

$$d_k^i = \left(k + \frac{1}{2} - \eta^i\right)\alpha^i, u_k^i = \left(k + \frac{1}{2} + \eta^i\right)\alpha^i. \quad (3.1)$$

Therefore U_k^a is a band of size $2\eta^a\alpha^a$ around the ask mid-tick grid value $(k + \frac{1}{2})\alpha^a$ and U_k^b is a band of size $2\eta^b\alpha^b$ around the bid mid-tick grid value $(k + \frac{1}{2})\alpha^b$. We call these bands the uncertainty zones. The zones on the bid and ask sides are characterized by the parameters η^b, η^a which control the width of the uncertainty zones. We will see in the next section how the fair bid and ask prices are deduced from the efficient price dynamics across the uncertainty zones. In particular, the larger η^i , the farther from the last traded price (on the bid or ask side) the efficient price has to be so that a price change occurs. The idea behind the model with uncertainty zones is that, in some sense, market participants feel more comfortable when the asset price is constant than when it is constantly moving. However, there are times when the transaction price has to change because they consider that the last traded price value is not reasonable anymore.

For sake of simplicity, we assume that transaction prices cannot jump by more than one tick. We also define the time series of bid and ask transaction times leading to a price change as $(\tau_j^b, \tau_j^a)_{j \geq 0}$. The last traded bid or ask price process is characterized by the couples of transaction times and transaction prices with price changes $(\tau_j; P_{\tau_j}^i)_{j \geq 0}$ where $P_{\tau_j}^i = S_{\tau_j}^{(\alpha^i)}$, the superscript (α^i) denoting the rounding to the nearest α^i .

Remark 18. *The trade price goes up by one tick when the efficient price process S crosses the uncertainty zone in the upward direction, and conversely it goes down by one tick when it crosses the uncertainty zone in the downward direction.*

One can actually show that the efficient price can be retrieved from transaction data using the equation

$$S_{\tau_j^i} = S_{\tau_j^i}^{(\alpha^i)} - \alpha^i \left(\frac{1}{2} - \eta^i \right) \operatorname{sgn} \left(S_{\tau_j^i}^{(\alpha^i)} - S_{\tau_{j-1}^i}^{(\alpha^i)} \right), \quad i \in \{a, b\}, j \in \mathbb{N}.$$

This formula is particularly useful in order to derive ultra high frequency estimators of volatility and covariation (see Robert and Rosenbaum [239]). The parameters η^i can be estimated very easily. Let $N_{\alpha^i, t}^{(a)}$ and $N_{\alpha^i, t}^{(c)}$ be respectively the number of alternations and continuations¹ of one tick over the period $[0, t]$. Then, an estimator of η^i over $[0, t]$ is given by

$$\hat{\eta}_{\alpha^i, t} = \frac{N_{\alpha^i, t}^{(c)}}{2N_{\alpha^i, t}^{(a)}}.$$

We refer to Robert and Rosenbaum [238, 239] for further details on these estimation procedures. In this chapter, we use the model with uncertainty zones for market-making purposes rather than for statistical estimation.

3.3 High frequency market-making under side-specific tick values and interaction with the exchange

3.3.1 The market-maker's problem

We consider a high frequency market maker acting on an asset whose efficient price S_t has the dynamics

$$dS_t = \sigma dW_t,$$

where $\sigma > 0$ denotes the volatility of the asset. He uses the model with uncertainty zones described earlier to materialize his views on the fair bid and ask prices. He increases (resp. decreases) his bid price if the efficient price is “sufficiently” higher (resp. lower) than his current fair bid price. The notion of “sufficiently” higher or lower is determined by the uncertainty zones parameters η^a, η^b , and the tick sizes α^a, α^b . If η^a is small (resp. large), the market-maker changes more (resp. less) frequently his ask price, and similarly for the bid price with η^b . This leads to the following definition of fair bid and ask prices of the market-maker S^a, S^b :²

$$\begin{aligned} S_t^a &= S_{t-}^a + \alpha^a \mathbf{1}_{\{S_t - S_{t-}^a > (\frac{1}{2} + \eta^a) \alpha^a\}} - \alpha^a \mathbf{1}_{\{S_t - S_{t-}^a < -(\frac{1}{2} + \eta^a) \alpha^a\}}, \\ S_t^b &= S_{t-}^b + \alpha^b \mathbf{1}_{\{S_t - S_{t-}^b > (\frac{1}{2} + \eta^b) \alpha^b\}} - \alpha^b \mathbf{1}_{\{S_t - S_{t-}^b < -(\frac{1}{2} + \eta^b) \alpha^b\}}. \end{aligned}$$

Thus the fair bid (resp. ask) is modified when the efficient price is close enough to a new tradable price on the tick grid with mesh α^b (resp. α^a).

¹An alternation/continuation corresponds to two consecutive price changes in the opposite/same direction.

²Note that we can have situations where the bid price is above the ask price. However, recall that S^a and S^b are only views about the fair bid and ask prices under the constraint that they have to lie on the tick grids. This does not imply an arbitrage opportunity as in this configuration the market-maker would not quote on at least one of the two sides.

Remark 19. Note that in the case $\alpha^a = \alpha^b, \eta^a = \eta^b$, the fair best bid is equal to the fair best ask. This means that at a given time, a buy or sell order would be at the same price. In this situation, in our stylized view, the market-maker would probably quote only on one side (bid or ask). It is consistent with the standard form of the model with uncertainty zones, where, at a given time, transactions can only happen only on one side of the market, depending on the location of the efficient price. Still, the market-maker collects the spread from transactions occurring at different times as it is the case in practice.

We assume a constant volume of transaction equal to one. The market-maker can choose to be present or not for a transaction at the bid (with a price S^b) or at the ask (with a price S^a). The corresponding cash process at terminal time T is given by

$$X_T = \int_0^T \left(S_t^a dN_t^a - S_t^b dN_t^b \right),$$

where the N_t^i represent the number of transactions on the bid or ask side between 0 and t . In this framework, the inventory of the market-maker is given by $Q_t = N_t^b - N_t^a \in \mathcal{Q} = [-\tilde{q}, \tilde{q}]$ where \tilde{q} is the risk limit of the market-maker. For $i \in \{a, b\}$, the dynamics of N_t^i is that of a point process with intensity

$$\lambda(\ell_t^i, Q_t) := \frac{\lambda \ell_t^i}{1 + (\kappa \alpha^i)^2} \mathbf{1}_{\{\phi(i)Q_t > -\tilde{q}\}}, \quad \phi(i) = \mathbf{1}_{\{i=a\}} - \mathbf{1}_{\{i=b\}}.$$

The process $\ell_t^i \in \{0, 1\}$ is the market-maker's control which lies in the set of \mathcal{F} – predictable processes with values in $\{0, 1\}$ denoted by \mathcal{L} . The parameter $\kappa > 0$ controls the sensitivity of the intensities to α^i , and $\lambda > 0$ is a scale parameter. When the market-maker does not want to be present on the bid (resp. ask side) at the price S^b (resp. S^a) he sets $\ell^b = 0$ (resp. $\ell^a = 0$) and conversely. In our large tick asset setting, the situation where the market-maker is not present is a simplified way to model the case where the market-maker's quote is higher than the best possible limit. At a given time $t \in [0, T]$, when $\ell_t^b = 0$ (resp. $\ell_t^a = 0$), the intensity of the point process N_t^b (resp. N_t^a) is equal to zero so that there are no incoming transactions. In addition to this, market-takers are more confident to send market orders when the tick size is small, as the market-maker has more flexibility to adjust his bid and ask prices.³ This explains the decreasing shape of the intensities of market order arrivals from market-takers with respect to the tick size. The chosen parametric form for the intensities ensures no degenerate behavior when the tick size gets close to zero. The parameter κ and the quadratic dependence of the intensity with respect to the tick size have no strong influence on the qualitative results obtained in Section 3.4.

The marked-to-market value of the market-maker's portfolio at time t is defined as $Q_t S_t$. His optimization problem writes

$$\sup_{\ell \in \mathcal{L}} \mathbb{E} \left[X_T + Q_T (S_T - A Q_T) - \phi \int_0^T Q_s^2 ds - \phi_- \int_0^T |Q_s|^2 \mathbf{1}_{Q_s < 0} ds \right], \quad (3.2)$$

where $\phi > 0$ represents the risk-aversion parameter of the market-maker, $\phi_- > 0$ is the additional risk aversion of the market-maker toward short position on $[0, T]$ and $A Q_T^2$, with $A > 0$, is a penalty term for the terminal inventory position regardless of its sign. In this setting, the market-maker wishes to hold a terminal inventory close to zero because of the quadratic penalty $A Q_T^2$. The term $\phi \int_0^T Q_s^2 ds$ penalizes long or short positions over the trading period. In the following, we will choose T large enough so that the final penalization has little impact on the results. Problem (3.2) can of course be rewritten as

$$\sup_{\ell \in \mathcal{L}} \mathbb{E} \left[Q_T (S_T - A Q_T) + \int_0^T \left(S_s^a \lambda(\ell_s^a) - S_s^b \lambda(\ell_s^b) - \phi Q_s^2 - \phi_- Q_s^2 \mathbf{1}_{Q_s < 0} \right) ds \right].$$

³When the tick size is smaller, the market-takers are more willing to trade. This does not necessarily lead to a higher number of orders as it depends on the market-maker's presence.

We define the corresponding value function h defined on the open set

$$\mathcal{D} = \left\{ (S^a, S^b, S) \in \alpha^a \mathbb{Z} \times \alpha^b \mathbb{Z} \times \mathbb{R} : -\left(\frac{1}{2} + \eta^a\right) \alpha^a < S - S^a < \left(\frac{1}{2} + \eta^a\right) \alpha^a \right. \\ \left. \text{and } -\left(\frac{1}{2} + \eta^b\right) \alpha^b < S - S^b < \left(\frac{1}{2} + \eta^b\right) \alpha^b \right\}$$

by

$$h(t, S^a, S^b, S, q) = \sup_{\ell \in \mathcal{L}_t} \mathbb{E}_{t, S^a, S^b, S, q} \left[Q_T (S_T - A Q_T) + \int_t^T \left(S_s^a \lambda(\ell_s^a) - S_s^b \lambda(\ell_s^b) \right. \right. \\ \left. \left. - \phi Q_s^2 - \phi_- Q_s^2 \mathbf{1}_{Q_s < 0} \right) ds \right], \quad (3.3)$$

where \mathcal{L}_t denotes the restriction of admissible controls to $[t, T]$ and $\mathbb{E}_{t, S^a, S^b, S, q}[\cdot] = \mathbb{E}[\cdot | S_t^a = S^a, S_t^b = S^b, S_t = S, Q_t = q]$. We define the boundary $\partial \mathcal{D}$ of \mathcal{D} as

$$\partial \mathcal{D} = \left\{ (S^a, S^b, S) \in \alpha^a \mathbb{Z} \times \alpha^b \mathbb{Z} \times \mathbb{R} : S - S^a = \pm \left(\frac{1}{2} + \eta^a\right) \alpha^a \text{ and/or } S - S^b = \pm \left(\frac{1}{2} + \eta^b\right) \alpha^b \right\},$$

and write $\bar{\mathcal{D}} = \mathcal{D} \cup \partial \mathcal{D}$. For given (S^a, S^b) , if $(S^a, S^b, S) \in \partial \mathcal{D}$, it means that S corresponds to an efficient price value that triggers a modification of the fair bid or ask price.

The Hamilton-Jacobi-Bellman equation associated to this stochastic control problem is given by

$$0 = \partial_t h(t, S^a, S^b, S, q) - \phi q^2 - \phi_- q^2 \mathbf{1}_{q < 0} + \frac{1}{2} \sigma^2 \partial_{SS} h(t, S^a, S^b, S, q) \\ + \frac{\lambda}{1 + (\kappa \alpha^a)^2} \max_{\ell^a \in \{0, 1\}} \left\{ \ell^a \left(S^a + h(t, S^a, S^b, S, q - \ell^a) - h(t, S^a, S^b, S, q) \right) \right\} \\ + \frac{\lambda}{1 + (\kappa \alpha^b)^2} \max_{\ell^b \in \{0, 1\}} \left\{ \ell^b \left((-S^b) + h(t, S^a, S^b, S, q + \ell^b) - h(t, S^a, S^b, S, q) \right) \right\}, \quad (3.4)$$

for $(t, S^a, S^b, S, q) \in [0, T) \times \mathcal{D} \times \mathcal{Q}$, with terminal condition

$$h(T, S^a, S^b, S, q) = q(S - Aq). \quad (3.5)$$

Let us consider the function h defined in 3.3. For $(t, S^a, S^b, S, q) \in [0, T) \times \partial \mathcal{D} \times \mathcal{Q}$, and $(t_n, S_n^a, S_n^b, S_n, q_n)$ a sequence in $[0, T) \times \mathcal{D} \times \mathcal{Q}$ which converges to (t, S^a, S^b, S, q) , we will show that $h(t_n, S_n^a, S_n^b, S_n, q_n)$ converges independently of the sequence and we denote by $h(t, S^a, S^b, S, q)$ its limit. On $[0, T) \times \partial \mathcal{D} \times \mathcal{Q}$, we will show the following boundary conditions (which we will naturally impose for the solution of 3.4):

$$0 = \mathbf{1}_{\{S - S^a = (\frac{1}{2} + \eta^a) \alpha^a, S - S^b < (\frac{1}{2} + \eta^b) \alpha^b\}} \left(h(t, S^a + \alpha^a, S^b, S, q) - h(t, S^a, S^b, S, q) \right) \\ + \mathbf{1}_{\{S - S^a < (\frac{1}{2} + \eta^a) \alpha^a, S - S^b = (\frac{1}{2} + \eta^b) \alpha^b\}} \left(h(t, S^a, S^b + \alpha^b, S, q) - h(t, S^a, S^b, S, q) \right) \\ + \mathbf{1}_{\{S - S^a = (\frac{1}{2} + \eta^a) \alpha^a, S - S^b = (\frac{1}{2} + \eta^b) \alpha^b\}} \left(h(t, S^a + \alpha^a, S^b + \alpha^b, S, q) - h(t, S^a, S^b, S, q) \right) \\ + \mathbf{1}_{\{S - S^a = -(\frac{1}{2} + \eta^a) \alpha^a, S - S^b > -(\frac{1}{2} + \eta^b) \alpha^b\}} \left(h(t, S^a - \alpha^a, S^b, S, q) - h(t, S^a, S^b, S, q) \right) \\ + \mathbf{1}_{\{S - S^a > -(\frac{1}{2} + \eta^a) \alpha^a, S - S^b = -(\frac{1}{2} + \eta^b) \alpha^b\}} \left(h(t, S^a, S^b - \alpha^b, S, q) - h(t, S^a, S^b, S, q) \right) \\ + \mathbf{1}_{\{S - S^a = -(\frac{1}{2} + \eta^a) \alpha^a, S - S^b = -(\frac{1}{2} + \eta^b) \alpha^b\}} \left(h(t, S^a - \alpha^a, S^b - \alpha^b, S, q) - h(t, S^a, S^b, S, q) \right). \quad (3.6)$$

In other words, the value function varies continuously when the efficient price leaves an uncertainty zone and the prices S^a and S^b are modified.⁴ In the following, we say that a function defined on $[0, T) \times \mathcal{D} \times \mathcal{Q}$ satisfies the continuity conditions if it satisfies (3.6).

The following proposition is of particular importance for the existence and uniqueness of a viscosity solution associated to the control problem of the market-maker.

Proposition 3.1. *The function h defined in Equation (3.3) is continuous on \mathcal{D} and satisfies the continuity conditions (3.6).*

The proof is given in Appendix 3.A.1 and relies on the specific structure of our model based on hitting times of a Brownian motion. We now state the main theorem of this chapter, whose proof is relegated to Appendix 3.A.2.

Theorem 3.1. *The value function h is the unique continuous viscosity solution to Equation (3.4) on $[0, T) \times \mathcal{D} \times \mathcal{Q}$ with terminal condition (3.5) and satisfying the continuity conditions.*

The value function depends on five variables. However, as (S^a, S^b) takes value in $\alpha^a \mathbb{Z} \times \alpha^b \mathbb{Z}$, it can essentially be reduced to three variables as we now explain. For any $(i, j) \in \mathbb{Z}^2$, we introduce the function $h^{i,j}$ defined on

$$[0, T] \times \mathcal{D}_{i,j} \times \mathcal{Q}$$

where

$$\mathcal{D}_{i,j} = \left(\alpha^a i - \left(\frac{1}{2} + \eta^a \right) \alpha^a, \alpha^a i + \left(\frac{1}{2} + \eta^a \right) \alpha^a \right) \cap \left(\alpha^b j - \left(\frac{1}{2} + \eta^b \right) \alpha^b, \alpha^b j + \left(\frac{1}{2} + \eta^b \right) \alpha^b \right),$$

by $h^{i,j}(t, S, q) = h(t, \alpha^a i, \alpha^b j, S, q)$. Then $h^{i,j}$ is the solution of the following HJB equation:

$$\begin{aligned} 0 = & \partial_t h^{i,j}(t, S, q) - \phi q^2 - \phi_-(q)_-^2 \mathbf{1}_{q < 0} + \frac{1}{2} \sigma^2 \partial_{SS} h^{i,j}(t, S, q) \\ & + \frac{\lambda}{1 + (\kappa \alpha^a)^2} \max_{\ell^a \in \{0,1\}} \left\{ \ell^a \left(\alpha^a i + h^{i,j}(t, S, q - \ell^a) - h^{i,j}(t, S, q) \right) \right\} \\ & + \frac{\lambda}{1 + (\kappa \alpha^b)^2} \max_{\ell^b \in \{0,1\}} \left\{ \ell^b \left(-\alpha^b j + h^{i,j}(t, S, q + \ell^b) - h^{i,j}(t, S, q) \right) \right\}, \end{aligned}$$

with terminal condition $h^{i,j}(T, S, q) = q(S - A\mathcal{Q})$ and natural Dirichlet boundary conditions for $S \in \partial \mathcal{D}_{i,j}$:

$$\begin{aligned} h^{i,j}(t, S, q) = & h^{i+1,j}(t, S, q) \mathbf{1}_{\{S - \alpha^a i = (\frac{1}{2} + \eta^a) \alpha^a, S - \alpha^b j < (\frac{1}{2} + \eta^b) \alpha^b\}} \\ & + h^{i,j+1}(t, S, q) \mathbf{1}_{\{S - \alpha^a i < (\frac{1}{2} + \eta^a) \alpha^a, S - \alpha^b j = (\frac{1}{2} + \eta^b) \alpha^b\}} \\ & + h^{i+1,j+1}(t, S, q) \mathbf{1}_{\{S - \alpha^a i = (\frac{1}{2} + \eta^a) \alpha^a, S - \alpha^b j = (\frac{1}{2} + \eta^b) \alpha^b\}} \\ & + h^{i-1,j}(t, S, q) \mathbf{1}_{\{S - \alpha^a i = -(\frac{1}{2} + \eta^a) \alpha^a, S - \alpha^b j > -(\frac{1}{2} + \eta^b) \alpha^b\}} \\ & + h^{i,j-1}(t, S, q) \mathbf{1}_{\{S - \alpha^a i > -(\frac{1}{2} + \eta^a) \alpha^a, S - \alpha^b j = -(\frac{1}{2} + \eta^b) \alpha^b\}} \\ & + h^{i-1,j-1}(t, S, q) \mathbf{1}_{\{S - \alpha^a i = -(\frac{1}{2} + \eta^a) \alpha^a, S - \alpha^b j = -(\frac{1}{2} + \eta^b) \alpha^b\}}. \end{aligned}$$

From this, we derive the optimal controls of the market-maker as

$$\begin{aligned} \ell^{*a}(t, i, j, S, q) &= \mathbf{1}_{\{\alpha^a i + h^{i,j}(t, S, q-1) - h^{i,j}(t, S, q) > 0\}}, \\ \ell^{*b}(t, i, j, S, q) &= \mathbf{1}_{\{-\alpha^b j + h^{i,j}(t, S, q+1) - h^{i,j}(t, S, q) > 0\}}. \end{aligned}$$

The practical interest of Theorem 3.1 is that it allows us to compute the value function and optimal controls based on a finite difference scheme. Examples of computations of the value function are given in Section 3.4 and Appendix 3.A.3. Having described the problem of the market-maker, we now turn to the optimization problem of the platform.

⁴Note that, as the terminal condition does not depend on S^a and S^b , it also satisfies this boundary condition on $\partial \mathcal{D}$.

3.3.2 The platform's problem

The market-maker acts on a platform whose goal is to maximize the number of market orders on $[0, T]$. The intensities of arrival of market orders are functions of ℓ^a, ℓ^b , which are themselves functions of α^a, α^b . We assume that the platform is risk-neutral and earns a fixed taker cost $c > 0$ for each market order.⁵ Therefore its optimization problem is defined as

$$\sup_{(\alpha^a, \alpha^b) \in \mathbb{R}_+^2} \mathbb{E}^{\ell^{*a}, \ell^{*b}} [X_T^p],$$

given the optimal controls (ℓ^{*a}, ℓ^{*b}) of the market-maker and $X_t^p = c(N_t^a + N_t^b)$.

It is easy to observe that this problem boils down to maximizing the function v defined below over \mathbb{R}_+^2 :

$$v(\alpha^a, \alpha^b) := \mathbb{E} \left[\int_0^T c\lambda \left(\frac{\ell^{*a}(t, S_t^a, S_t^b, S_t, q_t)}{1 + (\kappa\alpha^a)^2} + \frac{\ell^{*b}(t, S_t^a, S_t^b, S_t, q_t)}{1 + (\kappa\alpha^b)^2} \right) dt \right].$$

Here we clearly see the tradeoff of the platform. A small tick size α^a increase the term $(1 + (\kappa\alpha^a)^2)^{-1}$. This is because it attracts more buy market orders. However, the optimal control $\ell^{*,a}$ is more often equal to zero: the gain of the market-maker may be too small if he quotes at the price S^a , therefore he regularly sets $\ell^{*,a} = 0$. The problem is similar on the bid side. On the other hand, a large tick size increases the gain of the market-maker if a transaction occurs, but decreases the number of market orders sent by market-takers, hence decreasing the trading volume.

We study numerically this problem in the next section by computing the value of v on a two dimensional grid and finding its maximum.

3.4 Numerical results

In this section, we show from numerical experiments the benefits of side-specific tick values in terms of increase of their value function for both the market-maker and the platform. Also, we fix reference values η_0 and α_0 . From them, to choose the parameter η^i associated to a given tick size α^i we use a result from Dayri and Rosenbaum [91] which gives the new value of the parameter η^i in case of a change of tick size from α_0 to α^i . This formula writes

$$\eta^i = \eta_0 \sqrt{\frac{\alpha_0}{\alpha^i}}. \quad (3.7)$$

In the following, we only consider values of α^a and α^b such that the underlying remains a large tick asset both on the bid and ask sides, that is $\eta^a \leq \frac{1}{2}, \eta^b \leq \frac{1}{2}$.

For the first experiments, we set $T = 40s$, $\bar{q} = 5$, $\sigma = 0.01s^{-\frac{1}{2}}$, $A = 0.1$, $\kappa = 10$, $\phi = 0.005$, $\lambda = 4$, $\eta_0 = 0.3$ and $\alpha_0 = 0.01$ which correspond to reasonable values to model a liquid asset. In particular, taking a larger T has no significant impact on the dependence of our results on the tick sizes. To remain in the large tick regime, we investigate values of α^i satisfying $0.0045 \leq \alpha^i \leq 0.05$ for $i \in \{a, b\}$.

3.4.1 Similar tick values on both sides

In this section we investigate the case where $\alpha^a = \alpha^b$. We plot in Figure 3.1 the value functions of the market-maker and the exchange, respectively h and v , for various values of $\alpha = \alpha^a = \alpha^b$. We fix the efficient price $S = 10.5$, the inventory $q = 0$ and we only consider values of α so that $0.5/\alpha \in \mathbb{N}$.

⁵More complex fee schedules can be handled in this framework. We can for example add a component which is proportional to the amount of cash traded.




Figure 3.1: Value function h (on the left) and v (on the right) for $\phi_- = 0$ in blue, $\phi_- = 0.0005$ in orange, $\phi_- = 0.005$ in green, as a function of $\alpha = \alpha^a = \alpha^b$.

When $\phi_- = 0$, the value of the exchange reaches its maximum at $\alpha \simeq 0.012$. An increase of ϕ_- leads to a reduction of the number of transactions. However the optimal tick value for the exchange is not significantly modified.

The optimal tick value for the market-maker is larger than that of the exchange. This is because the exchange is only interested in attracting orders while the market-maker's gain per trade (not taking into account the inventory risk) is linear with respect to the tick value. The trade-off of the exchange is the following: on the one hand, he would like to implement a quite small tick value (to attract market orders) but on the other hand, he must ensure a reasonable presence of market-maker.

When ϕ_- increases, the value function of the market-maker decreases, for all tick values. This is no surprise since ϕ_- corresponds to an inventory penalization, hence reducing the market-maker's PnL.

In Figure 3.2, we subtract the value function when $\phi_- = 0$ to the other value functions displayed in Figure 3.1. We remark that for the market-maker, the larger the tick the more significant the penalization of short inventory in terms of value function. We observe the opposite phenomenon for the exchange: the difference is essentially slightly increasing with respect to α . In particular, we see a quite strong impact of the penalization on the value function of the exchange when the tick size is small.




Figure 3.2: Variation of the functions h and v (difference between $\phi_- = 0$ in blue, $\phi_- = 0.0005$ in orange, $\phi_- = 0.005$ in green, and $\phi_- = 0$) as a function of $\alpha = \alpha^a = \alpha^b$.

We now study the case of side-specific tick values.

3.4.2 Side-specific tick values: additional opportunities for the market-maker

We set $\alpha^b = 0.0124$ (optimal tick size in the non side-specific case) and let α^a vary. We plot the value functions of the market-maker and the exchange in Figure 3.3. Again we observe that both value




Figure 3.3: Value function h (on the left) and v (on the right) as functions of α^a , for $\alpha^b = 0.0124$, for $\phi_- = 0$ in blue, $\phi_- = 0.0005$ in orange, $\phi_- = 0.005$ in green.

functions are decreasing with respect to ϕ_- . From the point of view of the market-maker, having non side-specific tick values is sub-optimal, even in the case $\phi_- = 0$. This is because when the two tick values are different, it is possible for S^a to be greater than S^b and orders to arrive with the same intensities on both sides: the market-maker can collect the spread. It is not possible in the non side-specific case, where the market-maker can only pocket the spread from buy and sell orders at two different times. Side-specific tick values are also clearly beneficial for the exchange. The transaction flow increases for $\alpha^a > \alpha^b$ because of the good liquidity provided by the market-maker, and for $\alpha^a < \alpha^b$ because of the high number of incoming market orders.

Remark 20. *Remark that with shifted grids (same tick values on both sides but with a grid shifted compared to the other), those additional opportunities for the market-maker would remain. In section 3.4.3, we will see however, that, from the point of view of the exchange, side-specific tick values are much more interesting.*

3.4.3 Side-specific tick values: effect of ϕ_-

We plot the two-dimensional value functions of the market-maker and the exchange for side-specific tick values.

First we take $\phi_- = 0$ in Figure 3.4. We note that the opportunity for the market-maker mentioned above remains present for all tick values and that the value functions are symmetric around the axis $\alpha^b = \alpha^a$ (side-specific tick values are preferred). Furthermore, we see that the exchange prefers smaller tick values than the market-maker. The optimal values for the exchange lie on an anti-diagonal which goes from $(\alpha^a = 0.0045, \alpha^b = 0.025)$ to $(\alpha^a = 0.025, \alpha^b = 0.0045)$. On this line the number of transactions varies little. It seems however that the optimum is on the edges of the zone in which the asset remains large tick: the two couples (α^a, α^b) mentioned above.

If the tick values are too large the intensities of the market orders become too small and the number of transactions diminishes. If both ticks are too small, the market-maker does not trade much because the gain per trade becomes too little compared to the inventory cost (recall that the intensity of market orders is upper bounded). However, the case where one tick is quite small and the other is large is suitable for the market-maker: for example, if $\alpha^a < \alpha^b$ his strategy is to be long and liquidate his long position fast if needed thanks to the small value of α^a which ensures a large number of incoming market orders. This explains why the optimal tick values given by the exchange are side-specific and symmetric with respect to the axis $\alpha^a = \alpha^b$. More precisely, the choice of ticks $(\alpha^a = 0.0045, \alpha^b = 0.025)$ or $(\alpha^a = 0.025, \alpha^b = 0.0045)$ seems optimal.

We now plot in Figure 3.5 the value function for $\phi_- = 0.005$. This non-zero parameter implies a clear decrease of the value function of the market-maker and the reduction of the number of transactions. An important remark is that the value functions are no-longer symmetric around the axis $\alpha^b = \alpha^a$.




Figure 3.4: Value function h (on the left) and v (on the right) as functions of α^a and α^b , for $\phi_- = 0$.




Figure 3.5: Value function h (on the left) and v (on the right) as functions of α^a and α^b , for $\phi_- = 0.005$.

For clarity we plot in Figure 3.6 the difference of the value functions when $\phi_- = 0.005$ and when $\phi_- = 0$ as a function of α^a, α^b . We see that the added component is not symmetric regarding to




Figure 3.6: Difference between the value function h (on the left) and v (on the right) as functions of α^a and α^b , between the case $\phi_- = 0.005$ and the case $\phi_- = 0$.

the axis $\alpha^b = \alpha^a$ and both the market-maker and the exchange tend to prefer the case $\alpha^b > \alpha^a$. It is particularly clear for the market-maker's problem where the difference between the values at $(\alpha^a = 0.0045, \alpha^b = 0.025)$ and $(\alpha^a = 0.025, \alpha^b = 0.0045)$ is approximately 0.03 which is roughly 10%

of the value function. Indeed, as explained above, having α^b quite large and α^a rather small essentially ensures that the market-maker can maintain a positive inventory all along the trading trajectory: attractive PnL for incoming buy orders and possibility to quickly reduce a positive inventory. The right-hand side of Figure 3.6 is harder to interpret, as the changes are comparatively smaller and influenced more by the kink described in 3.4.2 and its analogues at $\alpha^b = 2\alpha^a$ and $\alpha^b = \frac{1}{2}\alpha^a$.

The exchange leans more towards a larger α^b compared to α^a . To see that more clearly, we fix $\alpha^a = 0.0045$ and plot in Figure 3.7 the value functions h and v , as functions of α^b , for different values of ϕ .




Figure 3.7: Value functions h and v for $\alpha^a = 0.0045$, as functions of α^b , for different values of ϕ .

The value function of the market-maker is increasing in α^b . This is the same phenomenon as already observed in Figure 3.5. The value of the exchange has a maximum which is increasing in ϕ_- : as the penalization gets more side-specific, the optimal tick values displayed by the exchange become more asymmetric. Indeed, for $\phi_- = 0$, the optimum is reached for $\alpha^b \simeq 0.024$, while for $\phi_- = 0.005$, it is obtained for $\alpha^b \simeq 0.034$. Note that a relevant tick value set by the exchange can compensate for his loss of value function due to an increase of ϕ_- . By choosing a new tick size optimally when going from $\phi_- = 0$ to $\phi_- = 0.005$, the loss in value function is of 7% only. Keeping $\alpha^b = 0.024$ would lead to a loss of 15%. Note that the compensation can be total for the market-maker (and even exceeds the loss) but is only partial for the exchange.

3.5 Conclusion

A suitable choice of tick values by the exchange is a subtle equilibrium. If the platform imposes the same tick value on the bid and ask sides, it has to be sufficiently large to ensure significant PnL per trade for the market-maker and therefore good liquidity provision, and sufficiently small to attract market orders from market-takers. When allowing for side-specific tick values with no constraint on short inventory, the optimal tick values for the exchange are of the form (α_1^*, α_2^*) or symmetrically (α_2^*, α_1^*) with $\alpha_1^* < \alpha_2^*$. In this case, the market-maker can take advantage from additional trading opportunities and increase his activity. The exchange benefits from this situation because of the higher number of trades on his platform. Moreover, when there is a penalty for short inventory positions of the market-maker, there is only one optimal couple of tick values. In this case, the market-maker and subsequently the exchange prefer $\alpha^b > \alpha^a$ and the difference between α^a and α^b at the optimum becomes larger. Finally, note that side-specific tick values could have subtle consequences in a multi-platform setting. This issue is left for further study, as well as the situation where market-takers are more strategic in their execution.

3.A Appendix

3.A.1 Proof of Proposition 3.1

First we prove the continuity of h on $\mathcal{D} \times [0, T)$.

Let $q \in \mathcal{Q}$, $t_1 \in [0, T)$, $(s^a, s^b, s_1) \in \mathcal{D}$. Note that $\{s \in \mathbb{R}, (s^a, s^b, s) \in \mathcal{D}\}$ is an open interval containing s_1 , which we denote by $(s^{\leftarrow}, s^{\rightarrow})$. If the process S_t starts from a point $s \in (s^{\leftarrow}, s^{\rightarrow})$ with $S_t^a = s^a$ and $S_t^b = s^b$, S_t^a and S_t^b will not jump as long as S_t stays in $(s^{\leftarrow}, s^{\rightarrow})$. We will prove that the function $(t, s) \in [0, T) \times (s^{\leftarrow}, s^{\rightarrow}) \mapsto h(t, s^a, s^b, s, q)$ is continuous at (t_1, s_1) .

We fix $\eta > 0$. There is a ball with positive diameter \mathcal{B} in $[0, T) \times (s^{\leftarrow}, s^{\rightarrow})$ centered on (t_1, s_1) and some $\epsilon > 0$ such that, if $(t_2, s_2) \in \mathcal{B}$, then

$$\mathbb{E}[\tau^1 - t_1 | S_{t_1} = s_1] < \eta, \quad \mathbb{E}[\tau^2 - t_2 | S_{t_2} = s_2] < \eta, \quad (3.8)$$

$$\mathbb{P}[\tau^1 < T | S_{t_1} = s_1] > 1 - \eta, \quad \mathbb{P}[\tau^2 < T | S_{t_2} = s_2] > 1 - \eta, \quad (3.9)$$

and⁶

$$\begin{aligned} \inf_{\ell \in \mathcal{L}} \mathbb{P} \left[\inf_{t_1 \leq s \leq \tau^1} Q_s = \sup_{t_1 \leq s \leq \tau^1} Q_s = q | S_{t_1} = s_1, Q_{t_1} = q \right] &> 1 - \eta, \\ \inf_{\ell \in \mathcal{L}} \mathbb{P} \left[\inf_{t_2 \leq s \leq \tau^2} Q_s = \sup_{t_2 \leq s \leq \tau^2} Q_s = q | S_{t_2} = s_2, Q_{t_2} = q \right] &> 1 - \eta, \end{aligned} \quad (3.10)$$

where we write

$$\begin{aligned} \tau^1 &= T \wedge \inf \{ t \geq t_1, S_t^{t_1, s_1} = (s_1 \vee s_2) + \epsilon \text{ or } S_t^{t_1, s_1} = (s_1 \wedge s_2) - \epsilon \}, \\ \tau^2 &= T \wedge \inf \{ t \geq t_2, S_t^{t_2, s_2} = (s_1 \vee s_2) + \epsilon \text{ or } S_t^{t_2, s_2} = (s_1 \wedge s_2) - \epsilon \}. \end{aligned}$$

The quantities τ^1 and τ^2 are stopping times such that $t_1 \leq \tau^1 \leq T$ a.s. and $t_2 \leq \tau^2 \leq T$ a.s. We impose

$$s^{\leftarrow} < (s_1 \wedge s_2) - \epsilon < s_1 < (s_1 \vee s_2) + \epsilon < s^{\rightarrow}, \quad s^{\leftarrow} < (s_1 \wedge s_2) - \epsilon < s_2 < (s_1 \vee s_2) + \epsilon < s^{\rightarrow}$$

for any $(t_2, s_2) \in \mathcal{B}$ by taking a smaller ball \mathcal{B} and a smaller ϵ if necessary. In particular, this tells us that if $(S_{t_1}, S_{t_1}^a, S_{t_1}^b) = (s_1, s^a, s^b)$, S_t^a does not jump between t_1 and τ^1 . Similarly, if $(S_{t_2}, S_{t_2}^a, S_{t_2}^b) = (s_2, s^a, s^b)$, S_t^b does not jump between between t_2 and τ^2 .

Let some arbitrary $(t_2, s_2) \in \mathcal{B}$ and τ^1 and τ^2 the associated stopping times. Using the dynamic programming principle, we obtain

$$h(t_1, s^a, s^b, s_1, q) = \sup_{\ell \in \mathcal{L}} \mathbb{E}_{s_1, s^a, s^b, q} \left[h(\tau^1, S_{\tau^1}^a, S_{\tau^1}^b, S_{\tau^1}, Q_{\tau^1}) + \int_{t_1}^{\tau^1} \left(-\phi Q_t^2 - \phi_- Q_t^2 \mathbf{1}_{Q_t < 0} \right) dt \right].$$

This can be rewritten as

$$h(t_1, s^a, s^b, s_1, q) = \sup_{\ell \in \mathcal{L}} \mathbb{E}_{s_1, q} \left[\sum_{\bar{q} \in \mathcal{Q}} h(\tau^1, s^a, s^b, S_{\tau^1}, \bar{q}) \mathbf{1}_{\{Q_{\tau^1} = \bar{q}\}} + \int_{t_1}^{\tau^1} \left(-\phi Q_t^2 - \phi_- Q_t^2 \mathbf{1}_{Q_t < 0} \right) \mathbf{1}_{\{Q_t = \bar{q}\}} dt \right].$$

Noticing from (3.3) that h is bounded on $[0, T) \times \{s^a\} \times \{s^b\} \times [s^{\leftarrow}, s^{\rightarrow}) \times \mathcal{Q}$, we deduce by (3.10) that

$$\left| h(t_1, s^a, s^b, s_1, q) - \sup_{\ell \in \mathcal{L}} \mathbb{E}_{s_1, q} \left[h(\tau^1, s^a, s^b, S_{\tau^1}, q) + \int_{t_1}^{\tau^1} \left(-\phi q^2 - \phi_-(q)_-^2 \right) dt \right] \right| \leq C\eta$$

⁶These inequalities can be attained independently of the control ℓ as S is independent from Q .

for a constant C which depends only on s^a , s^b , and $(q)_- = q^2 \mathbf{1}_{q < 0}$. The expectation above does not depend on the control ℓ , hence we drop the supremum and fix an arbitrary control $\ell = 0$. We denote by \mathbb{E}^0 the expectation under the probability measure given by this control. The expectation neither depends on the process Q_t , so we drop the conditioning with respect to Q_{t_1} .

This leads to

$$\left| h(t_1, s^a, s^b, s_1, q) - \mathbb{E}^0 \left[h(\tau^1, s^a, s^b, S_{\tau^1}, q) + \int_{t_1}^{\tau^1} \left(-\phi q^2 - \phi_-(q)_-^2 \right) dt \middle| S_{t_1} = s_1 \right] \right| \leq C\eta.$$

Similarly, starting from (t_2, s^a, s^b, s_2, q) with $(t_2, s_2) \in \mathcal{B}$, we get

$$\left| h(t_2, s^a, s^b, s_2, q) - \mathbb{E}^0 \left[h(\tau^2, s^a, s^b, S_{\tau^2}, q) + \int_{t_2}^{\tau^2} \left(-\phi q^2 - \phi_-(q)_-^2 \right) dt \middle| S_{t_2} = s_2 \right] \right| \leq C\eta,$$

and we deduce that

$$\begin{aligned} & |h(t_1, s^a, s^b, s_1, q) - h(t_2, s^a, s^b, s_2, q)| \\ & \leq \left| \mathbb{E}^0 \left[h(\tau^1, s^a, s^b, S_{\tau^1}, q) + \int_{t_1}^{\tau^1} \left(-\phi q^2 - \phi_-(q)_-^2 \right) dt \middle| S_{t_1} = s_1 \right] \right. \\ & \quad \left. - \mathbb{E}^0 \left[h(\tau^2, s^a, s^b, S_{\tau^2}, q) + \int_{t_2}^{\tau^2} \left(-\phi q^2 - \phi_-(q)_-^2 \right) dt \middle| S_{t_2} = s_2 \right] \right| + 2C\eta \\ & \leq \left| \mathbb{E}^0 \left[h(\tau^1, s^a, s^b, S_{\tau^1}, q) \middle| S_{t_1} = s_1 \right] - \mathbb{E}^0 \left[h(\tau^2, s^a, s^b, S_{\tau^2}, q) \middle| S_{t_2} = s_2 \right] \right| \\ & \quad + \left| \phi q^2 - \phi_-(q)_-^2 \right| \left(\mathbb{E}^0[\tau^1 - t_1 | S_{t_1} = s_1] + \mathbb{E}^0[\tau^2 - t_2 | S_{t_2} = s_2] \right) + 2C\eta. \end{aligned} \quad (3.11)$$

Using (3.8), we get

$$\left| \mathbb{E}^0[\tau^1 - t_1 | S_{t_1} = s_1] + \mathbb{E}^0[\tau^2 - t_2 | S_{t_2} = s_2] \right| < 2\eta. \quad (3.12)$$

Also, the conditional laws

$$\begin{aligned} & (\tau^1 | S_{t_1} = s_1, S_{\tau^1} = (s_1 \vee s_2) + \epsilon, \tau^1 < T), \quad (\tau^1 | S_{t_1} = s_1, S_{\tau^1} = (s_1 \wedge s_2) - \epsilon, \tau^1 < T), \\ & (\tau^2 | S_{t_2} = s_2, S_{\tau^2} = (s_1 \vee s_2) + \epsilon, \tau^2 < T), \quad (\tau^2 | S_{t_2} = s_2, S_{\tau^2} = (s_1 \wedge s_2) - \epsilon, \tau^2 < T), \end{aligned}$$

have bounded continuous densities, which we denote by $f^{1,+}$, $f^{1,-}$, $f^{2,+}$ and $f^{2,-}$ respectively (see for example Borodin and Salminen [52, Formula 3.0.6]). So, by decomposing the first term in (3.11) with respect to the values of S_{τ^1} and S_{τ^2} , we can write

$$\begin{aligned} & \left| \mathbb{E}^0 \left[h(\tau^1, s^a, s^b, S_{\tau^1}, q) \middle| S_{t_1} = s_1 \right] - \mathbb{E}^0 \left[h(\tau^2, s^a, s^b, S_{\tau^2}, q) \middle| S_{t_2} = s_2 \right] \right| \\ & \leq \left| \sum_{j \in \{+, -\}} \int_0^T h(t, s^a, s^b, s_j, q) \left(f^{1,j}(t) \mathbb{P}^0[S_{\tau^1} = s_j, \tau^1 < T | S_{t_1} = s_1] - f^{2,j}(t) \mathbb{P}^0[S_{\tau^2} = s_j, \tau^2 < T | S_{t_2} = s_2] \right) dt \right| \\ & \quad + \left| \mathbb{E}^0 \left[h(\tau^1, s^a, s^b, S_{\tau^1}, q) \mathbf{1}_{\{S_{\tau^1} \neq s_+, S_{\tau^1} \neq s_-\} \cup \{\tau^1 = T\}} \middle| S_{t_1} = s_1 \right] \right| \\ & \quad + \left| \mathbb{E}^0 \left[h(\tau^2, s^a, s^b, S_{\tau^2}, q) \mathbf{1}_{\{S_{\tau^2} \neq s_+, S_{\tau^2} \neq s_-\} \cup \{\tau^2 = T\}} \middle| S_{t_2} = s_2 \right] \right|, \end{aligned}$$

where $s_+ = s_1 \vee s_2 + \epsilon$ and $s_- = s_1 \wedge s_2 - \epsilon$. Remark that the event $S_{\tau^1} \neq s_+, S_{\tau^1} \neq s_-, S_{t_1} = s_1$ happens only if $\tau^1 = T$ so that $\mathbb{P}^0[\{S_{\tau^1} \neq s_+, S_{\tau^1} \neq s_-\} \cup \{\tau^1 = T\} | S_{t_1} = s_1] < \eta$ by (3.9). Similarly

$\mathbb{P}^0[\{S_{\tau^2} \neq s_+, S_{\tau^2} \neq s_-\} \cup \{\tau^2 = T\} | S_{t_2} = s_2] < \eta$ by (3.9). As a consequence, using again (3.11), (3.12) and the fact that h is bounded on $[0, T] \times \{s^a\} \times \{s^b\} \times [s^{\leftarrow}, s^{\rightarrow}] \times \mathcal{Q}$, we get

$$\begin{aligned} & |h(t_1, s^a, s^b, s_1, q) - h(t_2, s^a, s^b, s_2, q)| \\ & \leq \left| \sum_{j \in \{+, -\}} \int_0^T h(t, s^a, s^b, s_j, q) (f^{1,j}(t) \mathbb{P}^0[S_{\tau^1} = s_j, \tau^1 < T | S_{t_1} = s_1] - f^{2,j}(t) \mathbb{P}^0[S_{\tau^2} = s_j, \tau^2 < T | S_{t_2} = s_2]) dt \right| \\ & \quad + (2|\phi q^2 - \phi_-(q)_-| + 4C)\eta. \end{aligned}$$

Recall that the $f^{1,+}$, $f^{1,-}$, $f^{2,+}$ and $f^{2,-}$ depend on s_2 and t_2 . We have

$$\begin{aligned} & \left| \mathbb{P}^0[S_{\tau^2} = s_+, \tau^2 < T | S_{t_2} = s_2] f^{2,+} - \mathbb{P}^0[S_{\tau^1} = s_+, \tau^1 < T | S_{t_1} = s_1] f^{1,+} \right| \xrightarrow{(t_2, s_2) \rightarrow (t_1, s_1)} 0, \\ & \left| \mathbb{P}^0[S_{\tau^2} = s_-, \tau^2 < T | S_{t_2} = s_2] f^{2,-} - \mathbb{P}^0[S_{\tau^1} = s_-, \tau^1 < T | S_{t_1} = s_1] f^{1,-} \right| \xrightarrow{(t_2, s_2) \rightarrow (t_1, s_1)} 0 \end{aligned}$$

point-wise on $[0, T]$ directly by Borodin and Salminen [52, Formula 3.0.6, Appendix 11]. Having fixed ϵ and using again Borodin and Salminen [52, Formula 3.0.6, Appendix 11], we see that the above functions are uniformly bounded with respect to $(s_2, t_2) \in \mathcal{B}$. So, using that h is bounded on $[0, T] \times \{s^a\} \times \{s^b\} \times [s^{\leftarrow}, s^{\rightarrow}] \times \mathcal{Q}$, we can apply the dominated convergence theorem to deduce that

$$\begin{aligned} & \left| \sum_{j \in \{+, -\}} \int_0^T h(t, s^a, s^b, s_j, q) (f^{1,j}(t) \mathbb{P}^0[S_{\tau^1} = s_j, \tau^1 < T | S_{t_1} = s_1] - f^{2,j}(t) \mathbb{P}^0[S_{\tau^2} = s_j, \tau^2 < T | S_{t_2} = s_2]) dt \right| \\ & \xrightarrow{(t_2, s_2) \rightarrow (t_1, s_1)} 0. \end{aligned}$$

Thus we have shown that h is continuous at the point (t_1, S^a, S^b, s_1, q) . The case $t_1 = T$ is treated the same way.

The continuity conditions can be proved using the same lines: fixing $q \in \mathcal{Q}$, $t_1 \in [0, T)$ and $(S^a, S^b, s_1) \in \partial\mathcal{D}$, choosing (t_2, s_2) close enough to (t_1, s_1) and applying the dynamic programming principle between t_1 and τ^1 , and t_2 and τ^2 , for τ^1 and τ^2 two well-chosen stopping times (for example

$$\tau^1 = T \wedge \inf \left\{ t > t_1, S_t = s_1 + \epsilon \text{ or } S_t = \frac{s^{\leftarrow} \wedge s_1}{2} \right\}, \quad \tau^2 = T \wedge \inf \left\{ t > t_2, S_t = s_1 + \epsilon \text{ or } S_t = \frac{s^{\leftarrow} \wedge s_1}{2} \right\}.$$

with $\epsilon > 0$ small enough, for a boundary inducing an upward jump).

3.A.2 Proof of Theorem 3.1

We first prove that the value function of the market-maker's problem is indeed a viscosity solution of (3.4).

Proposition 3.2. *The value function h is a continuous viscosity solution on $[0, T) \times \mathcal{D} \times \mathcal{Q}$ of (3.4). Furthermore, $h(T, S^a, S^b, S, q) = q(S - Aq)$ for all $(S^a, S^b, S, q) \in \mathcal{D} \times \mathcal{Q}$, and*

$$\begin{aligned} h(t, S^a, S^b, S, q) &= \mathbf{1}_{\{S - S^a = (\frac{1}{2} + \eta^a)\alpha^a, S - S^b < (\frac{1}{2} + \eta^b)\alpha^b\}} h(t, S^a + \alpha^a, S^b, S, q) \\ & \quad + \mathbf{1}_{\{S - S^a < (\frac{1}{2} + \eta^a)\alpha^a, S - S^b = (\frac{1}{2} + \eta^b)\alpha^b\}} h(t, S^a, S^b + \alpha^b, S, q) \\ & \quad + \mathbf{1}_{\{S - S^a = (\frac{1}{2} + \eta^a)\alpha^a, S - S^b = (\frac{1}{2} + \eta^b)\alpha^b\}} h(t, S^a + \alpha^a, S^b + \alpha^b, S, q) \\ & \quad + \mathbf{1}_{\{S - S^a = -(\frac{1}{2} + \eta^a)\alpha^a, S - S^b > -(\frac{1}{2} + \eta^b)\alpha^b\}} h(t, S^a - \alpha^a, S^b, S, q) \\ & \quad + \mathbf{1}_{\{S - S^a > -(\frac{1}{2} + \eta^a)\alpha^a, S - S^b = -(\frac{1}{2} + \eta^b)\alpha^b\}} h(t, S^a, S^b - \alpha^b, S, q) \\ & \quad + \mathbf{1}_{\{S - S^a = -(\frac{1}{2} + \eta^a)\alpha^a, S - S^b = -(\frac{1}{2} + \eta^b)\alpha^b\}} h(t, S^a - \alpha^a, S^b - \alpha^b, S, q), \end{aligned}$$

for all $(t, S^a, S^b, S, q) \in [0, T) \times \partial\mathcal{D} \times \mathcal{Q}$.

Proof. Let $(\bar{S}^a, \bar{S}^b, \bar{q}) \in \alpha^a \mathbb{Z} \times \alpha^b \mathbb{Z} \times \mathcal{Q}$, and $(t_n, S_n)_{n \in \mathbb{N}} \in [0, T] \times \mathbb{R}$ be a sequence such that

$$\begin{aligned} t_n &\xrightarrow[n \rightarrow +\infty]{} \hat{t} \in [0, T], \\ S_n &\xrightarrow[n \rightarrow +\infty]{} \hat{S} \in \mathbb{R}, \\ h(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) &\xrightarrow[n \rightarrow +\infty]{} h(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}), \end{aligned}$$

with $(\bar{S}^a, \bar{S}^b, \hat{S}) \in \mathcal{D}$. Without loss of generality we can assume that $(\bar{S}^a, \bar{S}^b, S_n) \in \mathcal{D}$ for all $n \in \mathbb{N}$.

Let us first consider the case $\hat{t} = T$. Let us take two arbitrary controls $\ell_s^a = \ell_s^b = 0$, for all $s \in [0, T]$, then for all $n \in \mathbb{N}$ we have

$$h(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) \geq \mathbb{E}_{t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}} \left[Q_T(S_T - A Q_T) - \phi \int_{t_n}^T Q_s^2 ds - \phi_- \int_{t_n}^T Q_s^2 \mathbf{1}_{Q_s < 0} ds \right],$$

and by dominated convergence we can obtain

$$h(T, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) \geq \bar{q}(\hat{S} - A\bar{q}).$$

Now let us consider the case $\hat{t} < T$. Let $\varphi : [0, T] \times \mathcal{D} \times \mathcal{Q} \rightarrow \mathbb{R}$ be a continuous function, \mathcal{C}^1 in t , \mathcal{C}^2 in S and such that $0 = \min_{[0, T] \times \mathcal{D}} (h - \varphi) = (h - \varphi)(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q})$. We also assume that $h = \varphi$ only at the point $(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q})$. Let us assume that there exists $\eta > 0$ such that

$$\begin{aligned} 2\eta &\leq \partial_t \varphi(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) - \phi \bar{q}^2 - \phi_- \bar{q}^2 \mathbf{1}_{\bar{q} < 0} + \frac{1}{2} \sigma^2 \partial_{SS} \varphi(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) \\ &\quad + \frac{\lambda}{1 + (\kappa \alpha^a)^2} \max_{\ell^a \in \{0, 1\}} \left\{ \ell^a \left(\bar{S}^a + \varphi(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q} - \ell^a) - \varphi(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) \right) \right\} \\ &\quad + \frac{\lambda}{1 + (\kappa \alpha^b)^2} \max_{\ell^b \in \{0, 1\}} \left\{ \ell^b \left(-\bar{S}^b + \varphi(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q} + \ell^b) - \varphi(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) \right) \right\}. \end{aligned}$$

Then we must have

$$\begin{aligned} 0 &\leq \partial_t \varphi(t, \bar{S}^a, \bar{S}^b, S, \bar{q}) - \phi \bar{q}^2 - \phi_- \bar{q}^2 \mathbf{1}_{\bar{q} < 0} + \frac{1}{2} \sigma^2 \partial_{SS} \varphi(t, \bar{S}^a, \bar{S}^b, S, \bar{q}) \\ &\quad + \frac{\lambda}{1 + (\kappa \alpha^a)^2} \max_{\ell^a \in \{0, 1\}} \left\{ \ell^a \left(\bar{S}^a + \varphi(t, \bar{S}^a, \bar{S}^b, S, \bar{q} - \ell^a) - \varphi(t, \bar{S}^a, \bar{S}^b, S, \bar{q}) \right) \right\} \\ &\quad + \frac{\lambda}{1 + (\kappa \alpha^b)^2} \max_{\ell^b \in \{0, 1\}} \left\{ \ell^b \left(-\bar{S}^b + \varphi(t, \bar{S}^a, \bar{S}^b, S, \bar{q} + \ell^b) - \varphi(t, \bar{S}^a, \bar{S}^b, S, \bar{q}) \right) \right\}, \end{aligned}$$

for all $(t, S) \in B = ((\hat{t} - r, \hat{t} + r) \cap [0, T]) \times (\hat{S} - r, \hat{S} + r)$ for some $r > 0$. Without loss of generality, we can assume that B contains the sequence $(t_n, S_n)_n$ and that for all $(t, S) \in B$, we have $(\bar{S}^a, \bar{S}^b, S) \in \mathcal{D}$. We can choose the value of η such that

$$\varphi(t, \bar{S}^a, \bar{S}^b, S, \bar{q}) + \eta \leq h(t, \bar{S}^a, \bar{S}^b, S, \bar{q})$$

on $\partial_p B := \left(((\hat{t} - r, \hat{t} + r) \cap [0, T]) \times (\{\hat{S} - r\} \cup \{\hat{S} + r\}) \right) \cup (\{\hat{t} + r\} \times [\hat{S} - r, \hat{S} + r])$. We can also assume that

$$\varphi(t, S^a, S^b, S, q) + \eta \leq h(t, S^a, S^b, S, q),$$

for $(t, S^a, S^b, S, q) \in \tilde{B}$ with

$$\tilde{B} = \left\{ (t, \bar{S}^a, \bar{S}^b, S, q) \mid (t, S) \in B, q \in \{\bar{q} - 1, \bar{q} + 1\} \cap \mathcal{Q} \right\}.$$

We introduce the set

$$B_{\mathcal{D}} = \left\{ (t, \bar{S}^a, \bar{S}^b, S, \bar{q}) \mid (t, S) \in B \right\},$$

and set $\pi_n = \inf\{t \geq t_n \mid (t, S_t^a, S_t^b, S_t, q_t) \notin B_{\mathcal{D}}\}$ with $S_{t_n}^i = \bar{S}^i$, $q_{t_n} = \bar{q}$, $S_{t_n} = S_n$, where the processes are controlled by

$$\begin{aligned} \ell_t^a &= \mathbf{1}_{\{S_t^a + \varphi(t, S_t^a, S_t^b, S_t, q_t) - 1 - \varphi(t, S_t^a, S_t^b, S_t, q_{t-}) > 0\}}, \\ \ell_t^b &= \mathbf{1}_{\{-S_t^b + \varphi(t, S_t^a, S_t^b, S_t, q_{t-} + 1) - \varphi(t, S_t^a, S_t^b, S_t, q_{t-}) > 0\}}. \end{aligned}$$

Using Itô's formula and noting that S_t^a, S_t^b do not jump between t_n and π_n , we derive

$$\begin{aligned} \varphi(\pi_n, S_{\pi_n}^a, S_{\pi_n}^b, S_{\pi_n}, q_{\pi_n}) &= \varphi(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) + \int_{t_n}^{\pi_n} \left(\partial_t \varphi(t, S_t^a, S_t^b, S_t, q_t) + \frac{1}{2} \sigma^2 \partial_{SS} \varphi(t, S_t^a, S_t^b, S_t, q_t) \right) dt \\ &\quad + \int_{t_n}^{\pi_n} \lambda(\ell_t^a) (\varphi(t, S_t^a, S_t^b, S_t, q_{t-} - \ell_t^a) - \varphi(t, S_t^a, S_t^b, S_t, q_{t-})) dt \\ &\quad + \int_{t_n}^{\pi_n} \lambda(\ell_t^b) (\varphi(t, S_t^a, S_t^b, S_t, q_{t-} + \ell_t^b) - \varphi(t, S_t^a, S_t^b, S_t, q_{t-})) dt \\ &\quad + \int_{t_n}^{\pi_n} \sigma \partial_S \varphi(t, S_t^a, S_t^b, S_t, q_t) dW_t \\ &\quad + \int_{t_n}^{\pi_n} (\varphi(t, S_t^a, S_t^b, S_t, q_{t-} - \ell_t^a) - \varphi(t, S_t^a, S_t^b, S_t, q_{t-})) d\tilde{N}_t^a \\ &\quad + \int_{t_n}^{\pi_n} (\varphi(t, S_t^a, S_t^b, S_t, q_{t-} + \ell_t^b) - \varphi(t, S_t^a, S_t^b, S_t, q_{t-})) d\tilde{N}_t^b \\ &\geq \varphi(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) \\ &\quad - \int_{t_n}^{\pi_n} (S_t^a \lambda(\ell_t^a) - S_t^b \lambda(\ell_t^b) - \phi q_t^2 - \phi_- q_t^2 \mathbf{1}_{q_t < 0}) dt \\ &\quad + \int_{t_n}^{\pi_n} \sigma \partial_S \varphi(t, S_t^a, S_t^b, S_t, q_t) dW_t \\ &\quad + \int_{t_n}^{\pi_n} (\varphi(t, S_t^a, S_t^b, S_t, q_{t-} - \ell_t^a) - \varphi(t, S_t^a, S_t^b, S_t, q_{t-})) d\tilde{N}_t^a \\ &\quad + \int_{t_n}^{\pi_n} (\varphi(t, S_t^a, S_t^b, S_t, q_{t-} + \ell_t^b) - \varphi(t, S_t^a, S_t^b, S_t, q_{t-})) d\tilde{N}_t^b. \end{aligned}$$

Then by taking the expectation we get

$$\varphi(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) \leq \mathbb{E} \left[\varphi(\pi_n, S_{\pi_n}^a, S_{\pi_n}^b, S_{\pi_n}, q_{\pi_n}) + \int_{t_n}^{\pi_n} (S_t^a \lambda(\ell_t^a) - S_t^b \lambda(\ell_t^b) - \phi q_t^2 - \phi_- q_t^2 \mathbf{1}_{q_t < 0}) dt \right].$$

Thus

$$\varphi(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) \leq -\eta + \mathbb{E} \left[h(\pi_n, S_{\pi_n}^a, S_{\pi_n}^b, S_{\pi_n}, q_{\pi_n}) + \int_{t_n}^{\pi_n} (S_t^a \lambda(\ell_t^a) - S_t^b \lambda(\ell_t^b) - \phi q_t^2 - \phi_- q_t^2 \mathbf{1}_{q_t < 0}) dt \right].$$

As

$$\begin{aligned} \varphi(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) &\xrightarrow{n \rightarrow +\infty} \varphi(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) = h(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}), \\ h(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) &\xrightarrow{n \rightarrow +\infty} h(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}), \end{aligned}$$

there exists $n_0 \in \mathbb{N}$ such that for all $n \geq n_0$, $h(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) - \frac{\eta}{2} \leq \varphi(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q})$ and we deduce

$$h(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) \leq -\frac{\eta}{2} + \mathbb{E} \left[h(\pi_n, S_{\pi_n}^a, S_{\pi_n}^b, S_{\pi_n}, q_{\pi_n}) + \int_{t_n}^{\pi_n} (S_t^a \lambda(\ell_t^a) - S_t^b \lambda(\ell_t^b) - \phi q_t^2 - \phi_- q_t^2 \mathbf{1}_{q_t < 0}) dt \right],$$

which contradicts the dynamic programming principle. Therefore,

$$\begin{aligned} 0 &\geq \partial_t \varphi(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) - \phi \bar{q}^2 - \phi_- \bar{q}^2 \mathbf{1}_{\bar{q} < 0} + \frac{1}{2} \sigma^2 \partial_{SS} \varphi(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) \\ &\quad + \frac{\lambda}{1 + (\kappa \alpha^a)^2} \max_{\ell^a \in \{0,1\}} \left\{ \ell^a \left(\bar{S}^a + \varphi(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q} - \ell^a) - \varphi(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) \right) \right\} \\ &\quad + \frac{\lambda}{1 + (\kappa \alpha^b)^2} \max_{\ell^b \in \{0,1\}} \left\{ \ell^b \left(-\bar{S}^b + \varphi(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q} + \ell^b) - \varphi(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) \right) \right\}, \end{aligned}$$

and h is a viscosity supersolution of the HJB equation on $[0, T] \times \mathcal{D} \times \mathcal{Q}$.

The proof for the subsolution part is identical. \square

For a given $\rho > 0$, we introduce the function \tilde{h} such that

$$\tilde{h}(t, S^a, S^b, S, q) = e^{\rho t} h(t, S^a, S^b, S, q), \quad \text{for all } (t, S^a, S^b, S, q) \in [0, T] \times \mathcal{D} \times \mathcal{Q}.$$

Then \tilde{h} is a viscosity solution of the following HJB equation:

$$\begin{aligned} 0 &= -\rho \tilde{h}(t, S^a, S^b, S, q) + \partial_t \tilde{h}(t, S^a, S^b, S, q) - \phi q^2 - \phi_- q^2 \mathbf{1}_{q < 0} + \frac{1}{2} \sigma^2 \partial_{SS} \tilde{h} \\ &\quad + \frac{\lambda}{1 + (\kappa \alpha^a)^2} \max_{\ell^a \in \{0,1\}} \left\{ \ell^a \left(e^{\rho t} S^a + \tilde{h}(t, S^a, S^b, S, q - \ell^a) - \tilde{h}(t, S^a, S^b, S, q) \right) \right\} \\ &\quad + \frac{\lambda}{1 + (\kappa \alpha^b)^2} \max_{\ell^b \in \{0,1\}} \left\{ \ell^b \left(e^{\rho t} (-S^b) + \tilde{h}(t, S^a, S^b, S, q + \ell^b) - \tilde{h}(t, S^a, S^b, S, q) \right) \right\}, \end{aligned} \quad (3.13)$$

and we see that proving a maximum principle for (3.13) is equivalent to proving one for (3.4).

Definition 3.1. Let $U : [0, T] \times \mathcal{D} \times \mathcal{Q} \rightarrow \mathbb{R}$ be continuous with respect to (t, S) . For $(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) \in [0, T] \times \mathcal{D} \times \mathcal{Q}$, we say that $(y, p, A) \in \mathbb{R}^3$ is in the subset $\mathcal{P}^-U(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q})$ (resp. the superjet $\mathcal{P}^+U(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q})$) if

$$\begin{aligned} U(t, \bar{S}^a, \bar{S}^b, S, \bar{q}) &\geq U(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) + y(t - \hat{t}) + p(S - \hat{S}) + \frac{1}{2} A(S - \hat{S})^2 + o(|t - \hat{t}| + |S - \hat{S}|^2), \\ \left(\text{resp. } U(t, \bar{S}^a, \bar{S}^b, S, \bar{q}) \leq U(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) + y(t - \hat{t}) + p(S - \hat{S}) + \frac{1}{2} A(S - \hat{S})^2 + o(|t - \hat{t}| + |S - \hat{S}|^2) \right), \end{aligned}$$

for all (t, S) such that $(t, \bar{S}^a, \bar{S}^b, S, \bar{q}) \in [0, T] \times \mathcal{D} \times \mathcal{Q}$.

We also define $\bar{\mathcal{P}}^-U(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q})$ as the set of points $(y, p, A) \in \mathbb{R}^3$ such that there exists a sequence $(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}, y_n, p_n, A_n) \in [0, T] \times \mathcal{D} \times \mathcal{Q} \times \mathcal{P}^-U(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q})$ satisfying

$$(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}, y_n, p_n, A_n) \xrightarrow{n \rightarrow +\infty} (\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}, y, p, A).$$

The set $\bar{\mathcal{P}}^+U(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q})$ is defined similarly.

Let us recall one of the definitions of viscosity solutions which we are going to use for the proof of the uniqueness.

Lemma 3.1. A continuous function \tilde{U} is a viscosity supersolution (resp. subsolution) to (3.13) on $[0, T] \times \mathcal{D} \times \mathcal{Q}$ if and only if for all $(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) \in [0, T] \times \mathcal{D} \times \mathcal{Q}$ and all $(\hat{y}, \hat{p}, \hat{A}) \in \bar{\mathcal{P}}^-U(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q})$ (resp. $\bar{\mathcal{P}}^+U(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q})$), we have

$$\begin{aligned} &-\rho \tilde{U}(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) + \hat{y} - \phi \bar{q}^2 - \phi_- \bar{q}^2 \mathbf{1}_{\bar{q} < 0} + \frac{1}{2} \sigma^2 A \\ &\quad + \frac{\lambda}{1 + (\kappa \alpha^a)^2} \max_{\ell^a \in \{0,1\}} \left\{ \ell^a \left(e^{\rho t} \bar{S}^a + \tilde{U}(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q} - \ell^a) - \tilde{U}(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) \right) \right\} \\ &\quad + \frac{\lambda}{1 + (\kappa \alpha^b)^2} \max_{\ell^b \in \{0,1\}} \left\{ \ell^b \left(e^{\rho t} (-\bar{S}^b) + \tilde{U}(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q} + \ell^b) - \tilde{U}(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) \right) \right\} \leq 0. \end{aligned}$$

(resp. ≥ 0).

We refer to Bouchard [53] for a proof of this result. We can now state a maximum principle from which the uniqueness can be easily deduced.

Proposition 3.3. *Let U (resp. V) be a continuous viscosity supersolution (resp. subsolution) of (3.4) with polynomial growth on $[0, T] \times \mathcal{D} \times \mathcal{Q}$ and satisfying the continuity conditions (3.6). If $U \geq V$ on $\{T\} \times \mathcal{D} \times \mathcal{Q}$, then $U \geq V$ on $[0, T] \times \mathcal{D} \times \mathcal{Q}$.*

Proof. As before, we introduce the functions \tilde{U} and \tilde{V} such that

$$\tilde{U}(t, S^a, S^b, S, q) = e^{\rho t} U(t, S^a, S^b, S, q), \quad \text{and} \quad \tilde{V}(t, S^a, S^b, S, q) = e^{\rho t} V(t, S^a, S^b, S, q).$$

Then \tilde{U} and \tilde{V} are respectively viscosity supersolution and subsolution of Equation (3.13) on $[0, T] \times \mathcal{D} \times \mathcal{Q}$ with $\tilde{U} \geq \tilde{V}$ on $\{T\} \times \mathcal{D} \times \mathcal{Q}$. To prove the proposition, it is enough to prove that $\tilde{U} \geq \tilde{V}$ on $[0, T] \times \mathcal{D} \times \mathcal{Q}$. We proceed by contradiction. Let us assume that $\sup_{[0, T] \times \mathcal{D} \times \mathcal{Q}} \tilde{V} - \tilde{U} > 0$. Let $p \in \mathbb{N}^*$

such that

$$\lim_{\|S\|_2 \rightarrow +\infty} \sup_{\substack{t \in [0, T], q \in \mathcal{Q} \\ (S, S^a, S^b) \in \mathcal{D}}} \frac{|\tilde{U}(t, S^a, S^b, S, q)| + |\tilde{V}(t, S^a, S^b, S, q)|}{1 + \|S\|_2^{2p}} = 0,$$

where $\|\cdot\|_2$ is the Euclidian norm. Then there exists $(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) \in [0, T] \times \mathcal{D} \times \mathcal{Q}$ such that

$$\begin{aligned} 0 &< \tilde{V}(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) - \tilde{U}(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) - \phi(\hat{t}, \hat{S}, \hat{S}, \bar{q}) \\ &= \sup_{(t, S^a, S^b, S, q)} \tilde{V}(t, S^a, S^b, S, q) - \tilde{U}(t, S^a, S^b, S, q) - \phi(t, S, S, q), \end{aligned}$$

where

$$\phi(t, S, R, q) := \varepsilon e^{-\mu t} (1 + \|S\|_2^{2p} + \|R\|_2^{2p}),$$

with $\varepsilon > 0$, $\mu > 0$. The choice of the function ϕ allows us to look for a supremum in a bounded set with respect to (S, S^a, S^b) . Then the supremum is either reached for a point in $[0, T] \times \mathcal{D} \times \mathcal{Q}$ or on $[0, T] \times \partial \mathcal{D} \times \mathcal{Q}$ (recall that \mathcal{D} is open). But the continuity conditions tell us that if the supremum is reached on $[0, T] \times \partial \mathcal{D} \times \mathcal{Q}$, it is also reached in $[0, T] \times \mathcal{D} \times \mathcal{Q}$. Since $\tilde{U} \geq \tilde{V}$ on $\{T\} \times \mathcal{D} \times \mathcal{Q}$, it is clear that $\hat{t} < T$.

Then, for all $n \in \mathbb{N}^*$, we can find $(t_n, S_n, R_n) \in [0, T] \times \mathbb{R}^2$ such that $(\bar{S}^a, \bar{S}^b, S_n), (\bar{S}^a, \bar{S}^b, R_n) \in \mathcal{D}$ and

$$\begin{aligned} 0 &< \tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) - \tilde{U}(t_n, \bar{S}^a, \bar{S}^b, R_n, \bar{q}) \\ &\quad - \phi(t_n, S_n, R_n, \bar{q}) - n|S_n - R_n|^2 - (|t_n - \hat{t}|^2 + |S_n - \hat{S}|^4) \\ &= \sup_{(t, S, R)} \tilde{V}(t, \bar{S}^a, \bar{S}^b, S, \bar{q}) - \tilde{U}(t, \bar{S}^a, \bar{S}^b, R, \bar{q}) \\ &\quad - \phi(t, S, R, \bar{q}) - n|S - R|^2 - (|t - \hat{t}|^2 + |S - \hat{S}|^4). \end{aligned}$$

Then, we have

$$(t_n, S_n, R_n) \xrightarrow{n \rightarrow +\infty} (\hat{t}, \hat{S}, \hat{S}),$$

and

$$\begin{aligned} &\tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) - \tilde{U}(t_n, \bar{S}^a, \bar{S}^b, R_n, \bar{q}) - \phi(t_n, S_n, R_n) - n|S_n - R_n|^2 - (|t_n - \hat{t}|^2 + |S_n - \hat{S}|^4) \\ &\xrightarrow{n \rightarrow +\infty} \tilde{V}(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) - \tilde{U}(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) - \phi(\hat{t}, \hat{S}, \hat{S}). \end{aligned}$$

For $n \in \mathbb{N}^*$, let us write for $(t, S, R) \in [0, T] \times \mathbb{R}^2$

$$\varphi_n(t, S, R) := \phi(t, S, R) + n|S - R|^2 + |t - \hat{t}|^2 + |S - \hat{S}|^4.$$

Then Ishii's Lemma (see Barles and Imbert [36, Lemma 1], Crandall, Ishii, and Lions [83, Theorem 3.2]) guarantees that for any $\eta > 0$, we can find $(y_n^1, p_n^1, A_n^1) \in \bar{\mathcal{P}}^+ \tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q})$ and $(y_n^2, p_n^2, A_n^2) \in \bar{\mathcal{P}}^- \tilde{U}(t_n, \bar{S}^a, \bar{S}^b, R_n, \bar{q})$ such that

$$y_n^1 - y_n^2 = \partial_t \varphi_n(t_n, S_n, R_n), \quad (p_n^1, p_n^2) = (\partial_S \varphi_n, -\partial_R \varphi_n)(t_n, S_n, R_n),$$

and

$$\begin{pmatrix} A_n^1 & 0 \\ 0 & -A_n^2 \end{pmatrix} \leq H_{SR} \varphi_n(t_n, S_n, R_n) + \eta (H_{SR} \varphi_n(t_n, S_n, R_n))^2,$$

where $H_{SR} \varphi_n(t_n, \cdot, \cdot)$ denotes the Hessian of $\varphi_n(t_n, \cdot, \cdot)$. Applying Lemma 3.1, we get

$$\begin{aligned} \rho(\tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) - \tilde{U}(t_n, \bar{S}^a, \bar{S}^b, R_n, \bar{q})) &\leq y_n^1 - y_n^2 + \frac{1}{2} \sigma^2 (A_n^1 - A_n^2) \\ &+ \frac{\lambda}{1 + (\kappa \alpha^a)^2} \max_{\ell^a \in \{0,1\}} \left\{ \ell^a \left(e^{\rho t_n} \bar{S}^a + \tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q} - \ell^a) - \tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) \right) \right\} \\ &+ \frac{\lambda}{1 + (\kappa \alpha^b)^2} \max_{\ell^b \in \{0,1\}} \left\{ \ell^b \left(e^{\rho t_n} (-\bar{S}^b) + \tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q} + \ell^b) - \tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) \right) \right\} \\ &- \frac{\lambda}{1 + (\kappa \alpha^a)^2} \max_{\ell^a \in \{0,1\}} \left\{ \ell^a \left(e^{\rho t_n} \bar{S}^a + \tilde{U}(t_n, \bar{S}^a, \bar{S}^b, R_n, \bar{q} - \ell^a) - \tilde{U}(t_n, \bar{S}^a, \bar{S}^b, R_n, \bar{q}) \right) \right\} \\ &- \frac{\lambda}{1 + (\kappa \alpha^b)^2} \max_{\ell^b \in \{0,1\}} \left\{ \ell^b \left(e^{\rho t_n} (-\bar{S}^b) + \tilde{U}(t_n, \bar{S}^a, \bar{S}^b, R_n, \bar{q} + \ell^b) - \tilde{U}(t_n, \bar{S}^a, \bar{S}^b, R_n, \bar{q}) \right) \right\}. \end{aligned}$$

Moreover, we have

$$H_{SR} \varphi_n(t_n, S_n, R_n) = \begin{pmatrix} \partial_{SS}^2 \phi(t_n, S_n, R_n) + 2n + 12(S_n - \hat{S})^2 & \partial_{SR}^2 \phi(t_n, S_n, R_n) - 2n \\ \partial_{SR}^2 \phi(t_n, S_n, R_n) - 2n & \partial_{RR}^2 \phi(t_n, S_n, R_n) + 2n \end{pmatrix}.$$

It follows that

$$\begin{aligned} \rho(\tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) - \tilde{U}(t_n, \bar{S}^a, \bar{S}^b, R_n, \bar{q})) &\leq \partial_t \phi(t_n, S_n, R_n) + 2(t_n - \hat{t}) \\ &+ \frac{1}{2} \sigma^2 (\partial_{SS}^2 \phi(t_n, S_n, R_n) + \partial_{RR}^2 \phi(t_n, S_n, R_n) + 2\partial_{SR}^2 \phi(t_n, S_n, R_n) + 12(S_n - \hat{S})) + \eta C_n \\ &+ \frac{\lambda}{1 + (\kappa \alpha^a)^2} \max_{\ell^a \in \{0,1\}} \left\{ \ell^a \left(e^{\rho t_n} \bar{S}^a + \tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q} - \ell^a) - \tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) \right) \right\} \\ &+ \frac{\lambda}{1 + (\kappa \alpha^b)^2} \max_{\ell^b \in \{0,1\}} \left\{ \ell^b \left(e^{\rho t_n} (-\bar{S}^b) + \tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q} + \ell^b) - \tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) \right) \right\} \\ &- \frac{\lambda}{1 + (\kappa \alpha^a)^2} \max_{\ell^a \in \{0,1\}} \left\{ \ell^a \left(e^{\rho t_n} \bar{S}^a + \tilde{U}(t_n, \bar{S}^a, \bar{S}^b, R_n, \bar{q} - \ell^a) - \tilde{U}(t_n, \bar{S}^a, \bar{S}^b, R_n, \bar{q}) \right) \right\} \\ &- \frac{\lambda}{1 + (\kappa \alpha^b)^2} \max_{\ell^b \in \{0,1\}} \left\{ \ell^b \left(e^{\rho t_n} (-\bar{S}^b) + \tilde{U}(t_n, \bar{S}^a, \bar{S}^b, R_n, \bar{q} + \ell^b) - \tilde{U}(t_n, \bar{S}^a, \bar{S}^b, R_n, \bar{q}) \right) \right\}, \end{aligned}$$

where C_n does not depend on η . Therefore, as the maximums on the right-hand side are always positive, we deduce that for all $n \in \mathbb{N}^*$,

$$\begin{aligned} \rho(\tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) - \tilde{U}(t_n, \bar{S}^a, \bar{S}^b, R_n, \bar{q})) &\leq \partial_t \phi(t_n, S_n, R_n) + 2(t_n - \hat{t}) \\ &+ \frac{1}{2} \sigma^2 (\partial_{SS}^2 \phi(t_n, S_n, R_n) + \partial_{RR}^2 \phi(t_n, S_n, R_n) + 2\partial_{SR}^2 \phi(t_n, S_n, R_n) + 12(S_n - \hat{S})) \\ &+ \frac{\lambda}{1 + (\kappa \alpha^a)^2} \max_{\ell^a \in \{0,1\}} \left\{ \ell^a \left(e^{\rho t_n} \bar{S}^a + \tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q} - \ell^a) - \tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) \right) \right\} \\ &+ \frac{\lambda}{1 + (\kappa \alpha^b)^2} \max_{\ell^b \in \{0,1\}} \left\{ \ell^b \left(e^{\rho t_n} (-\bar{S}^b) + \tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q} + \ell^b) - \tilde{V}(t_n, \bar{S}^a, \bar{S}^b, S_n, \bar{q}) \right) \right\}. \end{aligned}$$

As \tilde{V} is continuous and $(t_n, S_n)_n$ converges to (\hat{t}, \hat{S}) , the last two terms are bounded from above by some constant M . Then by sending n to infinity, we get

$$\begin{aligned} \rho \left(\tilde{V}(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) - \tilde{U}(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) \right) &\leq \partial_t \phi(\hat{t}, \hat{S}, \hat{S}) \\ &+ \frac{1}{2} \sigma^2 \left(\partial_{SS}^2 \phi(\hat{t}, \hat{S}, \hat{S}) + \partial_{RR}^2 \phi(\hat{t}, \hat{S}, \hat{S}) + 2 \partial_{SR}^2 \phi(\hat{t}, \hat{S}, \hat{S}) \right) + M. \end{aligned}$$

For $\mu > 0$ large enough, the right-hand side is strictly negative, and as $\rho > 0$ we get

$$\tilde{V}(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) - \tilde{U}(\hat{t}, \bar{S}^a, \bar{S}^b, \hat{S}, \bar{q}) < 0,$$

hence the contradiction. \square

3.A.3 Effects of the uncertainty zones on h

We keep the same parameters as in Section 3.4 and take $\alpha^a = 0.01$ and $\alpha^b = 0.00625$. We plot the value function of the market-maker's problem (the function h) on some small range of values of S . Note that $S = 10.5$ is on both discrete grids. We distinguish 4 possible cases, depending on whether




Figure 3.8: Value function h of the market-maker for $q = 0$, as a function of S .

- $S^a = \alpha^a \lfloor S/\alpha^a \rfloor$ and $S^b = \alpha^b \lfloor S/\alpha^b \rfloor$ (green dots),
- $S^a = \alpha^a \lfloor S/\alpha^a \rfloor$ and $S^b = \alpha^b \lceil S/\alpha^b \rceil$ (red dash-dots),
- $S^a = \alpha^a \lceil S/\alpha^a \rceil$ and $S^b = \alpha^b \lfloor S/\alpha^b \rfloor$ (orange dash),
- $S^a = \alpha^a \lceil S/\alpha^a \rceil$ and $S^b = \alpha^b \lceil S/\alpha^b \rceil$ (blue solid).

Note that depending on the value of S , some of those cases can be excluded. The solid vertical red and black lines represent respectively the values on the ask ($\alpha^a \mathbb{Z}$) and the bid grid ($\alpha^b \mathbb{Z}$). The dotted vertical lines represent the limits of the uncertainty zones on each side.

In the uncertainty zones, the value function h depends non-trivially on S^a and S^b . Thanks to the continuity conditions at the boundaries of the uncertainty zones, we get a smooth behavior of h when S exits a zone. Remark that when $S \in [10 \pm ((\frac{1}{2} - \eta^a)\alpha^a) \wedge ((\frac{1}{2} - \eta^b)\alpha^b)]$, necessarily $S^a = S^b = 10$.

In our example, $\alpha^a > \alpha^b$ and $(\frac{1}{2} - \eta_a)\alpha^a > (\frac{1}{2} - \eta_b)\alpha^b$. So, if S is in $(10 + (\frac{1}{2} - \eta_b)\alpha^b, (\frac{1}{2} - \eta_a)\alpha^a)$, necessarily $S^a = 10$, but S^b can take either the value 10 or $10 + \alpha^b$ depending on whether S comes from higher prices or lower prices. This is why there are two curves in the interval $(10 + (\frac{1}{2} - \eta_b)\alpha^b, (\frac{1}{2} - \eta_a)\alpha^a)$. At $(\frac{1}{2} - \eta_a)\alpha^a$, two additional curves appear as S^a can also be two different values.

Chapter 4

How to design a derivatives market?

4.1 Introduction

Nowadays a typical role of an exchange is to give the possibility to investors to buy or sell financial products on electronic platforms, in sufficiently large quantity and at a reasonable price. Therefore exchanges have to set up their markets in a relevant way in order to achieve this goal. The issues related to market design cover a wide range of topics, from the microstructure of electronic trading platforms to the basic question of selecting the products that will be traded on the exchange.

Recently many papers have focused on the microstructural aspects of market design. For example the way of choosing an optimal tick size is addressed in Dayri and Rosenbaum [91], where the authors study the relations between tick size, volatility and bid-ask bounce frequency. In Budish, Cramton, and Shim [59], Jusselin, Mastrolia, and Rosenbaum [178], the relevance of continuous trading and its comparison with a frequent batch auction system is discussed, while market fragmentation is analyzed in Laruelle and Lehalle [192]. Macroscopic features have also been investigated, see for example Kalagnanam and Parkes [180], where different market structures are classified with respect to several criteria such as matching mechanism, information feedback and bid structure.

Most of the research on market design focuses on stock markets. However, even if exchanges concentrate a large part of their activities on simple products such as stocks or futures, many also offer to their clients the possibility to trade more complex financial instruments such as derivatives. Actually there is very few academic literature on derivatives market design, mostly addressing the relationship between stock and option markets. For example in Mayhew and Mihov [209] the authors investigate the factors influencing the selection of stocks for option listing. However, they neither question the optimality of those factors, nor search for more relevant ones. The papers dealing with market design can in fact be separated into two groups: the ones that review and try to understand market practice and those proposing a theoretical framework in order to help exchanges improve their market design. Surprisingly, to our knowledge, there is no paper of the last kind dealing with derivatives market. In this chapter we propose a first contribution in that direction.

We take the realistic point of view of an exchange who wants to organize, or reorganize, its derivatives market. We consider that the market is made of vanilla European options only, that we view as independent of the underlying. By this, we mean that we deal with options that are used as hedging instruments and whose prices are essentially fixed by supply and demand. Finally we suppose that the exchange has access to data allowing for the estimation of the distribution of options market demand. For example, if the exchange already has a derivatives market it can use its own data, otherwise that of other exchanges. We focus on two issues: selecting the options that are going to be traded and attracting liquidity on those options.

The first issue faced by the exchange is the choice of the derivatives offered to the clients. Obviously it is impossible for the exchange to propose all maturities and strikes on its platform. This would be very hard to manage from a technical point of view and it would be impossible to guarantee liquidity on each option. As the maturities are quite standardized, the main challenge relies in strikes selection satisfying clients needs. Therefore, we consider that the exchange's problem is to select n call options (or equivalently n strikes), with fixed maturity, with the aim of maximizing the clients satisfaction. We define a measurement of this satisfaction and write the exchange objective under the form of a quantization problems. We refer to Graf and Luschgy [132], Pagès, Pham, and Printems [228] for an introduction to quantization. Such approach allows the exchange to select automatically a set of options based only on the data at its disposal.

The next goal of the exchange is to attract liquidity on its platform in order to increase the amount of executed orders. To do so, one way is to use a make take fees system: the exchange typically associates a fee rebate to executed limit orders, while charging a transaction fee for market orders. This enables it to subsidize liquidity provision and tax liquidity consumption. In El Euch, Mastrolia, Rosenbaum, and Touzi [109] the authors design the optimal make take fees policy for a market with one market maker and a single undeying asset. This work has been extended in Chapter 2 to the case of multiple market makers. The general principle of the approach in El Euch, Mastrolia, Rosenbaum, and Touzi [109] and Chapter 2 is to consider that the exchange offers a contract to the market maker whose pay-off depends on the market order flow he generates. The problem of the exchange then boils down into designing the optimal contract in order to optimize the number of transactions.

However, in our setting the problem faced by the exchange is more complex to several extents. The main difference with the framework of El Euch, Mastrolia, Rosenbaum, and Touzi [109] and Chapter 2 is that the exchange has to manage several assets simultaneously, namely the different options quoted on the platform. In order to focus on this issue we assume that there is only one market maker setting bid and ask quotes for all available options. Another challenge for a derivatives exchange is the possible absence of quotations for far from the money options (or quotations with a too wide spread). Such issue arises essentially for commercial reasons. Indeed, an exchange does not wish to display to its clients a product with scarce liquidity. It wants to make sure that there is sufficient available volume on the market for the whole range of listed options. Therefore, the design of an optimal make take fees policy for options market must aim at providing incentives to the market maker to lower the spreads, notably for far from the money options.

To do so, we are inspired by El Euch, Mastrolia, Rosenbaum, and Touzi [109], using a principal-agent framework. The exchange (the principal) has to design a contract towards the market maker (the agent) that maximizes a certain utility that depends on the behavior of the market maker. The main point is that the market maker's behavior, here the quoted spread on every available option, cannot be dictated by the exchange and depends on the contract. For example if the contract offers high incentives for every executed ask market order, then it is likely that the ask price quoted by the market maker will be close to the mid price. Formally, for a given contract, the market maker determines its behavior by solving a stochastic control problem. Then in order to find the optimal contract, the exchange maximizes its expected utility over the set of admissible contracts, knowing the market maker's response to each contract. The chapter is organized as follows. In Section 4.2 we explain how an exchange can select the options that will be traded on its platform using only market data. Then in Section 4.3 we design the optimal contract that the exchange should offer to the market maker in order to maximize liquidity. Proofs and technical results are relegated to the Appendix.

4.2 Market driven selection of the listed options

In this section we build a method for the exchange to select the strikes that are going to be traded on its platform. This approach uses only data from trades volume reports and is based on a quantization algorithm. We illustrate this method by numerical experiments using data provided by Euronext.

4.2.1 How to choose the strikes in order to match market demand?

We consider European call options with strikes expressed in percentage of the spot price (in moneyness) and that the exchange wishes to select n strikes.¹ Choosing relevant strikes, the exchange's objective is to maximize the satisfaction of the investors. So, we focus in this section on the market taking side of the trading flow. Section 4.3 will be rather devoted to market makers. We measure the regret of a market taker associated to the execution of a market order as a function of the difference between the strike he would have ideally bought (or sold) and the strike he actually bought (or sold). More precisely, for a given maturity, consider strikes $K_1 < \dots < K_n$ that represent the options listed by the exchange. When a market taker wants to buy an option with strike K he sends a market order on the option whose strike is the nearest from K . Hence he buys (or sells) the option with strike K_i where i is such that

$$K_i = \arg \min_{1 \leq j \leq n} |K - K_j|.$$

We consider that the regret associated to this market order is $\rho(|K - K_i|)$ where ρ is an increasing function. Note that the regret of the market order can be written $\min_{1 \leq j \leq n} \rho(|K - K_j|)$. We finally assume that the strike K is randomly chosen according to the distribution \mathbb{P}^{mkt} . This probability measure represents the law of market demand. Thus the higher the demand for a given strike the higher the probability that K is close to this strike. The exchange can easily estimate the distribution \mathbb{P}^{mkt} using data from its own options market or from other exchanges. The average regret of a market order is therefore written

$$\mathbb{E}^{mkt} \left[\min_{1 \leq j \leq n} \rho(|K - K_j|) \right], \quad (4.1)$$

where \mathbb{E}^{mkt} denotes the expectation when $K \sim \mathbb{P}^{mkt}$. The problem of the exchange is then to find the n -uplet $(K_i)_{1 \leq i \leq n}$ that minimizes (4.1). Formally this corresponds to the following minimization problem:

$$\arg \min_{K_1 \leq \dots \leq K_n} \mathbb{E}^{mkt} \left[\min_{1 \leq j \leq n} \rho(|K - K_j|) \right]. \quad (4.2)$$

This type of optimization is classical in the field of signal or image processing and is called *quantization* problem. The main idea of quantization is to summarize the information contained in a complex probability measure into a uniform probability with finite support. As an example, it allows to compress a signal (or an image) by selecting among its spectrum a given number of frequencies that summarizes the signal with the smallest possible loss of information. For an introduction to quantization problem see Graf and Luschgy [132], Pagès, Pham, and Printems [228].

In this chapter we consider the quantization problem (4.2) when ρ is a power-law function of the form $\rho(x) = |x|^p$ with $p \geq 2$. The power-law function has the advantage to be symmetric and convex. Therefore greater errors are increasingly penalized. As a consequence we expect the solution of (4.2) to capture the features of the tails of \mathbb{P}^{mkt} . Moreover the greater p , the more large errors are penalized. Hence for a large p , the $(K_i)_{1 \leq i \leq n}$ solution of (4.2) are likely to be more spread towards large strikes and contain more extreme values of the distribution \mathbb{P}^{mkt} .

4.2.2 Solving the quantization problem

In this section we give some sufficient conditions that ensure that (4.2) has a unique solution. We also explain how (4.2) can be solved. To get existence of a solution to the problem (4.2) we need to make the following assumption.

¹We do not address here the problem of choosing the number of strikes to propose. This point is left for further research.

Assumption 4.1. *The probability \mathbb{P}^{mkt} is absolutely continuous with respect to the Lebesgue measure with density that is log-concave and compactly supported in $[0, \bar{K}]$, $\bar{K} > 0$.*

The assumption on the support of the probability is very reasonable since strikes between 0 and 200% of the spot price basically cover all the possible strikes of traded options. The log-concavity assumption is not really restrictive since it allows us to consider a wide class of probability distributions such as exponential type and Gaussian laws. It is shown in [132, Theorem I-5.1] that under Assumption 4.1, Problem (4.2) admits a unique non degenerate solution. The term non degenerate simply means that the optimal set of strikes satisfies $K_1 < \dots < K_n$.

We now present a way to approximate numerically the solution of (4.2). The idea behind the algorithm is that the solution $(K_i)_{1 \leq i \leq n}$ can be seen as the fixed point of a function. This provides us a numerical method to approximate the $(K_i)_{1 \leq i \leq n}$ that consists in iterating this function. This is known as the Lloyd's algorithm, which is a very intuitive approach that searches step by step the solution of (4.2). A very convenient aspect of this algorithm is that it is automatic and easy to implement.

Lloyd's algorithm starts with an initial set of strikes $(K_i)_{1 \leq i \leq n}$ and is made of three steps:

1. For any i , identify A_i the set of "wished" strikes that corresponds to market orders sent to the strike K_i . Equivalently A_i contains all the strikes K which are closer to K_i than from any other K_j

$$A_i = \left\{ K : i = \arg \min_{1 \leq j \leq n} |K - K_j| \right\}.$$

2. Set K'_i as the unique strike in A_i that minimizes the average regret of market orders sent with ideal strike in A_i . More precisely K'_i is given by

$$K'_i = \arg \min_{k \in A_i} \mathbb{E}^{mkt} [|K - k|^p \mathbf{1}_{K \in A_i}].$$

3. Go back to Step 1 with $(K_i)_{1 \leq i \leq n} = (K'_i)_{1 \leq i \leq n}$ (or stop if a certain stopping criterion is reached and consider $(K'_i)_{1 \leq i \leq n}$ as the approximate solution of (4.2)).

Lloyd's algorithm has a very clear interpretation in terms of selecting the optimal set of strikes: first it identifies the area "controlled" by the i -th strike and then improves the choice of the strikes. It is then intuitive that the solution of (4.2) is a fixed point of Lloyd's algorithm. The sets $(A_i)_{1 \leq i \leq n}$ form a covering of \mathbb{R}_+ that is often called the Voronoï tessellation associated to the $(K_i)_{1 \leq i \leq n}$. It is easy to show that, for Step 1

$$A_1 = [0, K_1], \quad A_n = [K_n, \bar{K}] \quad \text{and for } i \in \{2, \dots, n-1\} : A_i = \left[\frac{K_i + K_{i-1}}{2}, \frac{K_{i+1} + K_i}{2} \right].$$

A usual stopping criterion for Step 3 is when $(K'_i)_{1 \leq i \leq n}$ is too close from $(K_i)_{1 \leq i \leq n}$. More precisely the algorithm stops if $\sum_{i=1}^n |K'_i - K_i| < \varepsilon$, for a certain $\varepsilon > 0$. Note that, starting from a discrete valued \mathbb{P}^{mkt} (as will be the case here), when $p = 2$, Step 2 of Lloyd's algorithm boils down to compute the average realization of \mathbb{P}^{mkt} conditional on being in A_i . This can be obtained instantaneously. However when $p > 2$, Step 2 is not straightforward to compute in general. Yet the objective function being convex and taking the derivative with respect to k , a necessary and sufficient condition for k to be solution of Step 2 is $\mathbb{E}[|K - k|^{p-2}(K - k) \mathbf{1}_{K \in A_i}] = 0$ or equivalently

$$k = \frac{\mathbb{E}^{mkt}[K |K - k|^{p-2} \mathbf{1}_{K \in A_i}]}{\mathbb{E}^{mkt}[|K - k|^{p-2} \mathbf{1}_{K \in A_i}]}.$$

This characterizes the solution of Step 2 as a fixed point. Thus one usually replaces Step 2 by its iterative version:

$$K'_i = \frac{\mathbb{E}^{mkt}[K |K - K_i|^{p-2} \mathbf{1}_{K \in A_i}]}{\mathbb{E}^{mkt}[|K - K_i|^{p-2} \mathbf{1}_{K \in A_i}]}.$$

From now on, we call Lloyd's algorithm the initial algorithm where we replace Step 2 by its approximate version. We prove in Appendix 4.1 that $(K_i)_{1 \leq i \leq n}$ is solution of (4.2) if and only if it is a fixed point of the Lloyd's algorithm. The great strength of this method is that it is easy to implement, transparent, and completely automatic. Note also that if \mathbb{P}^{mkt} has a discrete support, say 10 strikes, then Lloyd's algorithm will not necessarily select those strikes as solution of (4.2). We now turn to numerical experiments illustrating the efficiency of our method.

4.2.3 Application

In this section we apply our methodology to market data. First we describe the data and then present our numerical results.

4.2.3.1 Description of the data

We use data from Euronext, one of the main stock and option exchanges in Europe. The dataset contains for every trading day from the 3-rd of December 2018 to the 24-th of May 2019 and for every available options the total number of trades (buy and sell) during the day. Our dataset is only made of transactions that occurred on the Euronext platform. In particular we neither use OTC data nor data from another exchange. We choose for our example the most standard call options in terms of underlying on Euronext, namely options on the CAC 40 index. We report in Table 4.1 the number of call options traded each month for different ranges of maturity and in Table 4.2 the number of call options traded each month for each strike.




Figure 4.1: Empirical distribution of traded option strikes (left). Quantile plot in log-scale of traded option maturities for the whole sample set (right).

In Figure 4.1, we display the empirical distribution of traded option strikes (for all maturities) and the quantile plot of the maturity distribution in log-scale. The distribution of the strikes is unimodal, concentrated near the money and skewed towards in the money strikes. In Figure 4.2, we provide the empirical distribution of traded options strikes for different ranges of maturity. We see that the distribution of the strikes depends on the maturity. In particular, the variance of the distribution is increasing with the maturity. The skewness towards in the money strikes is present for any maturity.

Maturity	December	January	February	March	April	May
$T \leq 1M$	135951	99202	96323	191357	161937	108491
$1M < T \leq 3M$	79016	61651	30371	117400	58914	121267
$3M < T \leq 6M$	10990	13279	15979	33901	11227	11779
$6M < T$	71977	30278	14197	17158	25354	21330

Table 4.1: Number of options traded by maturity and month.




Figure 4.2: Empirical distribution of the strikes for different maturities.

4.2.3.2 Numerical results

We now present our numerical results. Since the distribution of the strikes depends on the maturity and because short maturities are over-represented in our data, we split our dataset into four subsets depending on the maturity:

- maturity less than 1 month,
- maturity between 1 and 3 months,
- maturity between 3 and 6 months,
- maturity larger than 6 months.


For any of those subsets we approximate the solution of the quantization problem (4.2) using the Lloyd's algorithm for $n = 10$ and with stopping parameter $\varepsilon = 10^{-8}$. As initial value, we use n points $(K_i)_{1 \leq i \leq n}$ generated with uniform law between the 10-th and 90-th percentile of the dataset. In Figures 4.3 and 4.4 we plot a visualization of the quantization of the different sets obtained for $p = 2$ and $p = 8$.

The strikes selected by Lloyd's algorithm manage to reproduce some of the statistical properties of the demand distribution \mathbb{P}^{mkt} . In particular, for any range of maturity, the distribution of the $(K_i)_{1 \leq i \leq n}$ is skewed towards in the money strikes. Also the variance of the selected strikes is increasing with the maturity as for market data.

We observe that for $p = 8$ the strikes selected by the quantization method are more spread towards large strikes than for $p = 2$. This is not surprising since the penalization of large errors is increasing with p for the regret function $|\cdot|^p$. Therefore, as expected, the larger p , the more the solution of the quantization problem (4.2) contains extreme values of the distribution \mathbb{P}^{mkt} . We also note that the selected strikes for $p = 8$ exhibit some kind of redundancy: some of them are very close to each other. In practice, one would of course discard one of two strikes being very close (it may then be interesting to take a smaller n). For practical applications, the easiest approach is probably to use $p = 2$. With this choice, the Lloyd's algorithm is very fast and easy to implement. It also corresponds to the most documented case.

Strike (%)	December	January	February	March	April	May
20	0	0	0	0	55	10
30	1	1692	2	381	0	0
40	0	77	0	80	3	41
50	58	417	0	328	2031	1948
60	1933	152	31	323	691	2092
70	1402	1928	653	3837	2412	2956
80	12814	12952	3400	10118	14689	12147
90	113210	114463	10465	247877	184835	147362
100	159075	68747	130002	94714	50621	90528
110	5811	3586	12253	1766	83	2205
120	869	94	64	11	0	16
130	1	11	0	0	0	0
140	0	0	0	0	2012	1960
150	0	0	0	381	0	1602
160	1720	271	0	0	0	0
170	1040	20	0	0	0	0

Table 4.2: Number of options traded by strike and month.

Figure 4.3: Quantization of the option strikes using $p = 2$ and $\varepsilon = 10^{-8}$. Empirical distribution of traded strikes is plotted in blue or red. The dotted lines correspond to the optimal quantization of \mathbb{P}^{mkt} .

Finally we insist on the fact that when an exchange uses our methodology for strikes selection, it is interesting, if possible, to include transactions from other exchanges and from the OTC market in the dataset. This is because using only its own trade data may induce a bias in the strikes selection. For example if for some reasons clients of an exchange go on other venues to buy (or sell) out of the money options, then, in the exchange dataset, there will be very few transactions reported on out of the money options. This will lead to inaccuracies since the demand for out of the money options will be underestimated.

We now turn to the problem of providing incentives to the market maker to quote attractive spreads in order to attract liquidity towards the selected options.




Figure 4.4: Quantization of the option strikes using $p = 8$ and $\varepsilon = 10^{-8}$. Empirical distribution of traded strikes is plotted in blue or red. The dotted lines correspond to the optimal quantization of \mathbb{P}^{mkt} .

4.3 Incentive policy of the exchange

In this section, we assume that the exchange has already selected a list of options. The goal is to design a contract between the exchange and the market maker so that the latter receives incentives to provide suitable liquidity on all the options. We first describe the market and assumptions. In particular, due to the short time horizon we are working on, we can assume a Bachelier model for the underlying asset and constant delta for the options. Then, we introduce a class of tractable admissible contracts proposed to the market maker. These contracts are indexed on the transactions induced by the behavior of the market maker. We show that there is no loss of generality in considering such class of contracts. For a given contract, the market maker solves an optimization problem to deduce its optimal quotes for each option. Then, the exchange maximizes his expected utility over the set of admissible contracts, knowing the response of the market maker to a given contract.

The utility of the exchange is made of two parts: one component related to the actual Profit and Loss (PnL for short) due to transactions, and one aiming at ensuring that enough liquidity is constantly posted on every option. As explained in the introduction, this second component addresses commercial constraints in order to make the exchange competitive. In particular, our model is flexible and can be designed so that the exchange has more interest in reducing the spreads for far from the money options, although not very traded, than for near the money options. We derive explicitly the optimal incentives that should be offered, up to the resolution of a two-dimensional linear PDE.

We conclude this section with numerical results showing the impact of the incentive policy on the spread of the listed options. We emphasize that we adopt a financial engineering viewpoint in this section, that is our only goal is to provide a quantitative and operational answer to the lack of liquidity on certain options listed by the exchange.

4.3.1 The market

This section is devoted to the description of the market model. We consider a finite trading horizon time $T > 0$ and a probability space $(\Omega, \mathcal{F}, \mathbb{P}^0)$ under which all stochastic processes are defined. Following

Section 4.2, we work on a market where European call options with strike $k \in \mathcal{K} := \{K_1, \dots, K_n\}$ and maturity $\tau \in \mathcal{T} := \{T_1, \dots, T_m\}$ can be traded. We focus on call options but our results can be extended to put options in a straightforward manner.

Remark 21. *Here, the set of strike and maturities is unspecified. However, the problem would naturally be studied over the set derived using the quantization method in the previous section. This means that the Principal will design the contract proposed to the market-maker based on an exogeneous demand \mathbb{P}^{mkt} , calibrated on market data.*

The price of the underlying, observable by all market participants, has a dynamic given by

$$dS_t = \sigma dW_t, \quad (4.3)$$

where $\sigma > 0$ is the volatility of the asset and W is a one-dimensional Brownian motion. The choice of an arithmetic Brownian motion is motivated by the fact that we use a reasonably short time horizon T (less than one day). On such scale, Bachelier and Black-Scholes type dynamics are quite indistinguishable. Moreover, the Bachelier price has little probability to become negative.

In this model, we have in mind a time horizon T of one trading day. This is a reasonable assumption as the trading horizon of a market maker is usually one day, with the objective to end the trading day with a flat inventory. Therefore, the volatility σ must be interpreted as the daily volatility of the underlying asset. Moreover, the use of a stochastic volatility would be inefficient as on such time scale there is very little diffusion of the volatility process.

Assuming zero interest rate, we write the price at time t of the call option with maturity τ and strike k as $C_t^{k,\tau}$. Its dynamic is given by

$$dC_t^{k,\tau} = \sigma \Delta_t^{k,\tau} dW_t, \quad (4.4)$$

where $\Delta_t^{k,\tau} := \mathcal{N}(d_t)$ is the Bachelier delta of the call option $C_t^{k,\tau}$ at time t , $\mathcal{N}(\cdot)$ is the cumulative distribution function of the standard Gaussian law and $d_t := \frac{S_t - k}{\sigma\sqrt{\tau}}$.

Remark 22. *In this model, option prices are linear with respect to the underlying price, meaning that the value of the option comes from its Delta. We could consider the full Bachelier dynamics of a call option, which include the Gamma exposure, but as we will make use of the constant Delta assumption to obtain semi-explicit formulas, it is irrelevant to add such exposure (which moves way faster than the Delta) in the option dynamics if it is assumed to be constant.*

As we work over a short time horizon, the delta of the quoted options does not vary significantly. Hence, throughout the chapter, we assume it to be constant.

Assumption 4.2. *We assume that*

$$\Delta_t^{k,\tau} = \Delta^{k,\tau}.$$

This assumption, which can be relaxed, leads to technical simplifications. Note that in our problem setting, as the considered time horizon is one trading day, it is very reasonable to assume a constant delta (by taking the one at the opening of the market) and recalibrate it at the end of the day. This will lead to a different pay-off of the contract for the market maker in case of a significant price move from one day to another.

The market maker displays bid and ask quotes on the listed options. The market maker best bid price and best ask price at time t on the option with maturity τ and strike k are

$$P_t^{k,\tau,b} := C_t^{k,\tau} - \delta_t^{k,\tau,b}, \quad P_t^{k,\tau,a} = C_t^{k,\tau} + \delta_t^{k,\tau,a}, \quad t \in [0, T],$$

where the superscript b (resp. a) stands for bid (resp. ask). So we consider that the market maker

controls the spreads $\delta^{k,\tau} := (\delta^{k,\tau,a}, \delta^{k,\tau,b})$ on each option. The set of admissible controls for the market maker is therefore defined as

$$\mathcal{A} := \{(\delta_t)_{t \in [0,T]} = (\delta_t^{k,\tau,i})_{t \in [0,T]} : k \in \mathcal{K}, \tau \in \mathcal{T}, i \in \{a, b\}, \text{predictable}, |\delta_t^{k,\tau,i}| \leq \delta_\infty\}, \quad (4.5)$$

where $\delta_\infty > 0$ is a constant, assumed to be large enough to satisfy technical conditions (see Appendix 4.A.7). In practice it is of course not restrictive to assume that the spreads are bounded.

We now describe the dynamics of the market order flow. For every listed option, the arrival of ask (resp. bid) market orders is modeled by a point process $N^{k,\tau,a}$ (resp. $N^{k,\tau,b}$). We expect the intensity of buy (resp. sell) market order arrivals to be a decreasing function of both the spread quoted by the market maker $\delta^{k,\tau}$ and the transaction cost $f^{k,\tau}$ collected by the exchange. This has quite natural interpretation as a wider spread or higher fee decreases the number of transactions on the considered option. Moreover, we know from the literature (see Dayri and Rosenbaum [91], Madhavan, Richardson, and Roomans [203] and Wyart, Bouchaud, Kockelkoren, Potters, and Vettorazzo [271]) that the average number of trades per unit of time for single assets is a decreasing function of the ratio between spread and volatility. Assuming same kind of behavior for the options, this leads to the following form of the intensity function:

$$\lambda^{k,\tau}(\delta_t^{k,\tau,i}) := A \exp\left(-\frac{C}{\sigma}(\delta_t^{k,\tau,i} + f^{k,\tau})\right),$$

where A and C are positive constants that can be calibrated using market data, and $f^{k,\tau}$ represents the fee fixed by the exchange for each market order. Furthermore, we assume that all market orders are of unit size.

Remark 23. *In Equation (4.3), the efficient price is independent of the order flow. However, it is a direct extension to assume the dynamics*

$$dS_t = \sigma dW_t + \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \xi^{k,\tau} (dN_t^{k,\tau,a} - dN_t^{k,\tau,b}),$$

where $\xi^{k,\tau}$ represents the instantaneous impact of a trade on option $C^{k,\tau}$.

The parameter A represents the average number of orders sent on the option market, whereas C is the sensitivity of the intensities with respect to the spread. A direct generalization is set different order flow parameters $A^{k,\tau,i}$ for each side of each option. This helps to model the fact that there is on average more trading on at-the-money and in-the-money options, and that there is usually more option buyers than sellers. The main difficulty in our framework is that the market maker is dealing with multiple derivatives. If the market maker strategy depends on its inventory on each option, then the problem lies in dimension n , which becomes intricate for large n . However, we will see that we can circumvent this issue since in our case we can aggregate the risk factors related to the inventories through the delta weighted cumulated inventory:

$$\mathcal{Q}_t := \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \Delta^{k,\tau} Q_t^{k,\tau}, \quad (4.6)$$

where $Q_t^{k,\tau} := N_t^{k,\tau,b} - N_t^{k,\tau,a}$ is the number of options $C^{k,\tau}$ held by the market maker at time t . Each inventory is weighted by the corresponding Δ (see Section 4.3.2 for details). Thus, the quantity \mathcal{Q} represents the marked-to-market value of the market maker's portfolio. It therefore contains the market risk carried by the market maker. For example an out of the money option will account for a small part of the total risk, and conversely for in the money options. Finally we consider that the market maker has a critical absolute inventory $\bar{q} \in \mathbb{N}$. The intensity of the orders arrival is then

$$\lambda_t^{k,\tau,i} := \lambda^{k,\tau}(\delta_t^{k,\tau,i}) \mathbf{1}_{\{\phi(i)\mathcal{Q}_t > -\bar{q}\}} \text{ with } \phi(i) := \begin{cases} 1 & \text{if } i = a, \\ -1 & \text{if } i = b. \end{cases}$$

Remark 24. *Note that there is a direct link between the spread quoted by the market maker and his inventory process. Indeed a lower spread $\delta^{k,\tau,b}$ (resp. $\delta^{k,\tau,a}$) on the bid (resp. ask) side of the listed option $C^{k,\tau}$ increases the intensity of orders arrival $\lambda^{k,\tau,b}$ (resp. $\lambda^{k,\tau,a}$). This leads to an increase (resp. decrease) of the inventory process $Q^{k,\tau}$. In other words, the market maker skews his quotes depending on the level of its aggregated inventory.*

4.3.2 Market maker's problem and contract representation

In this section we exhibit the class of contracts used by the exchange. We also explain and solve the market maker's problem for any admissible contract.

The PnL of the market maker is defined as the sum of the cash earned from his executed orders and of the value of his inventory on each traded option. Thus, using that $\sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} Q_t^{k,\tau} \Delta^{k,\tau} S_t = Q_t S_t$, it writes

$$PL_t^\delta := \mathcal{W}_t^\delta + Q_t S_t, \quad (4.7)$$

where

$$\mathcal{W}_t^\delta := \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \int_0^t P_u^{k,\tau,a} dN_u^{k,\tau,a} - \int_0^t P_u^{k,\tau,b} dN_u^{k,\tau,b},$$

stands for his cash process at time $t \in [0, T]$. This expression shows the relevance of the variable Q for the market maker. It represents the volatility of the market maker's PnL with respect to the underlying price movements. Using (4.4), a direct integration by parts leads to the following form of the PnL process:

$$PL_t^\delta := \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \int_0^t \delta_u^{k,\tau,i} dN_u^{k,\tau,i} + Q_u dS_u.$$

Remark 25. *In this framework, we consider that there is no hedging of the market maker when he buy or sells an option. This can be added straightforwardly to the model in the spirit of Chapter 6 of this thesis, see Appendix 4.A.8. The results we obtain are not modified significantly.*

Moreover, the exchange offers to the market maker a contract ξ , namely an \mathcal{F}_T -measurable random variable, which is added to his PnL at the end of the trading period. This contract aims at incentivizing the market maker to reduce the spread quoted for each option. More details will be given in Section 4.3.3. The contract depends on all the transactions occurring between time 0 and time T , as well as on the efficient price moves.

Thus taking an exponential utility function, the market maker maximizes the following functional of his wealth:

$$V_{\text{MM}}(\xi) := \sup_{\delta \in \mathcal{A}} \mathbb{E}^\delta \left[- \exp \left(- \gamma (\xi + PL_T^\delta) \right) \right], \quad (4.8)$$

where $\gamma > 0$ denotes the market maker's risk aversion parameter and \mathbb{E}^δ the probability measure associated to a given control process $\delta \in \mathcal{A}$, see Appendix 4.A.2.2 for details. For the well-posedness of Equation (4.8), we need integrability conditions on the contract ξ , see Appendix 4.A.3 for details.

Finally we consider that the market maker accepts a contract ξ only if its associated optimal expected utility $V_{\text{MM}}(\xi)$ is above some fixed threshold $R < 0$. This threshold, called reservation utility of the agent, is the critical utility value under which the market maker has no interest in the contract. This quantity has to be taken into account carefully by the exchange before proposing a contract to the market makers.

We now introduce the class of contracts proposed to the market maker. Given $Y_0 > 0$, and predictable processes $Z := (Z^{C^{k,\tau}}, Z^{k,\tau,i})_{k \in \mathcal{K} \tau \in \mathcal{T} i \in \{a,b\}} \in \mathcal{Z}$ (see Appendix 4.A.3 for a definition of \mathcal{Z}), we introduce a special class of remuneration $\xi = Y_T^{Y_0, Z}$ of the form

$$\begin{aligned} Y_T^{Y_0, Z} := & Y_0 + \int_0^T \left(\sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} Z_r^{k,\tau,i} dN_r^{k,\tau,i} + Z_r^{C^{k,\tau}} dC_r^{k,\tau} \right) \\ & + \left(\frac{1}{2} \gamma \sigma^2 \left(\sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \Delta^{k,\tau} (Z_r^{C^{k,\tau}} + Q_r^{k,\tau}) \right)^2 - H(Z_r, Q_r) \right) dr, \end{aligned} \quad (4.9)$$

where for $(z, q) \in \mathbb{R}^{2 \times \#\mathcal{K} \times \#\mathcal{T}} \times \mathbb{R}$ with $z := (z^{k,\tau,i})_{\substack{(k,\tau) \in \mathcal{K} \times \mathcal{T} \\ i \in \{a,b\}}}$, the function H , called Hamiltonian of the market maker, is defined by²

$$H(z, Q) := \sup_{\delta \in \mathbb{R}^{2 \times \#\mathcal{K} \times \#\mathcal{T}}} h(\delta, z, Q)$$

with

$$h(\delta, z, Q) := \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \gamma^{-1} \left(1 - \exp \left(-\gamma (z^{k,\tau,i} + \delta^{k,\tau,i}) \right) \right) \lambda^{k,\tau} (\delta^{k,\tau,i}) \mathbf{1}_{\{\phi(i)Q > -\bar{q}\}}.$$

Actually, it turns out that it is enough to consider contracts of the form (4.9). More precisely, we show that any admissible contract (in the sense of the integrability conditions specified in Appendix 4.20), is of this form. We have the following lemma proved in Appendix 4.A.5.

Lemma 4.1. *Any contract ξ satisfying (4.20) has a unique representation $\xi = Y_T^{Y_0, Z}$ for some $(Y_0, Z) \in \mathbb{R} \times \mathcal{Z}$.*

Furthermore, the terms defining (4.9) have natural interpretation.

- The compensation Y_0 is calibrated by the exchange to ensure the reservation utility constraint with level R of the market maker.³
- The term $Z^{C^{k,\tau}}$ is the compensation given to the market maker with respect to the volatility risk induced by the option $C^{k,\tau}$.
- Each time a trade is executed on the ask (resp. bid) side for the option $C^{k,\tau}$, the market maker is compensated by the term $Z^{k,\tau,a}$ (resp. $Z^{k,\tau,b}$).
- The term $\frac{1}{2} \gamma \sigma^2 \left(\sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \Delta^{k,\tau} (Z^{C^{k,\tau}} + Q^{k,\tau}) \right)^2 - H(Z, Q)$ is a continuous coupon given to the market maker.

When the market maker remuneration is $Y^{Y_0, Z}$, its optimal response can be computed explicitly as a functional of Z .

Theorem 4.1. *For $\xi = Y^{Y_0, Z}$, the market maker utility is*

$$V_{\text{MM}}(Y_T^{Y_0, Z}) = -\exp(-\gamma Y_0),$$

associated to the optimal bid-ask policy $\hat{\delta}_t^{k,\tau,i}(\xi) := \Delta^i(Z_t^{k,\tau,i})$, where

$$\Delta^i(Z_t^{k,\tau,i}) := (-\delta_\infty) \vee \left(-Z_t^{k,\tau,i} + \frac{1}{\gamma} \log \left(1 + \frac{\sigma \gamma}{C} \right) \right) \wedge \delta_\infty \text{ for } (k, \tau, i) \in \mathcal{K} \times \mathcal{T} \times \{a, b\}. \quad (4.10)$$

Theorem 4.1 provides the optimal response of the market maker to any contract of the form (4.9), see Appendix 4.A.6 for the proof. Moreover from Equation (4.10), we get that the exchange can anticipate the optimal behavior of the market maker. It is therefore easy for the platform to compute its own utility for a given contract.

²This Hamiltonian term appears naturally when applying the dynamic programming principle for the market maker's problem.

³From Theorem 4.1, we see that taking $Y_0 = -\log(-R)$ ensures the reservation utility of the market maker.

4.3.3 Solving the exchange's problem

In this section we formalize the goal of the exchange and solve the problem of designing the optimal contract.

4.3.3.1 Description of the exchange's problem

We recall that the exchange has two objectives. The first one is to receive a high number of trades to collect the associated fees. The second is to have small spreads on its platform, in particular for far from the money options for which spreads are typically large. This is because the clients want to have sufficient liquidity on the whole list of options. In order to quantify the first objective, we introduce a weighted version of the total number of trades:

$$\mathcal{N}_t = \sum_{i=a,b} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} c^{k,\tau} N_t^{k,\tau,i},$$

where for any $(k, \tau) \in \mathcal{K} \times \mathcal{T}$, $c^{k,\tau} \geq 0$ represents the value attributed to a trade on the option $C^{k,\tau}$ by the exchange.⁴ Hence the more the exchange wants to attract liquidity on the option $C^{k,\tau}$, the higher $c^{k,\tau}$ has to be. If the considered option is very liquid (at the money options for example), the exchange may choose a rather small $c^{k,\tau}$.

To take into account the second objective, we consider the following quantity

$$\mathcal{L}_T^\delta := \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \int_0^T \omega (\delta_t^{k,\tau,i} - \delta_\infty^{k,\tau}) dN_t^{k,\tau,i}, \quad (4.11)$$

where $\omega \in [0, 1)$,⁵ and $\delta_\infty^{k,\tau}$ can be seen as a spread threshold the exchange would like to impose to the market maker. The more important the second objective for the exchange, the closer to one ω has to be chosen.

We thus consider that the exchange is looking for the contract ξ that maximizes the following quantity:

$$\mathbb{E}^{\hat{\delta}(\xi)} \left[- \exp \left(- \eta \left(\mathcal{N}_T - \mathcal{L}_T^{\delta(\xi)} - \xi \right) \right) \right], \quad (4.12)$$

where $\eta > 0$ is the risk aversion of the exchange and $\hat{\delta}(\xi)$ denotes the optimal response of the market maker given the contract ξ .

Remark 26. *We consider that the exchange does not have a reservation utility. For example, we can assume that R is so that $\xi = 0$ is an admissible contract, which would lead to positive PnL for the exchange.*

According to Lemma 4.1, we know that it is enough for the exchange to consider contracts of the form $Y_T^{Y_0, Z}$ with $(Y_0, Z) \in \mathbb{R} \times \mathcal{Z}$. So, (4.12) becomes

$$\mathbb{E}^{\hat{\delta}(Y^{Y_0, Z})} \left[- \exp \left(- \eta \left(\mathcal{N}_T - \mathcal{L}_T^{\delta(Y^{Y_0, Z})} - Y_T^{Y_0, Z} \right) \right) \right]. \quad (4.13)$$

Moreover for a contract of the form $Y^{Y_0, Z}$, from Theorem 4.1, the exchange knows the best response $\hat{\delta}(Y^{Y_0, Z})$ of the market maker. Indeed we recall that the optimal controls are given by

$$\hat{\delta}^{k,\tau,i}(Y^{Y_0, Z}) = \Delta^i(Z_t^{k,\tau,i}).$$

⁴One can for example take $c^{k,\tau} = f^{k,\tau}$. In this case, \mathcal{N}_T represents the total amount of fees collected by the exchange.

⁵The choice of $\omega \in [0, 1)$ is for technical reasons only.

It implies that

$$\mathcal{L}_T^{\hat{\delta}(Y^{Y_0,Z})} = \mathcal{L}_T^Z := \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \int_0^T \omega(\Delta^i(Z_t^{k,\tau,i}) - \delta_\infty^{k,\tau}) dN_t^{k,\tau,i}.$$

As in El Euch, Mastrolia, Rosenbaum, and Touzi [109], we notice that for a given contract $Y^{Y_0,Z}$, the market maker's optimal response does not depend on Y_0 . The exchange objective function (4.13) being decreasing in Y_0 , the maximization with respect to Y_0 is achieved at the level $\hat{Y}_0 = -\log(-R)$.⁶ Finally, the exchange problem becomes

$$V_0^E := \sup_{Z \in \mathcal{Z}} \mathbb{E}^{\Delta(Z)} \left[-\exp \left(-\eta \left(\mathcal{N}_T^Z - \mathcal{L}_T^Z - Y_T^{\hat{Y}_0,Z} \right) \right) \right]. \quad (4.14)$$

4.3.3.2 Stochastic control approach for the reduced exchange problem

In this section we solve the reduced exchange problem (4.14). We characterize the optimal contract components Z^* and explain how to compute them in practice. To solve this stochastic control problem, we study the associated Hamilton-Jacobi-Bellman (HJB for short) equation. This approach characterizes an optimal Z^* solving (4.14) under the form of a feedback function. The following result is proved in Appendix 4.A.7.

Theorem 4.2. *The maximization problem (4.14) admits a solution Z^* given by*

$$Z^{*k,\tau,i}(t, \mathcal{Q}_{t^-}) := \frac{1}{a-b} \log \left(\frac{bx_2 U(t, \mathcal{Q}_{t^-})}{ax_1^{k,\tau} U(t, \mathcal{Q}_{t^-} - \Delta^{k,\tau} \phi(i))} \right) \text{ and } Z^{*C^{k,\tau}}(t, Q_t^{k,\tau}) := -\frac{\gamma}{\gamma+\eta} Q_t^{k,\tau}, \quad (4.15)$$

for $(k, \tau, i) \in \mathcal{K} \times \mathcal{T} \times \{a, b\}$, where $a, b, (x_1^{k,\tau})_{k \in \mathcal{K}, \tau \in \mathcal{T}}$ and x_2 are constants defined in Appendix 4.A.7 and where $\tilde{U} := (-U)^{-\frac{C\gamma\eta}{\sigma\eta(1-\omega)}}$ is the unique solution of the following linear PDE on $[0, T] \times \mathbb{R}$:

$$\begin{cases} 0 = \partial_t \tilde{U}(t, \mathcal{Q}) - \tilde{U}(t, \mathcal{Q}) \frac{C\gamma\eta}{\gamma+\eta} \frac{\sigma}{2(1-\omega)} \mathcal{Q}^2 + \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \hat{C}^{k,\tau} \tilde{U}(t, \mathcal{Q} - \Delta^{k,\tau} \phi(i)) \mathbf{1}_{\{\phi(i)\mathcal{Q} > -\bar{q}\}}, \\ \tilde{U}(T, \mathcal{Q}) = 1, \end{cases} \quad (4.16)$$

where $\hat{C}^{k,\tau}$ are defined in Appendix 4.A.7.

Theorem 4.2 provides the incentives Z^* that maximize the exchange expected utility function, see Appendix 4.A.7 for the proof. The optimal contract is therefore given by

$$\begin{aligned} \xi^* = Y^{\hat{Y}_0, Z^*} = \hat{Y}_0 + \int_0^T \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \left(\sum_{i \in \{a,b\}} Z_r^{*k,\tau,i} dN_r^{k,\tau,i} + Z_r^{*C^{k,\tau}} dC_r^{k,\tau} \right) \\ + \int_0^T \left(\frac{1}{2} \gamma \sigma^2 \left(\sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \Delta^{k,\tau} (Z_r^{*C^{k,\tau}} + Q_r^{k,\tau}) \right)^2 - H(Z_r^*, \mathcal{Q}_r) \right) dr. \end{aligned} \quad (4.17)$$

We now provide some comments on the interpretation of the optimal incentives.

- The term $\int_0^T Z_u^{*C^{k,\tau}} dC_u^{k,\tau}$ in the optimal contract corresponds to part of the inventory risk process of the market maker $(Q_t^{k,\tau} C_t^{k,\tau})_{t \in [0, T]}$ that is supported by the exchange. As in El Euch, Mastrolia, Rosenbaum, and Touzi [109], the proportion of risk handled by the platform on each option is $\frac{\gamma}{\gamma+\eta}$. Hence, the more risk averse the exchange, the smallest this proportion.
- An application of Itô's formula gives the following approximation:

$$\log \left(\frac{\tilde{U}(t, \mathcal{Q})}{\tilde{U}(t, \mathcal{Q} - \Delta^{k,\tau} \phi(i))} \right) \approx \phi(i) 2 \frac{\sigma}{C} (T-t) \tilde{C} \Delta^{k,\tau} \mathcal{Q}, \quad (4.18)$$

where $\tilde{C} := \frac{C\gamma\eta}{\gamma+\eta} \frac{\sigma}{2(1-\omega)}$. Thus, when the aggregated inventory is highly positive, the exchange provides incentives to the market maker so that it attracts buy market orders and tries to dissuade him to accept more sell market orders, and conversely for a negative inventory.

⁶Note that $-\exp(-\hat{Y}_0) = R$.

- Numerically, we show that the incentive $Z^{*k,\tau,a}$ and $Z^{*k,\tau,b}$ given by (4.15) are increasing functions of the value $c^{k,\tau}$ that the principal associates to the option $C^{k,\tau}$. Hence, he logically provides higher incentives to an option he is more interested in.
- Although the principal manages a large number of listed options, we circumvent the curse of dimensionality by working with the aggregated inventory process. Note that the pay-off of the optimal contract depends only on t and \mathcal{Q} . Thus it is very easy to compute for the exchange at the end of the trading day.

In practice to implement the above methodology, one needs to compute the function \tilde{U} in order to design the optimal contract. A first way to do this is to use a classical finite difference scheme on the PDE (4.16). In Section 4.3.4 we use this technique for some numerical experiments on our method.

Moreover, as PDE (4.16) is linear, we can also resort to a probabilistic representation to compute \tilde{U} using a Monte-Carlo method. More precisely we have the following result which is a direct consequence of the Feynman-Kac formula.

Lemma 4.2. *We have the following representation:*

$$\tilde{U}(t, q) := \mathbb{E} \left[\exp \left(\int_t^T -\tilde{C}(\mathcal{Q}_s^{t,q})^2 + \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \bar{\lambda}_s^{k,\tau,i} ds \right) \right], \quad (4.19)$$

where

$$\mathcal{Q}_s^{t,q} = q + \int_t^s \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \Delta^{k,\tau} d(\bar{N}_u^{k,\tau,b} - \bar{N}_u^{k,\tau,a}),$$

where for any $(k, \tau) \in \mathcal{K} \times \mathcal{T}$ and $i = a$ or b , $\bar{N}^{k,\tau,i}$ is a point process with intensity $\bar{\lambda}_s^{k,\tau,i} := \hat{C}^{k,\tau} \mathbf{1}_{\{\phi(i)\mathcal{Q}_s^{t,q} > -\bar{q}\}}$, with $\hat{C}^{k,\tau}$ defined in Appendix 4.A.7.

The proof is in the same vein as [109, Proposition 4.1]. We now turn to numerical illustrations of our make take fees policy.

4.3.4 Numerical results

For numerical experiments, we consider three options which are characterized by their delta. We fix the following parameters: $A = 1.5s^{-1}$, $\sigma = C = 0.3s^{-1/2}$, $f^{k,\tau} = [0.5, 0.8, 0.8]$ the vector of fees, and $\delta_\infty^{k,\tau} = [2, 3, 3]$ the set of quotation thresholds. The first option is at the money, the second one is in the money and the third is out of the money, hence the following set of deltas $[0.5, 0.8, 0.2]$. Moreover, we take $\eta = 1$, $\gamma = 0.01$, $T = 100s$, $\bar{q} = 40$.

We analyze the impact of the penalty ω and the weight associated to each options $c^{k,\tau}$ in the value function of the exchange. In Figure 4.5, we display the average bid-ask spread at initial time on each option for $\omega = 0$, and $c^{k,\tau}$ being equal either to 0 or 0.1. We see that a higher $c^{k,\tau}$ leads to a decrease of the spread for the option $C^{k,\tau}$. This result is in line with the form of the incentives in Theorem 4.2. Indeed $Z^{*k,\tau,i}$ is an increasing function of $c^{k,\tau}$ and $\hat{\delta}^{k,\tau,i}$ is a decreasing function of $Z^{*k,\tau,i}$. Thus, increasing the interest of the principal for the option $C^{k,\tau}$ leads to a decrease of the spread proposed by the market maker on this option. This shows that the exchange has a direct control on each option he is interested in.

In Figure 4.6, we focus on the role of ω , equal to 0.1 on the spreads proposed by the market maker. As expected, a non-vanishing value of ω leads to a decrease of the spread for all the quoted options. This agrees with Theorem 4.2, where we see that the incentives are an increasing function of $\omega \in [0, 1)$. Thus, the exchange can influence the whole set of spreads proposed on the quoted options.




Figure 4.5: Evolution of the spread (in ticks) at initial time with respect to inventory, $\omega = 0$.




Figure 4.6: Evolution of the spread (in ticks) at initial time with respect to inventory, $\omega = 0.1$.

We conclude by showing in Figure 4.7 the behavior of the average spread with a higher ω , equal to 0.2. We obtain similar effects as in Figure 4.6, namely a decrease of the spread on all quoted options for a higher ω .




Figure 4.7: Evolution of the spread (in ticks) at initial time with respect to inventory, $\omega = 0.2$.

4.3.5 Conclusion

This work is, to our knowledge, the first to address the problem of designing a derivatives exchange, based solely on market data. In the first part, a simple market driven methodology enables us to choose which options the exchange should select to attract market takers. In the second part, we provide a make take fees policy between the exchange and the market maker which ensures a high quality of liquidity for the listed options.

4.A Appendix

4.A.1 Proof of the convergence of Lloyd's algorithm

According to Paragraph 5.2 in Graf and Luschgy [132], the set $(K_i)_{1 \leq i \leq n}$ is a solution of (4.2) if and only if for any i , A_i has positive Lebesgue measure and

$$\int_{A_i} |K_i - x|^{p-1} \text{sign}(x - K_i) \mathbb{P}^{mkt}(dx) = 0,$$

where sgin is the sign function. This is equivalent to

$$K_i = \frac{\int_{A_i} |K_i - x|^{p-2} x \mathbb{P}^{mkt}(dx)}{\int_{A_i} |K_i - x|^{p-2} \mathbb{P}^{mkt}(dx)} = \frac{\mathbb{E}^{mkt}[|K_i - K|^{p-2} K \mathbf{1}_{K \in A_i}]}{\mathbb{E}^{mkt}[|K_i - K|^{p-2} \mathbf{1}_{K \in A_i}]}.$$

Thus $(K_i)_{1 \leq i \leq N}$ is the solution of (4.2) if and only if it is a fixed point of Lloyd's algorithm.

We now give proofs and technical results for Section 4.3. They are mostly inspired by El Euch, Mastrolia, Rosenbaum, and Touzi [109]. However, for sake of completeness, we provide rigorous derivations.

4.A.2 Stochastic basis

4.A.2.1 Canonical process

In this section, we give an accurate definition of the probability space defined in Section 4.3.1. We consider a final horizon time $T > 0$ and the space $\Omega =: \Omega_c \times \Omega_d^{2 \times \#\mathcal{T} \times \#\mathcal{K}}$, with Ω_c the set of continuous functions from $[0, T]$ into \mathbb{R} and Ω_d the set of piecewise constant càdlàg functions from $[0, T]$ into \mathbb{N} . We consider Ω as a subspace of the Skorokhod space $\mathcal{D}([0, T], \mathbb{R}^{2 \times \#\mathcal{T} \times \#\mathcal{K} + 1})$ of càdlàg functions from $[0, T]$ into $\mathbb{R}^{2 \times \#\mathcal{T} \times \#\mathcal{K} + 1}$ and \mathcal{F} the trace Borel σ -algebra on Ω , where the topology is the one associated to the usual Skorokhod distance on $\mathcal{D}([0, T], \mathbb{R}^{2 \times \#\mathcal{T} \times \#\mathcal{K} + 1})$.

We define $(\mathcal{X}_t)_{t \in [0, T]} := (W_t, (N_t^{k, \tau, i})_{i=a, b; k \in \mathcal{K}; \tau \in \mathcal{T}})$ as the canonical process on Ω , that is for any $\omega = (w, n^{k, \tau, i}) \in \Omega$

$$W_t(\omega) = w(t), \quad N_t^{k, \tau, i}(\omega) = n^{k, \tau, i}(t).$$

4.A.2.2 Probability measure

We now properly define \mathbb{P}^0 and the associated change of measure. We set the probability \mathbb{P}^0 on (Ω, \mathcal{F}) such that under \mathbb{P}^0 , W , $N^{k, \tau, i}$ are independent, W is a one-dimensional Brownian motion and the $N^{k, \tau, i}$, $k \in \mathcal{K}$, $\tau \in \mathcal{T}$, $i \in \{a, b\}$ are Poisson processes with intensity $\lambda^{k, \tau, i}(0)$.⁷ Finally, we endow the space (Ω, \mathcal{F}) with the (\mathbb{P}^0 -completed) canonical filtration $\mathbb{F} := (\mathcal{F}_t)_{t \in [0, T]}$ generated by $(\mathcal{X}_t)_{t \in [0, T]}$.

By (4.5), the control process must be predictable and uniformly bounded. The last assumption is required to define the associated probability measure. So for $\delta \in \mathcal{A}$ we introduce the corresponding probability measure \mathbb{P}^δ under which $S_t = S_0 + \sigma W_t$ follows (4.3) and for $k \in \mathcal{K}$, $\tau \in \mathcal{T}$, $i \in \{a, b\}$ the

$$N_t^{\delta, k, \tau, i} := N_t^{k, \tau, i} - \int_0^t \lambda^{k, \tau}(\delta_r^{k, \tau, i}) \mathbf{1}_{\{\phi(i) Q_{r-} > -\bar{q}\}} dr$$

are martingales. This probability measure is defined by the corresponding Doléans-Dade exponential:

$$L_t^\delta := \exp \left(\sum_{i \in \{a, b\}} \sum_{(k, \tau) \in \mathcal{K} \times \mathcal{T}} \int_0^t \mathbf{1}_{\{\phi(i) Q_{r-} > -\bar{q}\}} \left(\log \left(\frac{\lambda^{k, \tau}(\delta_r^{k, \tau, i})}{A} \right) dN_r^{k, \tau, i} - \left(\lambda^{k, \tau}(\delta_r^{k, \tau, i}) - A \right) dr \right) \right),$$

⁷In other words, \mathbb{P}^0 is simply the product measure of the Wiener measure on Ω_c and the unique measure on $\Omega_d^{2 \times \#\mathcal{T} \times \#\mathcal{K}}$ that makes the canonical process an homogeneous Poisson process with the prescribed intensity.

which is a true martingale by the uniform boundedness of $\delta_t^{k,\tau,i}$.⁸ We can therefore define the Girsanov change of measure $\frac{d\mathbb{P}^\delta}{d\mathbb{P}^0} |_{\mathcal{F}_t} = L_t^\delta$, for all $t \in [0, T]$. In particular, all the probability measures \mathbb{P}^δ indexed by $\delta \in \mathcal{A}$ are equivalent. We shall write \mathbb{E}_t^δ for the conditional expectation with respect to \mathcal{F}_t under the probability measure \mathbb{P}^δ .

4.A.3 Well-posedness of the optimization problems

We give in this section the necessary integrability conditions ensuring that both exchange and market maker's problems are well defined. We consider the following assumptions:

$$\sup_{\delta \in \mathcal{A}} \mathbb{E}^\delta [\exp(-\gamma' \xi)] < +\infty, \text{ for some } \gamma' > \gamma, \quad \sup_{\delta \in \mathcal{A}} \mathbb{E}^\delta [\exp(\eta' \xi)] < +\infty, \text{ for some } \eta' > \eta. \quad (4.20)$$

Moreover, the next technical assumption is required in order to derive the best response of the market maker in Theorem 4.1:

$$\sup_{\delta \in \mathcal{A}} \mathbb{E}^\delta \left[\sup_{t \in [0, T]} \exp\left(-\gamma' Y_t^{0, Z}\right) \right] < +\infty, \text{ for some } \gamma' > \gamma. \quad (4.21)$$

Finally, we define \mathcal{Z} as the set of predictable processes $(Z_t)_{t \in [0, T]}$ such that Conditions (4.20) and (4.21) are satisfied. This is the set of admissible contract components of the exchange.

4.A.4 Dynamic programming principle

In the spirit of El Euch, Mastrolia, Rosenbaum, and Touzi [109], we provide a proof of a dynamic programming principle for the market maker's problem. Note that a same type of dynamic programming principle exists for the exchange's problem.

For any \mathbb{F} -predictable stopping time $\tau \in [t, T]$ and $\mu \in \mathcal{A}_\tau$, we define

$$J_T(\tau, \mu) = \mathbb{E}_\tau^\mu \left[- \exp \left(- \gamma \left(\xi + \sum_{i \in \{a, b\}} \sum_{(k, \tau) \in \mathcal{K} \times \mathcal{T}} \int_\tau^T \mu_u^{k, \tau, i} dN_u^{k, \tau, i} + Q_u^{k, \tau} dC_u^{k, \tau} \right) \right) \right]$$

where \mathcal{A}_τ denotes the restriction of \mathcal{A} to controls on $[\tau, T]$. We also define the set $\mathcal{J}_{\tau, T} = (J_T(\tau, \mu))_{\mu \in \mathcal{A}_\tau}$. The continuation utility of the market maker is defined for any \mathbb{F} -predictable stopping time τ by

$$V_\tau = \text{ess sup}_{\mu \in \mathcal{A}_\tau} J_T(\tau, \mu).$$

We first prove the following technical lemma.

Lemma 4.3. *Let τ be a \mathbb{F} -predictable stopping time with values in $[t, T]$. Then there exists an increasing sequence $(\mu^n)_{n \in \mathbb{N}}$ in \mathcal{A}_τ such that $V_\tau = \lim_{n \rightarrow +\infty} J_T(\tau, \mu^n)$.*

Proof. For $\mu, \mu' \in \mathcal{A}_\tau$ we define $\hat{\mu} := \mu \mathbf{1}_{\{J_T(\tau, \mu) \geq J_T(\tau, \mu')\}} + \mu' \mathbf{1}_{\{J_T(\tau, \mu) \leq J_T(\tau, \mu')\}}$. We have $\hat{\mu} \in \mathcal{A}_\tau$ and by definition of $\hat{\mu}$, $J_T(\tau, \hat{\mu}) \geq \max(J_T(\tau, \mu), J_T(\tau, \mu'))$. Thus $\mathcal{J}_{\tau, T}$ is increasing, and we obtain the same result as in El Euch, Mastrolia, Rosenbaum, and Touzi [109]. The conclusion follows. \square

We set

$$\mathcal{D}_{t, T}(\delta) := \exp \left(- \gamma \left(\sum_{i \in \{a, b\}} \sum_{(k, \tau) \in \mathcal{K} \times \mathcal{T}} \int_t^T \delta_u^{k, \tau, i} dN_u^{k, \tau, i} + Q_u^{k, \tau} dC_u^{k, \tau} \right) \right).$$

Given Lemma 4.3, we can now prove the dynamic programming principle associated to (4.8).

⁸The associated Novikov criterion is given in Sokol [253].

Lemma 4.4. *Let $t \in [0, T]$ and τ be an \mathbb{F} -predictable stopping time with values in $[t, T]$. Then*

$$V_t = \operatorname{ess\,sup}_{\delta \in \mathcal{A}} \mathbb{E}_t^\delta [-\mathcal{D}_{t,\tau}(\delta) V_\tau].$$

Proof. Let $t \in [0, T]$ and τ be a \mathbb{F} -predictable stopping time with values in $[t, T]$. First, by tower property, we have

$$V_t = \operatorname{ess\,sup}_{\delta \in \mathcal{A}} \mathbb{E}_t^\delta [-\mathcal{D}_{t,T}(\delta) \exp(-\gamma\xi)] = \operatorname{ess\,sup}_{\delta \in \mathcal{A}} \mathbb{E}_t^\delta [\mathcal{D}_{t,\tau}(\delta) \mathbb{E}_\tau^\delta [-\mathcal{D}_{\tau,T}(\delta) \exp(-\gamma\xi)]].$$

Then, Bayes rule yields

$$\begin{aligned} \mathbb{E}_\tau^\delta [-\mathcal{D}_{\tau,T}(\delta) \exp(-\gamma\xi)] &= \mathbb{E}_\tau^0 \left[-\frac{L_T^\delta}{L_\tau^\delta} \mathcal{D}_{\tau,T}(\delta) \exp(-\gamma\xi) \right] \leq \operatorname{ess\,sup}_{\delta \in \mathcal{A}} \mathbb{E}_\tau^\delta [\mathcal{D}_{\tau,T}(\delta) \exp(-\gamma\xi)] \\ &= V_\tau. \end{aligned}$$

Finally we obtain

$$V_t \leq \operatorname{ess\,sup}_{\mu \in \mathcal{A}} \mathbb{E}_t^\mu [V_\tau \mathcal{D}_{t,\tau}(\delta)].$$

We next prove the reverse inequality. Let $\delta \in \mathcal{A}$ and $\mu \in \mathcal{A}_\tau$. We define $(\delta \otimes_\tau \mu)_u = \delta_u \mathbf{1}_{\{0 \leq u \leq \tau\}} + \mu_u \mathbf{1}_{\{\tau \leq u \leq T\}}$. Then $\delta \otimes_\tau \mu \in \mathcal{A}$ and by tower property

$$V_t \geq \mathbb{E}_t^{\delta \otimes_\tau \mu} [-\mathcal{D}_{\tau,T}(\mu) \mathcal{D}_{t,\tau}(\delta) \exp(-\gamma\xi)] = \mathbb{E}_t^{\delta \otimes_\tau \mu} [\mathbb{E}_\tau^{\delta \otimes_\tau \mu} [-\mathcal{D}_{\tau,T}(\mu) \exp(-\gamma\xi)] \mathcal{D}_{t,\tau}(\delta)].$$

Using Bayes formula and noting that $\frac{L_T^{\delta \otimes_\tau \mu}}{L_\tau^{\delta \otimes_\tau \mu}} = \frac{L_T^\mu}{L_\tau^\mu}$, we have

$$\mathbb{E}_\tau^{\delta \otimes_\tau \mu} [-\mathcal{D}_{\tau,T}(\mu) \exp(-\gamma\xi)] = \mathbb{E}_\tau^0 \left[-\frac{L_T^\mu}{L_\tau^\mu} \mathcal{D}_{\tau,T}(\mu) \exp(-\gamma\xi) \right] = J_T(\tau, \mu).$$

This implies

$$V_t \geq \mathbb{E}_t^{\delta \otimes_\tau \mu} [\mathcal{D}_{t,\tau}(\delta) J_T(\tau, \mu)].$$

We can therefore use again Bayes rule and the fact that $\frac{L_T^{\delta \otimes_\tau \mu}}{L_t^{\delta \otimes_\tau \mu}} = \frac{L_T^\delta}{L_t^\delta}$ to obtain

$$\begin{aligned} V_t &\geq \mathbb{E}_t^0 \left[\frac{L_T^{\delta \otimes_\tau \mu}}{L_t^{\delta \otimes_\tau \mu}} \mathcal{D}_{t,\tau}(\delta) J_T(\tau, \mu) \right] = \mathbb{E}_t^0 \left[\mathbb{E}_\tau^0 \left[\frac{L_T^{\delta \otimes_\tau \mu}}{L_\tau^{\delta \otimes_\tau \mu}} \frac{L_\tau^{\delta \otimes_\tau \mu}}{L_t^{\delta \otimes_\tau \mu}} \mathcal{D}_{t,\tau}(\delta) J_T(\tau, \mu) \right] \right] \\ &= \mathbb{E}_t^0 \left[\mathbb{E}_\tau^0 \left[\frac{L_T^{\delta \otimes_\tau \mu}}{L_\tau^{\delta \otimes_\tau \mu}} \right] \frac{L_\tau^{\delta \otimes_\tau \mu}}{L_t^{\delta \otimes_\tau \mu}} \mathcal{D}_{t,\tau}(\delta) J_T(\tau, \mu) \right] \\ &= \mathbb{E}_t^0 \left[\frac{L_\tau^{\delta \otimes_\tau \mu}}{L_t^{\delta \otimes_\tau \mu}} \mathcal{D}_{t,\tau}(\delta) J_T(\tau, \mu) \right] \\ &= \mathbb{E}_t^\delta [\mathcal{D}_{t,\tau}(\delta) J_T(\tau, \mu)]. \end{aligned}$$

Since the previous inequality holds for any $\mu \in \mathcal{A}_\tau$, we deduce from monotone convergence theorem together with Lemma 4.3 that there exists a sequence $(\mu^n)_{n \in \mathbb{N}}$ of controls in \mathcal{A}_τ such that

$$V_t \geq \lim_{n \rightarrow +\infty} \mathbb{E}_t^\delta [\mathcal{D}_{t,\tau}(\delta) J_T(\tau, \mu^n)] = \mathbb{E}_t^\delta [\mathcal{D}_{t,\tau}(\delta) \lim_{n \rightarrow +\infty} J_T(\tau, \mu^n)] = \mathbb{E}_t^\delta [\mathcal{D}_{t,\tau}(\delta) V_\tau].$$

This concludes the proof. \square

The following technical lemma shows some integrability conditions on $\mathcal{D}_{t,T}$ and V_t .

Lemma 4.5. *For all $\delta \in \mathcal{A}$ and for a specific $\epsilon > 0$, we have*

$$\mathbb{E}^\delta \left[\sup_{t \in [0, T]} |V_t|^{1+\epsilon} \right] < +\infty, \quad \mathbb{E}^\delta \left[\sup_{(s,t) \in [0, T]^2} (D_{s,t}(\delta))^{1+\epsilon} \right] < +\infty.$$

Proof. We have for $t \in [0, T]$,

$$\begin{aligned} -V_t &= \operatorname{ess\,inf}_{\delta \in \mathcal{A}} \mathbb{E}_t^\delta \left[\exp \left(-\gamma \left(\xi + \sum_{i=a,b} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \int_\tau^T \mu_u^{k,\tau,i} dN_u^{k,\tau,i} + Q_u^{k,\tau} dC_u^{k,\tau} \right) \right) \right] \\ &\leq \mathbb{E}_t^\delta \left[e^{\gamma (\delta_\infty \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} N_T^{k,\tau,i} - \int_t^T Q_u dS_u)} \exp(-\gamma \xi) \right], \end{aligned}$$

where we used the fact that $N_T^{i,j} - N_t^{i,j} \leq N_T^{k,\tau,i}$, $Q_t \in [-\bar{q}, \bar{q}]$ and the control processes are uniformly bounded by δ_∞ . Moreover, for all $L > 0$,

$$\mathbb{E}_t^\delta \left[e^{-L \int_t^T Q_u dS_u} \right] \leq e^{\frac{L^2 \bar{q}^2 \sigma^2 T}{2}}.$$

Thus, using Holder's inequality, we have

$$\begin{aligned} -V_t &\leq \mathbb{E}_t^\delta \left[e^{\epsilon \gamma (\delta_\infty \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} N_T^{k,\tau,i} - \xi)} \right]^{\frac{1}{\epsilon}} \mathbb{E}_t^\delta \left[e^{-\gamma(1+\epsilon) \int_t^T Q_u dS_u} \right]^{\frac{1}{1+\epsilon}} \\ &\leq \mathbb{E}_t^\delta \left[e^{\epsilon \gamma (\delta_\infty \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} N_T^{k,\tau,i} - \xi)} \right]^{\frac{1}{\epsilon}} e^{\frac{(1+\epsilon)\gamma^2 \bar{q}^2 \sigma^2 T}{2}}. \end{aligned}$$

Then, we have

$$\mathbb{E}^\delta \left[\sup_{t \in [0, T]} (-V_t)^{1+\epsilon} \right] \leq e^{\frac{(1+\epsilon)^2 \gamma^2 \bar{q}^2 \sigma^2 T}{2}} \mathbb{E}^\delta \left[\sup_{t \in [0, T]} \mathbb{E}_t^\delta \left[e^{\epsilon \gamma (\delta_\infty \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} N_T^{k,\tau,i} - \xi)} \right]^{\frac{1+\epsilon}{\epsilon}} \right].$$

The term inside the conditional expectation is integrable⁹ and independent from $t \in [0, T]$ thus by Doob's inequality, we have

$$\mathbb{E}^\delta \left[\sup_{t \in [0, T]} (-V_t)^{1+\epsilon} \right] \leq C e^{\frac{(1+\epsilon)^2 \gamma^2 \bar{q}^2 \sigma^2 T}{2}} \mathbb{E}^\delta \left[e^{\gamma' (\delta_\infty \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} N_T^{k,\tau,i} - \xi)} \right],$$

where $C > 0$, $\gamma' = \gamma(1 + \epsilon)$. Thanks to Hölder's inequality, together with the boundedness of the intensities of the point processes $N^{k,\tau,i}$ and Condition (4.20), the right-hand side is bounded from above by a term independent of $t \in [0, T]$. Using the same arguments, we have

$$\mathbb{E}^\delta \left[\sup_{(s,t) \in [0, T]^2} (D_{s,t}(\delta))^{1+\epsilon} \right] \leq C' \mathbb{E}^\delta \left[e^{\gamma' (\delta_\infty \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} N_T^{k,\tau,i} + \frac{(1+\epsilon)\gamma^2 \bar{q}^2 \sigma^2 T}{2})} \right] < +\infty,$$

where $C' > 0$, using boundedness of the intensities of the point processes. The conclusion follows. \square

4.A.5 Proof of Lemma 4.1

We divide the proof into six steps.

Step 1: Derivation of the martingale representation.

For $\delta \in \mathcal{A}$, it follows from the dynamic programming principle of Lemma 4.4 that the process

$$U_t^\delta = V_t \mathcal{D}_{0,t}(\delta)$$

⁹Take $\epsilon > 1$ together with Condition (4.20) for example.

is of class (D) and defines a \mathbb{P}^δ -supermartingale for any $\delta \in \mathcal{A}$. By standard analysis, we may then consider it in its càdlàg version (by taking right limits along rationals). By the Doob-Meyer decomposition, we can write $U_t^\delta = M_t^\delta - A_t^\delta$ where M^δ is a \mathbb{P}^δ -martingale and $A_t^\delta = A_t^{\delta,c} + A_t^{\delta,d}$ is an integrable non-decreasing predictable process such that $A_0^{\delta,c} = A_0^{\delta,d} = 0$ with pathwise continuous component $A^{\delta,c}$ and with $A^{\delta,d}$ a piecewise constant predictable process.

From the martingale representation theorem under \mathbb{P}^δ , see Appendix A.1 in El Euch, Mastrolia, Rosenbaum, and Touzi [109], there exists $\tilde{Z}^\delta = (\tilde{Z}^{\delta,S}, \tilde{Z}^{\delta,k,\tau,i})_{k \in \mathcal{K}, \tau \in \mathcal{T}, i \in \{a,b\}}$ predictable, such that

$$M_t^\delta = V_0 + \int_0^t \tilde{Z}_r^{\delta,S} dS_r + \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \int_0^t \tilde{Z}_r^{\delta,k,\tau,i} dN_r^{\delta,k,\tau,i}.$$

Step 2: Boundedness of the value function.

We show that V is a negative process. In fact, thanks to the uniform boundedness of $\delta \in \mathcal{A}$, we have that

$$\frac{L_T^\delta}{L_t^\delta} \geq \alpha_{t,T} = \exp\left(-\sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \frac{k}{\sigma} N_T^{k,\tau,i} - 2 \times \#\mathcal{T} \times \#\mathcal{K} A e^{-\frac{kc_\infty}{\sigma}} (e^{\frac{k}{\sigma}} + 1)(T-t)\right),$$

where $c_\infty := \max_{k,\tau} c^{k,\tau}$. Therefore

$$V_t \leq \mathbb{E}_t^0 \left[-\alpha_{t,T} \exp\left(-\gamma(\delta_\infty \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} N_T^{k,\tau,i} + \int_t^T Q_u^{k,\tau} dC_u^{k,\tau})\right) e^{-\gamma\xi} \right] < 0.$$

Step 3: Identification of the coefficients (1/2).

Let Y be the process defined for any $t \in [0, T]$ by $V_t = -e^{-\gamma Y_t}$. As $A^{\delta,d}$ is a predictable point process and the jumps of $N^{k,\tau,i}, i \in \{a,b\}$ are totally inaccessible stopping times under \mathbb{P}^0 , we have $\langle N^{k,\tau,i}, A^{\delta,d} \rangle_t = 0$ a.s. Using Itô's formula, we obtain that

$$Y_T = \xi, \text{ and } dY_t = \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} Z_t^{k,\tau,i} dN_t^{k,\tau,i} + Z_t^S dS_t - dI_t - d\tilde{A}_t^d,$$

with

$$\begin{aligned} Z_t^{k,\tau,a} &= -\frac{1}{\gamma} \log\left(1 + \frac{\tilde{Z}_t^{\delta,k,\tau,a}}{U_{t^-}^\delta}\right) - \delta_t^{k,\tau,a}, \quad Z_t^{k,\tau,b} = -\frac{1}{\gamma} \log\left(1 + \frac{\tilde{Z}_t^{\delta,k,\tau,b}}{U_{t^-}^\delta}\right) - \delta_t^{k,\tau,b}, \\ Z_t^S &= -\frac{\tilde{Z}_t^{\delta,S}}{\gamma U_{t^-}^\delta} - \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} Q_{t^-}^{k,\tau} \Delta^{k,\tau}, \quad I_t = \int_0^t \left(\bar{h}(\delta_r, Z_r, Q_r) dr - \frac{1}{\gamma U_r^\delta} dA_r^{\delta,c}\right), \\ \bar{h}(\delta, Z_t, Q_t) &= h(\delta, Z_t, Q_t) - \frac{1}{2} \gamma \sigma^2 (Z_t^S)^2, \quad \tilde{A}_t^d = \frac{1}{\gamma} \sum_{s \leq t} \log\left(1 - \frac{\Delta A_s^{\delta,d}}{U_{s^-}^\delta}\right). \end{aligned}$$

In particular, the last relation between \tilde{A}^d and $A^{\delta,d}$ shows that $\Delta a_t \geq 0$ is independent of $\delta \in \mathcal{A}$, with $a_t = -\frac{A_t^{\delta,d}}{U_{t^-}^\delta}$ and abusing notations slightly, $\Delta a_t = -\frac{\Delta A_t^{\delta,d}}{U_{t^-}^\delta}$. In order to complete the proof, we argue in the subsequent steps that $Z \in \mathcal{Z}$ and that, for $t \in [0, T]$, $A_t^{\delta,d} = -\sum_{s \leq t} U_{s^-}^\delta \Delta a_s = 0$ so that $\tilde{A}_t^d = 0$ and $I_t = \int_0^t \bar{H}(Z_r, Q_r) dr$, where

$$\bar{H}(Z_t, Q_t) = H(Z_t, Q_t) - \frac{1}{2} \gamma \sigma^2 (Z_t^S)^2.$$

Step 4: Identification of the coefficients (2/2).

Since $V_T = -1$, we get that

$$\begin{aligned} 0 &= \sup_{\delta \in \mathcal{A}} \mathbb{E}^\delta [U_T^\delta] - V_0 = \sup_{\delta \in \mathcal{A}} \mathbb{E}^\delta [U_T^\delta - M_T^\delta] \\ &= \gamma \sup_{\delta \in \mathcal{A}} \mathbb{E}^0 \left[L_T^\delta \int_0^T U_{r-}^\delta \left(dI_r - \bar{h}(\delta, Z_r, Q_r) dr + \frac{da_r}{\gamma} \right) \right]. \end{aligned}$$

Moreover, the controls being uniformly bounded, we have

$$U_t^\delta \leq -\beta_t = V_t \exp \left(-\gamma \left(\delta_\infty \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} N_T^{k,\tau,i} + \int_0^t Q_u^{k,\tau} dC_u^{k,\tau} \right) \right) < 0.$$

Then, using $A^{\delta,d} \geq 0, U^\delta \leq 0$ and $dI_t - \bar{h}(\delta, Z_t, Q_t) dt \geq 0$, we obtain

$$\begin{aligned} 0 &\leq \sup_{\delta \in \mathcal{A}} \mathbb{E}^0 \left[\alpha_{0,T} \int_0^T -\beta_{r-} \left(dI_r - \bar{h}(\delta, Z_r, Q_r) dr + \frac{da_r}{\gamma} \right) \right] \\ &= -\mathbb{E}^0 \left[\alpha_{0,T} \int_0^T \beta_{r-} \left(dI_r - \bar{H}(Z_r, Q_r) dr + \frac{da_r}{\gamma} \right) \right]. \end{aligned}$$

The quantities $\alpha_{0,T} \int_0^T \beta_{r-} (dI_r - \bar{H}(Z_r, Q_r)) dr$ and $\alpha_{0,T} \int_0^T \beta_{r-} \frac{da_r}{\gamma}$ being non-negative random variables, the result follows.

Step 5: Admissibility of the process Z .

Using Lemma 4.5, we have that

$$\sup_{\delta \in \mathcal{A}} \mathbb{E}^\delta \left[\sup_{t \in [0, T]} \exp(-\gamma(p+1)Y_t) \right] < +\infty.$$

The conclusion follows using the fact that

$$\exp(-\gamma Y_t) = U_t^\delta \exp \left(\gamma \left(\sum_{i=a,b} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \int_0^t \delta_u^{k,\tau,i} dN_u^{k,\tau,i} + Q_u^{k,\tau} dC_u^{k,\tau} \right) \right).$$

Step 6: Uniqueness of the representation.

Let $(Y_0, Z), (Y'_0, Z') \in \mathbb{R} \times \mathcal{Z}$ be such that $\xi = Y_T^{Y_0, Z} = Y_T^{Y'_0, Z'}$. By following the lines of the verification argument in the proof of Theorem 4.1, we obtain the equality $Y_t^{Y_0, Z} = Y_t^{Y'_0, Z'}$ using the fact that the value of the continuation utility of the market maker satisfies

$$-e^{-\gamma Y_t^{Y_0, Z}} = -e^{-\gamma Y_t^{Y'_0, Z'}} = \operatorname{ess\,sup}_{\delta \in \mathcal{A}} \mathbb{E}_t^\delta \left[-e^{-\gamma(PL_T^\delta - PL_t^\delta + \xi)} \right].$$

This in turn implies that for $t \in [0, T]$ $Z_t^{k,\tau,i} dN_t^{k,\tau,i} = Z_t'^{k,\tau,i} dN_t^{k,\tau,i}$ and $Z_t^S \sigma^2 dt = Z_t'^S \sigma^2 dt = d(Y, S)_t$. Consequently, $(Y_0, Z) = (Y'_0, Z')$.

4.A.6 Proof of Theorem 4.1

Let $\xi = Y_T^{Y_0, Z}$ with $(Y_0, Z) \in \mathbb{R} \times \mathcal{Z}$. We first prove that for an arbitrary set of controls $\delta \in \mathcal{A}$, we have $J_{\text{MM}}(\delta, \xi) \leq -e^{-\gamma Y_0}$, where $J_{\text{MM}}(\delta, \xi)$ is such that $V_{\text{MM}}(\xi) = \sup_{\delta \in \mathcal{A}} J_{\text{MM}}(\delta, \xi)$. Then, we will see that this inequality is in fact an equality when the corresponding Hamiltonian $h(\delta, z, q)$ is maximized. Denote

$$\bar{Y}_t := Y_t^{Y_0, Z} + \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \int_0^t \delta_u^{k,\tau,i} dN_u^{k,\tau,i} + Q_u^{k,\tau} dC_u^{k,\tau},$$

with $t \in [0, T]$. A direct application of Itô's formula leads to

$$de^{-\gamma\bar{Y}_t} = \gamma e^{-\gamma\bar{Y}_t} \left(- \left(\sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} Q_t^{k,\tau} \Delta^{k,\tau} + Z_t^S \right) dS_t + (H(Z_t, \mathcal{Q}_t) - h(\delta, Z_t, \mathcal{Q}_t)) dt \right. \\ \left. - \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \gamma^{-1} \left(1 - \exp \left(-\gamma (Z_t^{k,\tau,i} + \delta_t^{k,\tau,i}) \right) \right) dN_t^{\delta,k,\tau,i} \right).$$

Thus, $e^{-\gamma\bar{Y}_\cdot}$ is a \mathbb{P}^δ -local submartingale. Thanks to Lemma 4.5, $(e^{-\gamma\bar{Y}_t})_{t \in [0, T]}$ is of class (D) hence is a true submartingale. Doob-Meyer decomposition theorem gives us that

$$\int_0^\cdot \gamma e^{-\gamma\bar{Y}_t} \left(- \left(\sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} Q_t^{k,\tau} \Delta^{k,\tau} + Z_t^S \right) dS_t - \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \gamma^{-1} \left(1 - \exp \left(-\gamma (Z_t^{k,\tau,i} + \delta_t^{k,\tau,i}) \right) \right) dN_t^{\delta,k,\tau,i} \right),$$

is a true martingale. This implies that

$$J_{\text{MM}}(\delta, \xi) = \mathbb{E}^\delta \left[-e^{-\gamma\bar{Y}_T} \right] \\ = -e^{-\gamma Y_0} - \mathbb{E}^\delta \left[\int_0^T \gamma e^{-\gamma\bar{Y}_t} (H(Z_t, \mathcal{Q}_t) - h(\delta, Z_t, \mathcal{Q}_t)) dt \right] \\ \leq -e^{-\gamma Y_0}.$$

In addition to this, the previous inequality becomes an equality if and only if δ is chosen as the maximizer of the Hamiltonian h thus leading to the optimal quotes provided in Theorem 4.1. So we deduce $J_{\text{MM}}(\delta, \xi) = -e^{-\gamma Y_0}$. Finally we have $V_{\text{MM}}(\xi) = -e^{-\gamma Y_0}$ with optimal response $(\hat{\delta}_t)_{t \in [0, T]}$.

4.A.7 Proof of Theorem 4.2

We define for any map $v : [0, T] \times \mathbb{Z}^{\#\mathcal{K} \times \#\mathcal{T}} \rightarrow (-\infty, 0)$, $x \in \mathbb{R}$, $(k, \tau) \in \mathcal{K} \times \mathcal{T}$ and $(t, q) \in [0, T] \times \mathbb{Z}^{\#\mathcal{K} \times \#\mathcal{T}}$

$$v(t, q \ominus_{K_i, T_j} x) := v(t, q^{K_1, T_1}, q^{K_1, T_2}, \dots, q^{K_i, T_{j+1}}, q^{K_i, T_j} - x, q^{K_i, T_{j+2}}, \dots, q^{K_n, T_m}).$$

The Hamilton-Jacobi-Bellman equation of the stochastic control problem (4.14) is given by

$$0 = \partial_t v(t, q) + \mathcal{H}_E(t, q, v(t, \cdot)), \quad v(T, q) = -1, \quad (4.22)$$

with

$$\mathcal{H}_E(t, q, v(t, \cdot)) = \sup_{z \in \mathcal{Z}} h_E(t, q, s, z, v(t, \cdot)),$$

$$h_E(t, q, s, z, v(t, \cdot)) = v(t, q) \left(\frac{\eta}{2} \gamma \sigma^2 \left(\sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \Delta^{k,\tau} (z^{C^{k,\tau}} + q^{k,\tau}) \right)^2 + \frac{\eta^2}{2} \sigma^2 \left(\sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \Delta^{k,\tau} z^{C^{k,\tau}} \right)^2 \right) \\ + \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} h_{k,\tau}^i \left(t, z^{k,\tau,i}, v(t, q), v(t, q \ominus_{k,\tau} \phi(i)) \right) \mathbf{1}_{\phi(i) \mathcal{Q} > -\bar{q}},$$

and

$$h_{k,\tau}^i(t, z, y, y') = (y' x_1^{k,\tau} e^{az} - y x_2 e^{bz}) O_{k,\tau},$$

where

$$x_1^{k,\tau} = e^{-\eta(c^{k,\tau} + \omega(\delta_\infty^{k,\tau} - \gamma^{-1} \log(1 + \frac{\sigma\gamma}{C}))}, \quad x_2 = \left(1 + \eta \frac{1 - (1 + \frac{\sigma\gamma}{C})^{-1}}{\gamma} \right), \quad O_{k,\tau} = \left(1 + \frac{\sigma\gamma}{C} \right)^{-\frac{C}{\gamma\sigma}} e^{-\frac{C}{\sigma} f^{k,\tau}},$$

and

$$a = \eta(1 - \omega) + \frac{C}{\sigma}, \quad b = \frac{C}{\sigma}.$$

Tedious but straightforward computations lead to the following optimizers:

$$\begin{aligned} z^{*k,\tau,i} &:= \frac{1}{a-b} \log \left(\frac{bx_2 v(t,q)}{ax_1^{k,\tau} v(t,q \ominus_{k,\tau} \phi(i))} \right), \\ z^{*C^{k,\tau}} &:= -\frac{\gamma}{\gamma + \eta} q^{k,\tau}. \end{aligned}$$

Note that from these computations, we get that this above optimization makes sense only if we assume that there exists δ_∞ large enough so that for $i = a$ or b , $k \in \mathcal{K}, \tau \in \mathcal{T}$ and any t, q :

$$\left| -z_t^{*k,\tau,i}(t,q) + \frac{1}{\gamma} \log \left(1 + \frac{\sigma\gamma}{C} \right) \right| < \delta_\infty. \quad (4.23)$$

We will check that we can make such choice at the end of the verification argument. Equation (4.22) is rewritten as

$$\begin{aligned} 0 &= \partial_t v(t,q) + v(t,q) \frac{\gamma\eta^2}{\gamma + \eta} \frac{\sigma^2}{2} \left(\sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \Delta^{k,\tau} q^{k,\tau} \right)^2 \\ &\quad - v(t,q) \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \tilde{C}^{k,\tau} \left(\frac{v(t,q)}{v(t,q \ominus_{k,\tau} \phi(i))} \right)^{\frac{C}{\sigma\eta(1-\omega)}} \mathbf{1}_{\phi(i)\mathcal{Q} > -\bar{q}}, \end{aligned} \quad (4.24)$$

where

$$\tilde{C}^{k,\tau} = x_2 \left(\frac{x_2}{x_1^{k,\tau}} \right)^{\frac{a}{a-b}} O_{k,\tau} \left(\left(\frac{b}{a} \right)^{\frac{b}{a-b}} - \left(\frac{b}{a} \right)^{\frac{a}{a-b}} \right) > 0.$$

We now make the ansatz $v(t,q) = u(t, \mathcal{Q})$. We derive the following PDE

$$0 = \partial_t u(t, \mathcal{Q}) + u(t, \mathcal{Q}) \frac{\gamma\eta^2}{\gamma + \eta} \frac{\sigma^2}{2} \mathcal{Q}^2 - u(t, \mathcal{Q}) \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \tilde{C}^{k,\tau} \left(\frac{u(t, \mathcal{Q})}{u(t, \mathcal{Q} - \Delta^{k,\tau} \phi(i))} \right)^{\frac{C}{\sigma\eta(1-\omega)}} \mathbf{1}_{\phi(i)\mathcal{Q} > -\bar{q}}, \quad (4.25)$$

with terminal condition $u(T, \mathcal{Q}) = -1$.

Using the classical change of variable $\tilde{u} := (-u)^{-\frac{C}{\sigma\eta(1-\omega)}}$, PDE (4.25) becomes

$$0 = \partial_t \tilde{u}(t, \mathcal{Q}) - \tilde{u}(t, \mathcal{Q}) \frac{C\gamma\eta}{\gamma + \eta} \frac{\sigma}{2(1-\omega)} \mathcal{Q}^2 + \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \hat{C}^{k,\tau} \tilde{u}(t, \mathcal{Q} - \Delta^{k,\tau} \phi(i)) \mathbf{1}_{\phi(i)\mathcal{Q} > -\bar{q}}, \quad (4.26)$$

where $\hat{C}^{k,\tau} := \tilde{C}^{k,\tau} \frac{C}{\sigma\eta(1-\omega)}$. Eventually Cauchy-Lipschitz theorem provides existence and uniqueness of a bounded solution to (4.26) and so to (4.24). For the verification argument, we first introduce a technical lemma.

Lemma 4.6. *Let $Z \in \mathcal{Z}, \xi = Y_T^{\hat{Y}_0, Z}$. We define*

$$K_t^Z := \exp \left(-\eta \left(\sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} c^{k,\tau} N_t^{k,\tau,i} - \int_0^t \omega (\Delta^i(Z_s^{k,\tau,i}) - \delta_\infty^{k,\tau}) dN_s^{k,\tau,i} - Y_t^{Y_0, Z} \right) \right), \quad t \in [0, T].$$

There exists $\epsilon > 0$ such that

$$\mathbb{E}^{\hat{\delta}(Z)} \left[\sup_{t \in [0, T]} |K_t^Z|^{1+\epsilon} \right] < +\infty,$$

where $\hat{\delta}(Z)$ is defined in Theorem 4.1.

The proof is borrowed from El Euch, Mastrolia, Rosenbaum, and Touzi [109]. We now verify that the unique solution v of Equation (4.22) coincides at any point $(0, Q_0)$ with the value v_0^E of the reduced problem (4.14). We also prove that in (4.14), the maximum is achieved for feedback controls issued from (4.15).

Using Itô's formula we get

$$\begin{aligned} d[v(t, Q_t)K_t^Z] &= K_t^Z \left(\left(h_E(t, Q_{t-}, S_t, Z_t, v(t, \cdot)) - \mathcal{H}_E(t, Q_{t-}, v(t, \cdot)) \right) dt \right. \\ &\quad + v(t, Q_{t-}) \eta \sum_{k=1}^{\mathcal{K}} \sum_{\tau=1}^{\mathcal{T}} Z_t^{C^{k,\tau}} dC_t^{k,\tau} \\ &\quad \left. + \sum_{i \in \{a,b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \left(e^{-\eta(c^{k,\tau} - Z_t^{k,\tau,i})} v(t, Q_{t-}^{k,\tau} - \phi(i)) - v(t, Q_{t-}) \right) dN_t^{\hat{\delta}(Z), k, \tau, i} \right). \end{aligned}$$

The process K^Z is of class (D) . Moreover v being uniformly bounded as a consequence of the Cauchy-Lipschitz theorem, the process $(v(t, Q_t)K_t^Z)_{t \in [0, T]}$ is a $\mathbb{P}^{\hat{\delta}(Z)}$ supermartingale of class (D) and the local martingale term in the above equation is a true martingale. Hence

$$v(0, Q_0) \geq \mathbb{E}^{\hat{\delta}(Z)}[v(T, Q_T)K_T^Z] = -\mathbb{E}^{\hat{\delta}(Z)}[K_T^Z]. \quad (4.27)$$

Since $Z \in \mathcal{Z}$ is arbitrary, we get

$$v(0, Q_0) \geq \sup_{Z \in \mathcal{Z}} -\mathbb{E}^{\hat{\delta}(Z)}[K_T^Z] = v_0^E.$$

The feedback form of Z , issued from (4.15), being bounded according to Equation (4.23), it is admissible. Considering the process Z^* , we get an equality instead of an inequality in the above equation. For consistency we now check that there does exist some constant δ_∞ such that (4.23) is satisfied. In the same vein as in Step 2 of the proof of Theorem 4.1, we can show that for any t and q , $v(t, q)$ is negative. Because of the compactness of the domain of v , the function is uniformly negative: we can find ε such that $v < -\varepsilon$ on $[0, T] \times \mathcal{D}$. Consequently $\log \left(\frac{v(t, q)}{v(t, q^{k,\tau} \ominus_{k,\tau} \phi(i))} \right)$ is uniformly bounded in i, k, τ, t and q . Thus we can always choose a δ_∞ satisfying (4.23).

4.A.8 Perfect and imperfect Delta-hedging

In the chapter, we assumed that the option market-maker does not hedge his positions with the underlying. We show in the present section how to relax this assumption. We first assume that the market-maker can ensure perfect Delta hedging under the constant Delta assumption. We define the Delta of the portfolio of options at time t as

$$\Delta_t := \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \Delta^{k,\tau} Q_t^{k,\tau}.$$

The cash process is now given by

$$d\mathcal{W}_t^\delta = \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \left(\delta_t^{k,\tau,a} dN_t^{k,\tau,a} + \delta_t^{k,\tau,b} dN_t^{k,\tau,b} - C_t^{k,\tau} dQ_t^{k,\tau} \right) + S_t d\Delta_t,$$

and the PnL of the market-maker at time t becomes

$$PL_t^\delta = \mathcal{W}_t^\delta - \Delta_t S_t + \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} Q_t^{k,\tau} C_t^{k,\tau}.$$

An application of Itô's formula gives

$$\begin{aligned}
dPL_t^\delta &= \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \left(\delta_t^{k,\tau,a} dN_t^{k,\tau,a} + \delta_t^{k,\tau,b} dN_t^{k,\tau,b} - C_t^{k,\tau} dQ_t^{k,\tau} \right) + S_t d\Delta_t - S_t d\Delta_t - \Delta_t dS_t \\
&\quad + \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \left(Q_t^{k,\tau} dC_t^{k,\tau} + C_t^{k,\tau} dQ_t^{k,\tau} \right) \\
&= \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \left(\delta_t^{k,\tau,a} dN_t^{k,\tau,a} + \delta_t^{k,\tau,b} dN_t^{k,\tau,b} \right) - \Delta_t dS_t + \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \left(Q_t^{k,\tau} dC_t^{k,\tau} \right) \\
&= \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \left(\delta_t^{k,\tau,a} dN_t^{k,\tau,a} + \delta_t^{k,\tau,b} dN_t^{k,\tau,b} \right).
\end{aligned}$$

Thus, the dynamics of the underlying asset disappears from the PnL process of the market-maker. This leads to the same form of contracts offered by the exchange, who simply sets $Z_t^{C^{k,\tau}} = 0$ for all $t \in [0, T]$.

Another possibility is to assume that the market-maker hedges his position on options by market-making activity on the stock. We define the market-maker's inventory on the stock process at time t as

$$Q_t^S = N_t^{S,b} - N_t^{S,a},$$

where $N_t^{S,a}, N_t^{S,b}$ are point processes of intensity

$$\lambda_t^{S,i} = \lambda^{S,i}(\delta_t^{S,i}) \mathbf{1}_{\{\phi(i)\bar{Q}_t > -\bar{q}\}}, \quad i \in \{a, b\},$$

where $\delta_t^{S,a}$ (resp. $\delta_t^{S,b}$) is the spread process of the market-maker on the ask (resp. bid) side of the stock,

$$\bar{Q}_t := Q_t + Q_t^S,$$

is the aggregated inventory on stocks and options, and

$$\lambda^S(\delta_t^{S,i}) = A \exp\left(-\frac{C}{\sigma}(\delta_t^{S,i} + f^S)\right),$$

with $f^S > 0$ is the fee fixed by the exchange for each market order on the underlying. Computations in the spirit of those of Lemma 4.1 show that every admissible contract offered to the market-maker is of the form

$$\begin{aligned}
\bar{Y}_T^{Y_0, Z} &:= \bar{Y}_0 + \int_0^T \left(\sum_{i \in \{a, b\}} \sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \bar{Z}_r^{k,\tau,i} dN_r^{k,\tau,i} + \bar{Z}_r^{S,i} dN_r^{S,i} + \bar{Z}_r^{C^{k,\tau}} dC_r^{k,\tau} \right) + \bar{Z}_r^S dS_r \\
&\quad + \left(\frac{1}{2} \gamma \sigma^2 \left(\sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \Delta^{k,\tau} (\bar{Z}_r^{C^{k,\tau}} + Q_r^{k,\tau}) + (\bar{Z}_r^S + Q_r^S) \right)^2 - \bar{H}(\bar{Z}_r, \bar{Q}_r) \right) dr,
\end{aligned} \tag{4.28}$$

where $\bar{Z} = (\bar{Z}^{C^{k,\tau}}, \bar{Z}^S, \bar{Z}^{k,\tau,i}, \bar{Z}^{S,i})$ are predictable processes controlled by the exchange and

$$\bar{H}(\bar{z}, \bar{Q}) := \sup_{\delta \in \mathbb{R}^{2 \times (\#\mathcal{K} \times \#\mathcal{T} + 1)}} h(\delta, \bar{z}, \bar{Q}),$$

with

$$\begin{aligned}
\bar{h}(\delta, \bar{z}, \bar{Q}) &:= \sum_{i \in \{a, b\}} \left(\sum_{(k,\tau) \in \mathcal{K} \times \mathcal{T}} \gamma^{-1} \left(1 - \exp\left(-\gamma(\bar{z}^{k,\tau,i} + \delta^{k,\tau,i})\right) \right) \lambda^{k,\tau}(\delta^{k,\tau,i}) \mathbf{1}_{\{\phi(i)\bar{Q} > -\bar{q}\}} \right. \\
&\quad \left. + \gamma^{-1} \left(1 - \exp\left(-\gamma(\bar{z}^{S,i} + \delta^{S,i})\right) \right) \lambda^S(\delta^{S,i}) \mathbf{1}_{\{\phi(i)\bar{Q} > -\bar{q}\}} \right).
\end{aligned}$$

Moreover, as we can aggregate the inventory of the market maker on the stock with the ones on the options, the problem of the exchange still boils down to the resolution of a two-dimensional HJB equation with different coefficients.

The two methods we presented in this Appendix show that we can incorporate easily Delta-hedging in the market-maker's behavior.

Chapter 5

Governmental incentives for green bonds investment

5.1 Introduction

Green bonds are fixed income products, issued by governments or companies to finance their debt. The only difference with the so-called conventional bonds is that they finance environmental or climate-related activities. Since its inception in 2007, the green bonds market has expanded rapidly to reach a total amount issued of \$100 billion in 2019. Corporate and finance companies issue more than 70% of the total amount of green bonds, whereas governments issue approximately 9% of this total, see, for example, the report of the Financial Stability Board [49] or the reports of OECD [225, 226]. The role of financial markets in promoting environmental policies via the green bonds is well documented in Park [229]. The characteristics of a bond to be defined as ‘green’ is given by the Green Bond Principles, which are ‘voluntary process guidelines that recommend transparency and disclosure, and promote integrity in the development of the Green Bond market by clarifying the approach for issuance of a Green Bond’, see the definition in the guidelines Association [25], published by the ICMA. These principles led the green bonds to become a standardized asset class, part of the traditional asset allocation. There is an important literature on the influence of green bonds on gas emissions and environmental ratings. In Flammer [118, 117], the author shows that the stock of a company responds positively to the announcement of green bond issues, and these issuances lead to an improvement of the environmental performance. The pricing and ownership of green bonds in the United States is studied in Baker, Bergstresser, Serafeim, and Wurgler [34], where the authors show in particular that green municipal bonds are issued at a premium to otherwise similar ordinary bonds. Similarly, the impact of corporate green bonds on the credit quality of the issuer and on the shareholders is well documented by Tang and Zhang [262]. In de Angelis, Tankov, and Zerbib [92], the authors show how green investments can help companies to reduce their greenhouse gas emissions by raising their cost of capital. In particular, they provide empirical evidence on the US markets that an increase of assets managed by green investors lead to a decrease of carbon emission by the companies.

The idea of financing renewable projects through green bonds is even more important since institutional investors, in particular pension funds and asset managers, have been considering the possibility of including sustainable environmental investments in their assets. As such, “sustainable investing” now accounts for more than one quarter of total assets under management (AUM) in the United States and more than half in Europe, see the report of the GSIA Alliance [6] for a detailed survey on the subject. The motivations of sustainable investing can be the search of higher alpha or lower risk (see Nilsson [223], Bauer and Smeets [37], Krüger [191]), or the will for a more socially responsible image (see Hong and Kacperczyk [160]). The two major practices in sustainable investing are exclusionary screening and environmental, social and governance (ESG) integration. Exclusionary screening involves

the exclusion of certain assets from the range of eligible investments on ethical grounds, such as the so-called sin stocks, while ESG integration involves under weighting assets with low ESG ratings and over weighting those with high ESG ratings. In Zerbib [275], the author builds a sustainable CAPM based on these two principles and shows how sustainable investing affects asset returns. Although the issuance of green corporate bonds has increased over the last years, the public sector accounts for two-thirds of the investments in sustainable energy infrastructure. This pleads in favour of a greater issuance of green bonds by public entities to finance their sustainable projects, which will be the focus of the present chapter.

However, there are still numerous issues related to the development of the green bond market: lack of green bond definition, framework, and transparency. In that regard, Zerbib [273, 274] investigates whether or not there is a yield premium for green bonds. The main result is that there exists a small negative premium: the yield of a green bond is lower than that of a conventional bond of comparable characteristics. This yield difference is mainly attributed to intangible asset creation, see for example Porter and Van der Linde [231], Ambec and Lanoie [19], or Brooks and Oikonomou [58]. The price difference between a green and a conventional bond is studied in Hachenberg and Schiereck [145], where the authors show that financial and corporate green bonds trade tighter than their conventional counterpart, and governmental bonds on the other hand trade marginally wider. Finally, Ekeland and Lefournier [105] relativize the use of green bonds to finance the ecological transition. As the green bond principles are by no means legally mandatory, and the investors are not necessarily motivated by the green transition, there are no intrinsic difference between a green bond and its conventional counterpart. An important aspect in order to avoid green-washing, that is when the investors use the funding obtained with the green bonds to finance non-sustainable projects, through green bonds is the issuer's reputation or green third-party verifications, as stated in Bachelet, Bechetti, and Manfredonia [28]. These studies show the several components which slow down the development of the green bonds market. It is therefore important to put in place practical solutions to overcome these constraints. Some mechanisms are already developed by the policy makers to facilitate the investment in this market.

Indeed, there are several types of incentives policy-makers can put in place to support green bond issuance, see Morel and Bordier [218], and Della Croce, Kaminker, and Stewart [94]: support for research and development (R&D), investment incentives (capital grants, loan guarantees and low-interest rate loans), policies which target the cost of investment in capital by hedging or mitigating risk, and tax incentives policies.¹ In particular, tax incentives are attractive from a cost-efficiency perspective, as they can provide a big boost to investment with a relatively low impact on public finances. In Agliardi and Agliardi [3], the authors show that governmental tax-based incentives play a significant role in scaling up the green bonds market. Finally, tax incentives (accelerated depreciation, tax credits, tax exemptions and rebates) can be provided either to the investor or to the issuer under the following forms.²

- Tax credit bonds: Bond investors receive tax credits instead of interest payments, so issuers do not pay coupon interests. Instead, they quarterly accrue phantom taxable income and tax credit equal to the amount of phantom income to holders, see Klein [185].
- Direct subsidy bonds: Bond issuers receive cash rebates from the government to subsidize their net interest payments. This type of incentives is mainly used by US municipalities, see for example Ang, Bhansali, and Xing [20].
- Tax-exempt bonds: Bond investors do not have to pay income tax on interest from the green bonds they hold (so issuer can get lower interest rate). This type of tax incentive is typically

¹For a complete survey of renewable energy promotion policies, we refer to Table 3 in Della Croce, Kaminker, and Stewart [94].

²The data provided below can be found at <https://www.climatebonds.net/policy/policy-areas/tax-incentives>.

applied to municipal bonds in the US market, see Calabrese and Ely [62] for a survey of the use of these tax-incentives.

All these incentives can be modelled as a function of the amount invested in green bonds. However, it should be clear that policy-makers cannot necessarily control directly the actions of the investor. This leads, for example, to the so-called ‘green-washing’ practice, when the investors use the funding obtained with the green bonds to finance non-sustainable projects, see Della Croce, Kaminker, and Stewart [94]. Moreover, the incentives are not dynamic in the sense that they do not depend on the evolution of market conditions (for example the price differences between green and conventional bonds). Thus, the incentives mechanism in the green bonds market is subject to a moral hazard component. In this chapter, we propose an alternative to tax incentives policy which is based on contract theory, and designed so as to increase the investment in green bonds. Moral hazard, whose related theory has been developed since the early 70’s, occurs when one person or entity (the Agent), is able to make decisions and/or take actions on behalf of, or that impact, another person or entity: the Principal. The classical continuous-time setting works as follow: the Principal hires an Agent to manage a ‘risky’ project, represented as a controlled stochastic differential equation. In exchange for the effort he puts into his work, the Agent receives a salary from the Principal which takes the form of a ‘contract’. The Principal’s goal is to offer a contract to the Agent allowing him to maximize its utility as a function of the terminal value of the project. The problem is addressed by solving a Stackelberg game, in two stages:

- (i) With a fixed contract, solve the problem of the Agent and obtain its optimal effort given a contract proposed by the Principal.
- (ii) Inject into the problem of the Principal the effort of better response of the Agent previously found, and solve the Principal’s problem, providing the optimal contract offered to the Agent.

Our goal is to propose a dynamic incentives model based on the prices and returns of green and conventional bonds issued by a government. We build a Principal-Agent model in which an investor (the Agent) runs a portfolio of green and conventional bonds. Without intervention of the government, the Agent has specific investment targets coming from his strategy. The policy-maker (the Principal) proposes incentives to the investor in order to achieve two objectives:

- (i) Increase the amount invested in green bonds according to a determined target;
- (ii) maximize the value of the portfolio of bonds issued by the government.

We show that without loss of utility for the government, we can consider incentives which take the form of stochastic integrals with respect to the portfolio process, the price of the bonds and their quadratic (co)variation. In order to propose tractable incentives for a possibly high number of bonds, we propose a form of contract that is based only on the dynamics of the portfolio process, the green bonds, an index of conventional bonds, and their respective quadratic (co)variations. In the case of deterministic short-term rates for the green and conventional bonds, both the Agent and the Principal’s problems can be solved by maximising deterministic functions with classical root-finding methods. When a one factor stochastic volatility model is used for short-term rates, we have to rely on stochastic control theory and determining the incentives of the policy-maker is equivalent to solve a high-dimensional, Hamilton-Jacobi-Bellman equation.

What we propose in this chapter is aimed to be used by governments as an alternative to the existing tax incentives, in order to increase the investment in green bonds. We summarize below the key features of our approach.

- The methodology we develop is completely tractable from a numerical point of view, thus the incentives can be designed on a large set of bonds.
- The remuneration we propose take into account the moral hazard between the investor and the government: the amount invested in the bonds is observed but not controlled by the government.

- The form of the optimal incentives is robust to model error: we show numerically that a more complex dynamics of the short-term rates of the bonds does not lead to an important loss in utility for the government, and causes minor variations in the form of the incentives.
- On a one-year horizon, the incentives show a rather constant behaviour. By using this, we show that the optimal incentives can be directly implemented with tradable financial products such as futures, log-contracts and variance swaps on the bonds.
- We compare our methodology with the current tax-incentives policy and show that, on a one-year period for a same target in green investments, our incentives policy leads to a higher value of the portfolio of bonds (15% to 20% on average).

In the numerical experiments, we provide general guidelines for the government to calibrate the model parameters, in particular the risk aversions, according to its objectives.

This chapter makes several contributions to the literature. First, to the best of our knowledge, it offers the first Principal-Agent framework to tackle the design of governmental incentives for green bonds. Contrary to articles like Zerbib [274] and Febi, Schäfer, Stephan, and Sun [115], where the authors provide a thorough descriptive analysis of the green bond market (risk premium, liquidity premium, ...) and examine the impact of green investing, we focus on answering a practical incentives problem from a quantitative viewpoint. The comparison with existing incentives policy on a set of French governmental bonds shows the benefits of our method for the government. The chapter contributes also to the Principal-Agent literature with volatility control, of which we give a brief overview.³ Contrary to the papers of Sung [259], Ou-Yang [227], the Principal observes the whole path of the controlled output process. Moreover, in our framework, moral hazard arises from unobservable sources of risk. In Lioui and Poncet [202], the authors consider a first-best problem with volatility control and assume that the agent has enough bargaining power to make the contract a linear function of the output and a benchmark risk factor. Another model is the one of Leung [199] where moral hazard with respect to the volatility arises because of the un-observability of the risk factors by the Principal and an exogenous source of risk multiplying the volatility of the Agent. These works are linked to the problem of ambiguity aversion on volatility and drift of the output process, see among others Chen and Sung [75], Hernández Santibáñez and Mastrolia [153], Mastrolia and Possamaï [205], Sung [261]. There is also a growing literature on the application of Principal-Agent with volatility control to the electricity market, see, for example, Élie, Hubert, Mastrolia, and Possamaï [111], Aïd, Possamaï, and Touzi [4]. Finally, we emphasize that the modelling framework of this chapter is directly inspired by the one in Cvitanić, Possamaï, and Touzi [88], where the authors consider the problem of delegated portfolio management and identify a family of admissible contracts for which the optimal agent's action is explicitly characterized. We extend this framework by allowing stochastic drift of the assets held by the Agent, and adapt it to our context.

The chapter is organized as follows. In 5.2, we present our framework and modelling assumptions. In 5.3, we solve the problems of the investor and the government with moral hazard and deterministic short rates. We present the numerical results in 5.4. Finally, we write in 5.A.1 the weak formulation of the control problem, while in 5.A.3 we solve the problem in the case of stochastic short rates.

5.2 Framework

Throughout the chapter, we work on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$ under which all stochastic processes are defined. We refer to 5.A.1 for the rigorous weak formulation of the problem, and we intend the present section to have a more accessible (and therefore more heuristic) flavour.

³This literature has been growing since the study of the well-posedness of second-order backward stochastic differential equations, see, for example, Possamaï, Tan, and Zhou [232], or Soner, Touzi, and Zhang [254]. A rigorous study of the Principal-Agent problem with volatility control in a general case can be found in Cvitanić, Possamaï, and Touzi [89].

We consider an investor wishing to develop his bonds' portfolio. He wants to acquire both green and conventional bonds issued by the same governmental entity or company with possible different amounts issued and different maturities. We assume that we are given a time horizon $T > 0$, and positive integers d^g and d^c . The investor manages, over the horizon $[0, T]$, d^g green bonds, d^c conventional bonds, and an index of conventional bonds of dynamics given by⁴

$$\begin{aligned} dP^g(t, T^g) &:= P^g(t, T^g) \circ \left((r^g(t) + \eta^g(t) \circ \sigma^g(t))dt + \text{diag}(\sigma^g(t))dW_t^g \right), \\ dP^c(t, T^c) &:= P^c(t, T^c) \circ \left((r^c(t) + \eta^c(t) \circ \sigma^c(t))dt + \text{diag}(\sigma^c(t))dW_t^c \right), \\ dI_t &:= I_t(\mu^I(t)dt + \sigma^I(t)dW_t^I). \end{aligned} \quad (5.1)$$

In the above equations, T^g is an \mathbb{R}^{d^g} -valued vector representing the maturities of each green bond and T^c is a \mathbb{R}^{d^c} -valued vector representing the maturities of each conventional bond. The functions $\mu^I : [0, T] \rightarrow \mathbb{R}$, $\sigma^I : [0, T] \rightarrow \mathbb{R}$ represent respectively the drift and volatility of the index of conventional of bonds $(I_t)_{t \in [0, T]}$. Similarly, the functions $r^g : [0, T] \rightarrow \mathbb{R}^{d^g}$, $r^c : [0, T] \rightarrow \mathbb{R}^{d^c}$ represent the vectors of short-term rate of the green and conventional bonds, and the functions $\eta^g : [0, T] \rightarrow \mathbb{R}^{d^g}$, $\eta^c : [0, T] \rightarrow \mathbb{R}^{d^c}$ represent the vectors of risk premia of the green and conventional bonds, while functions $\sigma^g : [0, T] \rightarrow \mathbb{R}^{d^g}$, $\sigma^c : [0, T] \rightarrow \mathbb{R}^{d^c}$ represent the vector of volatilities of the green and conventional bonds. The processes $(W_t^g)_{t \in [0, T]}$, $(W_t^c)_{t \in [0, T]}$, $(W_t^I)_{t \in [0, T]}$ are respectively \mathbb{R}^{d^g} , \mathbb{R}^{d^c} and \mathbb{R} -valued Brownian motions. Finally

$$W := \begin{pmatrix} W^g \\ W^c \\ W^I \end{pmatrix}$$

is an $\mathbb{R}^{d^g+d^c+1}$ -valued Brownian motion, whose co-variance structure is given by $d\langle W \rangle_t = \Sigma dt$, where

$$\Sigma \in \mathcal{M}_{d^g+d^c+1}(\mathbb{R}), \quad \Sigma := \begin{pmatrix} \Sigma^g & \Sigma^{g,c} & \Sigma^{g,I} \\ \Sigma^{g,c} & \Sigma^c & \Sigma^{c,I} \\ \Sigma^{g,I} & \Sigma^{c,I} & \Sigma^I \end{pmatrix},$$

with

$$\begin{aligned} \Sigma^g &\in \mathcal{M}_{d^g}(\mathbb{R}), \quad \Sigma_{i,j}^g := \rho_{i,j}^g \in [-1, 1] \text{ if } i \neq j, 1 \text{ otherwise, } (i, j) \in \{1, \dots, d^g\}^2, \\ \Sigma^c &\in \mathcal{M}_{d^c}(\mathbb{R}), \quad \Sigma_{i,j}^c := \rho_{i,j}^c \in [-1, 1] \text{ if } i \neq j, 1 \text{ otherwise, } (i, j) \in \{1, \dots, d^c\}^2, \\ \Sigma^{g,c} &\in \mathcal{M}_{d^g, d^c}(\mathbb{R}), \quad \Sigma_{i,j}^{g,c} := \rho_{i,j}^{g,c} \in [-1, 1], \quad (i, j) \in \{1, \dots, d^g\} \times \{1, \dots, d^c\}, \\ \Sigma^{g,I} &\in \mathbb{R}^{d^g}, \quad \Sigma_i^{g,I} := \rho_i^{g,I} \in [-1, 1], \quad i \in \{1, \dots, d^g\}, \\ \Sigma^{c,I} &\in \mathbb{R}^{d^c}, \quad \Sigma_i^{c,I} := \rho_i^{c,I} \in [-1, 1], \quad i \in \{1, \dots, d^c\}. \end{aligned}$$

Remark 27. *All these quantities are assumed to be deterministic, in order to derive a governmental incentive that is tractable for a large number of bonds. We will show in 5.A.3 that, at the expense of a higher computational cost and the use of stochastic control theory, one can also derive incentives for the investor when short-term rates are stochastic. In 5.4, we show numerically that the use of stochastic short-term rates for the green bonds does not impact qualitatively our results. In particular, when the short-term rates are driven by Ornstein-Uhlenbeck processes, the optimal investment policy in this case oscillates slightly around the one obtained with deterministic rates. Thus, the methodology we propose appears to be robust to model specification.*

⁴We define the index as an average of the dynamics of the conventional bonds. In practice, the investor may trade a large quantity of conventional bonds and only a couple of green bonds. Thus, we argue that it is more convenient for the government to index the remuneration proposed on an average dynamics of conventional bonds in order to have more granularity for the green bonds' incentives.

Throughout the chapter, we use the following technical assumption.

Assumption 5.1. *The functions r^g , r^c , η^c , η^g , σ^g , σ^c , μ^I , and σ^I are uniformly bounded on $[0, T]$.*

The investment policy is defined by a vector of control processes $\pi = (\pi_t^g, \pi_t^c, \pi_t^I)_{t \in [0, T]} \in \mathcal{A}$, representing the amount of money invested at time t , where

$$\mathcal{A} := \left\{ (\pi_t)_{t \in [0, T]} : K\text{-valued and } \mathbb{F}\text{-predictable processes} \right\}.$$

is the set of admissible control process, where $K := [\varepsilon, b_\infty]^{d^g} \times [\varepsilon, b_\infty]^{d^c} \times [\varepsilon, b_\infty]$, for some $0 < \varepsilon < b_\infty$ ⁵ and $\mathbb{F} := (\mathcal{F}_t)_{t \in [0, T]}$ is the natural filtration of the process (X, W) with X defined below. We define the dynamics of the vectors of returns on the bonds as

$$\begin{aligned} dR^g(t, T^g) &= (r^g(t) + \eta^g(t) \circ \sigma^g(t))dt + \text{diag}(\sigma^g(t))dW_t^g, \\ dR^c(t, T^c) &= (r^c(t) + \eta^c(t) \circ \sigma^c(t))dt + \text{diag}(\sigma^c(t))dW_t^c, \\ dR_t^I &= \mu^I(t)dt + \sigma^I(t)dW_t^I. \end{aligned}$$

For every $\pi \in \mathcal{A}$, one can define a probability measure \mathbb{P}^π ⁶ such that the dynamics of the value of portfolio of bonds is given by

$$dX_t := \pi_t^g \cdot dR^g(t, T^g) + \pi_t^c \cdot dR^c(t, T^c) + \pi_t^I dR_t^I.$$

We also denote by \mathbb{E}_t^π the conditional expectation under the probability measure \mathbb{P}^π with respect to \mathcal{F}_t for all $t \in [0, T]$. Throughout the investment period $[0, T]$, the investor wants to maintain his investment in bonds at some pre-defined levels, which can be seen as his investment profile. We introduce the vectors $\alpha = (\alpha^g, \alpha^c, \alpha^I) \in \mathbb{R}^{d^g} \times \mathbb{R}^{d^c} \times \mathbb{R}$ and the cost function $k : \mathbb{R}^{d^g} \times \mathbb{R}^{d^c} \times \mathbb{R} \rightarrow \mathbb{R}$, where for any $p := (p^g, p^c, p^I) \in \mathbb{R}^{d^g} \times \mathbb{R}^{d^c} \times \mathbb{R}$

$$k(p) := \frac{1}{2}\beta^g \cdot (p^g - \alpha^g)^2 + \frac{1}{2}\beta^c \cdot (p^c - \alpha^c)^2 + \frac{1}{2}\beta^I (p^I - \alpha^I)^2,$$

where $\beta := (\beta^g, \beta^c, \beta^I) \in \mathbb{R}^{d^g} \times \mathbb{R}^{d^c} \times \mathbb{R}$ are what we coin intensity vectors. For instance, at some time $t \in [0, T]$, the investor pays a cost to move the amount $(\pi_t^g)_i$ invested in the i -th green bond away from the initial target α_i^g , and this cost is equal to $\frac{1}{2}\beta_i^g ((\pi_t^g)_i - \alpha_i^g)^2$. Thus, $(\beta^g, \beta^c, \beta^I)$ represent the cost intensity of changing the investments of the agent: the higher these coefficients, the more incentives the investor will demand to change his investment profile.

In order to modify an investment policy $\pi \in \mathcal{A}$, the government proposes a remuneration to the investor. It takes the form of an \mathcal{F}_T -measurable random variable denoted by ξ , and we will see later that the form of remuneration considered is an indexation on the value of the portfolio of bonds as well as the sources of risk of each bond. The optimisation problem of the investor with CARA utility function writes, for a given contract provided by the government, as

$$V^A(\xi) := \sup_{\pi \in \mathcal{A}} \mathbb{E}^\pi \left[U_A \left(\xi - \int_0^T k(\pi_s) ds \right) \right], \quad U_A(x) := -\exp(-\gamma x),$$

where $\gamma > 0$ is his risk aversion parameter. To ensure that the control problem of the investor is non-degenerate, we impose the following integrability condition on the contracts

$$\sup_{\pi \in \mathcal{A}} \mathbb{E}^\pi \left[\exp(-\gamma' \xi) \right] < +\infty, \quad \text{for some } \gamma' > \gamma. \quad (5.2)$$

⁵We force the control processes to be strictly positive so that the density of the canonical process in 5.A.1 is invertible and we can define properly the weak formulation of the control problem. Practically, this simply means that the investor has to invest in the index, and in at least one of the conventional and one of the green bonds.

⁶See 5.A.1 for the weak formulation of the control problem, which explains how to construct \mathbb{P}^π .

Remark 28. *We emphasize here that the notion of price for a bond is meaningless as it is not quoted on the National Best Bid and Offer (NBBO): This is an OTC market where the liquidity is provided by one or several dealers. In particular, even though there is a quantity defined as the bond price on Bloomberg, it serves only as an indication as the dealers have no obligation to buy or sell at this price. However, especially in the case of treasury bonds, Futures on the bonds are listed on the Chicago Board Of Trade where the notion of price is meaningful. Thus, throughout the chapter, the notion of bond price must be thought as the price of a future on the considered bond.*

On the other hand, the government wishes to maximize the portfolio value of the bonds issued while increasing the amount invested in green bonds. Thus, he wants to maximize, on average, the quantity

$$X_T - \sum_{i=1}^{d^g} \int_0^T \kappa \left(G_i - (\hat{\pi}_t^g(\xi))_i \right)^2 dt,$$

where for $i \in \{1, \dots, d^g\}$, G_i is the investment target in the i -th green bond of the government entity, $\kappa > 0$ is the cost of moving away from the targets (G_1, \dots, G_{d^g}) and $\hat{\pi}(\xi)$ is a best response of the investor to a given contract ξ .⁷ We assume that the cost of moving away from the targets is the same for each green bond, meaning that the government does not have different preferences for each bond (this assumption can of course be relaxed). The government also subtracts from this quantity the contract ξ offered to the investor. Thus, his optimisation problem with CARA utility function writes

$$V_0^P = \sup_{\xi \in \mathcal{C}} \sup_{\hat{\pi} \in \mathcal{A}(\xi)} \mathbb{E}^{\hat{\pi}} \left[U_P \left(X_T - \sum_{i=1}^{d^g} \int_0^T \kappa \left(G_i - (\hat{\pi}_t^g(\xi))_i \right)^2 dt - \xi \right) \right], \quad U_P(x) = -\exp(-\nu x), \quad (5.3)$$

where $\nu > 0$ is the risk aversion parameter of the Principal,

$$\mathcal{A}(\xi) := \left\{ \hat{\pi} \in \mathcal{A} : V^A(\xi) = \mathbb{E}^{\hat{\pi}} \left[-\exp \left(-\gamma \left(\xi - \int_0^T k(\hat{\pi}_s) ds \right) \right) \right] \right\},$$

is the set of best-responses of the Agent to a given contract ξ and

$$\mathcal{C} = \left\{ \xi : \mathbb{R}\text{-valued, } \mathcal{F}_T\text{-measurable random variable such that } V^A(\xi) \geq R, \text{ and (5.2) is satisfied} \right\},$$

is the set of admissible contracts for the government, where $R < 0$ is the reservation utility of the investor: He will not accept to work for Principal (and accept the contract ξ) unless the contract is such that his expected utility is above R .

Remark 29. *We consider here that the reservation utility corresponds to the utility function of the investor in the case $\xi = 0$, that is*

$$R = V^A(0) = \sup_{\pi \in \mathcal{A}} \mathbb{E}^{\pi} \left[-\exp \left(\gamma \int_0^T k(\pi_s) ds \right) \right] = -1,$$

where the supremum is reached by choosing $\pi = (\alpha^g, \alpha^c, \alpha^I)$. We will see in the following section that the optimal contract proposed by the government will always saturate this constraint, that is the Principal will provide the Agent with the minimum reservation utility R he requires.

5.3 Solving the optimisation problem

5.3.1 The optimal contract

In this section, we derive the optimal governmental incentives proposed to the investor. As it would be unrealistic (and hardly tractable) to offer a compensation based on the whole universe of governmental

⁷We will see later that there might be several best responses of the Agent. Thus, following the tradition in the moral hazard literature, we assume that the Principal has enough bargaining power to be able to choose the best response of the Agent that maximizes his own utility.

bonds, we suggest a remuneration based on the green bonds, the value of the portfolio and an index of conventional bonds. This way, the contract is only indexed on $d^g + 2$ variables. The optimal incentives are obtained by maximising a deterministic function, which makes the problem easily tractable for a large number of green bonds. We begin this section with the definition of contractible and non-contractible variables.

Definition 5.1. *The set of contractible variables is defined as the \mathbb{R}^{d^g+2} -valued process*

$$B^{\text{obs}} := \begin{pmatrix} X \\ W^g \\ W^I \end{pmatrix}.$$

The set of non-contractible variables is defined as the \mathbb{R}^{d^c} -valued process $B^{\text{obs}} := W^c$, with the following dynamics

$$dB_t^{\text{obs}} := \mu^{\text{obs}}(t, \pi_t)dt + \Sigma^{\text{obs}}(t, \pi_t)dW_t, \quad dB_t^{\text{obs}} := \mu^{\text{obs}}dt + \Sigma^{\text{obs}}dW_t,$$

where $\mu^{\text{obs}} := (\mathbf{0}_{d^c,1})$, $\Sigma^{\text{obs}} := (\mathbf{0}_{d^c,d^g} \quad I_{d^c} \quad \mathbf{0}_{d^c,1})$, and the maps $\mu^{\text{obs}} : [0, T] \times \mathbb{R}^{d^g} \times \mathbb{R}^{d^c} \times \mathbb{R} \rightarrow \mathbb{R}^{d^g+2}$, as well as $\Sigma^{\text{obs}} : [0, T] \times \mathbb{R}^{d^g} \times \mathbb{R}^{d^c} \times \mathbb{R} \rightarrow \mathcal{M}_{d^g+2, d^g+d^c+1}(\mathbb{R})$ are defined for any $p := (p^g, p^c, p^I) \in \mathbb{R}^{d^g} \times \mathbb{R}^{d^c} \times \mathbb{R}$ and $t \in [0, T]$ by

$$\mu^{\text{obs}}(t, p) := \begin{pmatrix} p^g \cdot (r^g(t) + \eta^g(t) \circ \sigma^g(t)) + p^c \cdot (r^c(t) + \eta^c(t) \circ \sigma^c(t)) + p^I \mu^I(t) \\ \mathbf{0}_{d^g,1} \\ 0 \end{pmatrix},$$

$$\Sigma^{\text{obs}}(t, p) := \begin{pmatrix} (p^g \circ \sigma(t)^g)^\top & (p^c \circ \sigma(t)^c)^\top & p^I \sigma^I(t) \\ I_{d^g} & \mathbf{0}_{d^g, d^c} & \mathbf{0}_{d^g, 1} \\ \mathbf{0}_{1, d^g} & \mathbf{0}_{1, d^c} & 1 \end{pmatrix}.$$

Finding the optimal contract ξ in the optimisation problem (5.3) is an arduous task, as we search a solution in the space of \mathcal{F}_T -measurable random variables. However, see Cvitanić, Possamaï, and Touzi [89], it has been shown that without reducing the utility of the Principal, we can restrict our study to admissible contracts which have a specific form. In order to describe this result, we need first to introduce additional notations.

We define the quantities

$$B := \begin{pmatrix} B^{\text{obs}} \\ B^{\text{obs}} \end{pmatrix}, \quad \mu(t, p) := \begin{pmatrix} \mu^{\text{obs}}(t, p) \\ \mu^{\text{obs}} \end{pmatrix}, \quad \Sigma(t, p) := \begin{pmatrix} \Sigma^{\text{obs}}(t, p) \\ \Sigma^{\text{obs}} \end{pmatrix}, \quad (t, p) \in [0, T] \times \mathbb{R}^{d^g} \times \mathbb{R}^{d^c} \times \mathbb{R}.$$

We also will need to introduce the map $h : [0, T] \times \mathbb{R}^{d^g+d^c+2} \times \mathbb{S}_{d^g+d^c+2}(\mathbb{R}) \times K \rightarrow \mathbb{R}$, with for $(t, z, g, p) \in [0, T] \times \mathbb{R}^{d^g+d^c+2} \times \mathbb{S}_{d^g+d^c+2}(\mathbb{R}) \times K$,

$$h(t, z, g, p) = -k(p) + z \cdot \mu(t, p) + \frac{1}{2} \text{Tr}[g \Sigma(t, p) \Sigma(\Sigma(t, p))^\top],$$

and for all $(t, z, g) \in [0, T] \times \mathbb{R}^{d^g+d^c+2} \times \mathbb{S}_{d^g+d^c+2}(\mathbb{R})$,

$$\mathcal{O}(t, z, g) := \left\{ \hat{p} \in K : \hat{p} \in \underset{p \in K}{\text{argmax}} \{ h(t, z, g, p) \} \right\},$$

is the set of the maximizers of h with respect to its last variable, for (t, z, g) given. Following Schäl [245], there exists at least one Borel-measurable map $\hat{\pi} : [0, T] \times \mathbb{R}^{d^g+d^c+2} \times \mathbb{S}_{d^g+d^c+2}(\mathbb{R}) \rightarrow K$ such that for every $(t, z, g) \in [0, T] \times \mathbb{R}^{d^g+d^c+2} \times \mathbb{S}_{d^g+d^c+2}(\mathbb{R})$, $\hat{\pi}(t, z, g) \in \mathcal{O}(t, z, g)$. We denote by \mathcal{O} the corresponding set of all such maps.

Theorem 5.1. *Without reducing the utility of the Principal, we can restrict the study of admissible contracts to the set \mathcal{C}_1 where any $\xi \in \mathcal{C}_1 \subset \mathcal{C}$ is of the form $\xi = Y_T^{y_0, Z, \Gamma, \hat{\pi}}$ where for $t \in [0, T]$,*

$$Y_t^{y_0, Z, \Gamma, \hat{\pi}} := y_0 + \int_0^t Z_s \cdot dB_s + \frac{1}{2} \int_0^t \text{Tr}[(\Gamma_s + \gamma Z_s Z_s^\top) d\langle B \rangle_s] - \int_0^t h(s, Z_s, \Gamma_s, \hat{\pi}(s, Z_s, \Gamma_s)) ds, \quad (5.4)$$

where $y_0 \in \mathbb{R}$, $\hat{\pi} \in \mathcal{O}$ and (Z, Γ) are respectively $\mathbb{R}^{d^g + d^c + 2}$ - and $\mathbb{S}_{d^g + d^c + 2}(\mathbb{R})$ -valued, \mathbb{F} -predictable processes such that Condition (5.2) is satisfied for $Y_T^{y_0, Z, \Gamma, \hat{\pi}}$, and $V^A(Y_T^{y_0, Z, \Gamma, \hat{\pi}}) \geq U_A(y_0)$. We denote by \mathcal{ZG} the set of such processes, which is properly defined in (5.13). Moreover, we have

$$V^A(Y_T^{y_0, Z, \Gamma, \hat{\pi}}) = U_A(y_0), \quad \mathcal{A}(Y_T^{y_0, Z, \Gamma, \hat{\pi}}) = \left\{ (\hat{\pi}(t, Z_t, \Gamma_t))_{t \in [0, T]} : \hat{\pi} \in \mathcal{O}, (Z, \Gamma) \in \mathcal{ZG} \right\}.$$

The form of the admissible contracts we study deserves some remarks. The term $Z_t \cdot dB_t$ is a remuneration indexed linearly on the state variables. Contrary to the classical Principal-Agent problem where the agent controls the drift of the output process, see Sannikov [243] for example, the admissible contracts (5.4) are not only linear functions of the state variables but depend also linearly on their quadratic variation and covariation. This comes from the fact that by investing in the bonds, the investor controls directly the volatility of the portfolio process X . Using standard tools of static hedging, this contract can be replicated using futures, log-contracts and volatility products such as variance swaps, see 5.3.2.2 for details. In particular, this ensures that the contracts we recommend are practically implementable.

As stated at the beginning of this section, we wish to build an optimal contract based only on the green bonds, the portfolio process and the index of conventional bonds. In this regard, the form we obtained in 5.4 is too general, which is why we are now going to restrict our attention to a slightly smaller class of contracts. We thus define for any $(Z, \Gamma) \in \mathcal{ZG}$

$$Z_t := \begin{pmatrix} Z_t^{\text{obs}} \\ Z_t^{\text{obs}} \end{pmatrix}, \quad \Gamma := \begin{pmatrix} \Gamma_t^{\text{obs}} & \Gamma_t^{\text{obs, obs}} \\ \Gamma_t^{\text{obs, obs}} & \Gamma_t^{\text{obs}} \end{pmatrix},$$

where for Lebesgue-almost every $t \in [0, T]$

$$Z_t^{\text{obs}} \in \mathbb{R}^{d^g + 2}, \quad Z_t^{\text{obs}} \in \mathbb{R}^{d^c}, \quad \Gamma_t^{\text{obs}} \in \mathbb{S}_{d^g + 2}(\mathbb{R}), \quad \Gamma_t^{\text{obs}} \in \mathbb{S}_{d^c}(\mathbb{R}), \quad \Gamma_t^{\text{obs, obs}} \in \mathcal{M}_{d^g + 2, d^c}(\mathbb{R}).$$

We then consider a simplified Hamiltonian $h^{\text{obs}} : [0, T] \times \mathbb{R}^{d^g + 2} \times \mathbb{S}_{d^g + 2}(\mathbb{R}) \times K \rightarrow \mathbb{R}$ given by

$$h^{\text{obs}}(t, z^{\text{obs}}, g^{\text{obs}}, p) = -k(p) + z^{\text{obs}} \cdot \mu^{\text{obs}}(t, p) + \frac{1}{2} \text{Tr}[g^{\text{obs}} \Sigma^{\text{obs}}(t, p) \Sigma(\Sigma^{\text{obs}}(t, p)^\top)],$$

and for all $(t, z^{\text{obs}}, g^{\text{obs}}) \in [0, T] \times \mathbb{R}^{d^g + 2} \times \mathbb{S}_{d^g + 2}(\mathbb{R})$, we define

$$\mathcal{O}^{\text{obs}}(t, z^{\text{obs}}, g^{\text{obs}}) := \left\{ \hat{p} \in K : \hat{p} \in \underset{p \in K}{\text{argmax}} \{ h^{\text{obs}}(t, z^{\text{obs}}, g^{\text{obs}}, p) \} \right\}.$$

Following again Schäl [245], there exists at least one Borel-measurable map $\hat{\pi} : [0, T] \times \mathbb{R}^{d^g + 2} \times \mathbb{S}_{d^g + 2}(\mathbb{R}) \rightarrow K$ such that for every $(t, z^{\text{obs}}, g^{\text{obs}}) \in [0, T] \times \mathbb{R}^{d^g + 2} \times \mathbb{S}_{d^g + 2}(\mathbb{R})$, $\hat{\pi}(t, z^{\text{obs}}, g^{\text{obs}}) \in \mathcal{O}^{\text{obs}}(t, z^{\text{obs}}, g^{\text{obs}})$, and we let \mathcal{O}^{obs} be the corresponding set of all such maps.

We can now state precisely the class of contracts we are concerned with in this chapter.

Assumption 5.2. *We consider the subset of contracts*

$$\mathcal{C}_2 := \left\{ Y_T^{y_0, Z, \Gamma, \hat{\pi}} \in \mathcal{C}_1 : Z^{\text{obs}} = \mathbf{0}_{d^c}, \Gamma^{\text{obs}} = \mathbf{0}_{d^c, d^c}, \Gamma^{\text{obs, obs}} = \mathbf{0}_{d^g + 2, d^c} \right\}.$$

In particular, any $\xi \in \mathcal{C}_2$ is of the form $\xi = Y_T^{y_0, Z^{\text{obs}}, \Gamma^{\text{obs}}, \hat{\pi}}$, where for any $t \in [0, T]$,

$$\begin{aligned}
 Y_t^{y_0, Z^{\text{obs}}, \Gamma^{\text{obs}}, \hat{\pi}} := & y_0 + \int_0^t Z_s^{\text{obs}} \cdot dB_s^{\text{obs}} + \frac{1}{2} \text{Tr} \left[(\Gamma_s^{\text{obs}} + \gamma Z_s^{\text{obs}} (Z_s^{\text{obs}})^\top) d\langle B^{\text{obs}} \rangle_s \right] \\
 & - h^{\text{obs}} \left(s, Z_s^{\text{obs}}, \Gamma_s^{\text{obs}}, \hat{\pi}(s, Z_s^{\text{obs}}, \Gamma_s^{\text{obs}}) \right) ds,
 \end{aligned} \tag{5.5}$$

where $y_0 \geq 0$, $\hat{\pi} \in \mathcal{O}^{\text{obs}}$ and $(Z^{\text{obs}}, \Gamma^{\text{obs}}) \in \mathcal{ZG}^{\text{obs}}$ with

$$\mathcal{ZG}^{\text{obs}} := \left\{ (Z^{\text{obs}}, \Gamma^{\text{obs}}) : \mathbb{R}^{d^g+2} \times \mathbb{S}_{d^g+2}(\mathbb{R})\text{-valued, } \mathbb{F}\text{-predictable, s.t. } Y_T^{y_0, Z^{\text{obs}}, \Gamma^{\text{obs}}, \hat{\pi}} \in \mathcal{C}_2 \right\}.$$

The optimisation problem of the government that we now consider is⁸

$$\begin{aligned}
 \tilde{V}_0^P = \sup_{y_0 \geq 0} \sup_{(Z^{\text{obs}}, \Gamma^{\text{obs}}, \hat{\pi}) \in \mathcal{ZG}^{\text{obs}} \times \mathcal{O}^{\text{obs}}} & \mathbb{E}^{\hat{\pi}(Z, \Gamma)} \left[U_P \left(X_T - \sum_{i=1}^{d^g} \int_0^T \kappa \left(G_i - (\hat{\pi}^g(t, Z_t^{\text{obs}}, \Gamma_t^{\text{obs}}))_i \right)^2 dt \right. \right. \\
 & \left. \left. - Y_T^{y_0, Z^{\text{obs}}, \Gamma^{\text{obs}}, \hat{\pi}} \right) \right],
 \end{aligned} \tag{5.6}$$

This assumption allows us to consider more tractable contracts for a large portfolio of bonds, even if we consider less general contracts compared to (5.4). Moreover, as the objective of the government is to encourage the acquisition of green bonds, it is natural to consider a more granular contract with respect to the green bonds and to use only the index of conventional bonds as a representative contractible variable of this set of bonds. As we used only deterministic functions to model the risk premium, short-term rate and volatility processes, the optimal incentives of the government can be obtained by maximising a deterministic function, which leads to the following theorem.

Theorem 5.2 (Main result). *The optimal contract $\xi^* \in \mathcal{C}_2$ is given by*

$$\begin{aligned}
 \xi^* = Y_T^{0, z^{\star \text{obs}}, g^{\star \text{obs}}, \pi^*} = \int_0^T & z^{\star \text{obs}}(t) \cdot dB_t^{\text{obs}} + \frac{1}{2} \text{Tr} \left[(g^{\star \text{obs}}(t) + \gamma z^{\star \text{obs}}(t) (z^{\star \text{obs}}(t))^\top) d\langle B^{\text{obs}} \rangle_t \right] \\
 & - h^{\text{obs}} \left(t, z^{\star \text{obs}}(t), g^{\star \text{obs}}(t), \pi^*(t, z^{\star \text{obs}}(t), g^{\star \text{obs}}(t)) \right) dt,
 \end{aligned} \tag{5.7}$$

where for all $t \in [0, T]$, $z^{\star \text{obs}}(\cdot)$, $g^{\star \text{obs}}(\cdot)$, $\pi^*(\cdot, z^{\star \text{obs}}(\cdot), g^{\star \text{obs}}(\cdot))$ are deterministic functions of time, solving

$$\sup_{(z, g, \hat{\pi}) \in P \times \mathcal{O}^{\text{obs}}} \mathcal{H}(t, z, g, \hat{\pi}(t, z, g)), \tag{5.8}$$

where $P := \mathbb{R}^{d^g+2} \times \mathbb{S}_{d^g+2}(\mathbb{R})$ and $\mathcal{H} : [0, T] \times P \times K \rightarrow \mathbb{R}$ is given by

$$\begin{aligned}
 \mathcal{H}(t, z, g, p) := & - \sum_{i=1}^{d^g} (G_i - p_i)^2 - \frac{1}{2} \text{Tr} \left[(g + \gamma z z^\top) \Sigma^{\text{obs}}(t, p) \Sigma(\Sigma^{\text{obs}}(t, p)^\top) \right] \\
 & + h^{\text{obs}}(t, z, g, p) + \left(\mu^{\text{obs}}(t, p) \right)_1 - z^\top \mu^{\text{obs}}(t, p) \\
 & - \frac{1}{2} \nu^2 \left(\left(\Sigma^{\text{obs}}(t, p) \right)_{1,:} - z^\top \Sigma^{\text{obs}}(t, p) \right)^\top \Sigma \left(\left(\Sigma^{\text{obs}}(t, p) \right)_{1,:} - z^\top \Sigma^{\text{obs}}(t, p) \right).
 \end{aligned}$$

Moreover

$$\tilde{V}_0^P = U_P \left(\int_0^T \mathcal{H}(t, z^{\star \text{obs}}(t), g^{\star \text{obs}}(t), \pi^*(t, z^{\star \text{obs}}(t), g^{\star \text{obs}}(t))) dt \right).$$

Proof. The term in the exponential of the optimisation problem (5.6) is a linear function of y_0 hence the reservation utility of the investor is saturated using $y_0^* = 0$. Define for any martingale M the operator

⁸We use the notation $\mathbb{E}^{(\hat{\pi}(t, Z_t, \Gamma_t))_{t \in [0, T]}}[\cdot] =: \mathbb{E}^{\hat{\pi}(Z, \Gamma)}[\cdot]$

$$\mathcal{E}(M)_T := \exp\left(-\nu M_T + \frac{1}{2}\nu^2\langle M\rangle_T\right).$$

The government has now to solve

$$\begin{aligned} \sup_{(Z, \Gamma, \hat{\pi}) \in \mathcal{ZG}^{\text{obs}} \times \mathcal{O}^{\text{obs}}} \mathbb{E}^{\hat{\pi}(Z, \Gamma)} & \left[U_P \left(\int_0^T \left((\mu^{\text{obs}}(t, \hat{\pi}(t, Z_t, \Gamma_t)))_1 - \sum_{i=1}^{d_g} (G_i - \hat{\pi}_i(t, Z_t, \Gamma_t))^2 \right. \right. \right. \\ & - \frac{1}{2} \text{Tr} \left[\left(\Gamma(t, \hat{\pi}(t, Z_t, \Gamma_t)) + \gamma Z_t Z_t^\top \right) \Sigma^{\text{obs}}(t, \hat{\pi}(t, Z_t, \Gamma_t)) \Sigma \left(\Sigma^{\text{obs}}(t, \hat{\pi}(t, Z_t, \Gamma_t)) \right)^\top \right] \\ & \left. \left. \left. + h^{\text{obs}}(t, Z_t, \Gamma_t, \hat{\pi}(t, Z_t, \Gamma_t)) \right) dt \right) \right. \\ & \left. \times \exp \left(-\nu \int_0^T \left((\Sigma^{\text{obs}}(t, \hat{\pi}(t, Z_t, \Gamma_t)))_{0,:} - Z_t^\top \Sigma^{\text{obs}}(t, \hat{\pi}(t, Z_t, \Gamma_t)) \right) dW_t \right) \right]. \end{aligned}$$

We make appear the stochastic exponential so that the previous supremum becomes

$$\begin{aligned} \sup_{(Z, \Gamma, \hat{\pi}) \in \mathcal{ZG}^{\text{obs}} \times \mathcal{O}^{\text{obs}}} \mathbb{E}^{\hat{\pi}(Z, \Gamma)} & \left[U_P \left(\int_0^T \mathcal{H}(t, Z_t, \Gamma_t, \hat{\pi}(t, Z_t, \Gamma_t)) dt \right) \right. \\ & \left. \times \mathcal{E} \left(\int_0^\cdot \left((\Sigma^{\text{obs}}(t, \hat{\pi}(t, Z_t, \Gamma_t)))_{0,:} - Z_t^\top \Sigma^{\text{obs}}(t, \hat{\pi}(t, Z_t, \Gamma_t)) \right) dW_t \right)_T \right]. \end{aligned}$$

As the function $U_P(x)$ is increasing and the expectation of a stochastic exponential is bounded by one, we obtain

$$\tilde{V}_0^P \leq U_P \left(\int_0^T \sup_{(z, g, \hat{\pi}) \in P \times \mathcal{O}^{\text{obs}}} \mathcal{H}(t, z, g, \hat{\pi}(t, z, g)) dt \right).$$

We have

$$\mathcal{H}(t, z, g, \hat{\pi}(t, z, g)) \leq -\frac{1}{2} \text{Tr} \left[\gamma z z^\top \Sigma^{\text{obs}}(t, \hat{\pi}(t, z, g)) \Sigma \left(\Sigma^{\text{obs}}(t, \hat{\pi}(t, z, g)) \right)^\top \right] + \left(\mu^{\text{obs}}(t, \hat{\pi}(t, z, g)) \right)_1.$$

As $\hat{\pi}(t, z, g) < +\infty$ is uniformly bounded and strictly positive, Σ is definite positive and the components of Σ^{obs} are positive, we observe that when $\|z\|_2 + \|g\|_2 \rightarrow +\infty$, the first term goes to $-\infty$ while the second term is bounded. Therefore, the supremum on \mathcal{O}^{obs} cannot be attained for infinite values.

If we now choose the incentives $z^{*,\text{obs}}(t), g^{*,\text{obs}}(t), \pi^*(t, z^{*,\text{obs}}(t), g^{*,\text{obs}}(t))$ as the maximizers of \mathcal{H} , they are Borel-measurable deterministic functions of $t \in [0, T]$ thus belong to the set $\mathcal{ZG}^{\text{obs}}$ and are bounded on $[0, T]$, so that

$$\mathcal{E} \left(\int_0^\cdot \left((\Sigma^{\text{obs}}(t, \pi^*(t, z^{*,\text{obs}}(t), g^{*,\text{obs}}(t))))_{0,:} - z^{*,\text{obs}}(t)^\top \Sigma^{\text{obs}}(t, \pi^*(t, z^{*,\text{obs}}(t), g^{*,\text{obs}}(t))) \right) dW_t \right)_T,$$

is a \mathbb{P}^{π^*} -martingale and we obtain

$$\tilde{V}_0^P = U_P \left(\int_0^T \mathcal{H}(t, z^{*,\text{obs}}(t), g^{*,\text{obs}}(t), \pi^*(t, z^{*,\text{obs}}(t), g^{*,\text{obs}}(t))) dt \right).$$

□

Static maximisation (5.8) over $(z, g) \in P$ can easily be handled with classic root-finding algorithms for a large portfolio of green bonds.⁹ Before moving to the numerical experiments, we discuss the form and implementability of the optimal contract.

⁹In practice, we observe that for the set of parameters we choose for the numerical experiences, the function h^{obs} is strictly concave with respect to its last variable thus admits a unique maximizer $\hat{\pi}$.

5.3.2 Discussion

5.3.2.1 On the form of the optimal contracts

The contract consists of the following elements:

- The term Z_X^{obs} is a compensation given to the investor with respect to the risk associated to the evolution of the portfolio process. If $Z_X^{\text{obs}} > 0$ (resp. $Z_X^{\text{obs}} < 0$), the government encourages to increase (resp. decrease) the value of the portfolio: between two times $t_2 > t_1$, the investor receives approximately the amount $(Z_X^{\text{obs}})_{t_1}(X_{t_2} - X_{t_1})$.
- For $i \in \{1, \dots, d^g\}$, the term Z_i is a compensation given to the investor with respect to the volatility risk associated to the evolution of the i -th green bond price. Between two times $t_2 > t_1$, the investor receives approximately the amount $(Z_i^{\text{obs}})_{t_1}(W_{t_2}^i - W_{t_1}^i)$: if Z_i^{obs} is close to zero, the government does not give compensation with respect to the volatility of the i -th green bond and conversely for Z_i^{obs} far from zero. The intuition behind Z_i^{obs} is the same.
- The diagonal terms of Γ^{obs} are compensations with respect to the quadratic variation of the portfolio process and the risk sources of the green bonds and the index. For example, if $\Gamma_X^{\text{obs}} > 0$, the government provides remuneration to the investor for a high quadratic variation (which here can be thought of as volatility) of the portfolio process. If $\Gamma_X < 0$, the government penalizes a high volatility of the portfolio process.
- The non-diagonal terms of Γ^{obs} are compensations with respect to the quadratic covariation of the portfolio process and the risk sources of the green bonds and the index. For example, if $\Gamma_{X,i}^{\text{obs}} > 0$ for $i \in \{1, \dots, d^g\}$ the government provides remuneration to the investor for similar moves of the portfolio process and the i -th green bond. If $\Gamma_{X,i}^{\text{obs}} < 0$, the government encourages opposite moves of the portfolio process and the i -th green bond.
- The term $G^{\text{obs}}(t, Z^{\text{obs}}, \Gamma^{\text{obs}})$ is a continuous coupon that is given to the investor. It corresponds to the utility of the investor in the case $\xi = 0$.

For reasonable choices of parameters (α, β, G) , the supremum of h^{obs} and in (5.8) are strictly concave functions so that an optimizer is quickly found using root-finding algorithms. Note that the optimal contract is indexed on the portfolio process X the sources of risk coming from the green bonds W^g and the one coming from the index W^I . This can be reformulated as an indexing on X and the prices of the bonds. In this case we define

$$B^{\text{obs},p} := \begin{pmatrix} X \\ \log(P^g) \\ \log(P^I) \end{pmatrix}, \quad B^{\text{obs},p} := \log(P^c),$$

$$dB_t^{\text{obs},p} := \mu^{\text{obs},p}(t, \pi_t)dt + \Sigma^{\text{obs},p}(t, \pi_t)dW_t, \quad dB_t^{\text{obs},p} := \mu^{\text{obs},p}(t)dt + \Sigma^{\text{obs},p}(t)dW_t,$$

where

$$\mu^{\text{obs},p}(t, \pi) := \begin{pmatrix} \pi^g \cdot (r^g(t) + \eta^g(t) \circ \sigma^g(t)) + \pi^c \cdot (r^c(t) + \eta^c(t) \circ \sigma^c(t)) + \pi^I \mu^I(t) \\ r^g(t) + \eta^g(t) \circ \sigma^g(t) - (\sigma^g(t))^\top \Sigma^g \sigma^g(t) \\ \mu^I(t) - \frac{(\sigma^I(t))^2}{2} \end{pmatrix},$$

$$\Sigma^{\text{obs},p}(t, \pi) := \begin{pmatrix} (\pi^g \circ \sigma(t)^g)^\top & (\pi^c \circ \sigma(t)^c)^\top & \pi^I \sigma^I(t) \\ \text{diag}(\sigma^g(t)) & \mathbf{0}_{d^g, d^c} & \mathbf{0}_{d^g, 1} \\ \mathbf{0}_{1, d^g} & \mathbf{0}_{1, d^c} & \sigma^I(t) \end{pmatrix},$$

$$\mu^{\text{obs},p}(t) := (r^c(t) + \eta^c(t) \circ \sigma^c(t) - (\sigma^c(t))^\top \Sigma^c \sigma^c(t)), \quad \Sigma^{\text{obs},p}(t) := (\mathbf{0}_{d^c, d^g} \quad \text{diag}(\sigma^c(t)) \quad \mathbf{0}_{d^c, 1}).$$

This leads to minor changes in the computations and the optimal incentives.

5.3.2.2 On the practical implementation of the contract

We will show in the numerical section that the processes (π^*, Z^*, Γ^*) show a rather constant behaviour through the period $[0, T]$. Thus, the optimal contract does not need a frequent re-calibration throughout the year. This suggests the following approximation

$$\begin{aligned} \xi^* &\approx \xi_0^* + \bar{Z}^{\text{obs}} \cdot B_T^{\text{obs}} + \frac{1}{2} \text{Tr} \left[(\bar{\Gamma}^{\text{obs}} + \gamma \bar{Z}^{\text{obs}} (\bar{Z}^{\text{obs}})^\top) \langle B^{\text{obs}} \rangle_T \right] \\ &\quad - \int_0^T h^{\text{obs}} \left(t, \bar{Z}^{\text{obs}}, \bar{\Gamma}^{\text{obs}}, \pi^*(t, \bar{Z}^{\text{obs}}, \bar{\Gamma}^{\text{obs}}) \right) dt, \end{aligned} \quad (5.9)$$

where \bar{Z}^{obs} , and $\bar{\Gamma}^{\text{obs}}$ are constants corresponding the average of $z^{\text{obs}}(t)$, $g^{\text{obs}}(t)$ over $[0, T]$ defined by

$$\begin{aligned} \bar{Z}^{\text{obs}} &= (\bar{Z}_X^{\text{obs}}, \bar{Z}_1^{\text{obs}}, \dots, \bar{Z}_{d^g}^{\text{obs}}, \bar{Z}_I^{\text{obs}})^\top \in \mathbb{R}^{d^g+2}, \\ \bar{\Gamma}^{\text{obs}} &= \begin{pmatrix} \bar{\Gamma}_X^{\text{obs}} & \Gamma_{X,1}^{\text{obs}} & \dots & \Gamma_{X,d^g}^{\text{obs}} & \bar{\Gamma}_{X,I}^{\text{obs}} \\ \bar{\Gamma}_{X,1}^{\text{obs}} & \bar{\Gamma}_1^{\text{obs}} & \dots & \bar{\Gamma}_{1,d^g}^{\text{obs}} & \bar{\Gamma}_{1,I}^{\text{obs}} \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \vdots & \vdots & \vdots & \ddots & \Gamma_{d^g,I}^{\text{obs}} \\ \bar{\Gamma}_{X,I}^{\text{obs}} & \Gamma_{1,I}^{\text{obs}} & \dots & \Gamma_{d^g,I}^{\text{obs}} & \Gamma_I^{\text{obs}} \end{pmatrix} \in \mathbb{S}_{d^g+2}(\mathbb{R}). \end{aligned}$$

In order to provide a practical implementation of the contract, we propose a static replication of its payoff using financial instruments. First, note that the incentives \bar{Z}_X^{obs} , and $\bar{\Gamma}_X^{\text{obs}}$ are indexed on the holdings of the investor, thus do not need any replication using financial instruments. The portion $\bar{Z}^{\text{obs}} \cdot B_T^{\text{obs}}$ of the contract can be easily replicated using log-contracts. For example, for $i \in \{1, \dots, d^g\}$, we replicate $\bar{Z}_i^{\text{obs}} (B_T^{\text{obs}})_i$ using a long position of size Z_i^{obs} on a log-contract on the i -th green bond with maturity T . In this section, all the derivatives products will have a maturity equal to T .

The portion of the contract with respect to quadratic variation and covariation terms are more subtle to replicate. Define the matrix $\tilde{C} \in \mathbb{S}_{d^g+2}(\mathbb{R})$ whose coefficients are given by

$$\tilde{C}_{i,j} := \sum_{k=1}^{d^g+2} C_{i,k} \langle B_{k,j}^{\text{obs}} \rangle_T, \quad C_{i,j} := \bar{\Gamma}_{i,j}^{\text{obs}} + \gamma \bar{Z}_i^{\text{obs}} \bar{Z}_j^{\text{obs}}, \quad (i, j) \in \{1, \dots, d^g + 2\}.$$

Then, we can rewrite $\frac{1}{2} \text{Tr} \left[(\bar{\Gamma}^{\text{obs}} + \gamma \bar{Z}^{\text{obs}} (\bar{Z}^{\text{obs}})^\top) \langle B^{\text{obs}} \rangle_T \right] = \frac{1}{2} \sum_{i=1}^{d^g+2} \tilde{C}_{i,i}$. Following the reasoning of Carr and Lee [64], we note that the quadratic variations and covariations on the logarithm of the green bonds and the index of conventional bonds can be replicated statically using variance and covariance swaps on the bonds. Finally, the portfolio process is equivalent to holding $\pi^{*,g}$ green bonds, $\pi^{*,c}$ conventional bonds and $\pi^{*,I}$ index. Thus, the quadratic covariation between the portfolio process X and the bonds can be replicated using a linear combination of variance and covariance swaps. We are now in position to state the replication strategy for the implementation of the contract. The proof is an application of the no-arbitrage principle and Itô's formula on the logarithm of the bond prices.

Proposition 5.1. *The replication strategy on $[0, T]$ of the optimal contract in (5.9) is as follow:*

- For $i \in \{1, \dots, d^g\}$, hold a position of size \bar{Z}_i^{obs} in a log-contract on the i -th green bond.
- Hold a position of size \bar{Z}_I^{obs} in a log-contract on the index of conventional bonds.
- For $i \in \{2, \dots, d^g + 1\}$, $k \in \{2, \dots, d^g + 1\}$, hold a position of size $\frac{1}{2} C_{i,k}$ in a covariance swap between the $(i-1)$ -th and the $(k-1)$ -th green bonds.
- For $i = d^g + 2$, $k \in \{2, \dots, d^g + 1\}$, hold a position of size $\frac{1}{2} C_{i,k}$ in a covariance swap between the index of conventional bonds and the $(k-1)$ -th green bonds.

- For $i = k = d^g + 2$, hold a position of size $\frac{1}{2}\mathcal{C}_{i,i}$ in a variance swap on the index of conventional bonds.
- For $i = 1, k \in \{2, \dots, d^g + 1\}, l^g \in \{1, \dots, d^g\}, l^c \in \{1, \dots, d^c\}$, hold a position of size $\frac{1}{2}\mathcal{C}_{i,k}\pi_{l^g}^{*,g}$ in a (co)variance swap between the $(k-1)$ -th and the l^g -th green bonds, a position of size $\frac{1}{2}\mathcal{C}_{i,k}\pi_{l^c}^{*,c}$ between the $(k-1)$ -th green bond and the l^c -th conventional bonds, and a position of size $\frac{1}{2}\mathcal{C}_{i,k}\pi^{*,I}$ in a covariance swap between the index of conventional bonds and the $(k-1)$ -th green bond.
- For $i = 1, k = d^g + 2, l^g \in \{1, \dots, d^g\}, l^c \in \{1, \dots, d^c\}$, hold a position of size $\frac{1}{2}\mathcal{C}_{i,k}\pi_{l^g}^{*,g}$ in a covariance swap between the index of conventional bonds and the l^g -th green bond, a position of size $\frac{1}{2}\mathcal{C}_{i,k}\pi_{l^c}^{*,c}$ between the index of conventional bonds and the l^c -th conventional bond, and a position of size $\frac{1}{2}\mathcal{C}_{i,k}\pi^{*,I}$ in a variance swap on the index of conventional bonds.

The contract can be implemented practically only by using the value of the portfolio of bonds, log-contracts, variance and covariance swaps on the different bonds.

Remark 30. We would like to emphasize that, even though it is possible to replicate in practice the optimal contract using variance and covariance swaps on the government bonds, these derivatives might be highly illiquid on financial markets. However, it is possible to replicate these volatility derivatives using the log-contracts and the bonds. Indeed, a variance swap on a bond P_t (we omit to describe the type of bond for notational simplicity) of maturity T can be replicated by holding for all $t \in [0, T]$ one log-contract that pays $-2\log(P_T/P_0)$ and $2/P_t$ bonds P_t . A covariance swap on the bonds P_t^1 , and P_t^2 can be replicated by holding for all $t \in [0, T]$ one log-contract that pays $-2\log(P_T^1/P_0^1)$, one log-contract that pays $-2\log(P_T^2/P_0^2)$, short $\frac{1}{2}$ variance swap on P^1 , and short $\frac{1}{2}$ variance swap on P^2 , long $1/(P_t^1 P_t^2)$ bond $P_t^3 := P_t^1 P_t^2$. Thus, the optimal contract ξ in (5.7) can be implemented only using bond prices and log-contracts.

Finally, note that if vanilla options on the futures on the bonds are available on the market, one can use the Carr-Madan formula, see Carr and Madan [65] to replicate the log-contract payoffs in 30. Thus, the optimal contract in (5.9) can be implemented in practice in three different ways: using the bond prices, the portfolio process, the variance and covariance swaps on the bonds; using the bond prices, the portfolio process, and the log-contracts on the bonds; or using the bond prices, the portfolio process, and vanilla options on the bond prices.

5.4 Numerical results

In the current section, we provide numerical examples illustrating the efficiency of our incentives method.

5.4.1 Data, key results and remarks for the policy-maker

We illustrate our methodology on an example with real-world data. The dataset is composed of 3 French governmental bonds, one green bond and two conventional bonds with the following characteristics.

	Bloomberg Ticker	Valuation date	Maturity	Amount issued	Issue price	Coupon
Green bond	FRTR 1 3/4	24/01/2017	25/06/2039	27.375b	100.162	1.75
Conv. bond 1	FRTR 6	02/01/1994	25/10/2025	30.654b	95.29	6
Conv. bond 2	FRTR 4	09/03/2010	25/04/2060	16.000b	96.34	4

We also define the index of conventional bonds I_t as a geometric average of the conventional bonds, weighted by the amount issued. We perform the calibration using the daily prices of the bonds from 10/04/2019 to 10/04/2020 and the following affine parametrisation for short-term rates, volatilities and risk premiums:

$$\begin{aligned}
r^g(t) &= a^{r,g} + b^{r,g}(T^g - t), \quad \eta^g(t) = a^{\xi,g} + b^{\xi,g}(T^g - t), \quad \sigma^g(t) = a^{\sigma,g} + b^{\sigma,g}(T^g - t), \\
r^{1,c}(t) &= a_1^{r,c} + b_1^{r,c}(T^{1,c} - t), \quad \eta^{1,c}(t) = a_1^{\xi,c} + b_1^{\xi,c}(T^{1,c} - t), \quad \sigma^{1,c}(t) = a_1^{\sigma,c} + b_1^{\sigma,c}(T^{1,c} - t), \\
r^{2,c}(t) &= a_2^{r,c} + b_2^{r,c}(T^{2,c} - t), \quad \eta^{2,c}(t) = a_2^{\xi,c} + b_2^{\xi,c}(T^{2,c} - t), \quad \sigma^{2,c}(t) = a_2^{\sigma,c} + b_2^{\sigma,c}(T^{2,c} - t) \\
\mu^I(t) &= a^{\mu,I} + b^{\mu,I}(T^I - t), \quad \sigma^I(t) = a^{\sigma,I} + b^{\sigma,I}(T^I - t),
\end{aligned}$$

with $T^g = 19.73$, $T^{1,c} = 6.06$, $T^{2,c} = 40.58$, and $T^I = 18.29$.

In order to calibrate the dynamics of the bonds in (5.1) over the period, we use a classic least-square algorithm and we obtain the following set of parameters

$$\begin{aligned} a^{r,g} &= -0.07, & b^{r,g} &= 0.66, & a^{\xi,g} &= 0.38, & b^{\xi,g} &= 0.13, & a^{\sigma,g} &= 0.41, & b^{\sigma,g} &= 0.31, \\ a_1^{r,c} &= -0.05, & b_1^{r,c} &= -0.91, & a_1^{\xi,c} &= 0.01, & b_1^{\xi,c} &= 0.30, & a_1^{\sigma,c} &= 0.11, & b_1^{\sigma,c} &= 0.26, \\ a_2^{r,c} &= 0.28, & b_2^{r,c} &= 0.02, & a_2^{\xi,c} &= 0.12, & b_2^{\xi,c} &= -0.99, & a_2^{\sigma,c} &= 0.10, & b_2^{\sigma,c} &= -0.96, \\ a^{\mu,I} &= -0.01, & b^{\mu,I} &= 0.53, & a^{\sigma,I} &= 0.01, & b^{\sigma,I} &= 0.92, \end{aligned}$$

and the correlation matrix is given by

$$\Sigma = \begin{pmatrix} 1 & 0.2 & 0.8 & 0.8 \\ 0.2 & 1 & 0.2 & 0.7 \\ 0.8 & 0.2 & 1 & 0.7 \\ 0.8 & 0.7 & 0.7 & 1 \end{pmatrix}.$$

The time horizon of the investor and the government is equal to one year, i.e $T = 1$.

We define a so-called reference case, which is a reference to analyze the impact of our incentives policy. In this setting,

$$\nu = \gamma = 1, \quad G = \mathbf{0}_{dg}, \quad \kappa = 0, \quad \beta = (0.4, 0.4, 0.4, 0.4), \quad \alpha = (0.2, 0.2, 0.3, 0.5).$$

Thus, the investor and the government have the same risk aversion, and the government has no specific incentives to increase the investments in the green bond. The only objective of the government is to maximize the value of the portfolio of bonds. The investor has the same cost intensity for every bonds and wishes to invest more in the index and the second conventional bond compared to the green and the first conventional bond. This corresponds to a risk-averse investor who prefers a diversified portfolio of conventional bonds, and is reluctant to invest in the green bonds. Finally, the utility reservation of the investor is set equal to the his utility in the case $\xi = 0$.

We summarize the important empirical findings coming from the numerical results.

- The methodology we propose outperforms significantly the current tax-incentives policy: for a same result in terms of green investments, our methodology leads to a value of the portfolio process 15% to 20% higher.
- The optimal investment policy is robust to model specification: by using a one-factor model on the short-term rates of the green bond, we observe that the investor's strategy oscillates slightly around the one obtained with deterministic rates.
- The optimal controls show a rather constant behaviour throughout the year: The government does not have to frequently recalibrate the optimal contract.
- The government can increase the amount invested in the green bonds by the mean of G and κ . This decreases his utility as he must provides higher incentives to the investor.
- The most important incentive with respect to the contractible variables is Z_X^* : The government always encourage a higher value of the portfolio of bonds by setting $Z_X^* > 0$.
- When the government provides incentives to increase the investment in green bonds, he encourages higher variations of the value of the portfolio in order to compensate the amount given to the investor.
- At the expense of some substantial utility loss, the government can propose a contract indexed only on the contractible variables. This results in a higher incentive Z_X^* .

We also provide some general remarks for the policy-maker.

- The parameters (α, β) modelling the preferences of the investor should be calibrated using the historical data on the issuance of bonds. For example, for $i \in \{1, \dots, d^g\}$, the coefficient α_i should be equal to the historical amount invested in the green bond P_i^g , and β_i should be equal to the variance of the amount invested in this green bond throughout the year. Note however that one historical data on bonds with the same characteristic may not be available especially for countries with small amounts issued. Thus, the parameters (α, β) might be re-scaled depending on the maturity and the coupon of the newly issued bond: A risky investor such as a fixed-income hedge fund might increase his investment in the bond if it offers a higher coupon, whereas institutional investors such as pension funds will tend to buy bonds with a better rating.
- The risk-aversion parameter γ should be chosen such that, in the case $\xi = 0$ and with (α, β) chosen as explained previously, the optimal controls π^* correspond roughly to the historical positions of the investor.
- The risk-aversion parameter ν should be chosen heuristically such that the optimal contract offered to the investor bring the investments closer to the target G and the amount ξ^* offered by the government is reasonable. The terms ‘closer to’ and ‘reasonable’ have to be interpreted by the policy-maker in view of their own budget constraints and political objectives.
- In the case of a small number of bonds issued, the government can, for sake of simplicity, propose a contract indexed only on the value of the portfolio.

5.4.2 Reference case

5.4.2.1 Optimal controls and comparison with the no-contract case

In the absence of a contract, that is $\xi = 0$, the investor matches his investments $\pi^*(\xi)$ with the target α as he has no incentives to deviate. Thus, the optimal investments are given by $\pi^{*g}(0) = 0.2$, $\pi^{*c}(0) = (0.2, 0.3)$, $\pi^{*I}(0) = 0.5$. We can now analyze the influence of the contract on the behaviour of the investor. We first show in 5.1 the evolution of the optimal investment policy π^* and the optimal incentives Z^* , and Γ^* through time. One can see that, even if the risk premia, the short-term rates and volatility processes have a deterministic affine structure with respect to time, the processes (π^*, Z^*, Γ^*) show a rather constant behaviour through the year. Thus, the optimal contract does not need frequent recalibration through the year.

Compared to the case $\xi = 0$, we observe that the contract increases the investment in the green bond and the second conventional bond, while reducing the investment in the index and the first conventional bond. Given the dynamics of the bonds described previously, as well as the preferences of the investor, it is natural that he invests mostly in the index and the second conventional bond. As the green bond has a higher short-term rate and risk premium than the first conventional bond, the traders invests a higher part of his wealth in it.

The optimal incentives with respect to the sources of risk is as follow: the incentives with respect to the green bond and the index of conventional bonds are set to zero, whereas the incentive with respect to the value of the portfolio of bonds is strictly positive. Thus, the government provides incentives only to increase the value of the portfolio. We observe at the bottom of 5.1 the incentives with respect to the quadratic variations of the contractible variables. The government provides no incentives with respect to the quadratic variation of the index and the green bond while it encourages a high quadratic variation of the portfolio process. The incentives with respect to the quadratic covariations are as follow: the government penalizes a high covariation between the portfolio process and the index as well as between the green bonds and the index, while encouraging a high covariation between the portfolio and the green bond.




Figure 5.1: Optimal investment policy (upper left), optimal incentives Z^* (upper right) and Γ^* (bottom) as a function of time.

5.4.2.2 Trajectory simulation and portfolio value

To illustrate the benefits of the use of a contract, we plot in 5.2 some simulations of the evolution of the portfolio process over the year with and without contract (that is when $\xi = 0$). We observe that the portfolio process is higher when the government provides a contract to the investor. This is also illustrated in Figure 5.3 where we show the cumulated difference between the portfolio processes with and without contract, using 10000 simulations.

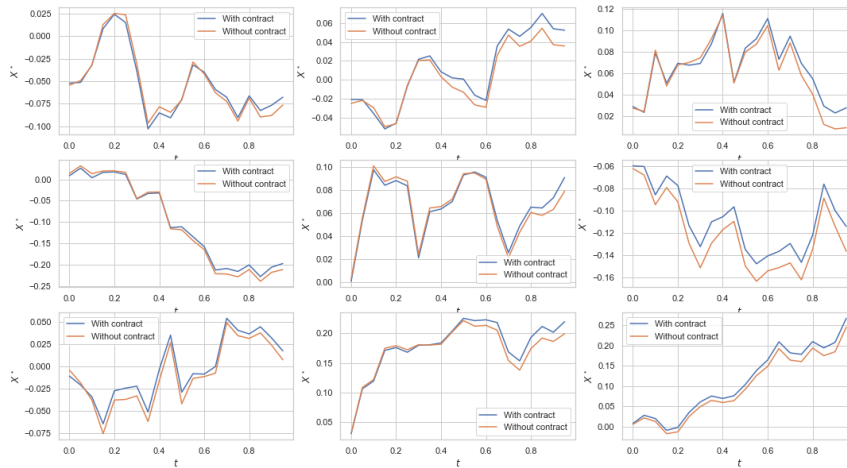


Figure 5.2: Some trajectories of the optimal portfolio process with and without contract.




Figure 5.3: Average absolute difference of portfolio value over time, for 10000 simulations.

5.4.2.3 Optimal contract with no indexation on quadratic variation

As the notion of incentives with respect to quadratic variation might not be easy to understand, we present in 5.4 the optimal investment and incentives when the government set $\Gamma = 0$.




Figure 5.4: Optimal investment policy (left) and optimal incentives Z^* (right) as a function of time.

Compared to 5.1, we observe that the government sets a higher incentive on the value of the portfolio, while the optimal investment policy is slightly higher on every asset, but not materially different compared to a framework with an optimal contract depending on both the dynamics and the quadratic variations of the contractible variables. Thus, for sake of simplicity, a government can build an optimal incentives scheme based only on the dynamics of the green bonds, the value of the portfolio and the index of conventional bonds.

5.4.2.4 Model robustness

We show that, using a more complex model for the short-term rates of the green bonds, the results are qualitatively the same. Using the methodology in 5.A.3, we assume that the short-term rate of the green bond is driven by a one-factor stochastic model, that is

$$dr_t^g = \theta^g(m^g - r_t^g)dt + \sigma^g dW_t^{g,r}, \quad (5.10)$$

where $W^{g,r}$ is a one-dimensional Brownian and $(\theta^g, m^g, \sigma^g) \in \mathbb{R}_+^3$. Using a least-square algorithm, a calibration on the short-term rate curve of the green bond gives the following parameters

$$\theta^g = 0.4, \quad m^g = 0.04, \quad \sigma^g = 0.02.$$

We show in 5.5 the optimal investment policy when the short-term rate of the green bond is driven by (5.10). This is obtained by solving the 4-dimensional HJB equation (5.18) using a fully implicit scheme and locally unidimensional methods on sparse grids.¹⁰ Note that the optimisation is much harder to complete since for every $\pi^*(t, z, g, r^g)$ we have to solve a 4-dimensional HJB equation and iterate until we find the optima (z^*, Γ^*) .




Figure 5.5: Optimal investment policy with stochastic rates.

¹⁰In particular, as the bond prices do not vary drastically during the year, we use 10 time steps, 40 space steps for the cash process, 10 for the stochastic rate and 20 for the risk factors of the green bond and index of conventional bonds.

We observe that the optimal policy oscillates around the values obtained in the case of deterministic short-term rates in 5.1. As the bonds are all positively but not perfectly correlated, a change of investment in the green bond induces a change of smaller magnitude in the other bonds. The magnitude of oscillation around the value with deterministic rates is not high, thus we observe same results from a qualitative point of view. As the use of stochastic rates can only be viable for a small portfolio of bonds, and as the difference of behaviour is negligible, we can argue that the use of deterministic short-term rates is more suited to practical applications.

5.4.2.5 Comparison with current tax-incentives policy

The purpose of the chapter is to show that a form of incentives based on the value of the portfolio and the prices of the bonds performs better than the current tax-incentives policy. As stated in the introduction, the incentives policy to increase investment in green bonds takes the form of tax credit or cash rebate, depending on the amount invested. Thus, in our Principal-Agent framework, it takes the following form

$$\xi = c \int_0^T \sum_{i=1}^{d^g} \pi_t^g dt,$$

where $c > 0$ is the amount of cash rebate or tax credit, controlled by the government. We choose c so that the amount invested in green bonds is the same as in 5.1. In 5.6, we plot the average relative difference between the cash processes of the government using our optimal policy and the actual tax-incentives. We observe that the difference increases with time, thus for a same result in terms of green investments our optimal contract increases its utility compared to the actual incentives policy.




Figure 5.6: Average relative difference (in %) of portfolio value over time for 10000 simulations.




Figure 5.7: Some trajectories of the optimal portfolio process with the optimal contract and with the tax-incentives policy (labeled ‘without contract’).

We also show in 5.7 some trajectories of the value of the portfolio process with the optimal contract and the optimal policy. We observe that the value of the portfolio process is always (slightly) higher in the presence of the optimal contract. In the next subsection we show that when the government wants to achieve a specific target in green investments, the difference between the two policies becomes larger.

5.4.3 Influence of the green target

5.4.3.1 Comparison with the reference case

We now study the impact of the incentives policy we propose when the government seeks to achieve a specific investment target in the green bond. We take $G = 3$, $\kappa = 0.8$ and present in 5.8 the new optimal controls of the investor and the government.




Figure 5.8: Optimal investment policy (upper left), optimal incentives Z^* (upper right) and Γ^* (bottom) as a function of time.

The behaviour of the investor is drastically different compared to 5.1. He now invests mostly in the green bond, while increasing the amount invested in the other assets. This comes from the fact that all assets are positively correlated so that the additional amount invested in the index is higher than the one invested in the first conventional bond. The government sets a higher incentive with respect to the value of the portfolio. The incentives with respect to the quadratic variation are now all positive and higher than in 5.1. While Γ_G and Γ_I are still set to zero, the incentive with respect to every covariations are now positive. In particular, Γ_{XI} has changed from -0.3 to 1 meaning that the government encourages a higher quadratic covariation between the portfolio and the index of conventional bonds. So as to maximize the value of the portfolio while giving higher incentives, the government encourages a higher variance of the portfolio process and positive co-variations between the portfolio and the bond prices.




Figure 5.9: Average difference of portfolio value over time, for 10000 simulations.

Note that, while the amount invested in the green bond is higher but not equal to the target of the government. As $\alpha^g = 0.2$, the government has to provide higher incentives to force the investor to shift his preferences toward a much higher investment in the green bond. As in the reference case, we show in 5.9 some simulations of the evolution of the portfolio process compared to the case without contract. We observe that the higher investment in green bonds leads to a higher average value of the portfolio process. Moreover due to the higher incentives on the quadratic variations, the portfolio process with the contract is more volatile, as it can be seen in 5.10.




Figure 5.10: Some trajectories of the optimal portfolio process with and without contract.

5.4.3.2 Comparison with the tax-incentives policy

We have seen in Figures 5.6 and 5.7 that without specific target in green investments, the optimal contract we propose leads to a higher value of the portfolio process compared to the tax-incentives policy. Here, we set the tax-incentives c so that the investor matches the investment in green bonds obtained with the optimal contract in 5.8.




Figure 5.11: Average relative difference (in %) of portfolio value over time (with the optimal contract and tax-incentives policy), for 10000 simulations.

We plot in 5.11, and 5.12 some trajectories and the average relative difference of cash processes obtained with the optimal contract and the tax-incentives policy. In this case, the relative differences of value are much higher compared to 5.6, and 5.7. Thus, if the government has a specific investment target in green bonds, the use of the optimal contract we propose guarantees a much higher value of the portfolio for a similar result than the tax-incentives policy.




Figure 5.12: Some trajectories of the optimal portfolio process with the optimal contract and with the tax-incentives policy (labeled ‘without contract’).

5.4.4 Sensitivity analysis

5.4.4.1 Influence of G , and κ

In 5.13, we show that reducing the value of κ makes the government target harder to achieve. In particular, we observe that the amount invested in all the assets has been reduced and especially the amount invested in the green bond. In this case, the government proposes a much higher incentive with respect to the dynamics of the portfolio compared to 5.8: as the investment target G is less important (because of a lower κ , he aims at maximising the value of the portfolio. Moreover, a high quadratic covariation between the green bond and the index is now penalized, while a high variance of the portfolio is encouraged in order to maximize its value.




Figure 5.13: Optimal investment policy (upper left), optimal incentives Z^* (upper right) and Γ^* (bottom) as a function of time.

In 5.14, we show that with the parameters $\kappa = 0.8$, $G = 1$, the investment target of the government can be reached more easily. In this case, the trader invest roughly the same amount in the the green bond and the second conventional bond. The government increases the incentive corresponding to the value of the portfolio compared to 5.8. Moreover, he encourages a high variance of the portfolio process while keeping the incentives Γ_G , Γ_I equal to zero.




Figure 5.14: Optimal investment policy (upper left), optimal incentives Z^* (upper right) and Γ^* (bottom) as a function of time.

5.4.4.2 Influence of α and β

We studied in the previous section the influence of the government's parameters, that is the target G and the cost intensity κ . We now show the influence of the targets $\alpha^g, \alpha^c, \alpha^I$ and the cost intensities $\beta^g, \beta^c, \beta^I$ of the investor. In 5.15, we place ourselves in the context of the reference case of 5.1, except that we set $\alpha^g = 0$. This means that the investor is not willing to put money in the green bond. Compared to 5.1, we see that in the absence of specific incentives for green investing, the investor effectively sets π^g equal to zero.




Figure 5.15: Optimal investment policy (upper left), optimal incentives Z^* (upper right) and Γ^* (bottom) as a function of time.

The other investment policies are slightly changed, as there is now more investment in the second conventional bond than in the index. As neither the government nor the investor are interested in the green bond, the government provides higher incentives Z_X in order to maximize the value of the portfolio. The incentive Γ_{XG} become negative while Γ_X becomes positive meaning that the government encourages opposite moves between the price of the green bond and the portfolio process. Moreover, Γ_{XI} becomes positive: the government encourages similar moves between the price of the index and the portfolio process. Finally, the incentives corresponding to the quadratic variation of the green bond and the index remain equal to zero.

In 5.16, we compare these results with the case $G = 3, \kappa = 0.8$ in order to show the influence of the contract when the investor and the government have very different investment targets. We observe that the amount invested in the green bond is clearly higher than in Figure 5.15 where the investor has $\alpha^g = 0$ but lower than in Figure 5.8 where the investor has $\alpha^g = 0.2$. The incentives with respect to the quadratic variations become positive meaning that the government encourages similar moves of all the contractible variables. In particular, compared to Figure 5.15, the government gives higher incentives toward similar moves of the portfolio value and the green bond.




Figure 5.16: Optimal investment policy (upper left), optimal incentives Z^* (upper right) and Γ^* (bottom) as a function of time.

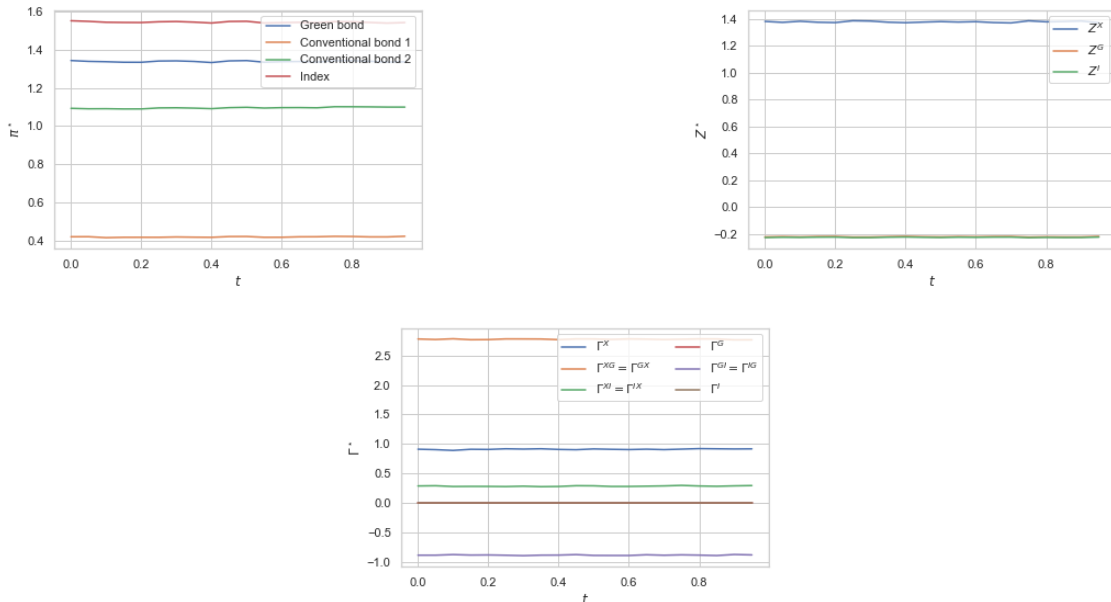


Figure 5.17: Optimal investment policy (upper left), optimal incentives Z^* (upper right) and Γ^* (bottom) as a function of time.

We conclude this section by showing in 5.17 the influence of the cost intensity. We take the same parameters as in 5.16 except that we set $\beta^g = 0.5$. As the intensity cost for moving the green bond target of the investor is higher than in 5.16, the optimal investment policy in the green bond is lower. The government sets a higher incentive Z_X to encourage a higher value of the portfolio. The incentives

with respect to quadratic variations are materially different compared to Figure 5.16. In particular, the government encourages opposite moves between the green bond and the index.

5.A Appendix

5.A.1 Weak formulation of the problem

We work on the canonical space \mathcal{Q} of continuous functions on $[0, T]$ with Borel algebra \mathcal{F} . The $(d^g + d^c + 2)$ -dimensional canonical process is

$$\mathcal{B} := \begin{pmatrix} X \\ W^g \\ W^c \\ W^I \end{pmatrix}$$

and $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$ is its natural filtration. We define \mathbb{P}_0 as the $(d^g + d^c + 1)$ -dimensional Wiener measure on \mathcal{Q} . Thus, \mathcal{B} is a $(d^g + d^c + 2)$ -dimensional Brownian motion where (W^g, W^c, W^I) has a correlation matrix Σ under \mathbb{P}_0 . We also define $\mathcal{M}(\Omega)$ as the set of probability measures on $(\mathcal{Q}, \mathcal{F}_T)$ and

$$\mathbb{H}^2(\mathbb{P}_0) := \left\{ (\pi_t)_{t \in [0, T]} : \mathcal{B}\text{-valued, } \mathbb{F}\text{-predictable processes such that } \mathbb{E}^{\mathbb{P}} \left[\int_0^T \|\pi_t\|_2^2 dt \right] < +\infty \right\}.$$

We consider the following family of processes, indexed by $\pi \in \mathbb{H}^2(\mathbb{P}_0)$

$$\mathcal{X}_t^\pi := \begin{pmatrix} \int_0^t \Sigma^{\text{obs}}(s, \pi_s) d\mathcal{B}_s \\ \int_0^t \Sigma^{\text{obs}} d\mathcal{B}_s \end{pmatrix},$$

and define the set \mathcal{P}_m as the set of probability measures $\mathbb{P}^\pi \in \mathcal{M}(\mathcal{Q})$ of the form

$$\mathbb{P}^\pi = \mathbb{P}_0 \circ (\mathcal{X}^\pi)^{-1}, \text{ for all } \pi \in \mathbb{H}^2(\mathbb{P}_0).$$

Thanks to Bichteler [46], we can define a pathwise version of the quadratic variation process $\langle \mathcal{B} \rangle$ and of its density process with respect to the Lebesgue measure $\hat{\alpha}_t := \frac{d\langle \mathcal{B} \rangle_t}{dt}$. As the processes $\pi \in \mathcal{A} \subset \mathbb{H}^2(\mathbb{P}_0)$ have all their coordinates strictly positive, the volatility of \mathcal{B} is invertible, which implies in particular that the process $W_t = \int_0^t \hat{\alpha}_s^{-\frac{1}{2}} d\mathcal{B}_s$ is an $\mathbb{R}^{d^g + d^c + 2}$ -valued, \mathbb{P} -Brownian motion with correlation matrix Σ for every $\mathbb{P} \in \mathcal{P}_m$. According to Soner, Touzi, and Zhang [255], there exists an $\mathbb{F}^{\mathcal{B}}$ -progressively measurable mapping $\beta_\pi : [0, T] \times \mathcal{Q} \rightarrow \mathbb{R}^{d^g + d^c + 2}$ such that

$$\mathcal{B} = \beta_\pi(\mathcal{X}^\pi), \mathbb{P}_0\text{-a.s.}, W = \beta_\pi(\mathcal{B}), \mathbb{P}^\pi\text{-a.s.}, \hat{\alpha}(\mathcal{B}) = \pi(\beta_\pi(\mathcal{B})), dt \otimes d\mathbb{P}^\pi\text{-a.e.}$$

In particular, the canonical process \mathcal{B} admits the following dynamics for all $\pi \in \mathcal{A}$

$$\mathcal{B}_t = \begin{pmatrix} \int_0^t \Sigma^{\text{obs}}(s, \pi(W \cdot)) dW_s \\ \int_0^t \Sigma^{\text{obs}} dW_s \end{pmatrix}, \mathbb{P}^\pi\text{-a.s.}$$

The first coordinate of the canonical process is the desired output process, the d^g next coordinates are the contractible sources of risk, that is the d^g green bonds and the index of conventional bond, and the last d^c coordinates are the non-contractible sources of risk. Then, we can introduce easily the drift of the output process by the means of Girsanov theorem. Denote

$$\frac{d\mathbb{Q}}{d\mathbb{P}^\pi} := \mathcal{E} \left(\int_0^\cdot \tilde{\Sigma}(s) dW_s \right)_T,$$

a change of measure independent of the control process π , where $\tilde{\Sigma} : [0, T] \longrightarrow \mathcal{M}_{d^g+d^c+2}(\mathbb{R})$ is such that

$$\tilde{\Sigma}(t) := \begin{pmatrix} \left(\frac{r^g(t) + \eta^g(t) \circ \sigma^g(t)}{\sigma^g(t)} \right)^\top & \left(\frac{r^c(t) + \eta^c(t) \circ \sigma^c(t)}{\sigma^c(t)} \right)^\top & \frac{\mu^I(t)}{\sigma^I(t)} \\ \mathbf{0}_{d^g+d^c+1, d^g} & \mathbf{0}_{d^g+d^c+1, d^c} & \mathbf{0}_{d^g+d^c+1, 1} \end{pmatrix}.$$

We finally obtain the desired dynamics for the output process and the $d^g + d^c + 1$ sources of risk.

5.A.2 Proof of Theorem 5.1

We can define the functions $\sigma : [0, T] \times K \longrightarrow \mathcal{M}_{d^g+d^c+2, d^g+d^c+1}(\mathbb{R})$, $\lambda : [0, T] \longrightarrow \mathbb{R}^{d^g+d^c+1}$ such that the set of contractible variables $(B_t)_{t \in [0, T]}$ can be rewritten for all $\pi \in \mathcal{A}$ as

$$dB_t = \sigma(t, \pi_t)(\lambda(t)dt + dW_t), \quad (5.11)$$

where for all $(t, p) \in [0, T] \times K$,

$$\sigma(t, p) := \begin{pmatrix} (p^g \sigma^g(t))^\top & (p^c \sigma^c(t))^\top & p^I \sigma^I(t) \\ \text{diag}(\sigma^g(t)) & \mathbf{0}_{d^g, d^c} & \mathbf{0}_{d^g, 1} \\ \mathbf{0}_{1, d^g} & \mathbf{0}_{1, d^c} & \sigma^I(t) \\ \mathbf{0}_{d^c, d^g} & \text{diag}(\sigma^c(t)) & \mathbf{0}_{d^c, 1} \end{pmatrix},$$

$$\lambda(t) := \left(\left(\frac{r^g(t) + \eta^g(t) \circ \sigma^g(t)}{\sigma^g(t)} \right)^\top \quad \left(\frac{r^c(t) + \eta^c(t) \circ \sigma^c(t)}{\sigma^c(t)} \right)^\top \quad \frac{\mu^I(t)}{\sigma^I(t)} \right)^\top,$$

Thanks to 5.1 and the definition of \mathcal{A} , the functions σ , and λ are bounded. As the function $\sigma(t, \pi)$ is continuous in time for some constant control process $\pi \in \mathcal{A}$, there always exists a weak solution to (5.11). Thanks to the boundedness of the function λ , we can use Girsanov's theorem which guarantees that every $\pi \in \mathcal{A}$ induces a weak solution for

$$B_t = B_0 + \int_0^t \sigma(s, \pi_s) dW'_s, \quad \frac{d\mathbb{P}'}{d\mathbb{P}} \Big|_{\mathcal{F}_T} = \mathcal{E} \left(\int_0^\cdot \lambda(s) \cdot dW_s \right)_T,$$

where W' is a \mathbb{P}' -Brownian motion.

The cost function $k : K \longrightarrow \mathbb{R}$ is measurable and bounded by boundedness of the elements of K . We introduce the norms

$$\|Z^e\|_{\mathbb{H}^p}^p = \sup_{\pi \in \mathcal{A}} \mathbb{E}^\pi \left[\left(\int_0^T |\tilde{\sigma}(t, \pi_t) Z_t|^2 dt \right)^{p/2} \right], \quad \|Y^e\|_{\mathbb{D}^p}^p = \sup_{\pi \in \mathcal{A}} \mathbb{E}^\pi \left[\sup_{t \in [0, T]} |Y_t|^p \right],$$

for any \mathbb{F} -predictable, $\mathbb{R}^{d^g+d^c+2}$ -valued process Z^e and \mathbb{R} -valued process Y^e , and for all $(t, p) \in [0, T] \times K$ $\tilde{\sigma} : [0, T] \times K \longrightarrow \mathcal{M}_{d^g+d^c+2}(\mathbb{R})$ is such that

$$\tilde{\sigma}^2(t, \pi_t) = \sigma(t, p) \sigma^\top(t, p).$$

We also define the functions $H^e : [0, T] \times \mathbb{R}^{d^g+d^c+2} \times \mathbb{S}_{d^g+d^c+2}(\mathbb{R}) \times \mathbb{R} \longrightarrow \mathbb{R}$ and $h^e : [0, T] \times \mathbb{R}^{d^g+d^c+2} \times \mathbb{S}_{d^g+d^c+2}(\mathbb{R}) \times \mathbb{R} \times K \longrightarrow \mathbb{R}$ as

$$H^e(t, z, g, y) := \sup_{p \in K} h^e(t, z, g, y, p)$$

$$h^e(t, z, g, y, p) := -\gamma k(p)y + z \cdot \sigma(t, p) \lambda(t) + \frac{1}{2} \text{Tr} \left[g \sigma(t, p) \Sigma(\sigma(t, p))^\top \right].$$

We introduce the set of so-called admissible incentives \mathcal{ZG}^e as the set of \mathbb{F} -predictable processes (Z^e, Γ^e) valued in $\mathbb{R}^{d^g+d^c+2} \times \mathbb{S}_{d^g+d^c+2}(\mathbb{R})$ such that

$$\|Z^e\|_{\mathbb{H}^p}^p + \|Y^{e, Z^e, \Gamma^e}\|_{\mathbb{D}^p}^p < +\infty, \quad (5.12)$$

for some $p > 1$ where for $y_0^e \in \mathbb{R}$,

$$Y_t^{e,y_0^e,Z^e,\Gamma^e} := y_0^e + \int_0^t Z_s^e dB_s + \frac{1}{2} \text{Tr} \left[\Gamma_s^e d\langle B \rangle_s \right] - H^e(s, Z_s^e, \Gamma_s^e, Y_s^{e,y_0^e,Z^e,\Gamma^e}) ds.$$

Condition (5.12) guarantees that the process $(Y_t^{e,y_0^e,Z^e,\Gamma^e})_{t \in [0,T]}$ is well defined: provided that the right-hand side integrals are well defined, and by noting that H^e is Lipschitz in its last variable (since the cost function k is bounded), $(Y_t^{e,y_0^e,Z^e,\Gamma^e})_{t \in [0,T]}$ is the unique solution of an ODE with random coefficient. Moreover, as K is a compact set and h^e is continuous with respect to its last variable, the supremum with respect to p is always attained. As $(Z^e, \Gamma^e) = (\mathbf{0}_{d^g+d^c+2}, \mathbf{0}_{d^g+d^c+2, d^g+d^c+2}) \in \mathcal{ZG}^e$, this set is non-empty and we are in the setting of Cvitanić, Possamaï, and Touzi [89]. Using [89, Proposition 3.3 and Theorem 3.6], we obtain that without reducing the utility of the Principal, any admissible contract admits the representation

$$U_A(\xi) = Y_T^{e,Z^e,\Gamma^e},$$

Define for all $t \in [0, T]$ the processes

$$Z_t := -\frac{Z_t^e}{\gamma Y_t^{e,y_0^e,Z^e,\Gamma^e}}, \quad \Gamma_t := -\frac{\Gamma_t^e}{\gamma Y_t^{e,y_0^e,Z^e,\Gamma^e}},$$

and

$$Y_t^{y_0,Z,\Gamma} = y_0 + \int_0^t Z_s dB_s + \frac{1}{2} \text{Tr} \left[(\Gamma_s + \gamma Z_s Z_s^\top) d\langle B \rangle_s \right] - H(s, Z_s, \Gamma_s) ds,$$

where $H : [0, T] \times \mathbb{R}^{d^g+d^c+2} \times \mathbb{S}_{d^g+d^c+2}(\mathbb{R}) \rightarrow \mathbb{R}$ is defined by $H(t, z, g) = \sup_{p \in K} h(t, z, g, p)$ and

$$\mathcal{ZG} := \left\{ (Z_t, \Gamma_t)_{t \in [0,T]} : \mathbb{R}^{d^g+d^c+2} \times \mathbb{S}_{d^g+d^c+2}(\mathbb{R})\text{-valued, } \mathbb{F}\text{-predictable processes s.t.} \right. \\ \left. \left(-\gamma Z_t U_A(Y_t^{y_0,Z,\Gamma}), -\gamma \Gamma_t U_A(Y_t^{y_0,Z,\Gamma}) \right)_{t \in [0,T]} \in \mathcal{ZG}^e \right\}. \quad (5.13)$$

An application of Itô's formula leads to $\xi = Y_T^{y_0,Z,\Gamma}$. Thus, we obtain the desired representation for admissible contracts and $V_A(Y_T^{y_0,Z,\Gamma}) = U_A(y_0)$. The characterisation of $\mathcal{A}(Y_T^{y_0,Z,\Gamma})$ is a direct consequence of [89, Proposition 3.3].

5.A.3 Green investments with stochastic interest rates

5.A.3.1 Framework

In the chapter, we considered a deterministic structure for the short-term rates. However, this omits some important stylized facts of the yield curve. In this section we show that at the expense of the use of stochastic control, the government can provide incentives based on short-term rates following a one factor stochastic model.

We now assume that the vectors of short rate dynamics of the green bonds are given by

$$dr_t^g := a^g(t, r_t^g) dt + \text{diag}(b^g) dW_t^{g,r}, \quad (5.14)$$

where $b^g \in \mathbb{R}_+^{d^g}$, $a^g : [0, T] \times \mathbb{R}^{d^g} \rightarrow \mathbb{R}^{d^g}$ and $W^{g,r}$ is a d^g -dimensional Brownian motion of correlation matrix $\Sigma^{g,r}$.

Remark 31. *For notational simplicity, we assume no dependence between the risk sources of the short-term rates and the ones of the bonds. Allowing such dependence is straightforward and does not lead to a higher dimension of the control problem.*

We contract only on the portfolio process, the risk factors of the green bonds and of the stochastic short-term rate of the green bonds, and the risk factor of the index of conventional bonds. The new sets of state variables are

$$B^{\text{obs},S} = \begin{pmatrix} X \\ W^g \\ r^g \\ W^I \end{pmatrix}, \quad B^{\text{obs},S} = W^c,$$

where the superscript S stands for stochastic, which can be written as

$$dB_t^{\text{obs},S} := \mu^{\text{obs},S}(t, \pi_t, r_t^g)dt + \Sigma^{\text{obs},S}(t, \pi_t)dW_t, \quad dB_t^{\text{obs},S} := \mu^{\text{obs},S}(t)dt + \Sigma^{\text{obs},S}(t)dW_t,$$

where for all $t \in [0, T], p = (p^g, p^c, p^I) \in \mathbb{R}^{d^g} \times \mathbb{R}^{d^c} \times \mathbb{R}, r^g \in \mathbb{R}^{d^g}$

$$W_t := \begin{pmatrix} W_t^g \\ W_t^{g,r} \\ W_t^I \\ W_t^c \end{pmatrix}, \quad \mu^{\text{obs},S}(t, p, r^g) := \begin{pmatrix} p^g \cdot (r^g + \eta^g(t) \circ \sigma^g(t)) + p^c \cdot (r^c(t) + \eta^c(t) \circ \sigma^c(t)) + p^I \mu^I(t) \\ \mathbf{0}_{d^g,1} \\ a^g(t, r^g) \\ 0 \end{pmatrix},$$

$$\Sigma^{\text{obs},S}(t, p) := \begin{pmatrix} (p^g \circ \sigma(t)^g)^\top & \mathbf{0}_{1,d^g} & p^I \sigma^I(t) & (p^c \circ \sigma(t)^c)^\top \\ I_{d^g} & \mathbf{0}_{d^g,d^g} & \mathbf{0}_{d^g,1} & \mathbf{0}_{d^g,d^c} \\ \mathbf{0}_{d^g,d^g} & \text{diag}(b^g) & \mathbf{0}_{d^g,1} & \mathbf{0}_{d^g,d^c} \\ \mathbf{0}_{1,d^g} & \mathbf{0}_{1,d^g} & 1 & \mathbf{0}_{1,d^c} \end{pmatrix},$$

$$\mu^{\text{obs},S}(t) = (\mathbf{0}_{d^c,1}), \quad \Sigma^{\text{obs},S}(t) = (\mathbf{0}_{d^c,d^g} \quad I_{d^c} \quad \mathbf{0}_{d^c,1} \quad \mathbf{0}_{d^c,d^g}).$$

We now specify the new set of admissible contracts that we consider for the incentives proposed by the government.

5.A.3.2 Representation of admissible contracts

Define \mathcal{C}^S as the set of admissible contracts in the case of stochastic short-term rates (the admissibility conditions are the same as for the set \mathcal{C}) and for any $\pi \in \mathcal{A}$ we introduce the following quantities

$$B^S := \begin{pmatrix} B^{\text{obs},S} \\ B^{\text{obs},S} \end{pmatrix}, \quad \mu^S(t, \pi) := \begin{pmatrix} \mu^{\text{obs},S}(t, \pi) \\ \mu^{\text{obs},S} \end{pmatrix}, \quad \Sigma^S(t, \pi) := \begin{pmatrix} \Sigma^{\text{obs},S}(t, \pi) \\ \Sigma^{\text{obs},S} \end{pmatrix}.$$

We define $h^S : [0, T] \times \mathbb{R}^{2d^g+d^c+2} \times \mathbb{S}_{2d^g+d^c+2}(\mathbb{R}) \times \mathbb{R}^{d^g} \times K \rightarrow \mathbb{R}$ such that

$$h^S(t, z, g, r^g, p) = -k(p) + z \cdot \mu^S(t, p, r^g) + \frac{1}{2} \text{Tr} \left[g \Sigma^S(t, p) \Sigma(\Sigma^S(t, p))^\top \right],$$

and for all $(t, z, g, r^g) \in [0, T] \times \mathbb{R}^{2d^g+d^c+2} \times \mathbb{S}_{2d^g+d^c+2}(\mathbb{R})$ we define

$$\mathcal{O}^S(t, z, g, r^g) := \left\{ \hat{p} \in K : \hat{p} \in \underset{p \in K}{\text{argmax}} h^S(t, z, g, r^g, p) \right\},$$

as the set of maximizers of h^S with respect to its last variable for (t, z, g, r^g) fixed. Following Schäl [245], there exists at least one Borel-measurable map $\hat{\pi} : [0, T] \times \mathbb{R}^{2d^g+d^c+2} \times \mathbb{S}_{2d^g+d^c+2}(\mathbb{R}) \times \mathbb{R}^{d^g} \rightarrow K$ such that for every $(t, z, g, r^g) \in [0, T] \times \mathbb{R}^{2d^g+d^c+2} \times \mathbb{S}_{2d^g+d^c+2}(\mathbb{R}) \times \mathbb{R}^{d^g}$, $\hat{\pi}(t, z, g, r^g) \in \mathcal{O}^S(t, z, g, r^g)$. We denote by \mathcal{O}^S the set of all such maps. By analogy with 5.1, the following theorem states the form of any admissible contracts in this setting.

Theorem 5.3. *Without reducing the utility of the Principal, we can restrict the study of admissible contracts to the set \mathcal{C}_1^S where any $\xi \in \mathcal{C}_1^S$ is of the form $\xi = Y_T^{y_0, Z^S, \Gamma^S, \hat{\pi}}$ where for all $t \in [0, T]$,*

$$Y_t^{y_0, Z^S, \Gamma^S, \hat{\pi}} := y_0 + \int_0^t Z_s^S \cdot dB_s + \frac{1}{2} \text{Tr} \left[(\Gamma_s^S + \gamma Z_s^S (Z_s^S)^\top) d\langle B^S \rangle_s \right] - h^S(s, Z_s^S, \Gamma_s^S, r_s^g, \hat{\pi}(s, Z_s, \Gamma_s, r_s^g)) ds, \quad (5.15)$$

where $\hat{\pi} \in \mathcal{O}^S$ and $(Z_t^S)_{t \in [0, T]}, (\Gamma_t^S)_{t \in [0, T]}$ are respectively $\mathbb{R}^{2d^g+2+d^c}$ and $\mathbb{S}_{2d^g+2+d^c}(\mathbb{R})$ -valued, \mathbb{F} -predictable processes satisfying similar conditions as the elements of \mathcal{ZG} . We denote the set of admissible incentives as \mathcal{ZG}^S . Moreover in the present case of stochastic rates for green bonds

$$V^A(Y_T^{y_0, Z^S, \Gamma^S, \hat{\pi}}) = U_A(y_0), \mathcal{A}(Y_T^{y_0, Z^S, \Gamma^S, \hat{\pi}}) = \left\{ (\hat{\pi}(t, Z_t^S, \Gamma_t^S, r_t^g))_{t \in [0, T]}, \hat{\pi} \in \mathcal{O}^S, (Z_t^S, \Gamma_t^S)_{t \in [0, T]} \in \mathcal{ZG}^S \right\}.$$

We now set

$$Z_t^S = \begin{pmatrix} Z_t^{\text{obs}, S} \\ Z_t^{\text{obs}, S} \end{pmatrix}, \Gamma_t^S = \begin{pmatrix} \Gamma_t^{\text{obs}, S} & \Gamma_t^{\text{obs}, \text{obs}, S} \\ \Gamma_t^{\text{obs}, \text{obs}, S} & \Gamma_t^{\text{obs}, S} \end{pmatrix},$$

where for all $t \in [0, T]$

$$Z_t^{\text{obs}, S} \in \mathbb{R}^{2d^g+2}, Z_t^{\text{obs}, S} \in \mathbb{R}^{d^c}, \Gamma_t^{\text{obs}, S} \in \mathbb{S}_{2d^g+2}(\mathbb{R}), \Gamma_t^{\text{obs}, S} \in \mathbb{S}_{d^c}(\mathbb{R}), \Gamma_t^{\text{obs}, \text{obs}, S} \in \mathcal{M}_{2d^g+2, d^c}(\mathbb{R}).$$

We define $h^{\text{obs}, S} : [0, T] \times \mathbb{R}^{2d^g+2} \times \mathbb{S}_{2d^g+2}(\mathbb{R}) \times \mathbb{R}^{d^g} \times K \rightarrow \mathbb{R}$ such that

$$h^{\text{obs}, S}(t, z^{\text{obs}, S}, g^{\text{obs}, S}, r^g, p) = -k(p) + z^{\text{obs}, S} \cdot \mu^{\text{obs}, S}(t, p, r^g) + \frac{1}{2} \text{Tr} \left[g^{\text{obs}, S} \Sigma^{\text{obs}, S}(t, p) \Sigma(\Sigma^{\text{obs}, S}(t, p))^\top \right],$$

and for all $(t, z^{\text{obs}, S}, g^{\text{obs}, S}, r^g) \in [0, T] \times \mathbb{R}^{2d^g+2} \times \mathbb{S}_{2d^g+2}(\mathbb{R}) \times \mathbb{R}^{d^g}$ we define

$$\mathcal{O}^{\text{obs}, S}(t, z^{\text{obs}, S}, g^{\text{obs}, S}, r^g) := \left\{ \hat{p} \in K : \hat{p} \in \underset{p \in K}{\text{argmax}} h^{\text{obs}, S}(t, z^{\text{obs}, S}, g^{\text{obs}, S}, r^g, p) \right\}.$$

Using again Schäl [245], there exists at least one Borel-measurable map $\hat{\pi} : [0, T] \times \mathbb{R}^{2d^g+2} \times \mathbb{S}_{2d^g+2}(\mathbb{R}) \times \mathbb{R}^{d^g} \rightarrow B$ such that for every $(t, z^{\text{obs}, S}, g^{\text{obs}, S}, r^g) \in [0, T] \times \mathbb{R}^{2d^g+2} \times \mathbb{S}_{2d^g+2}(\mathbb{R}) \times \mathbb{R}^{d^g}$, we have $\hat{\pi}(t, z^{\text{obs}, S}, g^{\text{obs}, S}, r^g) \in \mathcal{O}^{\text{obs}, S}(t, z^{\text{obs}, S}, g^{\text{obs}, S}, r^g)$ and $\mathcal{O}^{\text{obs}, S}$ denotes the set of all such maps. We consider the subset of admissible contracts

$$\mathcal{C}_2^S := \left\{ Y_T^{y_0, Z^S, \Gamma^S, \hat{\pi}} \in \mathcal{C}_1^S : Z^{\text{obs}, S} = \mathbf{0}_{d^c}, \Gamma^{\text{obs}, S} = \mathbf{0}_{d^c, d^c}, \Gamma^{\text{obs}, \text{obs}, S} = \mathbf{0}_{2d^g+2, d^c} \right\} \subset \mathcal{C}_1^S \subset \mathcal{C}^S,$$

where any contract in \mathcal{C}_2^S is of the form $Y_T^{y_0, Z^{\text{obs}, S}, \Gamma^{\text{obs}, S}, \hat{\pi}}$ where for all $t \in [0, T]$,

$$Y_t^{y_0, Z^{\text{obs}, S}, \Gamma^{\text{obs}, S}, \hat{\pi}} := y_0 + \int_0^t Z_s^{\text{obs}, S} \cdot dB_s^{\text{obs}, S} + \frac{1}{2} \text{Tr} \left[(\Gamma_s^{\text{obs}, S} + \gamma Z_s^{\text{obs}, S} (Z_s^{\text{obs}, S})^\top) d\langle B^{\text{obs}, S} \rangle_s \right] - h^{\text{obs}, S} \left(s, Z_s^{\text{obs}, S}, \Gamma_s^{\text{obs}, S}, r_s^g, \hat{\pi}(s, Z_s^{\text{obs}, S}, \Gamma_s^{\text{obs}, S}, r_s^g) \right) ds, \quad (5.16)$$

where $y_0 \geq 0$, $\hat{\pi} \in \mathcal{O}^{\text{obs}, S}$ and $(Z^{\text{obs}, S}, \Gamma^{\text{obs}, S}) \in \mathcal{ZG}^{\text{obs}, S}$ with

$$\mathcal{ZG}^{\text{obs}, S} := \left\{ (Z^{\text{obs}, S}, \Gamma^{\text{obs}, S}) : \mathbb{R}^{2d^g+2} \times \mathbb{S}_{2d^g+2}(\mathbb{R})\text{-valued, } \mathbb{F}\text{-predictable s.t. } Y_T^{y_0, Z^{\text{obs}, S}, \Gamma^{\text{obs}, S}, \hat{\pi}} \in \mathcal{C}_2^S \right\}.$$

We can now formulate the stochastic control problem faced by the government.

5.A.3.3 The Hamilton-Jacobi-Bellman equation

Let us define the process $(Q_t^{y_0, Z^{\text{obs}, S}, \Gamma^{\text{obs}, S}, \hat{\pi}})_{t \in [0, T]}$ where for all $(t, y_0, Z^{\text{obs}, S}, \Gamma^{\text{obs}, S}, \hat{\pi}) \in [0, T] \times \mathbb{R} \times \mathcal{ZG}^{\text{obs}, S} \times \mathcal{O}^{\text{obs}, S}$

$$Q_t^{y_0, Z^{\text{obs}, S}, \Gamma^{\text{obs}, S}, \hat{\pi}} := X_t - \int_0^t \sum_{i=1}^{d^g} (G_i - \hat{\pi}_i^g(s, Z_s^{\text{obs}, S}, \Gamma_s^{\text{obs}, S}, r_s^g))^2 ds - Y_t^{y_0, Z^{\text{obs}, S}, \Gamma^{\text{obs}, S}, \hat{\pi}}.$$

The optimisation problem of the government that we consider here is

$$\tilde{V}_0^P = \sup_{y_0 \geq 0} \sup_{(Z^{\text{obs},S}, \Gamma^{\text{obs},S}, \hat{\pi}) \in \mathcal{Z}^{\text{obs},S} \times \mathcal{O}^{\text{obs},S}} \mathbb{E}^{\hat{\pi}(Z^{\text{obs},S}, \Gamma^{\text{obs},S})} \left[- \exp \left(- \nu Q_T^{y_0, Z^{\text{obs},S}, \Gamma^{\text{obs},S}, \hat{\pi}} \right) \right]. \quad (5.17)$$

Due to the presence of state variables in the best response of the Agent, the optimal control of the Principal will no longer be deterministic, and we have to rely on the Hamilton-Jacobi-Bellman formulation of the stochastic control problem. First, we note that the supremum over y_0 is attained by setting $y_0 = 0$. Next, the state variables of the control problem are $(t, B_t^{\text{obs},S}, Q_t^{0, Z^{\text{obs},S}, \Gamma^{\text{obs},S}, \hat{\pi}})$ and as it is standard in control problems with CARA utility function, the last variable can be simplified. Define $P^S = \mathbb{R}^{2d^g+2} \times \mathbb{S}_{2d^g+2}(\mathbb{R})$, and the Hamiltonian $H^{\hat{\pi}} : [0, T] \times P^S \times \mathbb{R} \times P^S \rightarrow \mathbb{R}$

$$\begin{aligned} H^{\hat{\pi}}(t, z, g, u, u_b, u_{bb}) &:= \nu u \left(z \cdot \mu^{\text{obs},S}(t, \hat{\pi}(t, z, g, r^g), r^g) + \sum_{i=1}^{d^g} (G_i - \hat{\pi}_i^g(t, z, g, r^g))^2 \right. \\ &+ \frac{1}{2} \text{Tr} \left[(g + \gamma z z^\top) \Sigma^{\text{obs},S}(t, \hat{\pi}(t, z, g, r^g)) \left(\Sigma^{\text{obs},S}(t, \hat{\pi}(t, z, g, r^g)) \right)^\top \right] - h^{\text{obs},S}(t, z, g, r^g, \hat{\pi}(t, z, g, r^g)) \Big) \\ &+ \frac{1}{2} \nu^2 u \text{Tr} \left[z z^\top \Sigma^{\text{obs},S}(t, \hat{\pi}(t, z, g, r^g)) \left(\Sigma^{\text{obs},S}(t, \hat{\pi}(t, z, g, r^g)) \right)^\top \right] + u_b \cdot \mu^{\text{obs},S}(t, \hat{\pi}(t, z, g, r^g), r^g) \\ &+ \frac{1}{2} \text{Tr} \left[\Sigma^{\text{obs},S}(t, \hat{\pi}(t, z, g, r^g)) \left(\Sigma^{\text{obs},S}(t, \hat{\pi}(t, z, g, r^g)) \right)^\top u_{bb} \right]. \end{aligned}$$

The value function of the control problem of the Principal is solution of the following Hamilton-Jacobi-Bellman equation

$$\begin{cases} \partial_t U(t, b) + \sup_{(z, g, \hat{\pi}) \in P^S \times \mathcal{O}^{\text{obs},S}} H^{\hat{\pi}}(t, z, g, b, U, U_b, U_{bb}) = 0, \\ U(T, b) = -1, \end{cases} \quad (5.18)$$

where $U : [0, T] \times \mathbb{R}^{2d^g+2} \rightarrow \mathbb{R}$ and for all $(i, j) \in \{1, \dots, 2d^g + 2\}$, $(U_b)_i = \partial_{b_i} U$, $(U_{bb})_{i,j} = \partial_{b_i b_j} U$, in the sense that $\tilde{V}_0^P = U(0, b_0)$ where $B_0^{\text{obs},S} = b_0$ and $y_0 = 0$. Thus, the incentives provided to the investor are obtained up to the resolution of a $(2d^g + 2)$ -dimensional HJB equation. Although it provides greater flexibility on the modelling of short-term rates, this approach can only be applied to a small portfolio of bonds using classic numerical schemes on sparse grids.

Part II

Options market-making in high dimension

Chapter 6

Algorithmic market-making for options

6.1 Introduction

The electronification of financial markets started in the seventies with stock exchanges and now affects each and every asset class. For asset classes that are usually traded in a centralized way (stocks, futures, etc.), exchanges and other all-to-all trading platforms – based or not based on limit order books – are now fully automated. For assets that are still traded over the counter (OTC), electronification occurs through the introduction of new platforms, for instance single- and multi-dealer-to-client platforms.

This electronification is associated with a trend towards the automation of the trading process for many players of the financial industry: brokers, banks, but also systematic asset managers who often develop their own execution algorithms. For assets traded in centralized markets, trading automation is nowadays massive. For instance, in the cash equity world, a vast majority of the execution is now carried out using algorithms. For asset traded in dealer markets, the automation of the market-making process has been at the agenda for a few years and more and more banks are developing market-making algorithms for various asset classes (currencies, bonds, etc.).

In the academic literature, many market-making models have been proposed since the eighties. In the early literature on market-making, the two main references are the paper of Ho and Stoll Ho and Stoll [156] and the paper of Grossman and Miller Grossman and Miller [135]. Ho and Stoll introduced indeed a very relevant framework to tackle the main problem faced by market-makers: inventory management. Grossman and Miller, who were more interested in capturing the essence of liquidity, proposed a very simple model with 3 periods that encompassed both market-makers and final customers, enabled to understand what happens at equilibrium, and contributed to the important literature on the price formation process. If the latter paper belongs to a strand of literature that is extremely important to go beyond the simple Walrasian view of markets, it is of little help to build market-making algorithms. The former paper however, after more than 25 years, has paved the way to a recent mathematical literature on algorithmic market-making.

A seminal reference of the new literature on market-making is the paper of Avellaneda and Stoikov Avellaneda and Stoikov [26] who revived the dynamic approach proposed by Ho and Stoll. They indeed showed how the quoting and inventory management problems of market-makers could be addressed using the tools of stochastic optimal control. Since then, many models have been proposed, most of them to tackle the same problem of single-asset market-making as that of Avellaneda and Stoikov. For instance, Guéant, Lehalle, and Fernandez-Tapia provided in Guéant, Lehalle, and Fernandez-Tapia [142] a rigorous analysis of the Avellaneda-Stoikov stochastic optimal control problem and proved that the problem could be simplified into a system of linear ordinary differential equations (ODE) in the case of exponential intensity functions. Cartea, Jaimungal, and coauthors contributed a lot to the literature and added many features to the initial models: market impact, alpha signals, ambiguity aversion, etc.

(see Cartea, Donnelly, and Jaimungal [71], Cartea, Jaimungal, and Penalva [70], Cartea, Jaimungal, and Ricci [69]). They also considered a different objective function: a risk-adjusted expectation instead of a Von Neumann-Morgenstern expected utility.

The models proposed in the above papers all share the same characteristics: (i) they are agnostic with respect to the market structure, but are in fact more adapted to OTC markets or limit order book markets with small tick size, (ii) they only deal with single-asset market-making, and (iii) they do not deal with the market-making of options.

In fact, models have been specifically developed by Guilbaud and Pham (see Guilbaud and Pham [143, 144]) for assets traded through limit order books with large tick size (e.g. most stocks) and for assets traded on platforms with a pro-rata microstructure (e.g. some currency pairs). Interestingly, these models enable the use of aggressive orders by market-makers, which is – surprisingly – a standard behavior on equity markets (see Saliba [241]).

As far as multi-asset market-making is concerned, models have been developed recently to account for the correlation structure between asset price changes. Guéant extended to a multi-asset framework models *à la* Avellaneda-Stoikov and models *à la* Cartea-Jaimungal (see Guéant [137], Guéant [138], and Guéant and Lehalle [139]) and showed that the problem boils down, for general intensity functions, to solving a system of (*a priori* nonlinear) ODEs. The associated question of the numerical methods to approximate the solution of the equations characterizing the optimal quotes of a multi-asset market-maker is addressed in Bergault and Guéant [41] using a factorial approach and in Guéant and Manziuk [140] using reinforcement learning, both with applications to corporate bond markets.

Finally, as far as asset classes are concerned, there have been few attempts to address market-making problems outside of the cash world. market-making models for derivative contracts are indeed intrinsically more complicated because they must account for the strategies on both the market for the underlying asset and the market for the derivatives, and usually for numerous contracts (e.g. options for lots of strikes and maturities). Option market-making is only addressed in a paper by El Aoud and Abergel (see El Aoud and Abergel [106]) and in a paper by Stoikov and Sağlam Stoikov and Sağlam [258]. In the former, the authors consider a single-option market driven by a stochastic volatility model and assume that the position is always Δ -hedged. They provide optimal bid and ask quotes for the option and focus on the risk of model misspecification. In the latter, the authors consider three different settings, but all with only one option: (i) a market-maker in a complete market where continuous trading in the perfectly liquid underlying stock is allowed, (ii) a market-maker who may not trade continuously in the underlying stock, but rather sets bid and ask quotes in the option and the stock, and (iii) a market-maker in an incomplete market with residual risks due to stochastic volatility and overnight jumps.

In this chapter, we consider the case of a market-maker in charge of a book of options whose prices are driven by a stochastic volatility model. We assume that trading in continuous time can be carried out in the underlying asset so that the residual risk is only that of the Vega associated with the inventory. Using a constant-Vega approximation, we show that the problem of an option market-maker boils down to solving a low-dimensional functional equation of the Hamilton-Jacobi-Bellman type that can be tackled numerically using a simple Euler scheme along with interpolation techniques. In particular, in spite of the large number of assets, the market-making problem is tractable.

In Section 6.2, we describe the model and present the optimization problem of the option market-maker. In Section 6.3, we show how that problem can be simplified under the constant-Vega approximation. In particular, we show that solving the high-dimensional stochastic optimal control problem of the market-maker boils down to solving a low-dimensional functional equation. In Section 6.4, we consider the example of a book of options with several strikes and maturities and provide numerical results obtained through interpolation techniques and an explicit Euler scheme.

6.2 Description of the problem

We consider a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ with a filtration $(\mathcal{F}_t)_{t \in \mathbb{R}_+}$ satisfying the usual conditions. Throughout the chapter, we assume that all stochastic processes are defined on $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in \mathbb{R}_+}, \mathbb{P})$.

6.2.1 The market

We consider an asset whose price dynamics is described by a one-factor stochastic volatility model of the form

$$\begin{cases} dS_t = \mu S_t dt + \sqrt{\nu_t} S_t dW_t^S, \\ d\nu_t = a_{\mathbb{P}}(t, \nu_t) dt + \xi \sqrt{\nu_t} dW_t^{\nu}, \end{cases}$$

where $\mu \in \mathbb{R}, \xi \in \mathbb{R}^{*+}$, $(W_t^S, W_t^{\nu})_{t \in \mathbb{R}_+}$ is a couple of Brownian motions with quadratic covariation given by $\rho = \frac{d\langle W_t^S, W_t^{\nu} \rangle}{dt} \in (-1, 1)$, and $a_{\mathbb{P}}$ is such that the processes are well defined (in particular, we assume that the process $(\nu_t)_{t \in \mathbb{R}_+}$ stays positive almost surely).

Remark 32. A classical example for the function $a_{\mathbb{P}}$ is that of the Heston model (see Heston [154]), i.e. $a_{\mathbb{P}} : (t, \nu) \mapsto \kappa_{\mathbb{P}}(\theta_{\mathbb{P}} - \nu)$ where $\kappa_{\mathbb{P}}, \theta_{\mathbb{P}} \in \mathbb{R}^+$ satisfy the Feller condition $2\kappa_{\mathbb{P}}\theta_{\mathbb{P}} > \xi^2$.

Remark 33. For the sake of simplicity, we consider throughout this chapter a one-factor model where the instantaneous variance is the main variable of interest. Similar results could be obtained with a one-factor model focused on forward variances, such as the classical one-factor Bergomi model (see Bergomi [43, 45]). Moreover, it is noteworthy that our approach can easily be extended to two-factor stochastic volatility models such as the celebrated two-factor Bergomi model (see Bergomi [43, 45]), up to an increase – by 1 – in the dimension of the equation to solve.

Assuming interest rates are equal to 0, we introduce a risk-neutral / pricing probability measure¹ \mathbb{Q} under which the price and volatility processes become

$$\begin{cases} dS_t = \sqrt{\nu_t} S_t d\widehat{W}_t^S, \\ d\nu_t = a_{\mathbb{Q}}(t, \nu_t) dt + \xi \sqrt{\nu_t} d\widehat{W}_t^{\nu}, \end{cases}$$

where $(\widehat{W}_t^S, \widehat{W}_t^{\nu})_{t \in \mathbb{R}_+}$ is another couple of Brownian motions under \mathbb{Q} with quadratic covariation given by $\rho = \frac{d\langle \widehat{W}_t^S, \widehat{W}_t^{\nu} \rangle}{dt} \in (-1, 1)$, and where $a_{\mathbb{Q}}$ is such that the processes are well defined.

We consider $N \geq 1$ European options written on the above asset (hereafter, the underlying asset). For each $i \in \{1, \dots, N\}$, the option maturity date is denoted by T^i and we denote by $(\mathcal{O}_t^i)_{t \in [0, T^i]}$ the price process associated with the i -th option.

Remark 34. In applications, the options under consideration will always be call and/or put options. However, our setting enables to consider any European payoff.

In the above one-factor model, we know that for all $i \in \{1, \dots, N\}$, and all $t \in [0, T^i]$, $\mathcal{O}_t^i = O^i(t, S_t, \nu_t)$ where O^i is solution on $[0, T^i] \times \mathbb{R}_+^2$ of the following partial differential equation (PDE):

$$\begin{aligned} 0 = & \partial_t O^i(t, S, \nu) + a_{\mathbb{Q}}(t, \nu) \partial_{\nu} O^i(t, S, \nu) \\ & + \frac{1}{2} \nu S^2 \partial_{SS}^2 O^i(t, S, \nu) + \rho \xi \nu S \partial_{\nu S}^2 O^i(t, S, \nu) + \frac{1}{2} \xi^2 \nu \partial_{\nu\nu}^2 O^i(t, S, \nu). \end{aligned} \quad (6.1)$$

Remark 35. Options prices are also characterized by a terminal condition corresponding to the payoff. However, we will only consider short-term optimization problems for which the time horizon is before the maturity of all the options under consideration. Therefore, we shall never use the final condition associated with Equation (6.1).

¹For references, see for instance Gatheral [124].

6.2.2 The optimization problem of the market-maker

We consider an option market-maker in charge of providing bid and ask quotes for the N above options over the period $[0, T]$ where $T < \min_{i \in \{1, \dots, N\}} T^i$ (see Remark 35). For all $i \in \{1, \dots, N\}$ we denote by $\mathcal{O}_t^i - \delta_t^{i,b}(z)$ and $\mathcal{O}_t^i + \delta_t^{i,a}(z)$ the bid and ask prices proposed by the market-maker for a transaction of a number z of i -th options, where $(\delta_t^i(\cdot))_{t \in [0, T]} := (\delta_t^{i,b}(\cdot), \delta_t^{i,a}(\cdot))_{t \in [0, T]}$ is \mathbb{F} -predictable and bounded from below by a given constant δ_∞ .² Hereafter, we denote by \mathcal{A} the set of such admissible control processes. The dynamics of the inventory process $(q_t)_{t \in [0, T]} := (q_t^1, \dots, q_t^N)'_{t \in [0, T]}$ of the market-maker is given by

$$dq_t^i := \int_{\mathbb{R}_+^*} z (N^{i,b}(dt, dz) - N^{i,a}(dt, dz)), \quad i \in \{1, \dots, N\},$$

where $N^{i,b}(dt, dz), N^{i,a}(dt, dz)$ are two right-continuous \mathbb{R}_+^* -marked point processes, with almost surely no simultaneous jumps,³ modelling the transactions of the i -th option on the bid and ask side, whose respective intensity processes $(\lambda_t^{i,b}(dz))_{t \in \mathbb{R}_+}$ and $(\lambda_t^{i,a}(dz))_{t \in \mathbb{R}_+}$ are given by

$$\lambda_t^{i,b}(dz) := \Lambda^{i,b}(\delta_t^{i,b}(z)) \mathbf{1}_{\{q_{t-} + ze^i \in \mathcal{Q}\}} \mu^{i,b}(dz), \quad \lambda_t^{i,a}(dz) := \Lambda^{i,a}(\delta_t^{i,a}(z)) \mathbf{1}_{\{q_{t-} - ze^i \in \mathcal{Q}\}} \mu^{i,a}(dz)$$

with $(e^i)_{i \in \{1, \dots, N\}}$ the canonical basis of \mathbb{R}^N , \mathcal{Q} the set of authorized inventories⁴ for the market-maker, and $\mu^{i,b}, \mu^{i,a}$ two probability measures on \mathbb{R}_+^* , modelling the distributions of transaction sizes. For $i \in \{1, \dots, N\}$, $\Lambda^{i,b}, \Lambda^{i,a}$ are positive functions satisfying the following classical hypotheses (see Guéant [137, 138] for similar assumptions):

- $\Lambda^{i,b}$ and $\Lambda^{i,a}$ are twice continuously differentiable.
- $\Lambda^{i,b}$ and $\Lambda^{i,a}$ are strictly decreasing, with $\Lambda^{i,b'} < 0$ and $\Lambda^{i,a'} < 0$.
- $\lim_{\delta \rightarrow +\infty} \Lambda^{i,b}(\delta) = \lim_{\delta \rightarrow +\infty} \Lambda^{i,a}(\delta) = 0$.
- $\sup_{\delta \in \mathbb{R}} \frac{\Lambda^{i,b}(\delta) \Lambda^{i,b''}(\delta)}{(\Lambda^{i,b'}(\delta))^2} < 2$ and $\sup_{\delta \in \mathbb{R}} \frac{\Lambda^{i,a}(\delta) \Lambda^{i,a''}(\delta)}{(\Lambda^{i,a'}(\delta))^2} < 2$.

The above conditions are sufficiently general to allow for several relevant forms of intensities: the exponential intensities initially introduced in Avellaneda and Stoikov [26] and used in most of the literature, logistic intensities as in Bergault and Guéant [41], or many SU Johnson intensities as in Guéant and Manziuk [140].

In addition to quoting prices for the N options, the market-maker can buy and sell the underlying asset. We assume that the market for that asset is liquid enough to ensure a perfect Δ -hedging.

Remark 36. *In practice, for a portfolio that is not Vega-hedged, it is usually suboptimal to perfectly Δ -hedge the portfolio because of the correlation between the spot process and the instantaneous variance process. Nevertheless, we assume here for the sake of simplicity that Δ -hedging is carried out in continuous time. A study of the optimal position in the underlying asset and its consequence on our problem is carried out in Appendix 6.A.1.*

In what follows, we denote by $(\Delta_t)_{t \in [0, T]}$ the Δ of the portfolio:

$$\Delta_t := \sum_{i=1}^N \partial_S O^i(t, S_t, \nu_t) q_t^i \text{ for all } t \in [0, T].$$

²In applications, we always choose δ_∞ negative enough so that this lower bound is never binding.

³See Appendix 6.A.3 for more details on the construction of those processes.

⁴The frontier of this set defines the risk limits of the market-maker.

The resulting dynamics for the cash process $(X_t)_{t \in [0, T]}$ of the market-maker is:

$$dX_t := \sum_{i=1}^N \left(\int_{\mathbb{R}_+^*} z \left(\delta_t^{i,b}(z) N^{i,b}(dt, dz) + \delta_t^{i,a}(z) N^{i,a}(dt, dz) \right) - \mathcal{O}_t^i dq_t^i \right) + S_t d\Delta_t + d\langle \Delta, S \rangle_t.$$

We denote by $(V_t)_{t \in [0, T]}$ the process for the Mark-to-Market (MtM) value of the market-maker's portfolio (cash, shares, and options), i.e.,

$$V_t := X_t - \Delta_t S_t + \sum_{i=1}^N q_t^i \mathcal{O}_t^i.$$

The dynamics of that process is given by

$$\begin{aligned} dV_t &= dX_t - S_t d\Delta_t - \Delta_t dS_t - d\langle \Delta, S \rangle_t + \sum_{i=1}^N \mathcal{O}_t^i dq_t^i + \sum_{i=1}^N q_t^i d\mathcal{O}_t^i \\ &= \sum_{i=1}^N \int_{\mathbb{R}_+^*} z \left(\delta_t^{i,b}(z) N^{i,b}(dt, dz) + \delta_t^{i,a}(z) N^{i,a}(dt, dz) \right) + q_t^i d\mathcal{O}_t^i - \Delta_t dS_t \\ &= \sum_{i=1}^N \int_{\mathbb{R}_+^*} z \left(\delta_t^{i,a}(z) N^{i,a}(dt, dz) + \delta_t^{i,b}(z) N^{i,b}(dt, dz) \right) + q_t^i \partial_\nu O^i(t, S_t, \nu_t) (a_{\mathbb{P}}(t, \nu_t) - a_{\mathbb{Q}}(t, \nu_t)) dt \\ &\quad + \sqrt{\nu_t} \xi q_t^i \partial_\nu O^i(t, S_t, \nu_t) dW_t^\nu. \end{aligned}$$

For all $i \in \{1, \dots, N\}$, the Vega of the i -th option is defined as

$$\mathcal{V}_t^i := \partial_{\sqrt{\nu}} O^i(t, S_t, \nu_t) = 2\sqrt{\nu_t} \partial_\nu O^i(t, S_t, \nu_t) \text{ for all } t \in [0, T].$$

Hence, we can rewrite the dynamics of the portfolio as

$$dV_t = \sum_{i=1}^N \int_{\mathbb{R}_+^*} z \left(\delta_t^{i,b}(z) N_t^{i,b}(dt, dz) + \delta_t^{i,a}(z) N_t^{i,a}(dt, dz) \right) + q_t^i \mathcal{V}_t^i \frac{a_{\mathbb{P}}(t, \nu_t) - a_{\mathbb{Q}}(t, \nu_t)}{2\sqrt{\nu_t}} dt + \frac{\xi}{2} q_t^i \mathcal{V}_t^i dW_t^\nu.$$

Following the academic literature on market-making, we can consider two objective functions: As in the initial Avellaneda and Stoikov setting Avellaneda and Stoikov [26] (see also Guéant [137, 138], Guéant, Lehalle, and Fernandez-Tapia [142]), we can consider the following expected utility objective function:

$$\sup_{\delta \in \mathcal{A}} \mathbb{E} \left[- \exp(-\gamma V_T) \right],$$

where $\gamma > 0$ is the risk-aversion parameter of the market-maker. Instead, as in Cartea, Donnelly, and Jaimungal [71], Cartea, Jaimungal, and Penalva [70], Cartea, Jaimungal, and Ricci [69], but also in Guéant [138], we can consider a risk-adjusted expectation for the objective function, i.e.

$$\sup_{\delta \in \mathcal{A}} \mathbb{E} \left[V_T - \frac{\gamma}{2} \int_0^T \left(\sum_{i=1}^N \frac{\xi}{2} q_t^i \mathcal{V}_t^i \right)^2 dt \right].$$

The second objective function in our case can be written

$$\begin{aligned} \sup_{\delta \in \mathcal{A}} \mathbb{E} \left[\int_0^T \sum_{i=1}^N \left(\left(\sum_{j \in \{a,b\}} \int_{\mathbb{R}_+^*} z \delta_t^{i,j}(z) \Lambda^{i,j}(\delta_t^{i,j}(z)) \mathbf{1}_{\{q_t - \psi(j) z e^i \in \mathcal{Q}\}} \mu^{i,j}(dz) \right) \right. \right. \\ \left. \left. + q_t^i \mathcal{V}_t^i \frac{a_{\mathbb{P}}(t, \nu_t) - a_{\mathbb{Q}}(t, \nu_t)}{2\sqrt{\nu_t}} \right) dt - \frac{\gamma \xi^2}{8} \int_0^T \left(\sum_{i=1}^N q_t^i \mathcal{V}_t^i \right)^2 dt \right], \end{aligned} \tag{6.2}$$

where

$$\psi(j) := \begin{cases} +1, & \text{if } j = a, \\ -1, & \text{if } j = b. \end{cases}$$

These two objective functions are close to one other in practice. Guéant showed in Guéant [138] that they give similar optimal quotes in practical examples. Furthermore, in many cases, the expected utility framework with exponential utility function can be reduced to the maximization of the expected PnL minus a quadratic penalty of the above form, up to a change in the intensity functions (see Manziuk [204]).

In what follows, we consider the second framework. Therefore, we define the value function

$$u : (t, S, \nu, q) \in [0, T] \times \mathbb{R}_+^2 \times \mathcal{Q} \mapsto u(t, S, \nu, q),$$

associated with (6.2) as

$$\begin{aligned} u(t, S, \nu, q) = & \sup_{(\delta_s)_{s \in [t, T]} \in \mathcal{A}_t} \mathbb{E}_{(t, S, \nu, q)} \left[\int_t^T \sum_{i=1}^N \left(\left(\sum_{j \in \{a, b\}} \int_{\mathbb{R}_+^*} z \delta_s^{i, j}(z) \Lambda^{i, j}(\delta_s^{i, j}(z)) \mathbf{1}_{\{q_t - \psi(j) z e^i \in \mathcal{Q}\}} \mu^{i, j}(dz) \right) \right. \right. \\ & \left. \left. + q_s^i \mathcal{V}_s^i \frac{a_{\mathbb{P}}(s, \nu_s) - a_{\mathbb{Q}}(s, \nu_s)}{2\sqrt{\nu_s}} \right) ds - \frac{\gamma \xi^2}{8} \int_t^T \left(\sum_{i=1}^N q_s^i \mathcal{V}_s^i \right)^2 ds \right], \end{aligned} \quad (6.3)$$

where \mathcal{A}_t is the set of admissible controls defined on $[t, T]$.

6.2.3 Assumptions and approximations

The above stochastic optimal control problem can be addressed from a theoretical point of view using an approach similar to that of Guéant [138]. However, when it comes to approximating the optimal quotes a market-maker should set for the N options, classic numerical methods are of no help because the value function u has $N + 2$ variables (in addition to the time variable). In order to beat the curse of dimensionality and be able to approximate the solution of (6.3) we propose a method based on the following assumptions / approximations:

Assumption 6.1. *We approximate the Vega of each option over $[0, T]$ by its value at time $t = 0$, namely*

$$\mathcal{V}_t^i = \mathcal{V}_0^i =: \mathcal{V}^i \in \mathbb{R}, \text{ for all } i \in \{1, \dots, N\}.$$

Assumption 6.2. *We assume that the set of authorized inventories is associated with Vega risk limits, i.e.*

$$\mathcal{Q} = \left\{ q \in \mathbb{R}^N : \sum_{i=1}^N q^i \mathcal{V}^i \in [-\bar{\mathcal{V}}, \bar{\mathcal{V}}] \right\},$$

where $\bar{\mathcal{V}} \in \mathbb{R}_+^*$ is the Vega risk limit of the market-maker.

The first assumption is acceptable if T is not too large. This raises in fact the deep question of the reasonable value of T , as there is no natural choice for the horizon of the optimization problem. In practice, T has to be sufficient large to allow for several transactions in many options and small enough for the constant-Vega approximation to be relevant (and smaller than the maturities of the options). It is also noteworthy, although it is time-inconsistent, that one can use the proceed of the model (with the constant-Vega approximation) over a short period of time and then run the model again

with updated Vegas. This is a classical practice in applied optimal control when the parameters are estimated online.⁵

The second assumption states that risk limits are related to the only source of risk (as the portfolio is Δ -hedged). This is a natural assumption. The only drawback is that no risk limit can be set to individual options.

6.3 An approximate solution to the problem

6.3.1 Change of variables: beating the curse of dimensionality

Under the above assumptions, the $N + 2$ state variables can be replaced by only two: the instantaneous variance and the Vega of the portfolio. This portfolio Vega, defined by $\mathcal{V}_t^\pi := \sum_{i=1}^N q^i \mathcal{V}^i$, has the following dynamics

$$d\mathcal{V}_t^\pi = \sum_{i=1}^N \int_{\mathbb{R}_+^*} z \mathcal{V}^i (N^{i,b}(dt, dz) - N^{i,a}(dt, dz)).$$

It is clear indeed that the value function u verifies

$$\forall (t, S, \nu, q) \in [0, T] \times \mathbb{R}_+^2 \times \mathcal{Q}, u(t, S, \nu, q) = v\left(t, \nu, \sum_{i=1}^N q^i \mathcal{V}^i\right),$$

where

$$\begin{aligned} v(t, \nu, \mathcal{V}^\pi) = \sup_{(\delta_s) \in \mathcal{A}_t} \mathbb{E}_{(t, \nu, \mathcal{V}^\pi)} \left[\int_t^T \left(\left(\sum_{i=1}^N \sum_{j \in \{a, b\}} \int_{\mathbb{R}_+^*} z \delta_s^{i,j}(z) \Lambda^{i,j}(\delta_s^{i,j}(z)) \mathbf{1}_{\{|\mathcal{V}_s^\pi - \psi(j)z\mathcal{V}^i| \leq \bar{\nu}\}} \mu^{i,j}(dz) \right) \right. \right. \\ \left. \left. + \mathcal{V}_s^\pi \frac{a_{\mathbb{P}}(s, \nu_s) - a_{\mathbb{Q}}(s, \nu_s)}{2\sqrt{\nu_s}} - \frac{\gamma \xi^2}{8} \mathcal{V}_s^{\pi 2} \right) ds \right], \end{aligned} \quad (6.4)$$

where $\mathbb{E}_{(t, \nu, \mathcal{V}^\pi)}[\cdot] = \mathbb{E}[\cdot | \nu_t = \nu, \mathcal{V}_t^\pi = \mathcal{V}^\pi]$. In other words, the problem boils down, under the two above assumptions, to a low-dimensional optimal control problem where the two state variables are driven by $2N$ controlled point processes and a standard Brownian motion.

6.3.2 Hamilton-Jacobi-Bellman equation and optimal controls

$$\begin{aligned} 0 = \partial_t v(t, \nu, \mathcal{V}^\pi) + a_{\mathbb{P}}(t, \nu) \partial_\nu v(t, \nu, \mathcal{V}^\pi) + \frac{1}{2} \nu \xi^2 \partial_{\nu\nu}^2 v(t, \nu, \mathcal{V}^\pi) + \mathcal{V}^\pi \frac{a_{\mathbb{P}}(t, \nu) - a_{\mathbb{Q}}(t, \nu)}{2\sqrt{\nu}} - \frac{\gamma \xi^2}{8} \mathcal{V}^{\pi 2} \\ + \sum_{i=1}^N \sum_{j \in \{a, b\}} \int_{\mathbb{R}_+^*} z \mathbf{1}_{\{|\mathcal{V}^\pi - \psi(j)z\mathcal{V}^i| \leq \bar{\nu}\}} H^{i,j} \left(\frac{v(t, \nu, \mathcal{V}^\pi) - v(t, \nu, \mathcal{V}^\pi - \psi(j)z\mathcal{V}^i)}{z} \right) \mu^{i,j}(dz), \end{aligned} \quad (6.5)$$

with final condition $v(T, \nu, \mathcal{V}^\pi) = 0$, where

$$H^{i,j}(p) := \sup_{\delta^{i,j} \geq \delta_\infty} \Lambda^{i,j}(\delta^{i,j})(\delta^{i,j} - p), \quad i \in \{1, \dots, N\}, \quad j \in \{a, b\}.$$

We end up therefore with a low-dimensional functional equation of the Hamilton-Jacobi-Bellman type.

⁵In Appendix 6.A.2 we propose a method to relax the constant-Vega assumption. This method is based on a Taylor expansion around the constant-Vega case. The curse of dimensionality is tamed by the reduction of the problem to a Monte-Carlo simulation.

Once the value function is known, the optimal controls, which are the optimal mid-to-bid and ask-to-mid associated with the N options, are given by the following formula for $i \in \{1, \dots, N\}$, $j \in \{a, b\}$ (see Bergault and Guéant [41], Guéant [138]):

$$\delta_t^{i,j*}(z) = \max \left(\delta_\infty, (\Lambda^{i,j})^{-1} \left(-H^{i,j} \left(\frac{v(t, \nu_t, \mathcal{V}_{t-}^\pi) - v(t, \nu_t, \mathcal{V}_{t-}^\pi - \psi(j)z\mathcal{V}^i)}{z} \right) \right) \right).$$

Remark 37. In the case where $a_{\mathbb{P}} = a_{\mathbb{Q}}$, it is evident that v does not depend on ν . In that case $v(t, \nu, \mathcal{V}^\pi) = w(t, \mathcal{V}^\pi)$ where w is solution of the simpler Hamilton-Jacobi-Bellman

$$0 = \partial_t w(t, \mathcal{V}^\pi) - \frac{\gamma \xi^2}{8} \mathcal{V}^{\pi 2} + \sum_{i=1}^N \sum_{j \in \{a, b\}} \int_{\mathbb{R}_+^*} z \mathbf{1}_{\{|\mathcal{V}^\pi - \psi(j)z\mathcal{V}^i| \leq \bar{\nu}\}} H^{i,j} \left(\frac{w(t, \mathcal{V}^\pi) - w(t, \mathcal{V}^\pi - \psi(j)z\mathcal{V}^i)}{z} \right) \mu^{i,j}(dz),$$

with final condition $w(T, \mathcal{V}^\pi) = 0$.

6.4 Numerical results

6.4.1 Model parameters

In this section we consider a book of options and derive the optimal quotes using the above approach.

For this purpose, we consider an underlying stock with the following characteristics:

- Stock price at time $t = 0$: $S_0 = 10$ €.
- Instantaneous variance at time $t = 0$: $\nu_0 = 0.0225$ year⁻¹.
- Heston model with $a_{\mathbb{P}}(t, \nu) = \kappa_{\mathbb{P}}(\theta_{\mathbb{P}} - \nu)$ where $\kappa_{\mathbb{P}} = 2$ year⁻¹ and $\theta_{\mathbb{P}} = 0.04$ year⁻¹, and $a_{\mathbb{Q}}(t, \nu) = \kappa_{\mathbb{Q}}(\theta_{\mathbb{Q}} - \nu)$ where $\kappa_{\mathbb{Q}} = 3$ year⁻¹ and $\theta_{\mathbb{Q}} = 0.0225$ year⁻¹.
- Volatility of volatility parameter: $\xi = 0.2$ year⁻¹.
- Spot-variance correlation: $\rho = -0.5$.

We consider the case of a market-maker dealing with 20 European call options written on that stock where the strike×maturity couples are the elements $(K^i, T^i)_{i \in \{1, \dots, 20\}}$ of the set $\mathcal{K} \times \mathcal{T}$, where

$$\mathcal{K} = \{8 \text{ €}, 9 \text{ €}, 10 \text{ €}, 11 \text{ €}, 12 \text{ €}\} \quad \text{and} \quad \mathcal{T} = \{1 \text{ year}, 1.5 \text{ years}, 2 \text{ years}, 3 \text{ years}\}.$$

The associated implied volatility surface is plotted in Figure 6.1.⁶

The liquidity parameters of these options are the following:

- Intensity functions:

$$\Lambda^{i,j}(\delta) = \frac{\lambda^i}{1 + e^{\alpha + \frac{\beta}{\nu^i} \delta}}, \quad i \in \{1, \dots, N\}, \quad j \in \{a, b\}.$$

where $\lambda^i = \frac{252 \times 30}{1 + 0.7 \times |S_0 - K^i|}$ year⁻¹, $\alpha = 0.7$, and $\beta = 150$ year^{1/2}. The choice of λ^i corresponds to 30 requests per day for at-the-money options, and decreases to 12.5 for more in- and out-the-money options. The choice of α corresponds to a probability of $\frac{1}{1 + e^{0.7}} \approx 33\%$ to trade when the answered quote is the mid-price. The choice of β corresponds to a probability of $\frac{1}{1 + e^{-0.8}} \approx 69\%$ to trade when the answered quote corresponds to an implied volatility 1% better for the client and a probability of $\frac{1}{1 + e^{2.2}} \approx 10\%$ to trade when the answered quote corresponds to an implied volatility 1% worse for the client.

⁶This plot has been computed using 10^5 Monte-Carlo simulations for each option.

- Size of transactions: we assume constant sizes $z^i = \frac{5 \cdot 10^5}{O_0^i}$ contracts for option i . This corresponds approximately to 500000 € per transaction.⁷ The measures $\mu^{i,b}$ and $\mu^{i,a}$ are here Dirac masses.




Figure 6.1: Implied volatility surface associated with the above parameters.

Regarding the risk limits and the objective function, we consider the following:

- Vega risk limit: $\bar{\nu} = 10^7 \text{ €} \cdot \text{year}^{\frac{1}{2}}$.
- Time horizon given by $T = 0.0012 \text{ year}$ (i.e. 0.3 days). This short time horizon surprisingly ensures convergence towards stationary quotes at time $t = 0$ (see Figure 6.3 below).
- Risk aversion given by $\gamma = 1 \cdot 10^{-3} \text{ €}^{-1}$.

6.4.2 Optimal quotes

Using a monotone explicit Euler scheme with linear interpolation on a grid of size $180 \times 30 \times 40$, we approximate the value function solution to (6.5) (with Neumann conditions at the boundaries in ν) on the domain $[0, T] \times [0.0144, 0.0324] \times [-\bar{\nu}, \bar{\nu}]$. This value function is plotted in Figure 6.2.




Figure 6.2: Value function as a function of instantaneous variance and portfolio Vega.

From that value function, we deduce the optimal bid and ask quotes of the market-maker for each option as a function of the portfolio Vega. As mentioned above, we chose $T = 0.0008 \text{ year}$ (i.e. 0.2 days) – a choice that ensures convergence of the optimal quotes to their stationary values (see Figure 6.3).

⁷This is only an approximation as trade sizes are in number of options and option prices move.




Figure 6.3: Optimal mid-to-bid quotes as a function of time for option 1: $(K^1, T^1) = (8, 1) - \nu = 0.04$.

Focusing on the asymptotic values, we now present in Figures 6.4, 6.5, 6.6, 6.7, and 6.8, the optimal bid quotes as a function of the portfolio Vega for each strike and maturity. More precisely, as the options we consider can have very different prices, we consider instead of the optimal bid quotes themselves the ratio between each optimal mid-to-bid quote and the price (at time $t = 0$) of the corresponding option. In Figures 6.9, 6.10, 6.11, 6.12, and 6.13, we plot the same optimal bid quotes for the 20 options in terms of implied volatility (divided by the implied volatility at time $t = 0$).

The results are in line with what was expected: the mid-to-bid quotes increase with the portfolio Vega. The incentive to buy options with positive Vega decreases indeed with the Vega of the portfolio.




Figure 6.4: Optimal mid-to-bid quotes divided by option price as a function of the portfolio Vega for $K=8$.




Figure 6.5: Optimal mid-to-bid quotes divided by option price as a function of the portfolio Vega for $K=9$.




Figure 6.6: Optimal mid-to-bid quotes divided by option vega price as a function of the portfolio Vega for K=10.




Figure 6.7: Optimal mid-to-bid quotes divided by option vega price as a function of the portfolio Vega for K=11.




Figure 6.8: Optimal mid-to-bid quotes divided by option price as a function of the portfolio Vega for K=12.




Figure 6.9: Optimal (relative) bid implied volatility as a function of the portfolio Vega for K=8.




Figure 6.10: Optimal (relative) bid implied volatility as a function of the portfolio Vega for K=9.




Figure 6.11: Optimal (relative) bid implied volatility as a function of the portfolio Vega for K=10.




Figure 6.12: Optimal (relative) bid implied volatility as a function of the portfolio Vega for K=11.




Figure 6.13: Optimal (relative) bid implied volatility as a function of the portfolio Vega for K=12.

6.4.3 Conclusion

In this chapter, we tackled the problem of an option market-maker dealing with options on a single underlying asset.⁸ Using a constant-Vega approximation, we showed how to reduce the problem to a low-dimensional functional equation whose solution can easily be approximated using an explicit Euler scheme and linear interpolation. Furthermore, our method scales linearly in the number of options and can therefore be used with large books of options. Our method is illustrated by an example involving 20 European calls, but our model can be used with any European options.

6.A Appendix

6.A.1 An alternative to the Δ -hedging assumption

Throughout the chapter, we assumed that the market-maker ensured Δ -hedging. In this appendix, we show that this assumption can be relaxed without much change in the reasoning.

Let us introduce the process $(q_t^S)_{t \in [0, T]}$ representing the inventory of the market-maker in the underlying asset. The dynamics of the cash process of the market-maker $(X_t)_{t \in [0, T]}$ rewrites as

$$dX_t = \sum_{i=1}^N \left(\int_{\mathbb{R}_+^*} z \left(\delta_t^{i,b}(z) N_t^{i,b}(dt, dz) + \delta_t^{i,a}(z) N_t^{i,a}(dt, dz) \right) - \mathcal{O}_t^i dq_t^i \right) - q_t^S dS_t - d\langle q^S, S \rangle_t.$$

The Mark-to-Market value of the portfolio writes

$$V_t = X_t + q_t^S S_t + \sum_{i=1}^N q_t^i \mathcal{O}_t^i$$

and its dynamics is

$$\begin{aligned} dV_t &= \sum_{i=1}^N \left(\int_{\mathbb{R}_+^*} z \left(\delta_t^{i,b}(z) N_t^{i,b}(dt, dz) + \delta_t^{i,a}(z) N_t^{i,a}(dt, dz) \right) + q_t^i \mathcal{V}_t^i \frac{a_{\mathbb{P}}(t, \nu_t) - a_{\mathbb{Q}}(t, \nu_t)}{2\sqrt{\nu_t}} dt + \frac{\xi}{2} q_t^i \mathcal{V}_t^i dW_t^\nu \right) \\ &+ \sqrt{\nu_t} S_t \left(\sum_{i=1}^N q_t^i \partial_S \mathcal{O}^i(t, S_t, \nu_t) + q_t^S \right) dW_t^S. \end{aligned}$$

Denoting by $\Delta_t^\pi := \sum_{i=1}^N q_t^i \partial_S \mathcal{O}^i(t, S_t, \nu_t)$ the Δ of the market-maker's portfolio at time t , our mean-variance optimization problem becomes

$$\sup_{(\delta, q^S) \in \mathcal{A}'} \mathbb{E}[V_T] - \frac{\gamma}{2} \mathbb{V} \left[\int_0^T \frac{\xi}{2} \mathcal{V}_t^\pi dW_t^\nu + \sqrt{\nu_t} S_t (\Delta_t^\pi + q_t^S) dW_t^S \right],$$

where

$$\mathcal{A}' = \left\{ (\delta_t, q_t^S)_{t \in [0, T]} : \delta \text{ is an } \mathbb{R}^{2N}\text{-valued predictable process bounded from below by } \delta_\infty \right. \\ \left. \text{and } q^S \text{ is an } \mathbb{R}\text{-valued adapted process with } \mathbb{E} \left[\int_0^T \nu_t S_t^2 (\Delta_t^\pi + q_t^S)^2 dt \right] < +\infty \right\}.$$

⁸As noted while publishing this chapter in the form of an article, our method can easily be extended to the case of multiple underlying assets using the same method as in Bergault and Guéant [41] if the instantaneous variance processes of the different assets are driven by a few factors.

Noticing that

$$\mathbb{V}\left(\int_0^T \frac{\xi}{2} \mathcal{V}_t^\pi dW_t^\nu + \sqrt{\nu_t} S_t (\Delta_t^\pi + q_t^S) dW_t^S\right) = \mathbb{E}\left[\int_0^T \left(\frac{\xi^2}{4} \mathcal{V}_t^{\pi^2} + \nu_t S_t^2 (\Delta_t^\pi + q_t^S)^2 + \rho \xi \mathcal{V}_t^\pi \sqrt{\nu_t} S_t (\Delta_t^\pi + q_t^S)\right) dt\right],$$

we easily see that the variance term is minimized for $q^S = q^{S^*}$ where

$$\forall t \in [0, T], \quad q_t^{S^*} = -\Delta_t^\pi - \frac{\rho \xi \mathcal{V}_t^\pi}{2\sqrt{\nu_t} S_t},$$

and that its minimum value is

$$(1 - \rho^2) \int_0^T \frac{\xi^2}{4} \mathcal{V}_t^{\pi^2} dt.$$

Therefore, the optimization problem boils down to

$$\begin{aligned} \sup_{\delta \in \mathcal{A}} \mathbb{E} \left[\int_0^T \left(\left(\sum_{i=1}^N \sum_{j \in \{a,b\}} \int_{\mathbb{R}_+^*} z \delta_t^{i,j}(z) \Lambda^{i,j}(\delta_t^{i,j}(z)) \mathbf{1}_{|\mathcal{V}_t^\pi - \psi(j)z \mathcal{V}_t^i| \leq \bar{\nu}} \mu^{i,j}(dz) \right) \right. \right. \\ \left. \left. + \mathcal{V}_t^\pi \frac{a_{\mathbb{P}}(t, \nu_t) - a_{\mathbb{Q}}(t, \nu_t)}{2\sqrt{\nu_t}} - \frac{\gamma \xi^2}{8} (1 - \rho^2) \mathcal{V}_t^{\pi^2} \right) dt \right], \end{aligned}$$

and we recover the same optimization problem as in the body of the chapter, except that the risk aversion parameter is multiplied by $1 - \rho^2$ to account for the reduction of risk made possible by the optimal trading strategy in the underlying asset in presence of vol-spot correlation.

6.A.2 Beyond the constant-Vega assumption

In this appendix we propose a method to relax our main assumption: the constant-Vega approximation.

If, for all $i \in \{1, \dots, N\}$, the process $(\mathcal{V}_t^i)_{t \in [0, T]}$ stays close to its initial value \mathcal{V}_0^i , then it is reasonable to consider a perturbative approach around the constant-Vega approximation. In particular, instead of assuming that $\sum_{i=1}^N q_t^i \partial_{\sqrt{\nu}} O^i(t, S_t, \nu_t) \simeq \sum_{i=1}^N q_t^i \partial_{\sqrt{\nu}} O^i(0, S_0, \nu_0) = \mathcal{V}_t^\pi$, we consider the expansion

$$\sum_{i=1}^N q_t^i \partial_{\sqrt{\nu}} O^i(t, S_t, \nu_t) = \mathcal{V}_t^\pi + \varepsilon \mathcal{W}(t, S_t, \nu_t, q_t),$$

and we consider an expansion of the value function u of the following form:

$$u(t, S, \nu, q) = v\left(t, \nu, \sum_{i=1}^N q^i \partial_{\sqrt{\nu}} O^i(0, S_0, \nu_0)\right) + \varepsilon \varphi(t, S, \nu, q).$$

Assuming that $\mathcal{Q} = \mathbb{R}^N$ and noting that the Hamilton-Jacobi-Bellman equation associated with u is

$$\begin{aligned} 0 = & \partial_t u(t, S, \nu, q) + a_{\mathbb{P}}(t, \nu) \partial_\nu u(t, S, \nu, q) + \frac{1}{2} \nu S^2 \partial_{SS}^2 u(t, S, \nu, q) + \frac{1}{2} \nu \xi^2 \partial_{\nu\nu}^2 u(t, S, \nu, q) \\ & + \rho \nu S \xi \partial_{\nu S}^2 u(t, S, \nu, q) + \frac{a_{\mathbb{P}}(t, \nu) - a_{\mathbb{Q}}(t, \nu)}{2\sqrt{\nu}} \sum_{i=1}^N q^i \partial_{\sqrt{\nu}} O^i(t, S, \nu) - \frac{\gamma \xi^2}{8} \left(\sum_{i=1}^N q^i \partial_{\sqrt{\nu}} O^i(t, S, \nu) \right)^2 \\ & + \sum_{i=1}^N \sum_{j \in \{a,b\}} \int_{\mathbb{R}_+^*} z H^{i,j} \left(\frac{u(t, S, \nu, q) - u(t, S, \nu, q - \psi(j) z e^i)}{z} \right) \mu^{i,j}(dz), \end{aligned}$$

with terminal condition equal to 0, the first-order term in ε in the Taylor expansion gives

$$\begin{aligned}
0 &= \partial_t \varphi(t, S, \nu, q) + a_{\mathbb{P}}(t, \nu) \partial_\nu \varphi(t, S, \nu, q) + \frac{1}{2} \nu S^2 \partial_{SS}^2 \varphi(t, S, \nu, q) + \frac{1}{2} \nu \xi^2 \partial_{\nu\nu}^2 \varphi(t, S, \nu, q) \\
&+ \rho \nu S \xi \partial_{\nu S}^2 \varphi(t, s, \nu, q) + \frac{a_{\mathbb{P}}(t, \nu) - a_{\mathbb{Q}}(t, \nu)}{2\sqrt{\nu}} \mathcal{W}(t, S, \nu, q) - \frac{\gamma \xi^2}{4} \mathcal{W}(t, S, \nu, q) \sum_{i=1}^N q^i \mathcal{V}_0^i \\
&+ \sum_{i=1}^N \sum_{j \in \{a, b\}} \int_{\mathbb{R}_+^*} H^{i, j'} \left(\frac{v(t, \nu, \sum_{l=1}^N q^l \mathcal{V}_0^l) - v(t, \nu, \sum_{l=1}^N q^l \mathcal{V}_0^l - \psi(j) z \mathcal{V}_0^i)}{z} \right) \\
&\times (\varphi(t, S, \nu, q) - \varphi(t, S, \nu, q - \psi(j) z e^i)) \mu^{i, j}(dz),
\end{aligned}$$

with terminal condition equal to 0.

Although it involves $q \in \mathbb{R}^N$, this equation is linear and thus $\varphi(t, S, \nu, q)$ admits a Feynman-Kac representation that tames the curse of dimensionality for practical applications:

$$\varphi(t, S, \nu, q) = \mathbb{E}_{(t, S, \nu, q)} \left[\int_t^T \left(\frac{a_{\mathbb{P}}(s, \nu_s) - a_{\mathbb{Q}}(s, \nu_s)}{2\sqrt{\nu_s}} \mathcal{W}(s, S_s, \nu_s, q_s) - \frac{\gamma \xi^2}{4} \mathcal{W}(s, S_s, \nu_s, q_s) \sum_{i=1}^N q_s^i \mathcal{V}_0^i \right) ds \right],$$

where, for each $i \in \{1, \dots, N\}$, the processes $N^{i, b}$ and $N^{i, a}$ have respective intensities

$$\begin{aligned}
\tilde{\lambda}_t^{i, b}(dz) &= -H^{i, b'} \left(\frac{v(t, \nu_t, \sum_{l=1}^N q_{t-}^l \mathcal{V}_0^l) - v(t, \nu_t, \sum_{l=1}^N q_{t-}^l \mathcal{V}_0^l - \psi(j) z \mathcal{V}_0^i)}{z} \right) \mu^{i, b}(dz), \\
\tilde{\lambda}_t^{i, a}(dz) &= -H^{i, a'} \left(\frac{v(t, \nu_t, \sum_{l=1}^N q_{t-}^l \mathcal{V}_0^l) - v(t, \nu_t, \sum_{l=1}^N q_{t-}^l \mathcal{V}_0^l - \psi(j) z \mathcal{V}_0^i)}{z} \right) \mu^{i, a}(dz),
\end{aligned}$$

with as before $dq_t^i = \int_{\mathbb{R}_+^*} z (N^{i, b}(dt, dz) - N^{i, a}(dt, dz))$.

Subsequently, the function φ can be computed using a Monte-Carlo algorithm and quotes accounting for the variation of the Vegas can therefore be computed (to the first order in ε).

6.A.3 On the construction of the processes $N^{i, b}$ and $N^{i, a}$

Let us consider a new filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \in \mathbb{R}_+}, \tilde{\mathbb{P}})$. For the sake of simplicity, assume that there is only one option (the generalization is straightforward). Let us introduce \bar{N}^b and \bar{N}^a two independent compound Poisson processes of intensity 1 whose increments follow respectively the distributions $\mu^b(dz)$ and $\mu^a(dz)$ with support on \mathbb{R}_+^* . We denote by $\bar{N}^b(dt, dz)$ and $\bar{N}^a(dt, dz)$ the associated random measures. Let N^b and N^a be two processes, starting at 0, solutions of the coupled stochastic differential equation:

$$N_t^b = \int_0^t \int_{\mathbb{R}_+^*} \mathbf{1}_{\{N_{t-}^b - N_{t-}^a + z \in \mathcal{Q}\}} \bar{N}^b(dt, dz), \quad N_t^a = \int_0^t \int_{\mathbb{R}_+^*} \mathbf{1}_{\{N_{t-}^b - N_{t-}^a - z \in \mathcal{Q}\}} \bar{N}^a(dt, dz).$$

Then, under $\tilde{\mathbb{P}}$, N^b and N^a are two \mathbb{R}_+^* -marked point processes with respective intensity kernels

$$\lambda_t^b(dz) = \mathbf{1}_{\{q_{t-} + z \in \mathcal{Q}\}} \mu^b(dz) \quad \text{and} \quad \lambda_t^a(dz) = \mathbf{1}_{\{q_{t-} - z \in \mathcal{Q}\}} \mu^a(dz),$$

where $q_t := N_t^b - N_t^a$. We denote by $N^b(dt, dz)$ and $N^a(dt, dz)$ the associated random measures. For each $\delta \in \mathcal{A}$, we introduce the probability measure $\tilde{\mathbb{P}}^\delta$ given by the Radon-Nikodym derivative

$$\left. \frac{d\tilde{\mathbb{P}}^\delta}{d\tilde{\mathbb{P}}} \right|_{\mathcal{F}_t} = L_t^\delta, \tag{6.6}$$

where $(L_t^\delta)_{t \geq 0}$ is the unique solution of the stochastic differential equation

$$dL_t^\delta = L_t^\delta \left(\int_{\mathbb{R}_+^*} (\Lambda^b(\delta^b(t, z)) - 1) d\tilde{N}^b(dt, dz) + \int_{\mathbb{R}_+^*} (\Lambda^a(\delta^a(t, z)) - 1) d\tilde{N}^a(dt, dz) \right),$$

with $L_0^\delta = 1$, where \tilde{N}^b and \tilde{N}^a are the compensated processes associated with N^b and N^a respectively.

We then know from Girsanov theorem that under $\tilde{\mathbb{P}}^\delta$, the jump processes N^b and N^a have respective intensity kernels

$$\lambda_t^{\delta, b}(dz) = \Lambda^b(\delta^b(t, z)) \mathbf{1}_{\{q_{t-} + z \in \mathcal{Q}\}} \mu^b(dz) \quad \text{and} \quad \lambda_t^{\delta, a}(dz) = \Lambda^a(\delta^a(t, z)) \mathbf{1}_{\{q_{t-} - z \in \mathcal{Q}\}} \mu^a(dz),$$

as in the body of the chapter.

Chapter 7

An approximate solution for options market-making in high dimension

7.1 Introduction

After the electrification of delta-one trading, where high-frequency trading companies provide the vast majority of the liquidity on several thousands of assets, systematic options trading seems to be the next main challenge in quantitative trading. For assets listed in a central limit order book, as in the equity world, execution and market-making are carried out using algorithms. However, for less mature markets such as a great proportion of fixed income securities, systematic market-making activities are driven by request-for-quote (RFQ for short) systems: the client sends a request to obtain a buy or sell price, for a given quantity of a security, to one or several market-makers, who propose prices based on their current positions. Given the prices, the client accepts or refuses one or several transactions. On OTC markets, such as the corporate bonds market, the proportion of the volume traded with electronic market-makers is increasing.

The two primary references on optimal market-making are Grossman and Miller [135], Ho and Stoll [156]. In Grossman and Miller [135], the authors proposed a simple three-period economic model to understand the interaction between market-makers and market-takers and analyzed the equilibrium state. In Ho and Stoll [156], the authors studied the behavior of a market-maker facing a stochastic demand (thus an inventory risk) and obtained his optimal strategy using the stochastic optimal control theory. In the well-known paper of Avellaneda and Stoikov [26] inspired by this framework, they proposed a model applicable for market-making on the order-driven market at the high-frequency. However, due to the continuous nature of the market-maker's spreads, and the assumption that the underlying asset is a diffusion process, this model is more suited to quote-driven markets such as corporate bonds market.

By providing a rigorous analysis of the stochastic control problem of Avellaneda and Stoikov [26], the authors of Guéant, Lehalle, and Fernandez-Tapia [142] show, in the case of a CARA utility function, that the market-maker's problem boils is reduced to the resolution of a system of linear ordinary differential equations. A large part of the contribution to the market-making literature comes from works of Cartea and Jaimungal, who enriched the initial model by introducing alpha signals, ambiguity aversion, competition with other agents, see, for example, Cartea, Donnelly, and Jaimungal [71, 72], Cartea and Jaimungal [68], Cartea, Jaimungal, and Penalva [70]. In these works, they consider a risk-adjusted expectation maximization. As shown in Manziuk [204], the solution of such formulation can also be obtained through CARA utility maximization after a suitable intensity function transformation. More recently, multi-asset market-making, still on linear markets, has been addressed through reinforcement

learning techniques, see Guéant and Manziuk [140], and dimensionality reduction techniques, as in Bergault and Guéant [41].

Regardless of how rich is the part of academic literature considering linear markets, the part studying optimal market-making on options is far less extensive. A reasonable market-making model for options has to take into account a lot of stylized facts. First, option market-makers trade simultaneously derivatives and the corresponding underlying, which implies the construction of more complex trading strategies taking into account, for example, the Delta-Vega hedging. Consequently, one needs to impose a factorial stochastic volatility model, possibly with jumps, on the underlying asset. Second, option market-makers need to manage several thousands of positions, which lead to very high-dimensional problems that cannot be solved using classical numerical schemes. Even if machine learning techniques are used, involving, for example, deep reinforcement learning methods (see Guéant and Manziuk [140], E, Han, and Jentzen [103]), the computation time can still be an obstacle. The market-maker has to answer a request from a client in a given time, which can be insufficient to recalibrate the model if some parameter changes need to be applied (for example, the correlation structure). Finally, when dealing with short maturity options, the market-maker has to manage the positions individually to avoid sudden high exposure due to the Gamma of a specific position. This specificity prevents the use of some dimensionality reduction techniques.

In the existing academic literature, options market-making is addressed in El Aoud and Abergel [106], Stoikov and Sağlam [258] and Chapter 6 of this thesis. In Stoikov and Sağlam [258], the authors consider three different settings for a market-maker managing a single option and its underlying. The first setting is a complete market where continuous trading is allowed in a perfectly liquid underlying (in particular, there are no issues related to market liquidity). The second framework is a complete market with an illiquid underlying where the market-maker sets bid and ask quotes in the option and the stock. The last setting is an incomplete market with residual risks due to both stochastic volatility and overnight jumps in the stock price. In El Aoud and Abergel [106], the authors consider a market-maker in charge of a single option in a framework à la Avellaneda-Stoikov, where an underlying follows a one-factor stochastic volatility model, and the market-maker is always Delta-hedged. They provide analytically the bid and ask quotes for the option taking into account the risk of model misspecification. Finally, in Chapter 6 of this thesis, the authors consider a perfectly Delta-hedged market-maker in charge of a book of options with long maturities, whose prices are driven by a stochastic volatility model. The only risk factor comes from the Brownian motion driving the volatility of the underlying. Using a first-order approximation of the Vega of the portfolio, they show that the problem of an options market-maker boils down to a three-dimensional Hamilton-Jacobi-Bellman (HJB) equation, which can be solved using classical finite difference schemes. By linearizing the value function of the market-maker around the Vegas at the initial time, they provide a way to relax the constant Vega assumption. However, the disadvantage of this approach is its time-consumption due to the necessity to simulate inventory trajectories. Moreover, the constant Vega assumption, making the control problem time-inconsistent, is only valid for a market-maker in charge of long-dated options where possible jumps in the underlying do not influence the global risk position drastically. Finally, if one adds other Greeks such as Vanna and Vomma, the model becomes hardly tractable as the HJB equation is in dimension 5.

In this chapter, our goal is to propose a market-making algorithm that considers the three specificities mentioned above, more flexible and applicable in practice. To this end, we consider a market-maker in charge of a book of options on different underlyings. The assets follow a one-factor stochastic volatility model with jumps, and the Brownian motions driving the underlying and the volatility of each asset are correlated. We first consider the case of a perfectly Delta-hedged market-maker who manages his volatility Greeks, namely the Vega, the Vanna, and the Vomma, for all his positions. Inspired by Bergault, Evangelista, Guéant, and Vieira [42], we approximate the jump-diffusion HJB equation corresponding to the optimization problem of the market-maker with an elliptic Partial Differential

Equation (PDE for short). Using an ansatz quadratic in the inventories, we approximate the value function by a system of non-linear PDEs, which can be easily solved via classical numerical methods for a small number of assets. For a number of underlyings above two, we recast the ansatz by adding a non-local term, enabling the use of the Deep Galerkin method as in Hirsa and Fu [155] to solve the system of PDEs rapidly due to its simple non-linearity.

The method presented in this chapter has several advantages. First, contrary to El Aoud and Abergel [106] and similarly to Chapter 6 of this thesis, the market-maker can design trading strategies on a high number of options. Contrary to the approach of Chapter 6, the market-maker controls each position individually, which is particularly important for short-dated options that must be managed one by one. Moreover, it enables us to reproduce classic option market-making behavior where one option is hedged with another. Second, we allow continuous updates of the Greeks (Delta, Vega, Vanna, Vomma) of each option, and the dependence of the intensities of orders arrival on the dynamics of the underlying and its stochastic volatility. This is a major improvement compared to Chapter 6, as the quotes of the market-maker are adjusted dynamically with respect to the evolution of both an underlying and stochastic volatility, allowing the problem to be solved in a time-consistent way. Third, we can use a model for the underlying dynamics with an arbitrary number of factors without increasing the computation time. We show numerically how this algorithm outperforms the one in Chapter 6 in terms of average PnL for a portfolio of options, where Vegas vary significantly.

The chapter has the following structure: in Section 7.2, we present the framework of options market-making and the corresponding optimization problem faced by the market-maker. In Section 7.3, we show how to simplify the problem by approximating the value function. Finally, Section 7.4 is devoted to numerical experiments.

7.2 Framework

7.2.1 The option book

We consider a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$ where all stochastic processes are defined, and a time horizon $T > 0$. We consider $d > 1$ stocks with the following one factor stochastic volatility dynamics with jumps:

$$\begin{cases} dS_t^i = b_{\mathbb{P}}^i(t, S_t^i)dt + \sigma^i(t, S_t^i, \nu_t^i)dW_t^{i,S} + \int_{\mathbb{R}} Z^i(dt, dz), \\ d\nu_t^i = a_{\mathbb{P}}^i(t, \nu_t^i)dt + v_{\mathbb{P}}^i(t, \nu_t^i)dW_t^{i,\nu}, \end{cases} \quad (7.1)$$

where $(W_t^{i,S}, W_t^{i,\nu})_{t \in \mathbb{R}^+}$ is a couple of Brownian motions with quadratic covariation given by the coefficients $\rho^i = \frac{d\langle W_t^{i,S}, W_t^{i,\nu} \rangle}{dt} \in (-1, 1)$, and $a_{\mathbb{P}}^i, b_{\mathbb{P}}^i, v_{\mathbb{P}}^i, \sigma^i$ are such that the SDEs (7.1) admit a unique strong solution¹. The processes $Z^i(dt, dz)$ are marked point processes independent from the Brownian motions, with intensity kernels $\kappa_t^i(dz)$. We also assume that there exists covariance matrices $\Sigma^S, \Sigma^\nu \in \mathcal{M}_d(\mathbb{R})$ which correspond to the correlation structure of the stocks and the stochastic volatility in the option book. There also exists a risk-neutral probability measure \mathbb{Q} such that

$$\begin{cases} dS_t^i = \sigma^i(t, S_t^i, \nu_t^i)d\hat{W}_t^{i,S} + \int_{\mathbb{R}} Z^i(dt, dz), \\ d\nu_t^i = a_{\mathbb{Q}}^i(t, \nu_t^i)dt + v_{\mathbb{Q}}^i(t, \nu_t^i)d\hat{W}_t^{i,\nu}, \end{cases}$$

where $(\hat{W}_t^{i,S}, \hat{W}_t^{i,\nu}), i \in \{1, \dots, d\}$ are \mathbb{Q} -Brownian motions.

¹In particular, for the sake of readability, we assume that there is no correlation between the volatility process of an asset and the variations of another asset. This assumption can be directly relaxed.

Remark 38. As the reader will see in the following, by applying the ansatz detailed in Section 7.3, one can use a multi-factor stochastic volatility model for the underlying without increasing the complexity of the algorithm. For example, one can work with the well-known two-factor Bergomi model easily, see Bergomi [44, 45].

On every underlying $i \in \{1, \dots, d\}$ we consider a set of N^i European options $\mathcal{O}^{i,j}$ of maturity $T^{i,j}$, for $j \in \{1, \dots, N^i\}$. In the above one-factor model, we know that for all $(i, j) \in \{1, \dots, d\} \times \{1, \dots, N^i\}$, and all $t \in [0, T^{i,j}]$ such that $T < \min_{i,j} T^{i,j}$, $\mathcal{O}_t^{i,j} = O^{i,j}(t, S_t^i, \nu_t^i)$ where $O^{i,j}$ is a solution on $[0, T^{i,j}] \times \mathbb{R}_+^2$ of the following partial differential equation under the probability \mathbb{Q} :

$$\begin{aligned} 0 = & \partial_t O^{i,j}(t, S^i, \nu^i) + a_{\mathbb{Q}}^i(t, \nu^i) \partial_{\nu^i} O^{i,j}(t, S^i, \nu^i) + \frac{1}{2} (\sigma^i(t, S^i, \nu^i))^2 \partial_{S^i S^i}^2 O^{i,j}(t, S^i, \nu^i) \\ & + \rho^{i,i} \nu_{\mathbb{Q}}^i(t, \nu^i) \sigma^i(t, S^i, \nu^i) \partial_{\nu^i S^i}^2 O^{i,j}(t, S^i, \nu^i) + \frac{1}{2} (\nu_{\mathbb{Q}}^i(t, \nu^i))^2 \partial_{\nu^i \nu^i}^2 O^{i,j}(t, S^i, \nu^i) \\ & + \int_{\mathbb{R}} \left(O^{i,j}(t, S^i + \gamma^i(t, z), \nu^i) - O^{i,j}(t, S^i, \nu^i) \right) \kappa^i(dz). \end{aligned}$$

As the time horizon T is small compared to the maturity of the options (which can be from one day up to several years), the terminal condition of the PDEs does not have to be specified. In Section 7.4, numerical experiments are addressed using European call options but any other option with a path-independent payoff can be considered. We now define the market-maker's problem.

7.2.2 The market-maker's problem on OTC markets

We consider a market-maker in charge of providing bid and ask quotes for the $\sum_{i \in \{1, \dots, d\}} N^i$ options over the period $[0, T]$ where $T < \min_{i,j} T^{i,j}$. The bid and ask prices on the option $j \in \{1, \dots, N^i\}$ of stock $i \in \{1, \dots, d\}$ are defined, for transaction size z , by

$$P_t^{i,j,b} = \mathcal{O}_t^{i,j} - \delta_t^{i,j,b}(z), \quad P_t^{i,j,a} = \mathcal{O}_t^{i,j} + \delta_t^{i,j,a}(z),$$

where $(\delta_t^{i,j,b}(\cdot), \delta_t^{i,j,a}(\cdot)) \in \mathcal{A}$, where \mathcal{A} is the set of uniformly bounded \mathcal{F} -predictable processes. They represent the spread on the bid or ask side of the option $\mathcal{O}^{i,j}$. The number of transactions on these options are defined by marked point processes $N^{i,j,b}(dt, dz), N^{i,j,a}(dt, dz)$, with almost surely no simultaneous jumps, whose respective intensity processes are given by

$$\Lambda_t^{i,j,b}(S, \nu, dz) = \lambda^{i,j,b}(S, \nu, \delta_t^{i,j,b}(z)) \mu^{i,j,b}(dz), \quad \Lambda_t^{i,j,a}(S, \nu, dz) = \lambda^{i,j,a}(S, \nu, \delta_t^{i,j,a}(z)) \mu^{i,j,a}(dz).$$

The couples $(\mu^{i,j,b}, \mu^{i,j,a})$ are probability measures on \mathbb{R}_+^* modeling the distribution of transaction sizes for the options. Note that, in our framework, the intensities are allowed to depend on both the underlying and the stochastic volatility of the assets.

The market-maker manages his inventory process on each option, that is

$$dq_t^{i,j} = \int_{\mathbb{R}_+^*} z (N^{i,j,b}(dt, dz) - N^{i,j,a}(dt, dz)).$$

For the sake of simplicity, we represent the vector of inventories as follows:

$$q^\top = (q^{1,1}, \dots, q^{1,N^1}, \dots, q^{d,N^d}) \in \mathcal{M}_{\sum_{l=1}^d N^l, 1}(\mathbb{R}).$$

Assuming perfect Delta-hedging², the Δ of the portfolio on the i -th asset, $i \in \{1, \dots, d\}$, is given by

$$\Delta_t^i = \sum_{j \in \{1, \dots, N^i\}} \partial_{S^i} O^{i,j}(t, S_t^i, \nu_t^i) q_t^{i,j}, \quad \Delta_t = \sum_{i \in \{1, \dots, d\}} \Delta_t^i.$$

²This assumption can be relaxed by assuming that the market-maker acts on the stock market. This way, the mean-variance objective function will take into account the Delta of the portfolio.

The cash process of the market-maker at time t is defined by

$$dX_t = \sum_{(i,j) \in \{1, \dots, d\} \times \{1, \dots, N^i\}} \left(\int_{\mathbb{R}_+^*} z \left(\delta_t^{i,j,b}(z) N_t^{i,j,b}(dt, dz) + \delta_t^{i,j,a}(z) N_t^{i,j,a}(dt, dz) \right) - \mathcal{O}_t^{i,j} dq_t^{i,j} \right) + \sum_{i \in \{1, \dots, d\}} \left(S_t^i d\Delta_t^i + d\langle \Delta^i, S^i \rangle_t \right).$$

We finally define the Mark-to-Market value of the portfolio of the market-maker as

$$V_t = X_t - \sum_{i \in \{1, \dots, d\}} \Delta_t^i S_t^i + \sum_{(i,j) \in \{1, \dots, d\} \times \{1, \dots, N^i\}} q_t^{i,j} \mathcal{O}_t^{i,j}.$$

For all $(i, j) \in \{1, \dots, d\} \times \{1, \dots, N^i\}$, the Vega, the Vomma and the Vanna of the option $\mathcal{O}_t^{i,j}$ are defined as

$$\begin{aligned} \mathcal{V}_t^{i,j} &= \partial_{\sqrt{\nu^i}} \mathcal{O}^{i,j}(t, S_t^i, \nu_t^i) = 2\sqrt{\nu^i} \partial_{\nu^i} \mathcal{O}^{i,j}(t, S_t^i, \nu_t^i), \\ (\mathcal{VO})_t^{i,j} &= \partial_{\sqrt{\nu^i} \sqrt{\nu^i}} \mathcal{O}^{i,j}(t, S_t^i, \nu_t^i) = 4\nu^i \partial_{\nu^i \nu^i}^2 \mathcal{O}^{i,j}(t, S_t^i, \nu_t^i), \\ (\mathcal{VA})_t^{i,j} &= \partial_{S \sqrt{\nu^i}} \mathcal{O}^{i,j}(t, S_t^i, \nu_t^i) = 2\sqrt{\nu^i} \partial_{S \nu^i} \mathcal{O}^{i,j}(t, S_t^i, \nu_t^i). \end{aligned}$$

We also define the vectors $e^{i,j} \in \mathbb{R}^{\sum_{l=1}^d N^l}$ where $e_k^{i,j} = \mathbf{1}_{\{k = \sum_{l=1}^{i-1} N^l + j\}}$ and (e^1, \dots, e^d) as the canonical basis of \mathbb{R}^d . If we denote by $\Gamma_t^i = \frac{v_{\mathbb{P}}^i(t, \nu_t^i)}{2\sqrt{\nu_t^i}} \sum_{j \in \{1, \dots, N^i\}} q_t^{i,j} \mathcal{V}_t^{i,j}$, we can write the market-maker's problem as

$$\sup_{\delta \in \mathcal{A}} \mathbb{E} \left[V_T - \frac{\gamma}{2} \sum_{(i,k) \in \{1, \dots, d\}^2} \int_0^T \Gamma_t^i \Gamma_t^k \Sigma^{\nu, i, k} dt \right]. \quad (7.2)$$

Here we penalize the portfolio's total Vega. Any other penalization could be used, as long as it is quadratic in q . For example, this includes more complicated penalties linked to another position to hedge, or some target for the Greeks. We define the Hamiltonians

$$H^{i,j,a}(S, \nu, p) = \sup_{\delta} \lambda^{i,j,a}(S, \nu, \delta)(\delta - p), \quad H^{i,j,b}(S, \nu, p) = \sup_{\delta} \lambda^{i,j,b}(S, \nu, \delta)(\delta - p),$$

and the following processes $\mathcal{G}(t, S, \nu) \in \mathbb{R}^{\sum_{l=1}^d N^l}$ such that

$$\begin{aligned} \mathcal{G}_j(t, S, \nu) &= \mathcal{V}_t^{k_j, j - (\sum_{l=1}^{k_j-1} N^l)} \frac{a_{\mathbb{P}}^{k_j}(t, \nu^{k_j}) - a_{\mathbb{Q}}^{k_j}(t, \nu^{k_j})}{2\sqrt{\nu^{k_j}}} \\ &\quad + \rho^{k_j} (\mathcal{VA})_t^{k_j, j - (\sum_{l=1}^{k_j-1} N^l)} \frac{v_{\mathbb{P}}^{k_j}(t, \nu^{k_j}) - v_{\mathbb{Q}}^{k_j}(t, \nu^{k_j})}{2\sqrt{\nu^{k_j}}} \sigma^{k_j}(t, S^{k_j}, \nu^{k_j}) \\ &\quad + (\mathcal{VO})_t^{k_j, j - (\sum_{l=1}^{k_j-1} N^l)} \frac{(v_{\mathbb{P}}^{k_j}(t, \nu^{k_j}))^2 - (v_{\mathbb{Q}}^{k_j}(t, \nu^{k_j}))^2}{4\nu^{k_j}}, \end{aligned}$$

where $k_j = i$ if $j \in \{\sum_{l=1}^{i-1} N^l, \dots, \sum_{l=1}^i N^l\}$, for $i \in \{1, \dots, d\}$. We also define $\mathcal{R}(t, S, \nu) \in \mathcal{M}_{\sum_{l=1}^d N^l, d}(\mathbb{R})$ such that

$$\mathcal{R}_{j,i}(t, S, \nu) = \frac{v_{\mathbb{P}}^i(t, \nu^i)}{2\sqrt{\nu_t^i}} \mathcal{V}_t^{i, j - (\sum_{l=1}^{k_j-1} N^l)}, \quad \text{for } j \in \left\{ \sum_{l=1}^{i-1} N^l, \dots, \sum_{l=1}^i N^l \right\}, i \in \{1, \dots, d\},$$

and 0 otherwise. Finally, denote the diffusion part of the HJB equation as

$$\begin{aligned}
 \mathcal{L}(t, S, \nu, q, u) = & \sum_{i \in \{1, \dots, d\}} b_{\mathbb{P}}^i(t, S^i) \partial_{S^i} u(t, S, \nu, q) + \sum_{i \in \{1, \dots, d\}} a_{\mathbb{P}}^i(t, \nu^i) \partial_{\nu^i} u(t, S, \nu, q) \\
 & + \frac{1}{2} \sum_{(i,k) \in \{1, \dots, d\}^2} \partial_{S^i S^k} u(t, S, \nu, q) \sigma^i(t, S^i, \nu^i) \sigma^j(t, S^k, \nu^j) \Sigma^{S^i, k} \\
 & + \sum_{i \in \{1, \dots, d\}} \int_{\mathbb{R}} \kappa^i(dz) \left(u(t, S + e^i \gamma^i(t, z), \nu, q) - u(t, S, \nu, q) \right) \\
 & + \frac{1}{2} \sum_{(i,k) \in \{1, \dots, d\}^2} \partial_{\nu^i \nu^j} u(t, S, \nu, q) v_{\mathbb{P}}^i(t, \nu^i) v_{\mathbb{P}}^k(t, \nu^k) \Sigma^{\nu^i, k} \\
 & + \sum_{i \in \{1, \dots, d\}} \partial_{\nu^i S^i} u(t, S, \nu, q) \rho^i v_{\mathbb{P}}^i(t, \nu^i) \sigma^i(t, S^i, \nu^i).
 \end{aligned}$$

The HJB equation associated to (7.2) with compact notations is

$$\begin{aligned}
 0 = & \partial_t u(t, S, \nu, q) + \mathcal{L}(t, S, \nu, q, u) + q^\top \mathcal{G}(t, S, \nu) - \frac{\gamma}{2} q^\top \mathcal{R}(t, S, \nu) \Sigma^\nu \mathcal{R}^\top(t, S, \nu) q \\
 & + \sum_{(i,j) \in \{1, \dots, d\} \times \{1, \dots, N^i\}} \int_{\mathbb{R}_+} z H^{i,j,b} \left(S, \nu, \frac{u(t, S, \nu, q) - u(t, S, \nu, q + z e^{i,j})}{z} \right) \mu^{i,j,b}(dz) \\
 & + \sum_{(i,j) \in \{1, \dots, d\} \times \{1, \dots, N^i\}} \int_{\mathbb{R}_+} z H^{i,j,a} \left(S, \nu, \frac{u(t, S, \nu, q) - u(t, S, \nu, q - z e^{i,j})}{z} \right) \mu^{i,j,a}(dz).
 \end{aligned} \tag{7.3}$$

with terminal condition $u(T, S, \nu, q) = 0$. The proof of existence and uniqueness of a viscosity solution to (7.3) associated to the control problem (7.2) relies on classic arguments of second order viscosity solutions with jumps, see for example Barles and Imbert [36], Bergault and Guéant [41] and Chapter 3.

7.3 Solving the market-maker's problem with a system of non-linear PDEs

Equation (7.3) is intractable with classical numerical methods when dealing with several options on several underlyings. Notably, the method proposed in Chapter 6 of this thesis to overcome the constant Vega assumption requires Monte-Carlo simulations of high-dimensional inventory trajectories, which is very time-consuming.

In this section, inspired by Bergault, Evangelista, Guéant, and Vieira [42], we propose an approximation of the value function of the market-maker, quadratic with respect to the vector of inventories to reduce the dimensionality of the problem. A Taylor expansion at 0 on the third variable with respect to ϵ gives

$$\begin{aligned}
 & H^{i,j,b} \left(S, \nu, \frac{u(t, S, \nu, q) - u(t, S, \nu, q + \epsilon z e^{i,j})}{z} \right) + H^{i,j,a} \left(S, \nu, \frac{u(t, S, \nu, q) - u(t, S, \nu, q - \epsilon z e^{i,j})}{z} \right) \\
 & = H^{i,j,b}(S, \nu, 0) + H^{i,j,a}(S, \nu, 0) + \epsilon (H'^{i,j,a}(S, \nu, 0) - H'^{i,j,b}(S, \nu, 0)) \partial_q u(t, S, \nu, q) \\
 & \quad + \frac{\epsilon^2}{2} \left(H''^{i,j,a}(S, \nu, 0) (\partial_q u(t, S, \nu, q))^2 - z H'^{i,j,a}(S, \nu, 0) \partial_{qq} u(t, S, \nu, q) \right) \\
 & \quad + \frac{\epsilon^2}{2} \left(H''^{i,j,b}(S, \nu, 0) (\partial_q u(t, S, \nu, q))^2 - z H'^{i,j,b}(S, \nu, 0) \partial_{qq} u(t, S, \nu, q) \right) + o(\epsilon^3),
 \end{aligned}$$

and by taking $\epsilon = 1$, Equation (7.3) becomes

$$\begin{aligned}
0 &= \partial_t u(t, S, \nu, q) + \mathcal{L}(t, S, \nu, q, u) + q^\top \mathcal{G}(t, S, \nu) - \frac{\gamma}{2} q^\top \mathcal{R}(t, S, \nu) \Sigma^\nu \mathcal{R}^\top(t, S, \nu) q \\
&+ \sum_{(i,j) \in \{1, \dots, d\} \times \{1, \dots, N^i\}} \int_{\mathbb{R}_+} \left(H^{i,j,b}(S, \nu, 0) - H^{i,j,b}(S, \nu, 0) \partial_q u(t, S, \nu, q) \right. \\
&+ \left. \frac{1}{2} \left(H''^{i,j,b}(S, \nu, 0) (\partial_q u(t, S, \nu, q))^2 - z H^{i,j,b}(S, \nu, 0) \partial_{qq} u(t, S, \nu, q) \right) \right) \mu^{i,j,b}(dz) \\
&+ \sum_{(i,j) \in \{1, \dots, d\} \times \{1, \dots, N^i\}} \int_{\mathbb{R}_+} \left(H^{i,j,a}(S, \nu, 0) + H^{i,j,a}(S, \nu, 0) \partial_q u(t, S, \nu, q) \right. \\
&+ \left. \frac{1}{2} \left(H''^{i,j,a}(S, \nu, 0) (\partial_q u(t, S, \nu, q))^2 - z H^{i,j,a}(S, \nu, 0) \partial_{qq} u(t, S, \nu, q) \right) \right) \mu^{i,j,a}(dz).
\end{aligned}$$

In the following we will show how a simple ansatz, quadratic with respect to the vector of inventories, leads to significant simplifications. For the sake of the simplicity of the notation, assume that $H^{i,j,a} = H^{i,j,b} = H^{i,j}$ (extension to asymmetric intensities is straightforward). By setting

$$u(t, S, \nu, q) = \theta^0(t, S, \nu) + q^\top \theta^1(t, S, \nu) - q^\top \theta^2(t, S, \nu) q,$$

where $\theta^0 \in \mathbb{R}$, $\theta^1 \in \mathbb{R}^{\sum_{i=1}^d N^i}$, $\theta^2 \in \mathcal{M}_{\sum_{i=1}^d N^i}(\mathbb{R})$ are solutions of the following system of non-linear PDEs:

$$\left\{ \begin{aligned}
0 &= \partial_t \theta^0(t, S, \nu) + \bar{\mathcal{L}}(t, \theta^0, \nu, S) + 2 \sum_{(i,j) \in \{1, \dots, d\} \times \{1, \dots, N^i\}} H^{i,j}(S, \nu, 0) \\
&+ \int_{\mathbb{R}_+} \left(2z H^{i,j}(S, \nu, 0) \theta_{j,j}^2(t, S, \nu) + H''^{i,j}(S, \nu, 0) (\theta_j^1(t, S, \nu))^2 \right) \mu^{i,j}(dz) \\
0 &= \partial_t \theta^1(t, S, \nu) + \bar{\mathcal{L}}(t, \theta^1, \nu, S) + \mathcal{G}(t, S, \nu) + 4\theta^2(t, S, \nu) \text{diag}(H''(S, \nu, 0)) \theta^1(t, S, \nu) \\
0 &= \partial_t \theta^2(t, S, \nu) + \bar{\mathcal{L}}(t, \theta^2, \nu, S) - \frac{\gamma}{2} \mathcal{R}(t, S, \nu) \Sigma^\nu \mathcal{R}^\top(t, S, \nu) \\
&+ 4\theta^2(t, S, \nu) \text{diag}(H''(S, \nu, 0)) \theta^2(t, S, \nu),
\end{aligned} \right. \quad (7.4)$$

where

$$\begin{aligned}
\bar{\mathcal{L}}(t, \theta, \nu, S) &= \sum_{i \in \{1, \dots, N\}} b_{\mathbb{P}}^i(t, S^i) \partial_{S^i} \theta(t, S, \nu) + \sum_{i \in \{1, \dots, N\}} a_{\mathbb{P}}^i(t, \nu^i) \partial_{\nu^i} \theta(t, S, \nu) \\
&+ \frac{1}{2} \sum_{(i,k) \in \{1, \dots, d\}^2} (\partial_{S^i S^j} \theta(t, S, \nu) \sigma^i(t, S^i, \nu^i) \sigma^k(t, S^k, \nu^k) \Sigma^{S^i, i k} + \partial_{\nu^i \nu^k} \theta(t, S, \nu) v_{\mathbb{P}}^i(t, \nu^i) v_{\mathbb{P}}^k(t, \nu^k) \Sigma^{\nu^i, i, k}) \\
&+ \sum_{i \in \{1, \dots, d\}} \left(\partial_{\nu^i S^i} \theta(t, S, \nu) \rho^i v_{\mathbb{P}}^i(t, \nu^i) \sigma^i(t, S^i, \nu^i) + \int_{\mathbb{R}} \kappa^i(dz) \left(\theta(t, S + e^i \gamma^i(t, z), \nu) - \theta(t, S, \nu) \right) \right).
\end{aligned}$$

and $\theta^0(T, S, \nu) = 0$, $\theta^1(T, S, \nu) = \mathbf{0}_{\sum_{i=1}^d N^i, 1}$, $\theta^2(T, S, \nu) = \mathbf{0}_{\sum_{i=1}^d N^i}$. In system (7.4), one can note that the PDE with respect to θ^2 is independent from the two others, which reduces the overall complexity. It can easily be solved for a small number of underlyings and a large number of options using finite difference schemes. Note that a higher order expansion does not yield a polynomial solution. However, it is possible to truncate the high degree terms to obtain a polynomial solution. This does not lead to a significant change of the value function or the controls if the penalty term is at most quadratic.

We now show some numerical applications of the methodology.

7.4 Numerical results

To perform a comparison with respect to the existing methods, we first recall the methodology of Chapter 6 of this thesis. In this chapter, we considered a market-maker managing a book of options on a single underlying, and they suppose he is perfectly delta-hedged. We have the following set of market parameters:

- $d = 1, N^1 = N = 20$: there are 20 call options on a single underlying.
- Stock price at time $t = 0$: $S_0 = 100\text{€}$.
- Instantaneous variance at time $t = 0$: $\nu_0 = 0.04 \text{ year}^{-1}$.
- Heston model parameters: $b_{\mathbb{P}}(t, S) = \mu S$, $\sigma(t, S, \nu) = S\sqrt{\nu}$, $v_{\mathbb{P}}(t, \nu) = v_{\mathbb{Q}}(t, \nu) = \xi\sqrt{\nu}$, with $\xi = 0.7 \text{ year}^{-1}$.
- $a_{\mathbb{P}}(t, \nu) = \kappa_{\mathbb{P}}(\theta_{\mathbb{P}} - \nu)$, $a_{\mathbb{Q}}(t, \nu) = \kappa_{\mathbb{Q}}(\theta_{\mathbb{Q}} - \nu)$, with $\kappa_{\mathbb{P}} = \kappa_{\mathbb{Q}} = 2 \text{ year}^{-1}$, $\theta_{\mathbb{P}} = \theta_{\mathbb{Q}} = 0.04 \text{ year}^{-1}$.
- $Z(dt, dz) = 0$: there is no jump in the dynamics of the underlying.
- Spot-variance correlation: $\rho = -0.7$.

We consider the case of a market-maker dealing with 20 European call options written on that stock where the strike×maturity couples are the elements (K^j, T^j) , $j \in \{1, \dots, 20\}$ of the set $\mathcal{K} \times \mathcal{T}$, where

$$\mathcal{K} = \{97, 98, 99, 100\}, \quad \mathcal{T} = \{0.3 \text{ year}, 0.4 \text{ year}, 0.5 \text{ year}, 0.6 \text{ year}, 0.7 \text{ year}\}.$$

These market parameters provide the implied volatility surface as in Figure 7.1.




Figure 7.1: Implied volatility surface associated with the market parameters.

We consider mainly in-the-money options with maturity ranging from 3 to 6 months so that, due to the influence of both Vanna and Vomma, the Vega of the portfolio changes noticeably and the prices of options are non negligible.

We define the following intensity functions:

$$\Lambda^{j,a}(S, \nu, \delta) = \Lambda^{j,b}(S, \nu, \delta) = \frac{\lambda^j}{1 + \exp\left(\alpha + \frac{\beta}{\nu_i^j} \delta\right)},$$

for $j \in \{1, \dots, N\}$, where $\lambda^j = \frac{252 \times 50}{1 + 0.7 \times |S_0 - K^j|} \text{ year}^{-1}$, $\alpha = -0.7$, and $\beta = 10 \text{ year}^{\frac{1}{2}}$. The choice of λ^j corresponds to 50 requests per day for at-the-money options, and decreases to 13.2 for the most in-the-money options. The choice of α corresponds to a probability of $\frac{1}{1 + e^{-0.7}} \approx 66\%$ to trade when the answered quote is the mid-price (i.e $\delta = 0$). The choice of β corresponds to a probability of $\frac{1}{1 + e^{-0.8}} \approx 68\%$ to trade when the answered quote corresponds to an implied volatility 1% better for the client and a probability of $\frac{1}{1 + e^{-0.6}} \approx 64\%$ to trade when the answered quote corresponds to an implied volatility 1% worse for the client.

We assume transactions of constant size with $z^j = \frac{5 \times 10^5}{\mathcal{O}_0^j}$ contracts for option j , in other words, the measures $\mu^{j,b}$, $\mu^{j,a}$ are Dirac masses at z^j . This corresponds approximately to 500000€ per transaction. We finally set $T = 0.004 \text{ year}$ (i.e 1 day), and a risk aversion parameter $\gamma = 2 \cdot 10^{-5} \text{€}^{-1}$.

The HJB equation using the constant Vega assumption of Chapter 6 of this thesis is

$$0 = \partial_t u(t, \nu, \mathcal{V}^\pi) + a_{\mathbb{P}}(t, \nu) \partial_\nu u(t, \nu, \mathcal{V}^\pi) + \frac{1}{2} \nu \xi^2 \partial_{\nu\nu} u(t, \nu, \mathcal{V}^\pi) + \nu^\pi \frac{a_{\mathbb{P}}(t, \nu) - a_{\mathbb{Q}}(t, \nu)}{2\sqrt{\nu}} - \frac{\gamma \xi^2}{8} (\mathcal{V}^\pi)^2$$

$$+ \sum_{j \in \{1, \dots, N\}} z^j H^{j,b} \left(\frac{u(t, \nu, \mathcal{V}^\pi) - u(t, \nu, \mathcal{V}^\pi + z^j \mathcal{V}^j)}{z^j} \right) + \sum_{j \in \{1, \dots, N\}} z^j H^{j,a} \left(\frac{u(t, \nu, \mathcal{V}^\pi) - u(t, \nu, \mathcal{V}^\pi - z^j \mathcal{V}^j)}{z^j} \right),$$

with terminal condition $u(T, \nu, \mathcal{V}^\pi) = 0$, and

$$\mathcal{V}_t^\pi = \sum_{j \in \{1, \dots, N\}} z^j \mathcal{V}^j q_t^j, \quad H^{j,a/b}(p) = \sup_{\delta^{j,a/b}} \Lambda^{j,a/b}(\delta^{j,a/b})(\delta^{j,a/b} - p).$$

In the case where the Vegas are not constant, we use the following ansatz:

$$u(t, S, \nu, q) = \theta^0(t, S, \nu) + q^\top \theta^1(t, S, \nu) + q^\top \theta^2(t, S, \nu) q,$$

where $\theta^0 \in \mathbb{R}$, $\theta^1 \in \mathbb{R}^N$, $\theta^2 \in \mathcal{M}_N(\mathbb{R})$. Define

$$\tilde{\mathcal{L}}(t, S, \nu, \theta) = a_{\mathbb{P}}(t, \nu) \partial_\nu \theta(t, S, \nu) + \frac{1}{2} \nu \xi^2 \partial_{\nu\nu} \theta(t, S, \nu) + \frac{1}{2} \nu S^2 \partial_{SS} \theta(t, S, \nu) + \rho \nu S \xi \partial_{\nu S} \theta(t, S, \nu),$$

and assume symmetry of intensity functions, that is $H^{j,b} = H^{j,a} = H^j$, we obtain the following system of coupled PDEs:

$$\begin{cases} 0 = \partial_t \theta^0(t, S, \nu) + \tilde{\mathcal{L}}(t, S, \nu, \theta^0) + 2 \sum_{j \in \{1, \dots, N\}} H^j(S, \nu, 0) + 2 \sum_{j \in \{1, \dots, N\}} z^j H'^j(S, \nu, 0) \theta_{j,j}^2(t, S, \nu) \\ \quad + \sum_{j \in \{1, \dots, N\}} H''^j(S, \nu, 0) (\theta_j^1(t, S, \nu))^2, \\ 0 = \partial_t \theta^1(t, S, \nu) + \tilde{\mathcal{L}}(t, S, \nu, \theta^1) + \mathcal{V}_t \frac{a_{\mathbb{P}}(t, \nu) - a_{\mathbb{Q}}(t, \nu)}{2\sqrt{\nu}} + 4\theta^2(t, S, \nu) \text{diag}(H''(S, \nu, 0)) \theta^1(t, S, \nu), \\ 0 = \partial_t \theta^2(t, S, \nu) + \tilde{\mathcal{L}}(t, S, \nu, \theta^2) - \frac{\gamma \xi^2}{8} \text{diag}(\mathcal{V}_t) \frac{\mathbf{1}\mathbf{1}^\top}{N} \text{diag}(\mathcal{V}_t) + 4\theta^2(t, S, \nu) \text{diag}(H''(S, \nu, 0)) \theta^2(t, S, \nu), \end{cases} \quad (7.5)$$

where $\mathcal{V}_t = (\partial_{\sqrt{\nu}} O^1(t, S, \nu), \dots, \partial_{\sqrt{\nu}} O^N(t, S, \nu))^\top$, $\mathbf{1} = (1, \dots, 1)^\top \in \mathbb{R}^N$.

We first show in Figures 7.2 and 7.3 some plots of the value function obtained by solving (7.5).

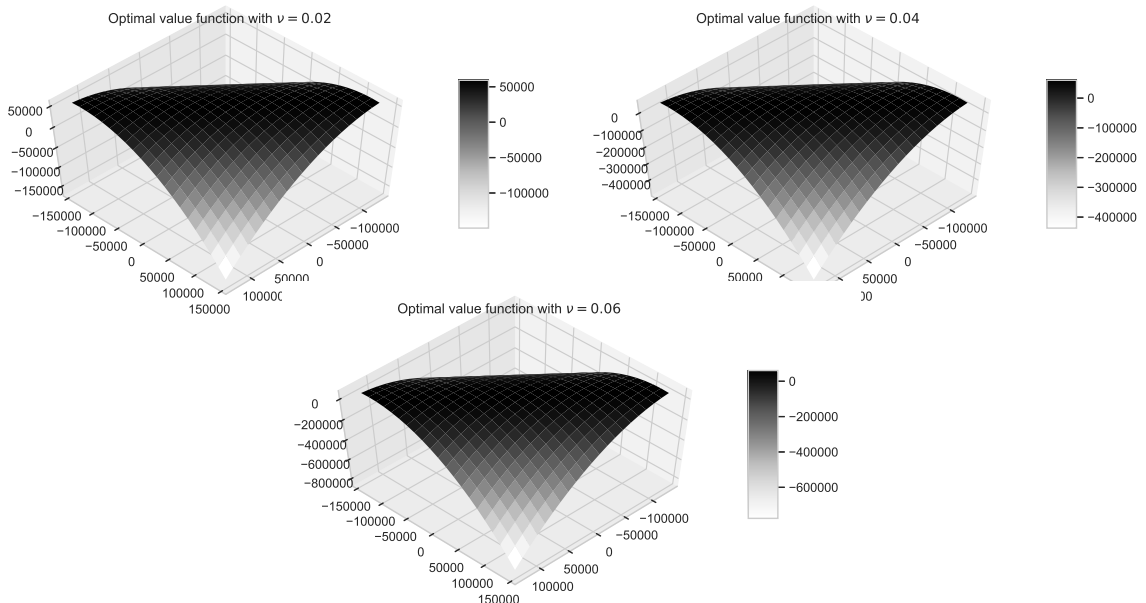


Figure 7.2: Value function for different inventories in (97, 0.3) and (98, 0.3) options, inventories in other options assumed to be equal 0, for different values of ν .

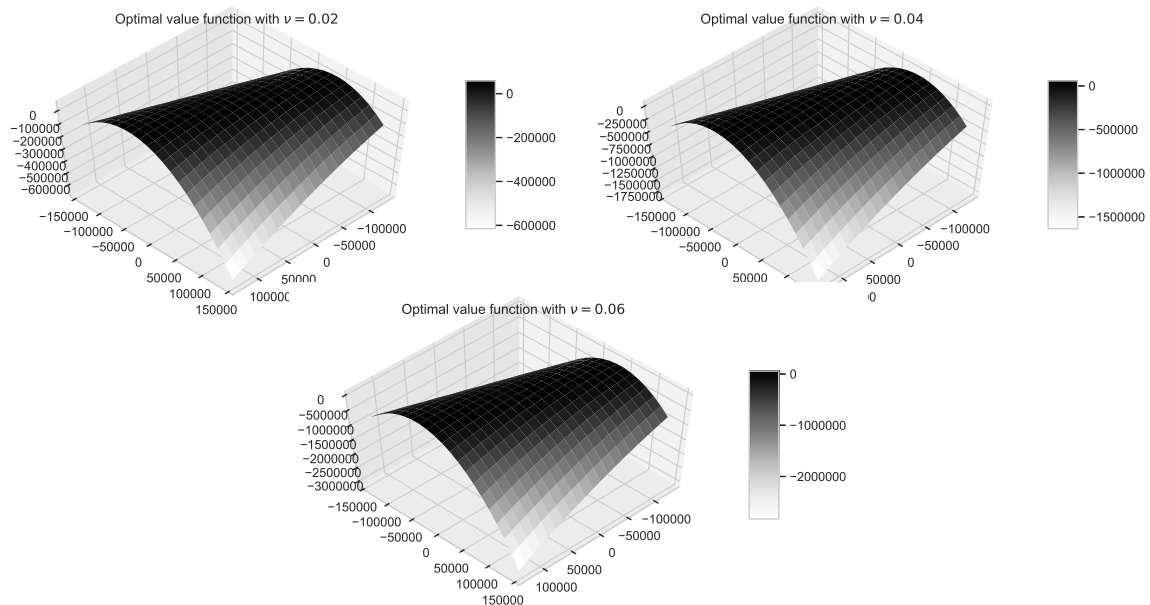


Figure 7.3: Value function for different inventories in $(97, 0.3)$ and $(100, 0.7)$ options, inventories in other options assumed to be equal 0, for different values of ν .

The value function often has higher values on the diagonals. The market-maker can compensate a long position in an option with a short position in another one. The values are noticeably lower for higher values of the volatility.

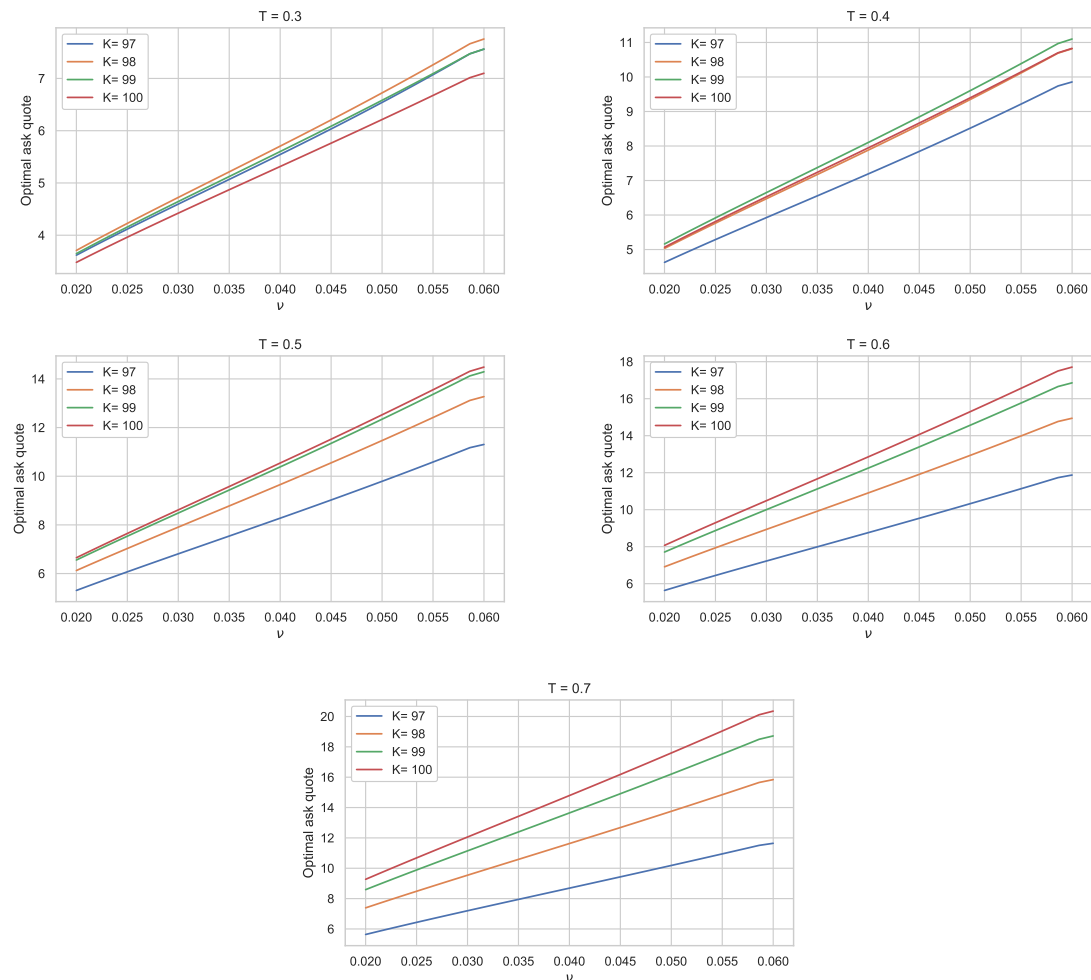


Figure 7.4: Optimal ask quotes with respect to ν for different options maturities.

We present in Figure 7.4 the evolution of the optimal ask quotes with respect to the stochastic volatility for the spot $S = 100$. We observe the usual increasing behavior of the optimal quotes with respect to both maturity and volatility of the underlying.

In Figure 7.5, we plot the evolution of the optimal ask quotes with respect to the underlying asset for the volatility $\nu = 0.04$.

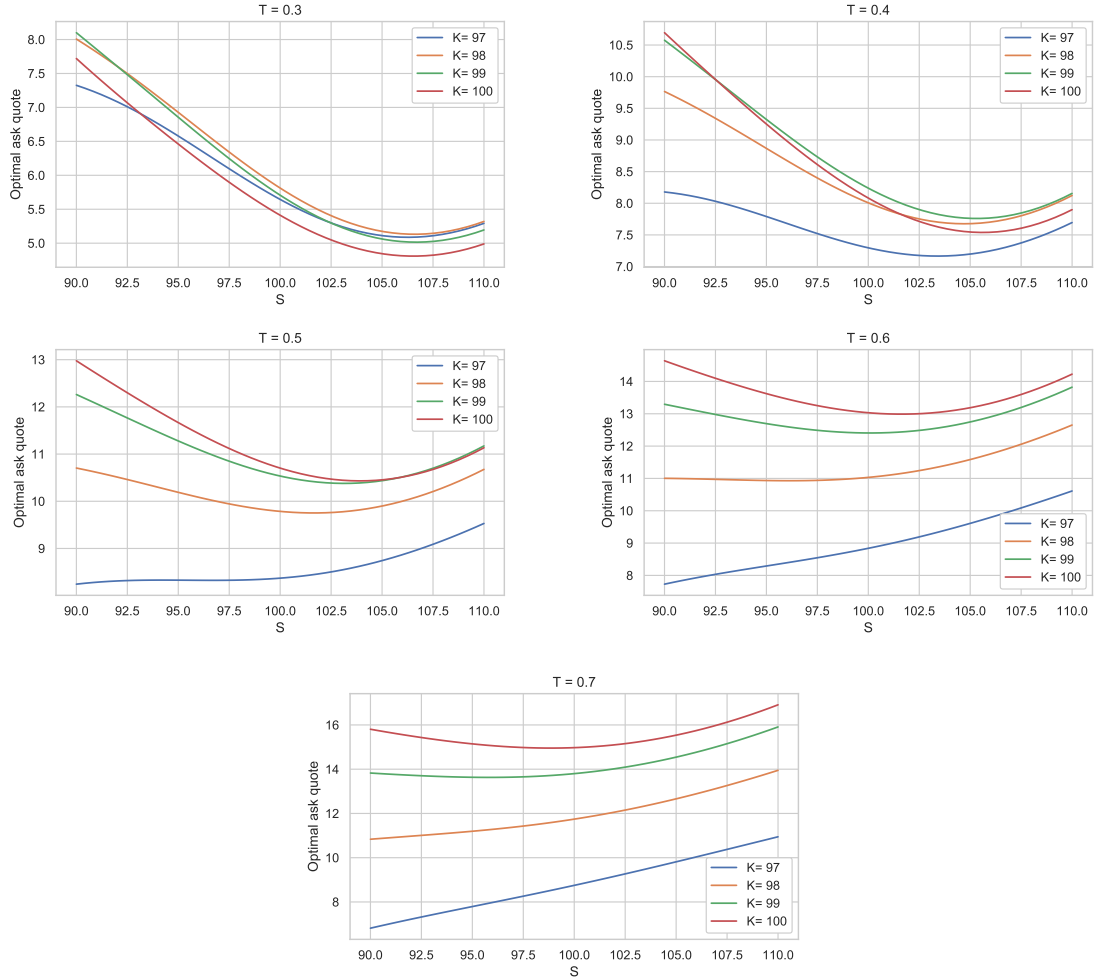


Figure 7.5: Optimal ask quotes with respect to S for different options maturities.

The behavior of the optimal quotes with respect to the strike depends on the expiry. We can see that the quotes are of the U-shaped nature, the quotes are decreasing in the spot price until some point depending on the strike and the expiry, and then become increasing. The inflection point decreases with the strike decreasing, and conversely for the expiry date. This way we can see that, for example, the quote for the option $(K, T) = (97, 0.7)$ is monotonously increasing in the spot price for the considered grid, which is fairly representative of the possible prices during one day. Conversely, for the option $(K, T) = (100, 0.3)$ the quote is decreasing for almost all values of the grid.

In Figure 7.6, we show the average PnL per request of the trader during the day over 1000 simulations, using the constant Greek approximation of Chapter 6 of this thesis and our algorithm.

At the beginning of the trading day, both methods yield a similar PnL per request. Notice that the PnL per request for the method with constant Greek approximation is slightly higher. Indeed the parameters at the beginning of the day correspond to the calibration parameters, and our algorithm is more conservative as it takes into account the risk that the underlying price could change. However, after roughly a tenth of the trading day the method with constant Greek approximations starts to

underperform our algorithm. This underperformance increases along the day as the constant Vega approximation becomes less accurate. On the contrary, with our method the PnL per request remains constant: there is no need for recalibration.




Figure 7.6: Average PnL per request over the trading day using constant and non-constant Greek approximations.

In Figure 7.7, we show one of 1000 simulation examples of the trajectories for the Vega of each option. We see that Vegas for this set of options are changing considerably during the day.




Figure 7.7: Example of Vega trajectories.




Figure 7.8: Cumulative distribution functions of the PnL over the trading day using both methods.

Finally, we present in Figure 7.8 the cumulative distribution function of the PnL of the trader using the constant Greek approximation of Chapter 6 of this thesis and our algorithm. We observe that the tail distribution of the PnL using our non-constant Greek approximation is higher compared to the method in Chapter 6.

7.A Appendix

7.A.1 The market-maker's problem for large number of underlyings

In this appendix, we present the system of low-dimensional PDEs analogous to (7.4) for more complex cases such as the market-making problem on several underlyings or the case where a number of different options' parameters is large (over one hundred).

We can rewrite the system of $(\sum_{i \in \{1, \dots, d\}} N^i)^2$ equations (7.4) on θ^2 as a set of d^2 equations by adding the strike and the maturity to the state variables. The same can be applied for the θ^1 equation. This way we obtain a smaller set of equations, though having more dimensions and some non-local terms.

Let $\mathbb{O}^i = \{(T^{i,j}, K^{i,j}), j \in \{1, \dots, N^i\}\}$ be the set of parameters of options on the underlying $i \in \{1, \dots, d\}$ and let us define $\hat{\theta}_i^1 : [0, T] \times \mathbb{R} \times \mathbb{R}_+ \times \mathbb{O}^i \rightarrow \mathbb{R}$ such that, for all $j \in \{1, \dots, N^i\}$,

$$\hat{\theta}_i^1(t, S, \nu, (T^{i,j}, K^{i,j})) = \theta_{\sum_{l=1}^{i-1} N^l + j}^1(t, S, \nu).$$

Similarly for $i_1, i_2 \in \{1, \dots, d\}$, define $\hat{\theta}_{i_1, i_2}^2 : [0, T] \times \mathbb{R} \times \mathbb{R}_+ \times \mathbb{O}^{i_1} \times \mathbb{O}^{i_2} \rightarrow \mathbb{R}$ such that, for any $j \in \{1, \dots, N^{i_1}\}$ and $l \in \{1, \dots, N^{i_2}\}$,

$$\hat{\theta}_{i_1, i_2}^2(t, S, \nu, (T^{i_1, j}, K^{i_1, j}), (T^{i_2, l}, K^{i_2, l})) = \theta_{\sum_{l=1}^{i_1-1} N^l + j, \sum_{l=1}^{i_2-1} N^l + l}^2(t, S, \nu).$$

Then the system of non-linear PDEs (7.4) can be rewritten as

$$\left\{ \begin{array}{l} 0 = \partial_t \theta^0(t, S, \nu) + \bar{\mathcal{L}}(t, S, \nu, \theta^0) + 2 \sum_{i \in \{1, \dots, d\}} \sum_{(T, K) \in \mathbb{O}^i} H^i(S, \nu, 0)(T, K) \\ \quad + 2 \sum_{i \in \{1, \dots, d\}} \sum_{(T, K) \in \mathbb{O}^i} \int_{\mathbb{R}_+} z H'^i(S, \nu, 0)(T, K) \hat{\theta}_{i, i}^2(t, S, \nu, (T, K), (T, K)) \mu^{i, (T, K)}(dz) \\ \quad + \sum_{i \in \{1, \dots, d\}} \sum_{(T, K) \in \mathbb{O}^i} H''^i(S, \nu, 0)(T, K) \left(\hat{\theta}_i^1(t, S, \nu, (T, K)) \right)^2, \\ 0 = \partial_t \hat{\theta}_i^1(t, S, \nu, (\mathcal{T}^1, \mathcal{K}^1)) + \bar{\mathcal{L}}^1(t, S, \nu, \hat{\theta}_i^1, (\mathcal{T}^1, \mathcal{K}^1)) + \mathcal{G}_i(t, S, \nu, (\mathcal{T}^1, \mathcal{K}^1)) \\ \quad + 4 \sum_{i_2 \in \{1, \dots, d\}} \sum_{(T, K) \in \mathbb{O}^{i_2}} \hat{\theta}_{i, i_2}^2(t, S, \nu, (\mathcal{T}^1, \mathcal{K}^1), (T, K)) H''^{i_2}(S, \nu, 0)(T, K) \hat{\theta}_{i_2}^1(t, S, \nu, (T, K)), \\ 0 = \partial_t \hat{\theta}_{i_1, i_2}^2(t, S, \nu, (\mathcal{T}^1, \mathcal{K}^1), (\mathcal{T}^2, \mathcal{K}^2)) + \bar{\mathcal{L}}^2(t, S, \nu, \hat{\theta}_{i_1, i_2}^2, (\mathcal{T}^1, \mathcal{K}^1), (\mathcal{T}^2, \mathcal{K}^2)) \\ \quad - \frac{\gamma}{2} \mathcal{R}_{i_1}(t, S, \nu, (\mathcal{T}^1, \mathcal{K}^1)) \Sigma^{\nu, i_1, i_2} \mathcal{R}_{i_2}(t, S, \nu, (\mathcal{T}^2, \mathcal{K}^2)) \\ \quad + 4 \sum_{i_3 \in \{1, \dots, d\}} \sum_{(T, K) \in \mathbb{O}^{i_3}} \hat{\theta}_{i_1, i_3}^2(t, S, \nu, (\mathcal{T}^1, \mathcal{K}^1), (T, K)) \hat{H}''^{i_3}(S, \nu, 0)(T, K) \hat{\theta}_{i_3, i_2}^2(t, S, \nu, (T, K), (\mathcal{T}^2, \mathcal{K}^2)), \end{array} \right.$$

where $((\mathcal{T}^1, \mathcal{K}^1), (\mathcal{T}^2, \mathcal{K}^2)) \in (\prod_{i \in \{1, \dots, d\}} \mathbb{O}^i)^2$ and, for $j \in \{1, \dots, N^{i_1}\}$, $l \in \{1, \dots, N^{i_2}\}$,

$$\begin{aligned} H^i(S, \nu, 0)(T^{i,j}, K^{i,j}) &= H^{i,j}(S, \nu, 0), \quad \mathcal{G}_i(t, S, \nu, (T^{i,j}, K^{i,j})) = \mathcal{G}(t, S, \nu)_{\sum_{l=1}^{i-1} N^l + j}, \\ \mathcal{R}_i(t, S, \nu, (T^{i,j}, K^{i,j})) &= \mathcal{R}(t, S, \nu)_{\sum_{l=1}^{i-1} N^l + j, i}, \quad \mu^{i, (T^{i,j}, K^{i,j})} = \mu^{i, j}, \\ \bar{\mathcal{L}}^1(t, S, \nu, \hat{\theta}_i^1, (T^{i,j}, K^{i,j})) &= \bar{\mathcal{L}}(t, S, \nu, \theta^1)_{\sum_{l=1}^{i-1} N^l + j}, \\ \bar{\mathcal{L}}^2(t, S, \nu, \hat{\theta}_{i_1, i_2}^2, (T^{i_1, j}, K^{i_1, j}), (T^{i_2, l}, K^{i_2, l})) &= \bar{\mathcal{L}}(t, S, \nu, \theta^2)_{\sum_{l=1}^{i_1-1} N^l + j, \sum_{l=1}^{i_2-1} N^l + l}. \end{aligned}$$

In particular, if $d = 1$, $\hat{\theta}^1$ and $\hat{\theta}^2$ are solutions of non-local PDEs in dimensions 5 and 7 respectively. The observed regularity of the solution with respect to the strike and expiry implies that the high-dimensional PDEs can be solved, for example, by a non-local variant of the Deep Galerkin Method, see Hirs and Fu [155], Sirignano and Spiliopoulos [251].

Chapter 8

High-frequency dynamics of the volatility surface

8.1 Introduction

In 1973, the seminal paper of Black and Scholes [48] has drastically changed the options markets and their operation. In practice, even if their model is not always applied directly, being replaced by newer and more complex ones, the notions introduced thanks to this work are still very much in use. One of these concepts is, of course, the implied volatility, which is defined as the volatility such that the market price of an option coincides with its Black-Scholes counterpart. By finding a mapping of strikes and maturities to corresponding implied volatilities, one could build a volatility surface, which became one of the options market's crucial tools.

Empirical evidence from options markets shows that the surface exhibits a non-constant behavior, such as skew, smile, and a term structure, see, for example, Derman [98], Derman, Kani, and Zou [99], Engle and Rosenberg [112]. In order to capture these features, the researchers and analysts started to develop more intricate models, mostly by changing the dynamics of the underlying asset: by introducing, for example, multiple factors and jumps to the diffusion. However, this way of modeling the volatility surface properties has two significant issues. First, the implied volatility surface can only be computed numerically in those frameworks, which is especially arduous for jump-diffusion models. Second, these models assume that the volatility surface dynamics are driven solely by the underlying; however, options markets have their supply-offer equilibrium, which cannot be described entirely by the underlying's dynamics, see, for example, Bakshi, Cao, and Chen [35]. By considering the options' Delta as the only source of risk, these models fail to accurately reproduce the surface's dynamics, which leads to frequent recalibrations. Thus, it seems more relevant to model the dynamics of the implied volatility surface directly.

Another challenge in the volatility surface's modeling is the time scale. As we can see from the example of stocks' prices, the price process's behavior differs at thin and coarse time scales, and we want to take it into account for the volatility surface modeling. The diffusion models are usually used to obtain realistic dynamics for the underlying's volatility at coarse time scale. However,

- The local volatility model reproduces the whole volatility surface at a given time but cannot reproduce the dynamics of the volatility observed in the market.
- Stochastic volatility models can accurately reproduce the volatility surface's dynamics but cannot fit both realized and implied volatility. In particular, the at-the-money skew of the volatility surface cannot be calibrated accurately using Brownian motion as a risk factor.

- The rough volatility models, whose recent development has been motivated by the seemingly universal rough behavior of financial assets, capture key features of the implied volatility surface and its dynamics, see Bayer, Friz, and Gatheral [38], El Euch, Gatheral, and Rosenbaum [108], Gatheral, Jaisson, and Rosenbaum [128].¹ However, as mentioned above, the calculation efforts to reproduce the volatility surface in this framework are considerably high.

The literature dealing with the implied volatility surface’s direct modeling at the daily or more coarse time scale usually uses the principal component analysis to extract the volatility surface’s main drivers, for example, Cont and Da Fonseca [80], Fengler, Härdle, and Mammen [116], Kamal and Gatheral [181], Skiadopoulos, Hodges, and Clewlow [252]. The resulting implied volatility surface is usually represented as a randomly fluctuating surface, whose deformation is driven by a small number of orthogonal random factors:

- The “level” factor corresponding to the global level of the whole surface of implied volatilities.
- The “calendar” factor corresponding to the skew of the volatility surface.
- The “butterfly” factor corresponding to the convexity of the volatility surface.

At the thin time scale, such as intraday, volatility surface dynamics models have been studied mainly from an empirical viewpoint, for example, in Dunis, Kellard, and Snaith [102], Koopman, Jungbacker, and Hol [187], Lee and Ryu [196], Mayhew [208], Wang and Wang [266]. In these articles, the authors use time-series models on intraday high-frequency implied volatility data to forecast implied and realized volatility behavior. To the best of our knowledge, there is no model aimed at reproducing the intraday dynamics of the volatility surface, which is, even more, the case for thinner time scales as the high-frequency scale.

Because of the massive amount of high-frequency data, it is problematic to obtain an accurate estimation of volatility and covariance between two or more assets over a given period. Moreover, trading times and price processes are discrete, which makes diffusion models hardly usable. To solve these two issues, we go on with a particular class of models, called tick-by-tick models, inspired by models of high-frequency asset prices using Hawkes processes. For example, in Bacry, Delattre, Hoffmann, and Muzy [31], Bacry and Muzy [30], the authors consider a Hawkes processes-based model, where they assume that all price jumps are of the same size, the microscopic price of the asset is the difference of the number of upward and downward jumps, that is

$$P_t = P_0 + N_t^+ - N_t^-,$$

where (N^+, N^-) is a bi-dimensional Hawkes process with intensity kernel

$$\phi(t) = \begin{pmatrix} \phi^{++}(t) & \phi^{+-}(t) \\ \phi^{-+}(t) & \phi^{--}(t) \end{pmatrix}.$$

Set ϕ is a set of endogenous sources of price moves: for example, ϕ^{+-} increases the intensity of upward price jumps after a downward price jump, creating a mean-reversion effect while ϕ^{++} creates a trending effect. This model reproduces some well-known stylized facts of the high-frequency price behavior (absence of arbitrage, long memory of the order flow, bid-ask symmetry, high degree of endogeneity of the markets) which can be incorporated into the Hawkes kernel. Moreover, it enables to obtain closed-form expressions for important statistical measures of high-frequency financial data, such as the signature plot or Epps effect. This model has been extended to the multi-asset case in Bacry, Delattre, Hoffmann, and Muzy [32], where the authors show the reproduction of the lead-lag effect between two assets, more general multi-asset case is studied in Tomas and Rosenbaum [263]. Finally, note that one important advantage of the Hawkes-based models is that their long-term behavior is consistent with financial assets’ roughness. At coarse time scale, the Hawkes price process converges to

¹The log-volatility process of an asset is well-approximated by a fractional Brownian motion with small Hurst parameter $H \approx 0.1$.

a stochastic rough volatility process. This property has raised interest in building microscopic models for market dynamics that reproduce rough volatility at a macroscopic scale, see Jaisson and Rosenbaum [174, 175], Jusselin and Rosenbaum [177], Tomas and Rosenbaum [263].

In this chapter, we propose a tick-by-tick model for the high-frequency dynamics of the volatility surface. Assuming $\#\mathcal{K}$ strikes and $\#\mathcal{T}$ maturities, the microscopic volatility surface is modeled as a $\#\mathcal{K} \times \#\mathcal{T}$ -dimensional process $(\sigma_t^{(k,\tau)})_{(k,\tau) \in \mathcal{K} \times \mathcal{T}}$, where

$$\sigma_t^{k,\tau} = N_t^{(k,\tau)+} - N_t^{(k,\tau)-},$$

and $(N_t^{(k,\tau)+}, N_t^{(k,\tau)-})_{(k,\tau) \in \mathcal{K} \times \mathcal{T}}$ is a $2 \times \#\mathcal{K} \times \#\mathcal{T}$ -dimensional Hawkes process. The processes $N_t^{(k,\tau)+}$ (resp. $N_t^{(k,\tau)-}$) count the number of upward (resp. downward) moves of the implied volatility of option with strike k and maturity τ . This modelling is especially well suited to Foreign Exchange (FX) options markets, where options are directly quoted in terms of implied volatility.²

We show how the coefficients of the Hawkes kernel govern the skew and convexity of the surface. Moreover, they can be parametrized in a way that no-arbitrage conditions of the volatility are satisfied. In addition, this parametrization of the volatility surface reduces the number of parameters to calibrate. We also study classic Hawkes kernels such as power-law, and more precisely, the kernel's parameters influence on the shape of the surface. The conditions provide practical ways to parametrize the volatility surface, reproducing most of the market's stylized facts.

To verify the model's properties at large time scales, we use the limit theorems as in Tomas and Rosenbaum [263] to show, for example, that our volatility surface behaves like a diffusion process. We obtain that a sum of orthogonal factors drives the volatility surface dynamics, whose volatility processes are rough. For example, the highest eigenvalue of ϕ is the "level" of the volatility surface. Hence, our results at the macroscopic scale are in the spirit of Cont and Da Fonseca [80], except that the risk factors driving the volatility surface come from the microscopic interactions between market participants. We show that our microscopic model can be used as a backtesting environment for trading and market-making strategies on options. At the macroscopic time scale, it can be used to compute these strategies' impact on the volatility surface.

The chapter is structured as follows. In Section 8.2, we introduce the model for the high-frequency dynamics of the implied volatility surface and provide sufficient conditions to ensure absence of arbitrage. In Section 8.3, we study the macroscopic behavior of the volatility surface. Finally, Section 8.4 is devoted to the numerical applications of the model to option market-making strategies.

8.2 Microscopic modelling of the volatility surface

8.2.1 Framework and no-arbitrage conditions

In this section, we first formalize our tick-by-tick model for the microscopic dynamics of the volatility surface. Let us consider a time horizon $T > 0$ and a set of options on an underlying asset with strikes $\{k_1, \dots, k_N\}$ and maturities $\{\tau_1, \dots, \tau_M\}$, $N, M \in \mathbb{N}^*$.³ At the microscopic scale, the point (k, τ) of the volatility surface is a piecewise constant process with upward and downward jumps defined by the following tick-by-tick model:

$$\sigma_t^{(k,\tau)} = N_t^{(k,\tau)+} - N_t^{(k,\tau)-}, \quad t \geq 0, (k, \tau) \in \{k_1, \dots, k_N\} \times \{\tau_1, \dots, \tau_M\} = \mathcal{K} \times \mathcal{T},$$

²It can also be used on markets where options are quoted in prices, by assuming for example sticky moneyness of the volatility surface.

³The strikes are expressed in log-moneyness, that is $k_i = \log(\frac{K_i}{S})$ where S is the price of the underlying asset and $K_i > 0$ for $i \in \{1, \dots, N\}$.

where $N_t^{(k,\tau)+}$ counts the number of upward volatility moves, and $N_t^{(k,\tau)-}$ counts the number of downward volatility moves. The vector $\mathbf{N} = (N^{(k,\tau)+}, N^{(k,\tau)-})_{(k,\tau) \in \mathcal{K} \times \mathcal{T}}$ is a $2 \times \#\mathcal{K} \times \#\mathcal{T}$ -dimensional Hawkes process with intensity

$$\boldsymbol{\lambda}_t = \boldsymbol{\mu}_t + \int_0^t \boldsymbol{\phi}(t-s) d\mathbf{N}_s.$$

Definition 8.1. *The microscopic volatility surface is defined by the $\mathcal{M}_{\#\mathcal{K} \times \#\mathcal{T}}(\mathbb{R})$ -valued process $\boldsymbol{\sigma}_t = (\sigma_t^{(k,\tau)})$ for $t \geq 0$, $(k, \tau) \in \mathcal{K} \times \mathcal{T}$.*

This definition is related to the notion of market impact: a net buy pressure on option (k, τ) leads to a higher implied volatility $\sigma^{(k,\tau)}$ and conversely for a net sell pressure. This tick-by-tick modelling is commonly used for price processes, see for example El Euch, Fukasawa, and Rosenbaum [107], Tomas and Rosenbaum [263]. As stated in the introduction, the use of a tick-by-tick model for the moves of the volatility surface is particularly well suited to FX options markets, where the options are directly quoted in terms of implied volatility. For other markets, as the price of an option is an increasing bijection of the implied volatility function, the impact on prices and implied volatility is qualitatively the same.

Remark 39. *Throughout the chapter, we assume implicitly that the volatility process of an option (k, τ) does not make unitary jumps but rather jumps equal to a fixed tick size. This tick size is assumed to be the same for all strikes and maturities for ease of notation. However, this assumption can be relaxed without a change of the key results.*

Let the vector $\boldsymbol{\mu} : \mathbb{R}_+ \rightarrow \mathbb{R}_+^{\#\mathcal{K} \times \#\mathcal{T}}$ denote the baseline (that is the intensity vector without self and mutual excitation) and $\boldsymbol{\phi} : \mathbb{R}_+ \rightarrow \mathcal{M}_{2 \times \#\mathcal{K} \times \#\mathcal{T}}(\mathbb{R})$ denote the kernel encoding the self- and cross-excitation of the different points of the volatility surface. These terms deserve some financial interpretations:

- The process $\mu^{(k,\tau)+}(t)$ (resp. $\mu^{(k,\tau)-}(t)$) is the exogenous source of upward (resp. downward) moves of the volatility surface at point (k, τ) . For example, the exogenous sources of jumps corresponding to at-the-money options have higher values than far-from-the-money options.
- The kernel $\boldsymbol{\phi}$ describes the endogenous influence of the past volatility moves to the current intensities, namely self- and cross-excitation. We denote by $\phi^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}}$ the influence of $N^{(\tilde{k},\tilde{\tau})\tilde{s}}$ on $N^{(k,\tau)s}$, $(s, \tilde{s}) \in \{+, -\}^2$. For example, the quantity $\phi^{(k_i, \tau_j)+, (k_i, \tau_j)+}$ has a trending self-exciting role: it increases the intensity of upward volatility jumps of option (k_i, τ_j) after an upward volatility jump of the same option (k_i, τ_j) , similarly for $\phi^{(k_i, \tau_j)-, (k_i, \tau_j)-}$. Or, for instance, for $(i_1, i_2), (j_1, j_2)$, the quantity $\phi^{(k_{i_1}, \tau_{j_1})+, (k_{i_2}, \tau_{j_2})-}$ has a mean-reverting cross-exciting role: it increases the intensity of upward volatility jumps of option (k_{i_1}, τ_{j_1}) after a downward volatility jump of the option (k_{i_2}, τ_{j_2}) .

The kernel $\boldsymbol{\phi}$ models the shape of the volatility surface at the microscopic scale. The following two examples aim at describing how the shape of the volatility surface is linked to the Hawkes kernel.

Example 8.1. Volatility skew. *Let us take a single slice of volatility,⁴ that is $\mathcal{T} = \{\tau\}$, and, for the sake of clarity, consider only two options $(k_1, \tau), (k_2, \tau)$. Let us assume a volatility skew, typical for equity markets, that is for all $t \geq 0$ and $k_1 \leq k_2$,*

$$\mathbb{E}[\sigma_t^{(k_1, \tau)}] \geq \mathbb{E}[\sigma_t^{(k_2, \tau)}].$$

Moreover, let us assume the following kernel structure for the 4-dimensional Hawkes process:

$$\boldsymbol{\phi} = \begin{pmatrix} 0 & 0 & \phi^{(k_1, \tau)+, (k_2, \tau)+} & \phi^{(k_1, \tau)+, (k_2, \tau)-} \\ 0 & 0 & \phi^{(k_1, \tau)-, (k_2, \tau)+} & \phi^{(k_1, \tau)-, (k_2, \tau)-} \\ \phi^{(k_2, \tau)+, (k_1, \tau)+} & \phi^{(k_2, \tau)+, (k_1, \tau)-} & 0 & 0 \\ \phi^{(k_2, \tau)-, (k_1, \tau)+} & \phi^{(k_2, \tau)-, (k_1, \tau)-} & 0 & 0 \end{pmatrix},$$

⁴In this chapter, we refer to the set of options having the same maturity as the slice of volatility.

where we drop the time dependence of the kernel coefficients for the sake of readability. In this example we assume no self- or cross-excitation of the volatility processes on the same option, that is

$$\phi^{(k_i, \tau) s_1, (k_i, \tau) s_2} = 0, \quad i \in \{1, 2\}, s_i \in \{+, -\},$$

but only cross-excitation from option (k_1, τ) to option (k_2, τ) and conversely. Assume a high number of buy orders on option (k_2, τ) . This increase of the counting process $N^{(k_2, \tau)+}$ increases the level of volatility $\sigma^{(k_2, \tau)}$. However, it also impacts the left side of the volatility skew because of the coefficients $\phi^{(k_1, \tau)-, (k_2, \tau)+}$, $\phi^{(k_1, \tau)+, (k_2, \tau)+}$. The slice of volatility can be impacted in the following ways:

- If $\phi^{(k_1, \tau)-, (k_2, \tau)+} > \phi^{(k_1, \tau)+, (k_2, \tau)+}$, then, on average, an increase of the volatility for option (k_2, τ) decreases the volatility for option (k_1, τ) so that the volatility slice for maturity τ flattens.
- If $\phi^{(k_1, \tau)-, (k_2, \tau)+} < \phi^{(k_1, \tau)+, (k_2, \tau)+}$, then, on average, an increase of the volatility for option (k_2, τ) increases the volatility for option (k_1, τ) :
 - If the buy orders on option (k_2, τ) lead to more buy orders on option (k_1, τ) , then the volatility skew steepens.
 - If the buy orders on option (k_2, τ) lead to less buy orders on option (k_1, τ) , then the volatility skew flattens.

These cases are represented in Figure 8.1 below.




Figure 8.1: Evolution of the slice of volatility when $\phi^{(k_1, \tau)-, (k_2, \tau)+} > \phi^{(k_1, \tau)+, (k_2, \tau)+}$ (left), $\phi^{(k_1, \tau)-, (k_2, \tau)+} < \phi^{(k_1, \tau)+, (k_2, \tau)+}$ and there is less buy orders on option (k_2, τ) (center), $\phi^{(k_1, \tau)-, (k_2, \tau)+} < \phi^{(k_1, \tau)+, (k_2, \tau)+}$ and there is more buy orders on option (k_2, τ) (right)

Example 8.2. Volatility convexity. Let us take a slice of volatility, that is $\mathcal{T} = \{\tau\}$, and consider only three options $(k_1, \tau), (k_2, \tau), (k_3, \tau)$. Assume a volatility smile, as in FX markets, for example, such that for all $t \geq 0$,

$$\mathbb{E}[\sigma_t^{(k_1, \tau)}] \geq \mathbb{E}[\sigma_t^{(k_2, \tau)}], \quad \mathbb{E}[\sigma_t^{(k_2, \tau)}] \leq \mathbb{E}[\sigma_t^{(k_3, \tau)}]$$

for $k_1 \leq k_2 \leq k_3$. Thus, we consider that k_1 is an in-the-money strike, k_2 is an at-the-money strike and k_3 is an out-of-the-money strike. Moreover, assume the following kernel structure for the 6-dimensional Hawkes process:

$$\phi = \begin{pmatrix} 0 & 0 & 0 & 0 & \phi^{(k_1, \tau)+, (k_3, \tau)+} & \phi^{(k_1, \tau)+, (k_3, \tau)-} \\ 0 & 0 & 0 & 0 & \phi^{(k_1, \tau)-, (k_3, \tau)+} & \phi^{(k_1, \tau)-, (k_3, \tau)-} \\ 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 \\ \phi^{(k_3, \tau)+, (k_1, \tau)+} & \phi^{(k_3, \tau)+, (k_1, \tau)-} & 0 & 0 & 0 & 0 \\ \phi^{(k_3, \tau)-, (k_1, \tau)+} & \phi^{(k_3, \tau)-, (k_1, \tau)-} & 0 & 0 & 0 & 0 \end{pmatrix}.$$

In this example we assume no self- and cross-excitation of the volatility processes on the same options, and no self or cross-excitation coming from the at-the-money option (k_2, τ) . The only cross-excitation

comes from the out-of-the-money to the in-the-money option and conversely. The at-the-money part of the smile is not-affected by far-from-the-money transactions for sake of clarity of this example. We assume a high number of sell orders on (k_3, τ) .

- If $\phi^{(k_1, \tau)^-, (k_3, \tau)^-} > \phi^{(k_1, \tau)^+, (k_3, \tau)^-}$, the convexity of this slice of volatility decreases on average.
- If $\phi^{(k_1, \tau)^-, (k_3, \tau)^-} < \phi^{(k_1, \tau)^+, (k_3, \tau)^-}$, the volatility smile becomes (on average) left skewed.

Finally, assume a high number of buy orders on the out-of-the-money option (k_3, τ) . This increase of the counting process $N^{(k_3, \tau)^+}$ increases the level of volatility $\sigma^{(k_3, \tau)}$, and it also impacts the left side of the volatility smile.

- If $\phi^{(k_1, \tau)^-, (k_3, \tau)^+} < \phi^{(k_1, \tau)^+, (k_3, \tau)^+}$, left and right side of the volatility smile rise (on average) so that the convexity of this slice of volatility increases.
- If $\phi^{(k_1, \tau)^-, (k_3, \tau)^+} > \phi^{(k_1, \tau)^+, (k_3, \tau)^+}$, the volatility smile becomes right-skewed.

These four cases are represented in Figure 8.2 below.




Figure 8.2: Evolution of the slice of volatility with sell orders on (k_3, τ) and $\phi^{(k_1, \tau)^-, (k_3, \tau)^-} > \phi^{(k_1, \tau)^+, (k_3, \tau)^-}$ (upper left), $\phi^{(k_1, \tau)^-, (k_3, \tau)^-} < \phi^{(k_1, \tau)^+, (k_3, \tau)^-}$ (upper right). Evolution of the slice of volatility with buy orders on (k_3, τ) and $\phi^{(k_1, \tau)^-, (k_3, \tau)^+} < \phi^{(k_1, \tau)^+, (k_3, \tau)^+}$ (lower left), $\phi^{(k_1, \tau)^-, (k_3, \tau)^+} > \phi^{(k_1, \tau)^+, (k_3, \tau)^+}$ (lower right).

These two simple examples show that a suitable parametrization for a multidimensional Hawkes process can reproduce shapes of a tick-by-tick volatility surface. However, to ensure meaningful results, it is crucial to enforce no-arbitrage conditions on the volatility surface at the microscopic scale. Following Gatheral and Jacquier [125], we consider that a volatility surface is arbitrage-free if it is free of calendar spread and butterfly arbitrages. We now explain these two notions and show how to incorporate these properties into our framework.

A calendar spread arbitrage is an implied volatility arbitrage consisting in buying and selling two calls (or puts) on the same underlying with the same strike and different maturities. As the time value of

the option with a shorter maturity decreases quicker than the one with a longer maturity, the latter should have a higher price than the former. Thus, there is no arbitrage opportunity in a calendar spread strategy if European call option prices are monotonous with respect to the maturity. Assuming that dividends are proportional to the price of the underlying, a necessary and sufficient condition for the volatility surface to be free of calendar spread arbitrage is the following.

Proposition 8.1. *Let us define the total implied variance at time t for an option (k, τ) as $\varpi_t(k, \tau) = (\sigma_t^{(k, \tau)})^2 \tau$. The microscopic volatility surface is free of calendar spread arbitrage if and only if*

$$\partial_\tau \varpi_t(k, \tau) \geq 0 \text{ for all } (t, k, \tau) \in [0, T] \times \mathcal{K} \times \mathcal{T}.$$

The proof can be found in Gatheral and Jacquier [125]. This condition cannot be applied directly to our model as we work with a discrete implied volatility. To incorporate the no calendar spread arbitrage condition in the Hawkes kernel, we rewrite the condition of Proposition 8.1 as

$$\frac{(\sigma_t^{(k, \tau_i)})^2 \tau_i - (\sigma_t^{(k, \tau_j)})^2 \tau_j}{\tau_i - \tau_j} \geq 0, \quad \text{for all } k \in \mathcal{K}, (\tau_i, \tau_j) \in \mathcal{T}^2.$$

We recall that, for a given Hawkes process $N^{(k, \tau)s}$ with intensity $\lambda^{(k, \tau)s}$, $s \in \{+, -\}$, we have the following equalities:

$$\begin{aligned} \mathbb{E}[N_t^{(k, \tau)s}] &= \int_0^t \mathbb{E}[\lambda_u^{(k, \tau)s}] du, \\ \mathbb{E}[\lambda_t^{(k, \tau)s}] &= \mu^{(k, \tau)s} + \sum_{(\tilde{k}, \tilde{\tau})} \sum_{\tilde{s} \in \{+, -\}} \int_0^t \phi^{(k, \tau)s, (\tilde{k}, \tilde{\tau})\tilde{s}}(t - u) \mathbb{E}[\lambda_u^{(\tilde{k}, \tilde{\tau})\tilde{s}}] du, \end{aligned} \quad (8.1)$$

where, for sake of simplicity, we assume that the exogenous source of jumps μ is constant. Assuming $\tau_i \geq \tau_j$ and using the definition of the expectation of a Hawkes process, for no statistical arbitrage for calendar spread we obtain

$$\mathbb{E}[\mu^{(k, \tau_i)+} - \mu^{(k, \tau_i)-}] + \int_0^t \mathbb{E}[\lambda_s^{(k, \tau_i)+} - \lambda_s^{(k, \tau_i)-}] ds \geq \sqrt{\frac{\tau_j}{\tau_i}} \left(\mathbb{E}[\mu^{(k, \tau_j)+} - \mu^{(k, \tau_j)-}] + \int_0^t \mathbb{E}[\lambda_s^{(k, \tau_j)+} - \lambda_s^{(k, \tau_j)-}] ds \right). \quad (8.2)$$

Using Equation (8.1), we can deduce that to satisfy the above condition we can directly impose the following constraints on each coefficient of the intensities:

$$\begin{aligned} \phi^{(k, \tau_i)+, \cdot} &= \beta^{\phi, +} \sqrt{\frac{\tau_j}{\tau_i}} \phi^{(k, \tau_j)+, \cdot}, & \phi^{(k, \tau_i)-, \cdot} &= \beta^{\phi, +} \sqrt{\frac{\tau_j}{\tau_i}} \phi^{(k, \tau_j)-, \cdot}, \\ \mu^{(k, \tau_i)+} &= \beta^{\mu, +} \sqrt{\frac{\tau_j}{\tau_i}} \mu^{(k, \tau_j)+}, & \mu^{(k, \tau_i)-} &= \beta^{\mu, -} \sqrt{\frac{\tau_j}{\tau_i}} \mu^{(k, \tau_j)-}, \end{aligned} \quad (8.3)$$

where $(\beta^{\mu, +}, \beta^{\mu, -}, \beta^{\phi, +}, \beta^{\phi, -}) \in (1, +\infty)^4$. This leads to the following corollary.

Corollary 8.1. *The microscopic volatility surface $(\sigma_t)_{t \geq 0}$ is free of statistical calendar spread arbitrage if condition (8.3) is satisfied.*

We emphasize that Condition (8.3) is a sufficient but not necessary condition to ensure absence of statistical calendar spread arbitrage. Indeed, different kernel parametrizations are possible, notably by imposing the conditions on $\psi = \sum_{l \geq 1} \phi^{*l}$, where ϕ^{*l} is the l -th convolution product of ϕ with itself. Moreover, the coefficients $(\beta^{\mu, +}, \beta^{\mu, -}, \beta^{\phi, +}, \beta^{\phi, -})$ can be time- or strike-dependent for each coefficient of the Hawkes kernel.

Condition (8.3) leads to a maturity parametrization for the Hawkes kernel, thus reduces the number of parameters to calibrate. For example, let us assume $\beta^{\mu, +} = \beta^{\mu, -} = \beta^{\phi, +} = \beta^{\phi, -} = 1$, so that the

no calendar spread arbitrage constraint (8.2) becomes an equality. This way, we have only $2 \times \#\mathcal{K}$ coefficients of the Hawkes kernel to calibrate instead of $2 \times \#\mathcal{K} \times \#\mathcal{T}$.

Another condition for the volatility surface to be arbitrage-free is the absence of butterfly arbitrage. In financial markets, a butterfly is a strategy that consists in: buying one call option with strike k_1 and one call option with strike k_3 , and selling two call options with strike k_2 with the same maturity τ , and $k_1 < k_2 < k_3, k_2 - k_1 = k_3 - k_2$. If we denote by $C(k, \tau)$ the price of a call option with strike k and maturity τ , the payoff of the butterfly strategy is

$$BF_{(k_1, k_2, k_3), \tau} = C(k_1, \tau) - 2C(k_2, \tau) + C(k_3, \tau).$$

Using the results of Breeden and Litzenberger [55], we know that if the price of a vanilla option is available, one can retrieve the density of the underlying f_S via the formula

$$f_S(k, \tau) = \partial_{kk} C(k, \tau), \text{ for all } (k, \tau) \in \mathcal{K} \times \mathcal{T}.$$

Without conditions on call option prices (and therefore on the implied volatility), f_S in practice is not necessarily a density. The butterfly strategy is related to the function f_S in the following way: for a set of strikes (k_1, \dots, k_N) and a maturity τ , $f_S(k_i, \tau) \approx \frac{C(k_{i-1}, \tau) - 2C(k_i, \tau) + C(k_{i+1}, \tau)}{(k_i - k_{i-1})^2}$ for $i \in \{1, \dots, N - 1\}$. Therefore, the absence of butterfly arbitrage states that f_S must stay positive and integrate to one. The following proposition, whose proof can be found in Gatheral and Jacquier [125], provides necessary and sufficient conditions for the volatility surface to be free of butterfly arbitrage.

Proposition 8.2. *Let us define for a slice τ the following quantities at time $t \in [0, T]$*

$$d_t(k, \tau) = -\frac{k}{\sqrt{\varpi_t(k, \tau)}} + \frac{\sqrt{\varpi_t(k, \tau)}}{2},$$

$$g_t(k, \tau) = \left(1 - \frac{\partial_k \varpi_t(k, \tau)}{2\varpi_t(k, \tau)}\right)^2 - \frac{\partial_k \varpi_t(k, \tau)^2}{4} \left(\frac{1}{\varpi_t(k, \tau)} + \frac{1}{4}\right) + \frac{\partial_{kk} \varpi_t(k, \tau)}{2}.$$

A slice is free of butterfly arbitrage if and only if $g_t(k, \tau) \geq 0$ and $\lim_{k \rightarrow +\infty} d_t(k, \tau) = -\infty$.

This condition is more complicated to express in terms of the Hawkes kernel. The limit over $d(k, \tau)$ when $k \rightarrow +\infty$ implies that the density of the underlying integrates to one, and can be rewritten as

$$(\sigma^{(k, \tau)})^2 < \frac{2k}{\tau},$$

for large k , which is highly similar to the Roger-Lee moment formula for tail behavior of the volatility surface, see Lee [197]. For technical convenience, we can impose the stricter condition

$$\sigma^{(k_i, \tau)} = \beta \sqrt{\frac{k_i}{k_j}} \sigma^{(k_j, \tau)}, \quad \beta \in (0, 1),$$

and using the results of (8.1), we can deduce that this relationship is satisfied by imposing the following condition on the coefficients of the Hawkes kernel:

$$\phi^{(k_i, \tau)+, \cdot} = \beta \sqrt{\frac{k_i}{k_j}} \phi^{(k_j, \tau)+, \cdot}, \quad \phi^{(k_i, \tau)-, \cdot} = \beta \sqrt{\frac{k_i}{k_j}} \phi^{(k_j, \tau)-, \cdot},$$

$$\mu^{(k_i, \tau)+} = \beta \sqrt{\frac{k_i}{k_j}} \mu^{(k_j, \tau)+}, \quad \mu^{(k_i, \tau)-} = \beta \sqrt{\frac{k_i}{k_j}} \mu^{(k_j, \tau)-},$$
(8.4)

for large $(k_i, k_j) \in \mathcal{K}^2, \tau \in \mathcal{T}$. In practice, $\beta \in (0, 1)$ must be linked to the number of finite moments of the underlying asset, see Lee [197]. The condition $g_t(k, \tau) \geq 0$, which guarantees the positivity of the density of the underlying asset, is harder to include. We propose in the following section an alternative to this condition which ensures at-the-money convexity of the volatility slices. This alternative is also more suited in our case of discrete volatility surfaces.

8.2.2 Parametrization of the volatility surface

8.2.2.1 Three-points volatility surface

An important advantage of this kind of microscopic volatility modeling is that it is easy to control right and left skew and convexity of the slices of volatility. We choose one slice τ and we note that, by using no calendar spread arbitrage condition (8.3), it is sufficient to parametrize only one slice of volatility. We take three strikes of the options $-25\Delta P$, 50Δ and $25\Delta C$, that is

- the put option $-25\Delta P$ has a strike such that its delta is equal to -0.25 ,
- the call option 50Δ has a strike such that its delta is equal to 0.5 ,
- the call option $25\Delta C$ has a strike such that its delta is equal to 0.25 .

We denote by $\sigma^{-25\Delta P, \tau}$, $\sigma^{50\Delta, \tau}$, $\sigma^{25\Delta C, \tau}$ the volatilities associated to these options. The implied volatilities of the so-called risk-reversal and butterfly strategies, which are in fact measures of the skew and the convexity of the slice, with maturity τ are given by

$$BF_{25\Delta, \tau} = \frac{\sigma^{25\Delta C, \tau} + \sigma^{-25\Delta P, \tau}}{2} - \sigma^{50\Delta, \tau}, \quad RR_{25\Delta, \tau} = \sigma^{25\Delta C, \tau} - \sigma^{-25\Delta P, \tau}.$$

In particular, the convexity of the volatility slice is verified if and only if

$$\frac{\mathbb{E}[\sigma_t^{25\Delta C, \tau}] + \mathbb{E}[\sigma_t^{-25\Delta P, \tau}]}{2} \geq \mathbb{E}[\sigma_t^{50\Delta, \tau}].$$

If this condition is violated, one can find an arbitrage opportunity by creating a butterfly product whose price is negative. It can be replaced by a more restrictive condition:

$$\frac{\mathbb{E}[\sigma_t^{25\Delta C, \tau}] + \mathbb{E}[\sigma_t^{-25\Delta P, \tau}]}{2} = \beta^B \mathbb{E}[\sigma_t^{50\Delta, \tau}], \quad \beta^B \in [1, +\infty).$$

By imposing conditions directly on the coefficients of the Hawkes kernel, we obtain

$$\begin{aligned} & \frac{1}{2} \sum_{(\tilde{k}, \tilde{\tau})} \sum_{s \in \{+, -\}} (\phi^{(25\Delta C, \tau)+, (\tilde{k}, \tilde{\tau})s} + \phi^{(-25\Delta P, \tau)+, (\tilde{k}, \tilde{\tau})s} - \phi^{(25\Delta C, \tau)-, (\tilde{k}, \tilde{\tau})s} - \phi^{(-25\Delta P, \tau)-, (\tilde{k}, \tilde{\tau})s}) \\ & = \beta^B \sum_{(\tilde{k}, \tilde{\tau})} \sum_{s \in \{+, -\}} \phi^{(50\Delta, \tau)+, (\tilde{k}, \tilde{\tau})s}, \end{aligned}$$

which leads to the following natural condition to ensure convexity:

$$\begin{aligned} & \frac{\phi^{(25\Delta C, \tau)+, (\tilde{k}, \tilde{\tau})s} + \phi^{(-25\Delta P, \tau)+, (\tilde{k}, \tilde{\tau})s}}{2} = \beta^B \phi^{(50\Delta, \tau)+, (\tilde{k}, \tilde{\tau})s}, \\ & \frac{\phi^{(25\Delta C, \tau)-, (\tilde{k}, \tilde{\tau})s} + \phi^{(-25\Delta P, \tau)-, (\tilde{k}, \tilde{\tau})s}}{2} = \beta^B \phi^{(50\Delta, \tau)-, (\tilde{k}, \tilde{\tau})s}. \end{aligned} \tag{8.5}$$

If $\beta^B = 1$, the at-the-money volatility surface is flat. Increasing β^B leads to a higher at-the-money convexity. Combining this formula with the previous results leads to the following theorem.

Theorem 8.1 (No-arbitrage of the volatility surface). *We say that the three-points microscopic volatility surface is arbitrage free if*

$$\phi \in \mathcal{NA}_3 = \{ \phi \in \mathcal{M}_{2 \times \#\mathcal{K} \times \#\mathcal{T}}(\mathbb{R}), \text{ s.t. (8.3), (8.4) and (8.5) are satisfied} \}.$$

To summarize:

- Condition (8.3) guarantees the absence of calendar spread arbitrage, and reduces the dimension of the Hawkes kernel from $2 \times \#\mathcal{K} \times \#\mathcal{T}$ to $2 \times \#\mathcal{K}$. It states that, for a fixed strike, the intensity kernel for a different maturity is simply the initial intensity kernel scaled by the square root of the quotient of the two considered maturities.

- Condition (8.4) guarantees no-arbitrage for the wings of the volatility surface, ensuring that on average, the behavior of the volatility for a large strike is of \sqrt{k} order of magnitude, where strike k is expressed in log-moneyness.
- Condition (8.5) guarantees convexity of the at-the-money strike which, for slices composed of three strikes, is sufficient to ensure no butterfly arbitrage.

The simplicity of the formulae allows to give simple conditions on the kernel parameters so that the volatility slices are right or left skewed. Consider the following general power-law kernel

$$\phi^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}}(t) = \frac{\alpha^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}}}{(1+t)^{1+\gamma^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}}}}, \quad (8.6)$$

with $(k, \tilde{k}) \in \{-25\Delta P, 50\Delta, 25\Delta C\}^2$, $(\tau, \tilde{\tau}) \in \mathcal{T}^2$, $(s, \tilde{s}) \in \{+, -\}^2$ and assume stationary increments of the Hawkes processes. We have, see Bacry, Mastromatteo, and Muzy [33], for example,

$$\bar{\lambda} = \mathbb{E}[\lambda] = (I - \tilde{\phi})^{-1} \mu,$$

where $\tilde{\phi}^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}} = \frac{\alpha^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}}}{\gamma^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}}}$. Therefore, on average, $RR_{25\Delta} > 0$ if for every $\tau \in \mathcal{T}$, $\bar{\lambda}^{(25\Delta C, \tau)+} - \bar{\lambda}^{(25\Delta C, \tau)-} > \bar{\lambda}^{(-25\Delta P, \tau)+} - \bar{\lambda}^{(-25\Delta P, \tau)-}$ and conversely. This condition can be relaxed in several ways in the spirit of (8.4), thus we obtain the following corollary.

Corollary 8.2.

1. The three-points volatility surface is right (resp. left) skewed if for every $\tau \in \mathcal{T}$

$$\bar{\lambda}^{(25\Delta C, \tau)+} - \bar{\lambda}^{(25\Delta C, \tau)-} = \beta^{RR} (\bar{\lambda}^{(-25\Delta P, \tau)+} - \bar{\lambda}^{(-25\Delta P, \tau)-}), \quad \beta^{RR} > 1 \text{ (resp. } \beta^{RR} \in (0, 1)).$$

In particular, high β^{RR} increases the right skewness of the volatility surface and conversely for β^{RR} close to zero.

2. If β^B defined in Equation (8.5) is close to one, the at-the-money volatility surface is flat, while increasing β^B leads to a higher at-the-money convexity.

8.2.2.2 Five-points volatility surface

The volatility surface from the previous section takes into account, for each slice of maturity, the three most liquid traded points denoted by the $-25\Delta P, 50\Delta, 25\Delta C$ strikes. We can obviously extend this modeling by including for every slice of maturity τ the volatilities $\sigma^{10\Delta C, \tau}, \sigma^{-10\Delta P, \tau}$ corresponding to call and put with strikes such that their delta are equal to 0.10 and -0.10 respectively. We can now define the volatilities of the 10Δ risk reversal and butterfly strategies as

$$BF_{10\Delta, \tau} = \frac{\sigma^{10\Delta C, \tau} + \sigma^{-10\Delta P, \tau}}{2} - \sigma^{50\Delta, \tau}, \quad RR_{10\Delta, \tau} = \sigma^{10\Delta C, \tau} - \sigma^{-10\Delta P, \tau}.$$

To ensure convexity of the volatility slice, we must have $BF_{10\Delta, \tau} > 0$ which, by analogy with the previous section, leads to the following conditions

$$\begin{aligned} \frac{\phi^{(10\Delta C, \tau)+, (\tilde{k}, \tilde{\tau})s} + \phi^{(-10\Delta P, \tau)+, (\tilde{k}, \tilde{\tau})s}}{2} &= \beta^{\tilde{B}} \phi^{(50\Delta, \tau)+, (\tilde{k}, \tilde{\tau})s}, \\ \frac{\phi^{(10\Delta C, \tau)-, (\tilde{k}, \tilde{\tau})s} + \phi^{(-10\Delta P, \tau)-, (\tilde{k}, \tilde{\tau})s}}{2} &= \beta^{\tilde{B}} \phi^{(50\Delta, \tau)-, (\tilde{k}, \tilde{\tau})s}, \end{aligned} \quad (8.7)$$

for $\beta^{\tilde{B}} > 1$, which controls the convexity at extreme strikes of the volatility surface. In particular, if $\beta^B < \beta^{\tilde{B}}$, the convexity of the volatility slices increases at extreme strikes. Conversely, if $\beta^{\tilde{B}}$ is close to β^B , the volatility slice flattens at extreme strikes. This leads to the following theorem, which is the counterpart of Theorem 8.1 for the five-points volatility surface.

Theorem 8.2 (No-arbitrage of the volatility surface). *We say that the five-points microscopic volatility surface is arbitrage free if*

$$\phi \in \mathcal{NA}_5 = \{\phi \in \mathcal{M}_{2 \times \#\mathcal{K} \times \#\mathcal{T}}(\mathbb{R}) : (8.3), (8.4), (8.5) \text{ and } (8.7) \text{ are satisfied}\}.$$

It is straightforward to find conditions on the coefficients of the kernels $\phi \in \mathcal{NA}$ to control the skewness and convexity of the volatility surface. Let us consider the power-law kernel (8.6), for $(\tau, \tilde{\tau}) \in \mathcal{T}^2$, $(s, \tilde{s}) \in \{+, -\}^2$, $(k, \tilde{k}) \in \{-10\Delta P, -25\Delta P, 50\Delta, 25\Delta C, 10\Delta C\}^2$, and assume stationary increments of the Hawkes processes, so that we have

$$\bar{\lambda} = \mathbb{E}[\lambda] = (I - \tilde{\phi})^{-1} \mu, \quad (8.8)$$

where $\tilde{\phi}^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}} = \frac{\alpha^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}}}{\gamma^{(k,\tau)s,(\tilde{k},\tilde{\tau})\tilde{s}}}$, and (8.8) is in closed-form. By analogy with Corollary 8.2, we derive the following results.

Corollary 8.3.

1. *The five-points volatility surface is right (resp. left) skewed if for every $\tau \in \mathcal{T}$*

$$\begin{aligned} \bar{\lambda}^{(25\Delta C, \tau)+} - \bar{\lambda}^{(25\Delta C, \tau)-} &= \beta_{25\Delta}^{RR} (\bar{\lambda}^{(-25\Delta P, \tau)+} - \bar{\lambda}^{(-25\Delta P, \tau)-}), \quad \beta_{25\Delta}^{RR} > 1 \text{ (resp. } \beta_{25\Delta}^{RR} \in (0, 1)), \\ \bar{\lambda}^{(10\Delta C, \tau)+} - \bar{\lambda}^{(10\Delta C, \tau)-} &= \beta_{10\Delta}^{RR} (\bar{\lambda}^{(-10\Delta P, \tau)+} - \bar{\lambda}^{(-10\Delta P, \tau)-}), \quad \beta_{10\Delta}^{RR} > 1 \text{ (resp. } \beta_{10\Delta}^{RR} \in (0, 1)). \end{aligned}$$

2. *A slice $\tau \in \mathcal{T}$ exhibits a smile (resp. skew) if $RR_{10\Delta, \tau} > RR_{25\Delta, \tau}$ (resp. $<$), which is satisfied if*

$$(\beta_{10\Delta}^{RR} - 1)(\bar{\lambda}^{(-10\Delta P, \tau)+} - \bar{\lambda}^{(-10\Delta P, \tau)-}) > (\beta_{25\Delta}^{RR} - 1)(\bar{\lambda}^{(-25\Delta P, \tau)+} - \bar{\lambda}^{(-25\Delta P, \tau)-}),$$

and a sufficient condition is

$$(\beta_{10\Delta}^{RR} - 1)(\bar{\lambda}^{(-10\Delta P, \tau)+} - \bar{\lambda}^{(-10\Delta P, \tau)-}) = \beta^{SS} (\beta_{25\Delta}^{RR} - 1)(\bar{\lambda}^{(-25\Delta P, \tau)+} - \bar{\lambda}^{(-25\Delta P, \tau)-})$$

with $\beta^{SS} \in (1, +\infty)$ (resp. $\beta^{SS} \in (0, 1)$). In particular, a high β^{SS} leads to a steeper smile and conversely.

3. *The convexity of the five-points volatility surface increases (resp. decreases) for extreme strikes if, for every $\tau \in \mathcal{T}$, $BF_{10\Delta, \tau} > BF_{25\Delta, \tau}$ (resp. $BF_{10\Delta, \tau} < BF_{25\Delta, \tau}$). A sufficient condition is*

$$\beta^{\tilde{B}} > \beta^B \text{ (resp. } \beta^{\tilde{B}} < \beta^B).$$

8.3 Macroscopic limit of the volatility surface

We now show the macroscopic behavior of the volatility dynamics. As the absence of calendar spread arbitrage offers a simple parametrization compared to other conditions, in this section, we only assume the absence of calendar spread arbitrage, that is, Condition (8.3) is verified and becomes an equality. Using no butterfly arbitrage would change the Hawkes kernel's coefficients with respect to the strikes, but not the qualitative nature of our results. We adapt the framework used in Tomas and Rosenbaum [263]. For every result stated in this section, we refer the reader to Tomas and Rosenbaum [263] for corresponding proofs.

We set $T > 0$ and take, in a basis O independent of both T and $t \in [0, T]$, a sequence of triangularizable intensity kernels $\phi^{\mathbf{T}}(t) \in \mathcal{M}_{2 \times \#\mathcal{K} \times \#\mathcal{T}}(\mathbb{R})$ with $n_c > 0$ non-zero eigenvalues. Using the block matrix notation, the Hawkes kernel can be written as

$$\phi^{\mathbf{T}}(t) = \mathbf{O} \begin{pmatrix} \mathbf{A}^{\mathbf{T}}(t) & 0 \\ \mathbf{B}^{\mathbf{T}}(t) & \mathbf{C}^{\mathbf{T}}(t) \end{pmatrix} \mathbf{O}^{-1},$$

with $\mathbf{A}^{\mathbf{T}} : \mathbb{R}_+ \rightarrow \mathcal{M}_{n_c}(\mathbb{R})$, $\mathbf{B}^{\mathbf{T}} : \mathbb{R}_+ \rightarrow \mathcal{M}_{2 \times \#\mathcal{K} \times \mathcal{T} - n_c, n_c}(\mathbb{R})$, $\mathbf{C}^{\mathbf{T}} : \mathbb{R}_+ \rightarrow \mathcal{M}_{2 \times \#\mathcal{K} \times \mathcal{T} - n_c}(\mathbb{R})$. We impose the following limit condition on the Hawkes kernel:

$$\mathbf{T}^\alpha \left(\mathbf{I} - \int_0^\infty \mathbf{A}^{\mathbf{T}} \right) \xrightarrow{\mathbf{T} \rightarrow +\infty} \mathbf{K},$$

where \mathbf{K} is an invertible matrix and $\alpha > 0$ is the Hurst exponent of the volatility surface. This condition is in fact the saturation of the stability condition of the Hawkes kernel. We finally add the heavy-tail condition on $\mathbf{A} = \lim_{\mathbf{T} \rightarrow +\infty} \mathbf{A}^{\mathbf{T}}$, that is

$$\alpha x^\alpha \int_x^\infty \mathbf{A}(s) ds \xrightarrow{x \rightarrow +\infty} \mathbf{M},$$

with \mathbf{M} being an invertible matrix.

For $t \in [0, T]$, we define the following rescaled processes:

$$\mathbf{X}_t^{\mathbf{T}} = \frac{1}{T^{2\alpha}} \mathbf{N}_{t\mathbf{T}}^{\mathbf{T}}, \quad \mathbf{Y}_t^{\mathbf{T}} = \frac{1}{T^{2\alpha}} \int_0^{t\mathbf{T}} \boldsymbol{\lambda}_s ds, \quad \mathbf{Z}_t^{\mathbf{T}} = T^\alpha (\mathbf{X}_t^{\mathbf{T}} - \mathbf{Y}_t^{\mathbf{T}}) = \mathbf{M}_{t\mathbf{T}}^{\mathbf{T}}, \quad \boldsymbol{\sigma}_t^{\mathbf{T}} = \frac{1}{T^{2\alpha}} (\mathbf{N}_{t\mathbf{T}}^{\mathbf{T}+} - \mathbf{N}_{t\mathbf{T}}^{\mathbf{T}-}).$$

Particularly, $\boldsymbol{\sigma}^{\mathbf{T}}$ is the rescaled microscopic volatility surface. Throughout the section, we assume the following:

Assumption 8.1. For all $(\tilde{k}, \tilde{\tau}) \in \mathcal{K} \times \mathcal{T}$, and $(s, \tilde{s}) \in \{+, -\}^2$,

$$\phi^{(k, \tau)s, (\tilde{k}, \tilde{\tau})\tilde{s}}(t) = \sqrt{\frac{\tau}{\tilde{\tau}}} \phi^{k, \tilde{k}}(t),$$

meaning that the self $(++, --)$ and cross $(+-, -+)$ exciting terms are the same.

This assumption can be relaxed and is made only to provide meaningful results for some specific forms of ϕ .

8.3.1 Limiting processes in the separable, one-dimensional case

In this section, we consider the case of the separable Hawkes kernel. We assume

$$\phi^{k, \tilde{k}}(t) = \tilde{z}(k) \tilde{z}(\tilde{k}) \varphi(t),$$

where $\varphi : \mathbb{R}_+ \rightarrow \mathbb{R}_+$, $\tilde{z} : \mathcal{K} \rightarrow \mathbb{R}_+$ represent the overall intensity for changes of the implied volatility surface depending on time and strike respectively.

Remark 40. In the context of a power-law kernel (8.6), this form of kernel coefficients means that there is a single power coefficient $\gamma^{(k, \tau)s, (\tilde{k}, \tilde{\tau})\tilde{s}} = \gamma > 0$, and only the scaling coefficients $\alpha^{(k, \tau)s, (\tilde{k}, \tilde{\tau})\tilde{s}}$ can differ from one strike to another.

Therefore, ϕ is a rank one matrix defined by

$$\boldsymbol{\phi} = \sqrt{\tau} \varphi \begin{pmatrix} \frac{1}{\sqrt{\tau}} \tilde{z} \tilde{z}^\top & \frac{1}{\sqrt{\tau}} \tilde{z} \tilde{z}^\top & \dots & \frac{1}{\sqrt{\tau}} \tilde{z} \tilde{z}^\top \\ \frac{1}{\sqrt{\tau_2}} \tilde{z} \tilde{z}^\top & \frac{1}{\sqrt{\tau_2}} \tilde{z} \tilde{z}^\top & \dots & \frac{1}{\sqrt{\tau_2}} \tilde{z} \tilde{z}^\top \\ \vdots & \vdots & \ddots & \vdots \\ \frac{1}{\sqrt{\tau_M}} \tilde{z} \tilde{z}^\top & \frac{1}{\sqrt{\tau_M}} \tilde{z} \tilde{z}^\top & \dots & \frac{1}{\sqrt{\tau_M}} \tilde{z} \tilde{z}^\top \end{pmatrix}.$$

This matrix has a single non-zero eigenvalue $\|\tilde{z}\|_2^2 \left(\frac{1}{\sqrt{\tau}} + \frac{1}{\sqrt{\tau_1}} + \dots + \frac{1}{\sqrt{\tau_M}} \right)$ associated to the eigenvector $v = \frac{1}{\|\tilde{z}\|_2 \sqrt{\frac{1}{\sqrt{\tau}} + \frac{1}{\sqrt{\tau_2}} + \dots + \frac{1}{\sqrt{\tau_M}}}} \left(\frac{1}{\sqrt{\tau}} \tilde{z}, \frac{1}{\sqrt{\tau_2}} \tilde{z}, \dots, \frac{1}{\sqrt{\tau_M}} \tilde{z} \right)$. In order for the Hawkes process to be nearly-unstable we impose

$$\|\varphi\|_1 = \frac{1}{\|\tilde{z}\|_2^2 \left(1 + \frac{\sqrt{\tau}}{\sqrt{\tau_2}} + \dots + \frac{\sqrt{\tau}}{\tau_M} \right)}.$$

Then, by using the results of El Euch, Fukasawa, and Rosenbaum [107], we have that at the limit, the dynamics of the volatility is given by $\sigma_t = F_t v$, where

$$F_t = C \int_0^t \sqrt{V_s} dW_s,$$

$$V_t = \frac{1}{\Gamma(\alpha)} \int_0^t (t-s)^{\alpha-1} (\theta - V_s) ds + \frac{1}{\Gamma(\alpha)} \int_0^t (t-s)^{\alpha-1} \lambda \sqrt{V_s} dZ_s,$$

$(C, \theta, \lambda) \in \mathbb{R}_+^3$ and (W, Z) is a 2-dimensional Brownian motion. Note that the value of the coefficients in the macroscopic volatility process depends crucially on the rescaling. In Section 8.3.3, we provide a characterization of these coefficients in the general case of a triangularizable intensity kernel with the specific rescaling described at the beginning of the section. In the present case, the macroscopic volatility surface is driven by a single factor of risk F_t whose dynamics is that of a rough Heston model. This corresponds to the so-called ‘‘market’’ or ‘‘level’’ factor described in Cont and Da Fonseca [80]. All the components of the eigenvector v are positive. Thus, when factor F_t increases, all the implied volatilities increase and conversely.

Finally, note that the volatility of the macroscopic implied volatility is a rough process. This property does not contradict the studies on the volatility’s roughness always treating the realized volatility because, in this chapter, we model the implied volatility, which has no reason to be rough.

8.3.2 Limiting process in the semi-separable, factor case

In this section, we consider the case of the semi-separable Hawkes kernel. We assume that

$$\phi^{k, \bar{k}}(t) = \sum_{i=1}^r \tilde{z}_i(k) \tilde{z}_i(\bar{k})^\top \varphi_i(t),$$

where each factor has an associated kernel in time to reflect lead-lag effects in the overall surface.

Remark 41. *In the context of a power-law kernel (8.6), this form of kernel coefficients means that there is a single power coefficient $\gamma_i^{(k, \tau)s, (\bar{k}, \bar{\tau})\bar{s}} = \gamma_i > 0$ for every of the $F_t^i, i \in \{1, \dots, r\}$ factors and the scaling coefficients $\alpha^{(k, \tau)s, (\bar{k}, \bar{\tau})\bar{s}}$ can differ depending on the strike and the considered factor $F_t^i, i \in \{1, \dots, r\}$.*

We assume that $\tilde{z}_i, i \in \{1, \dots, r\}$ are orthogonal, i.e. $\sum_k \tilde{z}_i(k) \tilde{z}_j(\bar{k}) = 0$ if $i \neq j$. Therefore, ϕ is a rank r matrix

$$\phi = \sqrt{\tau} \sum_{i=1}^r \varphi_i \begin{pmatrix} \frac{1}{\sqrt{\tau}} \tilde{z}_i \tilde{z}_i^\top & \frac{1}{\sqrt{\tau}} \tilde{z}_i \tilde{z}_i^\top & \dots & \frac{1}{\sqrt{\tau}} \tilde{z}_i \tilde{z}_i^\top \\ \frac{1}{\sqrt{\tau_2}} \tilde{z}_i \tilde{z}_i^\top & \frac{1}{\sqrt{\tau_2}} \tilde{z}_i \tilde{z}_i^\top & \dots & \frac{1}{\sqrt{\tau_2}} \tilde{z}_i \tilde{z}_i^\top \\ \vdots & \vdots & \ddots & \vdots \\ \frac{1}{\sqrt{\tau_M}} \tilde{z}_i \tilde{z}_i^\top & \frac{1}{\sqrt{\tau_M}} \tilde{z}_i \tilde{z}_i^\top & \dots & \frac{1}{\sqrt{\tau_M}} \tilde{z}_i \tilde{z}_i^\top \end{pmatrix}.$$

Thus, the eigenvectors of the Hawkes kernel are $v_i = \frac{1}{\|\tilde{z}_i\|_2 \sqrt{\frac{1}{\sqrt{\tau}} + \frac{1}{\sqrt{\tau_2}} + \dots + \frac{1}{\sqrt{\tau_M}}}} (\frac{1}{\tau} \tilde{z}_i, \frac{1}{\sqrt{\tau_2}} \tilde{z}_i, \dots, \frac{1}{\sqrt{\tau_M}} \tilde{z}_i)$ with associated eigenvalues $\|z_i\|_2^2 (\frac{1}{\sqrt{\tau}} + \frac{1}{\sqrt{\tau_2}} + \dots + \frac{1}{\sqrt{\tau_M}})$. They form an orthonormal basis as the \tilde{z}_i are orthogonal. Therefore, the criticality condition for the eigenvector v_i is:

$$\|\varphi_i\|_1 = \frac{1}{\|\tilde{z}_i\|_2^2 (1 + \frac{\sqrt{\tau}}{\sqrt{\tau_2}} + \dots + \frac{\sqrt{\tau}}{\sqrt{\tau_M}})}.$$

Then, at the limit, we obtain $\sigma_t = \sum_{i=1}^r v_i F_t^i$ where

$$F_t^i = C^i \int_0^t \sqrt{V_s^i} dW_s^i,$$

$$V_t^i = \frac{1}{\Gamma(\alpha)} \int_0^t (t-s)^{\alpha-1} (\theta^i - V_s^i) ds + \frac{1}{\Gamma(\alpha)} \int_0^t (t-s)^{\alpha-1} \lambda^i \sqrt{V_s^i} dZ_s^i,$$

with $(C^i, \theta^i, \lambda^i) \in \mathbb{R}_+^3$ for all $i \in \{1, \dots, r\}$ and $(W^i, Z^i)_{i \in \{1, \dots, r\}}$ being a $2r$ -dimensional Brownian motion. The macroscopic limit of the volatility surface dynamics is given by a sum of risk factors having rough volatility.

In this setting, it is easy to parametrize the microscopic kernel to recover the classic “level-skew-convexity” behavior of the volatility surface. Assume that $r = 3$, that is, there are three factors of risk driving the macroscopic volatility surface.

- By setting $\tilde{z}_1(k) > 0$ for all $k \in \mathcal{K}$, the first factor corresponds to the “level” factor of the implied volatility surface: when this factor increases, all the implied volatilities increase and conversely. A possible parametrization is then

$$\tilde{z}_1(k) = \nu k,$$

with $\nu > 0$, close to zero: the eigenvector components increase linearly with respect to the moneyness, meaning that the “level” factor’s impact is slightly higher for out-of-the-money calls.

- By setting $\tilde{z}_2(k) > 0$ (resp. $\tilde{z}_2(k) < 0$) for $k > 1$ (resp. $k < 1$), the second factor corresponds to the “calendar” factor of the implied volatility surface: when this factor increases, the implied volatility of out-of-the-money calls increase, while those of out-of-the-money puts decrease. A possible parametrization is then

$$\tilde{z}_2(k) = c_1 \left(\frac{1}{1 + e^{-c_2(k-1)}} - 0.5 \right),$$

with $c_1, c_2 > 0$. The parameter c_1 controls the scale of the factor. The parameter c_2 controls the steepness of the factor’s change around the money: when c_2 increases, $\tilde{z}_2(k)$ changes its sign quicker around $k = 1$ and conversely.

- By setting $\tilde{z}_3(k)$ parabolic with respect to k , minimized at $k = 1$, the third factor F_t^3 corresponds to the “butterfly” factor of the implied volatility surface: a variation of this factor leads to a change of convexity of the volatility surface and a downward sloping term structure (which is already incorporated into the intensity kernel through the no calendar spread arbitrage condition (8.3)). A possible parametrization is then

$$\tilde{z}_3(k) = c_3(k-1)^2 - c_4,$$

with $c_3, c_4 > 0$. The parameter c_3 scales the factor and the parameter c_4 controls the factor’s level.

8.3.3 General case

In this section, we assume no specific form on $\phi^{k, \tilde{k}}(t) = \tilde{Z}(k, \tilde{k}, t)$. This covers the case of power-law kernels where

$$\phi^{k, \tilde{k}}(t) = \frac{\alpha^{k, \tilde{k}}}{(1+t)^{1+\gamma^{k, \tilde{k}}}},$$

with $(\alpha^{k, \tilde{k}}, \gamma^{k, \tilde{k}}) \in \mathbb{R}_+^{*2}$. We can rewrite the kernel in the following basis:

$$\phi(t) = \mathbf{O} \begin{pmatrix} \mathbf{A}(t) & 0 \\ \mathbf{B}(t) & \mathbf{C}(t) \end{pmatrix} \mathbf{O}^{-1},$$

where

$$\mathbf{O} = \begin{pmatrix} \frac{1}{\sqrt{\tau}} I & \frac{1}{\sqrt{\tau}} I & \dots & 0 \\ \frac{1}{\sqrt{\tau_2}} I & 0 & \dots & 0 \\ \vdots & & & \\ \frac{1}{\sqrt{\tau_M}} I & -\frac{1}{\sqrt{\tau_M}} I & \dots & -\frac{1}{\sqrt{\tau_M}} I \end{pmatrix} \mathbf{G} = \begin{pmatrix} \mathbf{O}_{11} & \mathbf{O}_{12} \\ \mathbf{O}_{21} & \mathbf{O}_{22} \end{pmatrix},$$

and $\mathbf{G} := \text{diag}(\tilde{G}, \dots, \tilde{G})$, where \tilde{G} is the basis change which makes \tilde{Z} tridiagonal for all t . Let

$$\begin{aligned} \mathbf{O}^{-1} &= \begin{pmatrix} \mathbf{O}_{11}^{-1} & \mathbf{O}_{12}^{-1} \\ \mathbf{O}_{21}^{-1} & \mathbf{O}_{22}^{-1} \end{pmatrix}, \\ \Theta^1 &= \left(\mathbf{O}_{11} + \mathbf{O}_{12} \left(\mathbf{I} - \int_0^\infty \mathbf{C} \right)^{-1} \int_0^\infty \mathbf{B} \right) \mathbf{K}^{-1}, \\ \Theta^2 &= \left(\mathbf{O}_{21} + \mathbf{O}_{22} \left(\mathbf{I} - \int_0^\infty \mathbf{C} \right)^{-1} \int_0^\infty \mathbf{B} \right) \mathbf{K}^{-1}, \\ \theta_0 &= \begin{pmatrix} \mathbf{O}_{11}^{-1} & \mathbf{0} \\ \mathbf{0} & \mathbf{O}_{22}^{-1} \end{pmatrix} \boldsymbol{\mu}, \\ \Lambda &= \frac{\alpha}{\Gamma(1-\alpha)} \mathbf{K} \mathbf{M}^{-1}, \end{aligned}$$

and assume that ϕ has $n_c > 0$ non-zero eigenvalues. Then, for any limit point $(\mathbf{X}, \mathbf{Y}, \mathbf{Z})$ of the sequence $(\mathbf{X}^T, \mathbf{Y}^T, \mathbf{Z}^T)$, which is C -tight for the Skorokhod topology, there exists a positive process \mathbf{V} and a $2 \times \#\mathcal{K} \times \#\mathcal{T}$ dimensional Brownian motion \mathbf{B} such that

$$\mathbf{X}_t = \int_0^t \mathbf{V}_s ds, \quad \mathbf{Z}_t = \int_0^t \text{diag}(\sqrt{\mathbf{V}_s}) d\mathbf{B}_s.$$

Moreover there exists a process $\tilde{\mathbf{V}}$ of Holder regularity $\alpha - \frac{1}{2} - \epsilon$ for any $\epsilon > 0$ such that $\Theta^1 \tilde{\mathbf{V}} = (\mathbf{V}^1, \dots, \mathbf{V}^{n_c})$, $\Theta^2 \tilde{\mathbf{V}} = (\mathbf{V}^{n_c+1}, \dots, \mathbf{V}^{2 \times \#\mathcal{K} \times \#\mathcal{T}})$ and for all $t \in [0, 1]$,

$$\begin{aligned} \tilde{\mathbf{V}}_t &= \frac{1}{\Gamma(\alpha)} \Lambda \int_0^t (t-s)^{\alpha-1} (\theta_0 - \tilde{\mathbf{V}}_s) ds \\ &\quad + \frac{1}{\Gamma(\alpha)} \Lambda \int_0^t (t-s)^{\alpha-1} \mathbf{O}_{11}^{(-1)} \text{diag}(\sqrt{\Theta^1 \tilde{\mathbf{V}}_s}) d\mathbf{W}_s^1 \\ &\quad + \frac{1}{\Gamma(\alpha)} \Lambda \int_0^t (t-s)^{\alpha-1} \mathbf{O}_{12}^{(-1)} \text{diag}(\sqrt{\Theta^2 \tilde{\mathbf{V}}_s}) d\mathbf{W}_s^2, \end{aligned}$$

where $\mathbf{W}^1 := (B^1, \dots, B^{n_c})$, $\mathbf{W}^2 := (B^{n_c+1}, \dots, B^{2 \times \#\mathcal{K} \times \#\mathcal{T}})$, Θ^1 , Θ^2 , $\mathbf{O}_{11}^{(-1)}$, $\mathbf{O}_{12}^{(-1)}$, θ_0 do not depend on the chosen basis. Finally, any limit point $\boldsymbol{\sigma}$ of the rescaled volatility surface $\boldsymbol{\sigma}^T$ satisfies

$$\boldsymbol{\sigma}_t = (\mathbf{I} + \Delta) \mathbf{Q}^T \left(\int_0^t \text{diag}(\sqrt{\mathbf{V}_s}) d\mathbf{B}_s + \int_0^t \boldsymbol{\mu}_s ds \right), \quad (8.9)$$

where

$$\begin{aligned} \Delta_{(k,\tau),(\bar{k},\bar{\tau})} &= \lim_{T \rightarrow \infty} \left(\left\| \psi_{(\bar{k},\bar{\tau})+, (k,\tau)+}^T \right\|_1 - \left\| \psi_{(\bar{k},\bar{\tau})-, (k,\tau)+}^T \right\|_1 \right), \\ \mathbf{Q} &= (e_1 - e_2 | \dots | e_{2 \times \#\mathcal{K} \times \#\mathcal{T}-1} - e_{2 \times \#\mathcal{K} \times \#\mathcal{T}}), \end{aligned}$$

and $(e_i)_{1 \leq i \leq 2 \times \#\mathcal{K} \times \#\mathcal{T}}$ is the canonical basis of $\mathbb{R}^{2 \times \#\mathcal{K} \times \#\mathcal{T}}$.

We refer to Tomas and Rosenbaum [263] for the proof of this theorem. In this setting, the macroscopic volatility surface is driven by $2 \times \mathcal{K} \times \mathcal{T}$ sources of risk having rough volatilities. The first n_c sources of risk \mathbf{W}^1 come from the non-zero eigenvalues of the Hawkes kernel ϕ whereas the $2 \times \mathcal{K} \times \mathcal{T} - n_c$ others come from the zero-valued eigenvalues. This result contrasts with previous sections where the intensity kernel was diagonalizable, and the only sources of risk corresponded to the non-zero eigenvalues.

8.4 Applications

8.4.1 Backtesting of option market-making strategies

8.4.1.1 Description

Option market-making is known to be the new “hot” topic in the industry. Indeed, it accounts for a considerable part of the trading revenue for proprietary shops and hedge funds. In the academic literature, option market-making models have been studied in El Aoud and Abergel [106], Stoikov and Sağlam [258] and in Chapter 6 and Chapter 7 of this thesis. In El Aoud and Abergel [106], Stoikov and Sağlam [258], the authors consider a single-option market driven by a stochastic volatility model and assume that the position is always Δ -hedged. They provide optimal bid and ask quotes for the options and focus either on the risk of model misspecification or the residual risk due to the presence of stochastic volatility. In Chapter 6 and Chapter 7, the authors propose several models to deal with the market-making of a large portfolio of options. The authors make a constant Greek assumption in the former, which is well suited for market-making on long-dated options: the only variable of interest is the aggregated portfolio, where each position is weighted by its Vega. In the latter, the authors make an approximation directly on the form of the value function of the market-maker’s optimization problem. Thus, the Greeks are allowed to vary with time, spot, and stochastic volatility, and one can manage each position individually, even for a very large portfolio.

These models enable to design market-making strategies on options. However, they all make the assumption that the number of buy (resp. sell) filled limit orders on option (k, τ) on $[0, t]$ is a point process $N^{(k, \tau)+}$ (resp. $N^{(k, \tau)-}$) without self- or cross-exciting properties. In other words, the intensity of $N^{(k, \tau)+}$ is only a function of the time t , the spot S , the stochastic volatility ν and the spread $\delta^{(k, \tau)+}$ quoted by the market-maker on the ask side of option (k, τ) . If we assume $\mathbb{P} = \mathbb{Q}$, that is the risk-neutral measure is calibrated only using historical data, the intensities are functions only of time and spread, and are denoted by

$$\Lambda_t^{(k, \tau)s}(\delta_t^{(k, \tau)s}), \quad (k, \tau) \in \mathcal{K} \times \mathcal{T}, s \in \{+, -\}.$$

In this case, it is particularly easy to backtest the market-making strategy obtained in Chapter 6 and Chapter 7. We denote the inventory on option (k, τ) by $q^{(k, \tau)}$ and the optimal strategy of the market-maker on option (k, τ) by

$$\delta^{*(k, \tau)+}(t, q_t), \delta^{*(k, \tau)-}(t, q_t), \quad (8.10)$$

where $q_t = (q_t^{(k, \tau)})_{(k, \tau) \in \mathcal{K} \times \mathcal{T}}$.

In order to backtest this optimal strategy, we simulate the number of filled buy (resp. sell) limit orders on option (k, τ) using a point process $\bar{N}^{(k, \tau)+}$ (resp. $\bar{N}^{(k, \tau)-}$) with intensity

$$\bar{\lambda}^{(k, \tau)s}(t, \delta^{*(k, \tau)s}(t, q_t)) = \lambda_t^{(k, \tau)s} \Lambda_t^{(k, \tau)s}(\delta^{*(k, \tau)s}(t, q_t)), \quad (8.11)$$

where

$$\lambda_t^{(k, \tau)s} = \mu^{(k, \tau)s} + \sum_{(\tilde{k}, \tilde{\tau})} \sum_{\tilde{s} \in \{+, -\}} \int_0^t \phi^{(k, \tau)s, (\tilde{k}, \tilde{\tau})\tilde{s}}(t - u) dN_u^{(\tilde{k}, \tilde{\tau})\tilde{s}},$$

and $q_t^{(k, \tau)} = \bar{N}_t^{(k, \tau)+} - \bar{N}_t^{(k, \tau)-}$. Thanks to Equation (8.11), we incorporate the strategy of the market-maker in the intensities of the Hawkes processes which represent the market flow on each option. Thus, we can backtest the trading strategy in a more realistic framework.

8.4.1.2 Numerical results

We conduct a backtest of the optimal market-making strategy in the context of the three point volatility surface described previously. We choose a single slice of maturity $\mathcal{T} = \{\tau\} = 1$ week and three options of strike (expressed in log-moneyness)

$$k_1 = 0.98, \quad k_2 = 1, \quad k_3 = 1.02.$$

The influence of the market-maker's quotes on the intensity of orders' arrival are given by

$$\Lambda_t^{(k,\tau)+}(\delta) = \Lambda_t^{(k,\tau)-}(\delta) = \frac{\lambda^{k,\tau}}{1 + \exp\left(\alpha + \frac{\beta}{\nu^{(k,\tau)}}\delta\right)},$$

where $\lambda^{k,\tau} = \frac{252 \times 70}{1 + 8 \times |k-1|}$, $\alpha = -0.7$, $\beta = 10$. The choice of $\lambda^{k,\tau}$ corresponds to 70 requests per day for the at-the-money option, and decreases to 27 for far-from-the-money options. The choice of α corresponds to a probability of $\frac{1}{1+e^{-0.8}} \approx 66\%$ to trade when the answered quote is the mid-price ($\delta = 0$). The choice of β corresponds to a probability of $\frac{1}{1+e^{-0.8}} \approx 68\%$ to trade when the answered quote corresponds to an implied volatility 1% better for the client and a probability of $\frac{1}{1+e^{-0.8}} \approx 64\%$ to trade when the answered quote corresponds to an implied volatility 1% worse for the client. We assume transactions of unitary volume, a trading horizon of $T = 1$ day and a risk-aversion parameter of 2×10^{-5} for the market-maker.⁵ Finally, we set $\mu^{(k,\tau)s} = 1$ for all $(k, \tau) \in \mathcal{K} \times \mathcal{T}$, $s \in \{+, -\}$ and the Hawkes kernel is

$$\begin{aligned} \phi^{(k_1,\tau)s,(k_2,\tau)\tilde{s}} &= \phi^{(k_2,\tau)s,(k_1,\tau)\tilde{s}} = \phi^{\text{ITM-ATM}}, & \phi^{(k_1,\tau)s,(k_3,\tau)\tilde{s}} &= \phi^{(k_3,\tau)s,(k_1,\tau)\tilde{s}} = \phi^{\text{OTM-ATM}}, \\ \phi^{(k_2,\tau)s,(k_3,\tau)\tilde{s}} &= \phi^{(k_3,\tau)s,(k_2,\tau)\tilde{s}} = 0, & \phi^{(k_1,\tau)s,(k_1,\tau)\tilde{s}} &= \phi^{\text{ITM}}, \\ \phi^{(k_2,\tau)s,(k_2,\tau)\tilde{s}} &= \phi^{\text{ATM}}, & \phi^{(k_3,\tau)s,(k_3,\tau)\tilde{s}} &= \phi^{\text{OTM}}, \end{aligned}$$

for all $(s, \tilde{s}) \in \{+, -\}$, and for all $t \in [0, T]$,

$$\begin{aligned} \phi^{\text{ITM-ATM}}(t) &= \frac{\alpha^{\text{ITM-ATM}}}{(1+t)^{1+\gamma^{\text{ITM-ATM}}}}, & \alpha^{\text{ITM-ATM}} &= 0.18, \gamma^{\text{ITM-ATM}} = 0.15, \\ \phi^{\text{OTM-ATM}}(t) &= \frac{\alpha^{\text{OTM-ATM}}}{(1+t)^{1+\gamma^{\text{OTM-ATM}}}}, & \alpha^{\text{OTM-ATM}} &= 0.13, \gamma^{\text{OTM-ATM}} = 0.15, \\ \phi^{\text{ITM}}(t) &= \frac{\alpha^{\text{ITM}}}{(1+t)^{1+\gamma^{\text{ITM}}}}, & \alpha^{\text{ITM}} &= 0.48, \gamma^{\text{ITM}} = 0.08, \\ \phi^{\text{ATM}}(t) &= \frac{\alpha^{\text{ATM}}}{(1+t)^{1+\gamma^{\text{ATM}}}}, & \alpha^{\text{ATM}} &= 0.52, \gamma^{\text{ATM}} = 0.08, \\ \phi^{\text{OTM}}(t) &= \frac{\alpha^{\text{OTM}}}{(1+t)^{1+\gamma^{\text{OTM}}}}, & \alpha^{\text{OTM}} &= 0.14, \gamma^{\text{OTM}} = 0.08. \end{aligned}$$

We assume power law intensity kernels and no cross-excitation between in-the-money and out-of-the-money options. For sake of simplicity, we assume symmetry between buy and sell self and cross intensities. Moreover

- $\phi^{\text{ITM-ATM}} > \phi^{\text{OTM-ATM}}$: there is more cross excitation between the in-the-money and at-the-money options than between the out-of-the-money and at-the-money options.
- $\phi^{\text{ATM}} > \phi^{\text{ITM}} > \phi^{\text{OTM}}$: the self excitation of buy or sell orders is the highest on at-the money options and the lowest on out-of-the-money options.

⁵The extension to multiple transaction sizes is straightforward using marked point processes, see for example Chapter 6.




Figure 8.3: PnL with respect to time (in seconds) under the Poisson and Hawkes assumption.

We show in Figure 8.3 the Profit and Loss (PnL) obtained using the optimal strategy (8.10) when we perform a backtest with the assumption that the intensity of orders' arrival is driven by Hawkes processes compared to the one obtained with Poisson processes assumption (corresponding to the case $\phi = 0$).

8.4.2 Market impact curves

8.4.2.1 Description

Another important application of our framework is the derivation of market impact curves for a given trading strategy at the microscopic scale. We take again the example of option market-making with optimal trading strategy given by (8.10). We use the rescaling results of Section (8.3.3) on the volatility surface:

$$(\bar{\sigma}_t^{(k,\tau)})_{t \in [0, T]} = (\bar{N}_t^{(k,\tau)+} - \bar{N}_t^{(k,\tau)-})_{t \in [0, T]},$$

where $\bar{N}_t^{(k,\tau)+}, \bar{N}_t^{(k,\tau)-}$ have intensities defined by (8.11), to obtain an impacted volatility surface $(\bar{\sigma}_t)_{t \in [0, T]}$ in the form of (8.9). Thus, we can compute the market impact of the market-making strategy (8.10) on the volatility surface at time $t \in [0, T]$ as

$$\text{MI}(t) = \mathbb{E}[\|\bar{\sigma}_t - \sigma_t\|_1], \quad (8.12)$$

where σ_t corresponds to the rescaled volatility surface in absence of trading activity, meaning that the Hawkes processes driving the corresponding microscopic volatility surface have intensity

$$\lambda_t^{(k,\tau)s} = \mu^{(k,\tau)s} + \sum_{(\tilde{k}, \tilde{\tau})} \sum_{\tilde{s} \in \{+, -\}} \int_0^t \phi^{(k,\tau)s, (\tilde{k}, \tilde{\tau})\tilde{s}}(t-u) dN_u^{(\tilde{k}, \tilde{\tau})\tilde{s}}.$$

For a given trading strategy, we can therefore compute the cumulated impact on the whole volatility surface using (8.12). We can also compute the market impact of the trading strategy on a specific point (k, τ) of the surface, that is:

$$\text{MI}^{(k,\tau)}(t) = \mathbb{E}[\bar{\sigma}_t^{(k,\tau)} - \sigma_t^{(k,\tau)}].$$

Thus, depending on the trading strategy, we are able to compare its influence on specific parts of the volatility surface. It is of particular interest for a desk of systematic option trading wishing to estimate the PnL of the strategies.

8.4.2.2 Numerical results

In the classical optimal trading or optimal market-making models, the price processes of stock or options are assumed to be independent from the trading activity. One extension, see Guéant, Lehalle, and Fernandez-Tapia [142], consists in adding a linear temporary impact, that is the price of an option (k, τ) is given by

$$d\tilde{C}_t^{k,\tau} = dC^{k,\tau} + \xi^{k,\tau} (d\bar{N}_t^{(k,\tau)+} - d\bar{N}_t^{(k,\tau)-}), \quad \xi^{k,\tau} > 0,$$

where the dynamics of $C^{k,\tau}$ is the one of a stochastic volatility model, see Chapter 6. In the case of Poisson process, that is $\lambda_t^{(k,\tau)s} = 0$ for all $(k, \tau) \in \mathcal{K} \times \mathcal{T}, s \in \{+, -\}$, there is no cross-impact: for example, buy orders on option (k_1, τ) do not change the price of option (k_2, τ) . We use the same model parameters as in the previous section, and set $\xi^{k,\tau} = \frac{5 \cdot 10^{-4}}{C_0^{k,\tau}}$ for all $k \in \mathcal{K}$, meaning that a buy (resp. sell) order increases (resp. decreases) the price of the option by five basis points. We compute the cross-impact curves of option k_1 on the prices of options k_2 and k_3 .

We show in Figure 8.4 the results in the case of Poisson and Hawkes intensities. The Poisson processes are unable to reproduce the important stylized fact of options markets that buying or selling an option impacts the whole volatility surface. As $\phi^{\text{OTM-ATM}} < \phi^{\text{ITM-ATM}}$, the impact on option k_2 of trading option k_1 is obviously higher than on option k_3 .




Figure 8.4: Cross impact curves (in basis points) with respect to time (in seconds) in the Poisson and Hawkes cases.

8.5 Conclusion

We designed a Hawkes-based model for the dynamics of the implied volatility at the microscopic scale. The Hawkes kernel coefficients control the skew and convexity of the volatility surface, and we present sufficient conditions to ensure the absence of arbitrage. These conditions give simple parametrization of the kernel and therefore reduce the number of coefficients to estimate. At the scaling limit, we use some existing results to show that the macroscopic volatility surface dynamics are a sum of risk factors having rough volatility. We show how to parametrize the Hawkes kernel to recover the classic “level-skew-convexity” behavior of the volatility surface at the macroscopic limit. Finally, we conduct a backtest of systematic option market-making strategies using our framework and compute the associated market impact on the volatility surface.

This work opens doors for several extensions: introduction of interactions between the moves of the spot and the volatility surface and its macroscopic extension, study of the scaling limit of the volatility surface when accounting for different Hurst exponents for different options’ moneyness. These extensions would lead to new no-arbitrage conditions on the volatility surface.

Part III

Optimal trading in high dimension

Chapter 9

Adaptive trading strategies across liquidity pools

9.1 Introduction

A vast majority of quantitative trading strategies are based on cross-platform arbitrage. These strategies involve cross-listed stocks, that are assets traded on two or more liquidity venues. In Jain and Jain [173], the authors investigate the prices of cross-listed stocks in different venues and provide evidence of price deviations for the majority of the 600 cross-listed stocks they studied. In Alsayed and McGroarty [16], the authors highlight mispricing that can exist between a domestic stock and its ADR (American Deposit Receipt) counterpart. The study conducted in Werner and Kleidon [267] for US-UK cross-listed stocks shows that markets for cross-listed securities are among the most heavily arbitrated. In particular, higher potential of arbitrage can be exploited for cross-listed stocks from emerging markets, see Rabinovitch, Silva, and Susmel [235].

Usually, the trader builds an execution curve targeting, for example, an Implementation Shortfall or volume-weighted average price (VWAP). Then, he buys or sells shares of the asset following the execution curve by sending limit and market orders to the different venues. But how to find the best splitting of orders between the venues? The trader splits his orders depending mainly on the imbalance and spread of the different venues, which of course depend on each other. For example, a higher imbalance on the ask side of one venue can indicate a potential imminent price change and a lower probability to have an ask limit order executed, so it may be more profitable to send the order to another venue. The problem of optimal trading across liquidity pools has been treated, for instance, in Almgren and Harts [11], Cont and Kukanov [81], Laruelle, Lehalle, and Pagès [193, 194]. In Almgren and Harts [11], the authors develop a dynamic estimate of the hidden liquidity present on several venues and use this information to make order splitting decisions (Smart Order Routing). The paper Cont and Kukanov [81] solve a general order placement problem and provide explicit solution for the optimal split between limit and market orders on different venues. Finally, in Laruelle, Lehalle, and Pagès [193, 194], the authors build a stochastic algorithm to find the optimal splitting between liquidity pools, including dark venues.

Building a good model for optimal trading cross-listed assets requires to take into account the cross-dependence between the imbalance and spread of each venue, as well as the probability and the proportion of execution of limit orders. However, the quality of the model will mainly rely on the estimation of the market parameters. If one assumes constant parameters over the trading period, he believes in the quality of his parameters' estimation. In this case, his strategy is not robust to changes in price dynamics or the platforms' behavior. For example, the trading period may occur when another market participant is executing a buy (or sell) metaorder on one or several venues. This participant

will consume the vast majority of the liquidity available on the sell (or buy) side of those venues. If the trader does not adjust his market parameters, the algorithm will keep sending limit orders on the platform where the metaorder is being split, without an execution opportunity. That is why it is essential to update model parameter estimations with new information obtained from observing the market dynamics. Here, we treat the updates in a Bayesian manner.

In this chapter, we formulate the problem of a trader dealing in a stock, listed on several venues, by placing limit and market orders. Trader's activity can be formulated as a stochastic control problem. The controls are the splitting of volumes between limit and market orders on each venue, and the limits chosen by the trader. The optima are obtained from a classical Hamilton-Jacobi-Bellman (HJB) quasi-variational inequality, which, for a parsimonious model, can be easily solved by grid methods.

Then we propose a Bayesian update of each market parameter, decoupled from the control problem. One of the advantages of this method is the simplicity of the formulae for each parameter's posterior estimate. In particular, we do not need to use Markov chain Monte-Carlo. This method's choice comes from the fact that updating the market parameters continuously in the control problem increases the number of state variables drastically, leading to high computation time. A continuous Bayesian update would require first to compute the conditional expectation of the value function given the market parameters and then to integrate it over their posterior distribution. This last integration brings multiple non-linearities in the equation, making this fully Bayesian control problem hard to solve numerically.

The proposed Bayesian procedure is easier to apply in practice: we divide the trading period into time slices of about several seconds up to a few minutes long, assuming that market conditions do not vary drastically throughout the slice. For each slice, we keep track of all the market events. Specifically, on each venue, we count the number of executed limit orders, the executed proportions (for example, 50% or 100% of the order volume) given the couple spread-imbalance on each venue at the time of the execution. We also keep track of the price dynamics. At the end of each slice, we update our view on the market parameters and recompute the optimal trading strategy for the next slice. This application of the Bayesian updates on slices of execution is time-inconsistent. However, we see it as a first step toward a more integrated Bayesian learning framework for cross-listed trading. By using finite difference schemes or deep reinforcement learning methods (which could also be mixed) for high-dimensional PDE resolution, we can compute in a couple of minutes the optimal trading strategy on a slice given new market conditions.

This chapter aims at giving a useful and applicable model for practitioners who work on cross-trading strategies. For a quantitative firm, the control model is flexible enough to reproduce the main stylized facts about the market and to design trading strategies taking into account real signals. Moreover, the procedure for Bayesian updates of market parameters in the control problem enables to reevaluate the optimal strategy when the market conditions may differ from the prior empirical estimation of the trader.

The chapter has the following structure: in Section 9.2, we describe the framework for cross-platform trading and formulate the trader's optimization problem. In Section 9.2.2, we derive the Hamilton-Jacobi-Bellman quasi-variational inequality (HJBQVI) associated with the trader's optimal trading problem. We introduce a change of variable to reduce the dimensionality of the problem and prove the existence and uniqueness of the viscosity solution of the initial HJBQVI in 9.A.1. In Section 9.3, we first define the conjugate Bayesian update of all market parameters. Then, we describe the update procedure in practice and its link to the control problem of the trader. Section 9.4 is dedicated to some extensions of the model and their impact on the dimensionality of the resulting HJBQVI. We devote Section 9.5 to numerical results, for the sake of clarity of interpretations considered in the case of limit orders only. Finally in Appendix 9.A.3, we present an application of the Bayesian update of the market parameters to the problem of an OTC market-maker.

9.2 Optimal trading on several liquidity pools

The model presented in this section is a generalization of the classic optimal trading framework, developed notably in Avellaneda and Stoikov [26], Guéant, Lehalle, and Fernandez-Tapia [142], Guilbaud and Pham [143], Stoikov and Sağlam [258] and in the reference books Cartea, Jaimungal, and Penalva [70], Guéant [137], to the case of several liquidity venues.

9.2.1 Framework

We consider a trader acting on N liquidity platforms operating with limit order books over time interval $[0, T]$. We define a filtered probability space $(\Omega, \mathcal{F}, \mathbb{P})$ on which all stochastic processes are defined and $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$. He trades continuously on each venue by sending limit and market orders. For $n \in \{1, \dots, N\}$, the n -th venue is characterized by the following continuous-time Markov chains:

- the bid-ask spread process $(\psi_t^n)_{t \in [0, T]}$ taking values in the state space $\bar{\psi}^n = \{\delta^n, \dots, J\delta^n\}$,
- the imbalance process $(I_t^n)_{t \in [0, T]}$ taking values in the state space $\bar{I}^n = \{I_1^n, \dots, I_K^n\}$,

where $J, K \in \mathbb{N}$ denote the number of possible spreads and imbalances respectively and δ^n stands for the tick size of the n -th venue. We define the sets $\Psi = \{\Psi_1, \dots, \Psi_{\#\Psi}\}, \mathcal{I} = \{\mathcal{I}_1, \dots, \mathcal{I}_{\#\mathcal{I}}\}$ of disjoint intervals, representing different market regimes of interest in terms of spreads and imbalances.

Example 9.1. Assume for all $n \in \{1, \dots, N\}$ that $\delta^n = \delta$. The set $\Psi = \{\delta, \{2\delta, 3\delta\}, \{4\delta, 5\delta\}\}$ denotes three spread regimes: low (one tick), medium (two or three ticks), and high (four or five ticks).

Example 9.2. Assume for all $n \in \{1, \dots, N\}$ and $k \in \{1, \dots, K\}$ that $I_k^n = I_k$. In this case the set $\mathcal{I} = \{[-1, -0.66], (-0.66, -0.33], (-0.33, 0.33], (0.33, 0.66], (0.66, 1]\}$ denotes five regimes of imbalance: low (-33% to 33%), medium on the ask (resp. bid) from 33% to 66% (resp. from -66% to -33%) and high on the ask (resp. bid) from 66% to 100% (resp. from -100% to -66%).

Whenever the spread and the imbalance of each venue enter the state $\mathbf{k} = (\mathbf{k}^\psi, \mathbf{k}^I) \in \mathcal{K}$ where $\mathcal{K} = \prod_{n=1}^N \bar{\psi}^n \times \prod_{n=1}^N \bar{I}^n$, they remain in this state for a time exponentially distributed with mean $\frac{1}{\nu_{\mathbf{k}}}$. We define a transition matrix $\mathbf{P} = (p_{\mathbf{k}\mathbf{k}'})$, $(\mathbf{k}, \mathbf{k}') \in \mathcal{K}$, and corresponding intensity vectors $\nu = (\nu_{\mathbf{k}})_{\mathbf{k}}^\top$. We assume that $p_{\mathbf{k}\mathbf{k}} = 0$, meaning that we cannot come to the same state twice in a row. The infinitesimal generator of the processes can be written as

$$r_{\mathbf{k}\mathbf{k}'} = \nu_{\mathbf{k}} p_{\mathbf{k}\mathbf{k}'}, \quad \text{if } \mathbf{k} \neq \mathbf{k}', \quad r_{\mathbf{k}\mathbf{k}} = - \sum_{\mathbf{k}' \neq \mathbf{k}} r_{\mathbf{k}\mathbf{k}'} = -\nu_{\mathbf{k}}, \quad \text{otherwise.}$$

Remark 42. This general formulation allows us a full coupling between the spread and imbalance of all venues. If one wants a more parsimonious model, the following simplifications could be made. When the spread (imbalance) of the n -th venue enters the state k , it remains there for an exponentially distributed time with mean $(\nu_k^{n,\psi})^{-1}$ ($(\nu_k^{n,I})^{-1}$ for the imbalance). Therefore, we define a transition matrix $\mathbf{P}^{n,\psi} = (p_{kk'}^{n,\psi})$, $n \in \{1, \dots, N\}, (k, k') \in \bar{\psi}^n$ such that $p_{kk}^{n,\psi} = 0$, and corresponding intensity vectors $\nu^{n,\psi} = (\nu_1^{n,\psi}, \dots, \nu_K^{n,\psi})^\top$. Similarly we define a transition matrix $\mathbf{P}^{n,I}$ for the imbalance. Then, the infinitesimal generator of the processes can be written as

$$r_{kk'}^{n,\psi} = \nu_k^{n,\psi} p_{kk'}^{n,\psi}, \quad \text{if } k \neq k', \quad r_{kk}^{n,\psi} = - \sum_{k' \neq k} r_{kk'}^{n,\psi} = -\nu_k^{n,\psi}, \quad \text{otherwise.}$$

This framework will be used in Section 9.5, where we present the numerical results.

In what follows, the trader designs his strategy on the ask side of the market (optimal liquidation problem). The extension to trading on both sides of the market is straightforward and does not cause an increase in the problem's dimensionality. The number of, possibly partially, filled ask orders in the venue n is modeled by a Cox process denoted by $N^n, n \in \{1, \dots, N\}$ with intensities $\lambda^n(\psi_t, I_t, p_t^n, \ell_t)$ where $p_t^n \in Q_\psi^n$ represent the limit at which the trader sends a limit order of size ℓ_t^n , and

$$Q_\psi^n = \{0, 1\} \text{ if } \psi^n = \delta^n, \text{ and } \{-1, 0, 1\} \text{ otherwise,}$$

$$\mathcal{A} = \left\{ (\ell_t)_{t \in [0, T]} : \mathbb{F}\text{-predictable and for all } t \in [0, T], 0 \leq \sum_{n=1}^N \ell_t^n \leq q_t \right\},$$

where $(q_t)_{t \in [0, T]}$ is defined in Equation (9.1). Practically for $n \in \{1, \dots, N\}$, when the spread is equal to the tick size, the trader can post at the first best limit ($p^n = 0$) or the second best limit (if $p^n = 1$). When the spread is equal to two ticks or more, the trader can either create a new best limit ($p^n = -1$) or post at the best or the second best limit as previously. The arrival intensity of a buy market order at time t on the venue $n \in \{1, \dots, N\}$ at the limit $p \in Q_\psi^n$, given a couple $(\psi_t, I_t) = \mathbf{m}$ of spread and imbalance on each venue, is equal to $\lambda^{n, \mathbf{m}, p} > 0$. When the trader posts limit orders of volume ℓ_t^n on the n -th venue for $n \in \{1, \dots, N\}$, the probability that it is executed is equal to $f^\lambda(\ell_t)$, where $f^\lambda(\cdot) \in [0, 1]$ is a continuously differentiable function, decreasing with respect to each of its coordinate. Therefore, the arrival intensity of an ask market order filling the buy limit order of the trader on the n -th venue at the limit p_t^n , given spread and imbalance (ψ_t, I_t) is a multi-regime function defined by

$$\lambda^n(\psi_t, I_t, p_t^n, \ell_t) = f^\lambda(\ell_t) \sum_{\mathbf{m} \in \mathcal{M}, p \in Q_\psi^n} \lambda^{n, \mathbf{m}, p} \mathbf{1}_{\{(\psi_t, I_t) \in \mathbf{m}, p_t^n = p\}},$$

where $\mathcal{M} = \Psi^N \times \mathcal{I}^N$. Moreover, we allow for partial execution, the fact of which we represent by random variables $\epsilon_t^n \in [0, 1]$. The proportion of executed volume for limit orders in each venue depends on the spread and the imbalance in all N venues, as well as the volume and the limit of the order chosen by the trader. We assume a categorical distribution with $R > 0$ different execution proportions $\omega^r, r \in \{1, \dots, R\}$ for each venue with $\mathbb{P}(\epsilon_t^n = \omega^r) = \rho^{n, r}(\psi_t, I_t, p_t^n, \ell_t)$, where

$$\rho^{n, r}(\psi_t, I_t, p_t^n, \ell_t) = f^\rho(\ell_t) \sum_{\mathbf{m} \in \mathcal{M}, p \in Q_\psi^n} \rho^{n, \mathbf{m}, p, r} \mathbf{1}_{\{(\psi_t, I_t) \in \mathbf{m}, p_t^n = p\}},$$

where $f^\rho(\cdot)$ is a continuously differentiable function, decreasing with respect to each of its coordinate.

Remark 43. *The estimation of this kind of parameters for executed proportions can be quite intricate in practice. To simplify, one can assume that $\rho^{n, r}(\psi_t, I_t, p_t^n, \ell_t) = \rho^{n, r} \in [0, 1]$. In practice, this means that there are different execution proportion probabilities inherent by each venue, depending on its toxicity.*

Finally, we allow for the execution of market orders (denoted by a point process $(J_t^n)_{t \in [0, T]}$) on each venue of size $(m_t^n)_{t \in [0, T]} \in [0, \bar{m}]$ where $\bar{m} > 0$ and $J_t^n = J_{t-}^n + 1$. We assume that market orders are always fully executed.

The cash process of the trader at time $t \in [0, T]$ is

$$dX_t = \sum_{n=1}^N \left(\ell_t^n \left(S_t + \frac{\psi_t^n}{2} + p_t^n \delta^n \right) \epsilon_t^n dN_t^n + m_t^n \left(S_t - \frac{\psi_t^n}{2} \right) dJ_t^n \right),$$

where

$$dS_t = \mu dt + \sigma dW_t, \quad (\mu, \sigma) \in \mathbb{R} \times \mathbb{R}^+,$$

is the dynamics of the mid-price process. The inventory process of the trader at time $t \in [0, T]$ is defined by

$$q_t = q_0 - \sum_{n=1}^N \left(\int_0^t \ell_u^n \epsilon_u^n dN_u^n + \int_0^t m_u^n dJ_u^n \right). \quad (9.1)$$

We also assume that the trader has a pre-computed trading curve q^* that he wants to follow (Almgren-Chriss trading curve or VWAP strategy, for example). Then the trader's optimization problem is

$$\sup_{(p,\ell,m) \in Q_\psi \times \mathcal{A} \times [0,\bar{m}]^N} \mathbb{E} \left[X_T + q_T S_T - \int_0^T g(q_t - q_t^*) dt \right], \quad (9.2)$$

where the function g penalizes deviation from the pre-computed optimal trading curve.

9.2.2 The Hamilton-Jacobi-Bellman quasi-variational inequality

The HJBQVI associated with the optimization problem of the trader (9.2) is the following:

$$\begin{aligned} 0 = \min & \left\{ -\partial_t u(t, x, q, S, \psi, I) + g(q - q_t^*) - \mu \partial_S u - \frac{1}{2} \sigma^2 \partial_{SS} u \right. \\ & - \sum_{\mathbf{k} \in \mathcal{K}} r_{(\psi, I), (\mathbf{k}^\psi, \mathbf{k}^I)} (u(t, x, q, S, \mathbf{k}^\psi, \mathbf{k}^I) - u(t, x, q, S, \psi, I)) \\ & - \sup_{p \in Q_\psi, \ell \in \mathcal{A}} \sum_{n=1}^N \lambda^n(\psi, I, p^n, \ell) \mathbb{E} \left[u \left(t, x + \epsilon^n \ell^n \left(S + \frac{\psi^n}{2} + p^n \delta^n \right), q - \ell^n \epsilon^n, S, \psi, I \right) \right. \\ & \left. \left. - u(t, x, q, S, \psi, I) \right]; \sum_{n=1}^N u(t, x, q, S, \psi, I) - \sup_{m^n \in [0, \bar{m}]} u \left(t, x + m^n \left(S - \frac{\psi^n}{2} \right), q - m^n, S, \psi, I \right) \right\}, \end{aligned} \quad (9.3)$$

with terminal condition

$$u(t, x, q, S, \psi, I) = x + qS,$$

where $\psi = (\psi^1, \dots, \psi^N)$, $I = (I^1, \dots, I^N)$. The expectation in (9.3) is taken over the variables ϵ^n , $n \in \{1, \dots, N\}$. We prove the following theorem in 9.A.1:

Theorem 9.1. *There exists a unique viscosity solution to the HJBQVI (9.3), which coincides with the value function of the control problem of the trader (9.1).*

The proof of existence and uniqueness of the viscosity solution mainly relies on adaptations of the theory of the second order viscosity solution with jumps, see Barles and Imbert [36], for example.

The value function has to be linear with respect to the cash process and the mark-to-market value of the trader's inventory due to the form of the terminal condition. Therefore we use the following ansatz:

$$u(t, x, q, S, \psi, I) = x + qS + v(t, q, \psi, I).$$

The HJBQVI then becomes a system of ODEs with $2N + 1$ state variables:

$$\begin{aligned} 0 = \min & \left\{ -\partial_t v(t, q, \psi, I) + g(q - q_t^*) - \mu q \right. \\ & - \sum_{\mathbf{k} \in \mathcal{K}} r_{(\psi, I), (\mathbf{k}^\psi, \mathbf{k}^I)} (v(t, q, \mathbf{k}^\psi, \mathbf{k}^I) - v(t, q, \psi, I)) \\ & - \sup_{p \in Q_\psi, \ell \in \mathcal{A}} \sum_{n=1}^N \lambda^n(\psi, I, p^n, \ell) \mathbb{E} \left[\epsilon^n \ell^n \left(\frac{\psi^n}{2} + p^n \delta^n \right) + v(t, q - \ell^n \epsilon^n, \psi, I) \right. \\ & \left. \left. - v(t, q, \psi, I) \right]; \sum_{n=1}^N v(t, q, \psi, I) - \sup_{m^n \in [0, \bar{m}]} -m^n \frac{\psi^n}{2} + v(t, q - m^n, \psi, I) \right\}, \end{aligned} \quad (9.4)$$

with terminal condition $v(T, q, \psi, I) = 0$.

Conditionally on the market parameters such as the transition matrix of both the spread and the imbalance processes, the drift and volatility of the underlying asset and the execution proportion probabilities, solving Equation (9.4) is done using simple finite difference schemes and the optimal splitting of volumes as well as the optimal limits can be computed in advance.

If one wants to incorporate directly Bayesian learning of the parameters in the control problem, the result would be a very high number of state variables, which makes the problem intractable in practice. For example, if we want to update continuously the value of the processes λ^n for $n \in \{1, \dots, N\}$ we need to add the counting processes $(N_t^n)_{t \in [0, T]}$ to the state variables, which increases the dimension of the HJBQVI (9.4) by N . What we propose in the following section is a practical way to update the market parameters according to trader's observations in a Bayesian way. This method, which is performed separately from the optimization procedure, allows to update, at the end of a slice, the trading strategy according to changing market conditions.

9.3 Adaptive trading strategies with Bayesian update

The framework presented in the above section allows to choose generic parametric forms for the state variables prior distributions (transition matrix of spreads and imbalances, intensities of orders' arrival on each venue) suitable to the use of conjugate Bayesian updates.

9.3.1 Bayesian update of the model parameters

In this section, we present the conjugate Bayesian update of the market parameters and how to choose the prior distributions.

9.3.1.1 Update of the intensities

Let us recall the form of the intensities for counterpart market orders' arrival:

$$\lambda^n(\psi_t, I_t, p_t^n, \ell_t) = f^\lambda(\ell_t) \sum_{\mathbf{m} \in \mathcal{M}, p \in Q_\psi^n} \lambda^{n, \mathbf{m}, p} \mathbf{1}_{\{(\psi_t, I_t) \in \mathbf{m}, p_t^n = p\}}.$$

In the vast majority of optimal liquidation models, the probability of execution $\lambda^{n, \mathbf{m}, p}$ is estimated empirically. We propose to put a prior law $\Gamma(\alpha^{n, \mathbf{m}, p}, \beta^{n, \mathbf{m}, p})$ on the arrival rate, and to update a prior belief at the end of each slice of execution. The parameters $\alpha^{n, \mathbf{m}, p}, \beta^{n, \mathbf{m}, p}$ are chosen by the trader according to his vision of the market before he starts to trade. Up to time $t \in [0, T]$ the trader observes the processes

$$N_t^{n, \mathbf{m}, p} = \int_0^t \mathbf{1}_{\{(\psi_s, I_s) \in \mathbf{m}, p_s^n = p\}} dN_s^n,$$

which represent the number of executed orders on each venue for every spread-imbalance zone \mathbf{m} . The posterior distribution of $\lambda^{n, \mathbf{m}, p}$ for $n \in \{1, \dots, N\}$ is then given by

$$\lambda^{n, \mathbf{m}, p} | N_t^{n, \mathbf{m}, p} \sim \Gamma(\alpha^{n, \mathbf{m}, p} + N_t^{n, \mathbf{m}, p}, \beta^{n, \mathbf{m}, p} + \int_0^t f^\lambda(\ell_s) ds),$$

and at time t , our best estimate of the filling ratio becomes

$$\lambda^{n, \mathbf{m}, p}(t, N_t^{n, \mathbf{m}, p}, \ell_t) = \mathbb{E} \left[\lambda^{n, \mathbf{m}, p} | N_t^{n, \mathbf{m}, p} \right] = \frac{\alpha^{n, \mathbf{m}, p} + N_t^{n, \mathbf{m}, p}}{\beta^{n, \mathbf{m}, p} + \int_0^t f^\lambda(\ell_s) ds}.$$

The posterior estimate of the intensity $\lambda^n(\psi_t, I_t, p_t^n, \ell_t)$ becomes

$$\hat{\lambda}^n(\psi_t, I_t, p_t^n, \ell_t) = f^\lambda(\ell_t) \sum_{\mathbf{m} \in \mathcal{M}, p \in Q_\psi^n} \frac{\alpha^{n, \mathbf{m}, p} + N_t^{n, \mathbf{m}, p}}{\beta^{n, \mathbf{m}, p} + \int_0^t f^\lambda(\ell_s) ds} \mathbf{1}_{\{(\psi_t, I_t) \in \mathbf{m}, p_t^n = p\}}.$$

As the convergence of the prior parameters toward the true market specification follows from the central limit theorem, the convergence rate equals to $\frac{1}{\sqrt{o^{\mathbf{m}}}}$ where $o^{\mathbf{m}}$ is the number of observations of filled limit orders on the spread-imbalance zone \mathbf{m} . If we consider even a quite parsimonious model, for example two venues, two regimes of spread and three regimes of imbalance, we have $\#\mathcal{M} = 36$ different zones. This means that we need a sufficiently large amount of observations (large number of executed orders) to get an accurate approximation of the market behavior.

If the trader anticipates that the number of observations he will have is not adequate to obtain a suitable approximation of the “true” market parameters (in the case of a mid to low frequency strategy with only a few number of trades throughout the day), he might choose at the beginning the couples $(\alpha^{n,\mathbf{m},p}, \beta^{n,\mathbf{m},p})$ such that $\frac{\alpha^{n,\mathbf{m},p}}{\beta^{n,\mathbf{m},p}} \gg \frac{N_t^{n,\mathbf{m},p}}{\int_0^t f^\lambda(\ell_s) ds}$. That way, his prior will not be sensitive to a small number of observations, and with sufficient number of observations the prior will have less influence and the estimation will be less biased.

9.3.1.2 Update of the executed proportion

We propose to use the Dirichlet prior distribution on the executed proportion parameters so that $\rho^{n,\mathbf{m},p} \sim \text{Dirichlet}(\alpha^{\epsilon,n,\mathbf{m},p})$ where $\alpha^{\epsilon,n,\mathbf{m},p} = (\alpha^{\epsilon,n,\mathbf{m},p,1}, \dots, \alpha^{\epsilon,n,\mathbf{m},p,R})$ for all $(n, \mathbf{m}, p, r) \in \{1, \dots, N\} \times \mathcal{M} \times Q_\psi \times \{1, \dots, R\}$. Given observations of ϵ_t^n , the executed proportion parameters have Dirichlet posterior distribution

$$\rho^{n,\mathbf{m},p} \sim \text{Dirichlet}(\alpha^{\epsilon,n,\mathbf{m},p} + c_t^{n,\mathbf{m},p}),$$

where $c_t^{n,\mathbf{m},p} = (c_t^{n,\mathbf{m},p,1}, \dots, c_t^{n,\mathbf{m},p,R})$ and $c_t^{n,\mathbf{m},p,r} = \sum_{s \leq t} \mathbf{1}_{\{\epsilon_s^n = \omega^r, (\psi_s, I_s) \in \mathbf{m}, p_s^n = p, N_s^n - N_{s-}^n = 1\}}$ is the number of observations before time t in zone \mathbf{m} for a limit p in the venue n . Therefore, the ϵ_t^i have the following posterior distribution:

$$\hat{\rho}^{n,r}(\psi_t, I_t, p_t^n, \ell_t) = f^p(\ell_t) \sum_{\mathbf{m} \in \mathcal{M}, p \in Q_\psi} \frac{\alpha^{\epsilon,n,\mathbf{m},p,r} + c_t^{n,\mathbf{m},p,r}}{\sum_{r=1}^R (\alpha^{\epsilon,n,\mathbf{m},p,r} + c_t^{n,\mathbf{m},p,r})} \mathbf{1}_{\{(\psi_t, I_t) \in \mathbf{m}, p^n = p\}}.$$

This Bayesian update is linked to the filling of limit orders of the trader: the proportion executed is updated only if the limit order is (partially) executed. If one chooses the parametrization independent of the spread-imbalance zones and the order volume, that is execution proportion depends only on the venue, the speed of convergence is much faster as the same amount of gathered information is used to update a much smaller number of parameters. Using this more parsimonious parametrization the trader can rely on the observations more than on his prior.

9.3.1.3 Update of the characteristics of the venues

We observe the states of the Markov chains $\psi_{t_d}, I_{t_d}, d \in \{0, \dots, D\}$ and the times t_d of the $D > 0$ transitions. The likelihood function for the spread and the imbalance processes is

$$\begin{aligned} \mathcal{L}(\mathbf{P}, \nu | \psi_{t \leq t_D}, I_{t \leq t_D}) &= \prod_{d=1}^D \nu_{t_{d-1}} \exp(-\nu_{t_{d-1}}(t_d - t_{d-1})) p_{(\psi_{t_{d-1}}, I_{t_{d-1}})(\psi_{t_d}, I_{t_d})} \\ &\propto \prod_{\mathbf{k} \in \mathcal{K}} (\nu_{\mathbf{k}})^{n_{\mathbf{k}}} \exp(-\nu_{\mathbf{k}} T_{\mathbf{k}}) \prod_{\mathbf{k}' \in \mathcal{K}} (p_{\mathbf{k}\mathbf{k}'})^{n_{\mathbf{k}\mathbf{k}'}} \end{aligned}$$

where $n_{\mathbf{k}\mathbf{k}'}$ is the number of observed transitions from state \mathbf{k} to \mathbf{k}' for $(\mathbf{k}, \mathbf{k}') \in \mathcal{K}$, $T_{\mathbf{k}}$ is the total time spent in state \mathbf{k} , and $n_{\mathbf{k}} = \sum_{\mathbf{k}' \in \mathcal{K}} n_{\mathbf{k}\mathbf{k}'}$ is the total number of transitions out of state \mathbf{k} .

Given independent prior distributions for \mathbf{P}, ν , the posterior distributions will also be independent. We can carry out Bayesian inference separately on the probability matrix and the intensity vectors of

the Markov chains. We assume the following priors:

$$\nu_{\mathbf{k}} \sim \Gamma(a_{\mathbf{k}}, b_{\mathbf{k}}), \quad \mathbf{p}_{\mathbf{k}} = (p_{\mathbf{k}\mathbf{k}'}')_{\mathbf{k}' \in \mathcal{K}} \sim \text{Dirichlet}(\alpha_{\mathbf{k}}), \quad \text{where } \alpha_{\mathbf{k}} = (\alpha_{\mathbf{k}\mathbf{k}'}')_{\mathbf{k}' \in \mathcal{K}}.$$

Given these conjugate priors, our best estimators of $\nu_{\mathbf{k}}, \mathbf{p}_{\mathbf{k}}$ are

$$\hat{\nu}_{\mathbf{k}} = \frac{a_{\mathbf{k}} + n_{\mathbf{k}\cdot} - 1}{b_{\mathbf{k}} + T_{\mathbf{k}}}, \quad \hat{p}_{\mathbf{k}\mathbf{k}'} = \frac{\alpha_{\mathbf{k}\mathbf{k}'} + n_{\mathbf{k}\mathbf{k}'}}{\sum_{\mathbf{l} \neq \mathbf{k}} (\alpha_{\mathbf{k}\mathbf{l}} + n_{\mathbf{k}\mathbf{l}})}.$$

Then the posterior transition matrix is

$$\hat{r}_{\mathbf{k}\mathbf{k}'} = \hat{\nu}_{\mathbf{k}} \hat{p}_{\mathbf{k}\mathbf{k}'}, \quad \mathbf{k} \neq \mathbf{k}', \quad \hat{r}_{\mathbf{k}\mathbf{k}} = -\hat{\nu}_{\mathbf{k}} \text{ otherwise.}$$

This update aims at finding the “true” behavior of the imbalance and spread processes of each venue. This is of particular importance if an event (for instance, an announcement or news) happens in the market. More specifically, if one event occurs in a particular platform (if a metaorder is executed in one specific platform, for example), this helps to discriminate one venue from the others and to redirect the orders to the less toxic liquidity platforms. Given the large number of observations (transitions from one state of imbalance or spread to another occur fast), the trader does not necessarily need to be confident about his prior distributions.

Remark 44. *If one wants to use a more parsimonious model as in Remark 42, the same methodology applies. In particular for $k \in \bar{\psi}^n$, we assume the following prior:*

$$\nu_k^{n,\psi} \sim \Gamma(a_k^{n,\psi}, b_k^{n,\psi}), \quad \mathbf{p}_k^{n,\psi} = (p_{kk'}^{n,\psi})_{k' \in \bar{\psi}^n} \sim \text{Dirichlet}(\alpha_k^{n,\psi}), \quad \text{where } \alpha_k^{n,\psi} = (\alpha_{kk'}^{n,\psi})_{k' \in \bar{\psi}^n}.$$

Given these conjugates priors, our best estimators of $\nu_k^{n,\psi}, \mathbf{p}_k^{n,\psi}$ are

$$\hat{\nu}_k^{n,\psi} = \frac{a_k^{n,\psi} + n_{k\cdot}^{n,\psi} - 1}{b_k^{n,\psi} + T_k^{n,\psi}}, \quad \hat{p}_{kk'}^{n,\psi} = \frac{\alpha_{kk'}^{n,\psi} + n_{kk'}^{n,\psi}}{\sum_{\mathbf{l} \neq k} (\alpha_{kl}^{n,\psi} + n_{kl}^{n,\psi})}.$$

The posterior transition matrix is given by

$$\hat{r}_{kk'}^{n,\psi} = \hat{\nu}_k^{n,\psi} \hat{p}_{kk'}^{n,\psi}, \quad k \neq k', \quad \hat{r}_{kk}^{n,\psi} = -\hat{\nu}_k^{n,\psi} \text{ otherwise.}$$

Similar formulae apply for $\nu_k^{n,I}, \mathbf{p}_k^{n,I}$.

9.3.1.4 Update of the mid-price

We recall that the price process has the following dynamics:

$$dS_t = \mu dt + \sigma dW_t,$$

so that $(S_t - S_0 | \mu, \sigma) \sim \mathcal{N}(\mu t, \sigma^2 t)$. We assume that the couple (μ, σ^2) follows a Normal-Inverse-Gamma prior distribution $NIG(\mu_0, \nu, \alpha^s, \beta^s)$, where $(\mu_0, \nu, \alpha^s, \beta^s) \in \mathbb{R} \times \mathbb{R}_+^3$. Therefore the posterior distribution has the following form:

$$(\mu, \sigma^2 | S_t - S_0) \sim NIG\left(\frac{(S_t - S_0) + \mu_0 \nu}{\nu + t}, \nu + t, \alpha^s + \frac{t}{2}, \beta^s + \frac{t\nu}{\nu + t} \frac{(S_t - S_0 - \mu_0)^2}{2}\right).$$

Given our observations of the stock price up to time t , the best approximation of the drift and volatility are given by

$$\mu(t, S_t) = \mathbb{E}[\mu | S_t - S_0] = \frac{(S_t - S_0) + \mu_0 \nu}{\nu + t}, \quad \sigma^2(t, S_t) = \mathbb{E}[\sigma^2 | S_t - S_0] = \frac{\beta^s + \frac{t\nu}{\nu + t} \frac{(S_t - S_0 - \mu_0)^2}{2}}{\alpha^s + \frac{t}{2} - 1}.$$

The volatility σ does not appear explicitly in the HJBQVI (9.3). However, it is taken into account when the trader computes his trading curve q^* .

In the case where the trader is confident with his estimation of σ , one can use a Normal prior distribution on μ such that $\mu \sim \mathcal{N}(\mu_0, \nu^2)$. Then, the best approximation of the drift is given by

$$\mu(t, S_t) = \mathbb{E}[\mu | S_t - S_0] = \frac{\mu_0 \sigma^2 + \nu^2 (S_t - S_0)}{\sigma^2 + \nu^2 t}. \quad (9.5)$$

If the trader firmly believes in the a priori parameter estimation, he can set ν close to 0 so that he mostly relies on his prior. On the contrary, if he sets ν high enough, his estimation comes mostly from market information. Given the large amount of data coming from the market (each time step corresponding to one new observation), convergence to the real value of the drift is fast.

Remark 45. *One can argue about the use of a frequentist estimator of the model parameters, which would actually lead to quite similar formulae. However the original problem, that is continuous update of market parameters in the control problem, is of Bayesian nature. Moreover, in our approach, the formulae for posterior distribution of market parameters are as explicit as in the frequentist approach.*

9.3.2 Algorithm description

We now present the use of the Bayesian updates in order to obtain adaptive trading strategies in practice. We emphasize that the procedure is decoupled from the optimization problem (9.4), so that we do not perform Bayesian optimization but rather a Bayesian update of the parameters of an optimization problem.

Number of time steps is an important parameter of the optimization problem because its choice is a trade-off between computation time and computation precision. To address this problem, we use the trading algorithm with fixed market parameters over a short period of time (a couple of seconds up to a few minutes), which we call a slice. Let us consider $\mathcal{V} > 0$ slices $\mathcal{T}_v = [T_v, T_{v+1}]$, $v = 0, \dots, \mathcal{V} - 1$, such that $T_0 = 0, T_{\mathcal{V}} = T$. We define for each slice $v \in \mathcal{V}$ a set of market parameters

$$\theta_v^m = (r, \rho^n, \lambda^{n, \mathbf{m}, p}, \mu, \sigma)_{\{n \in \{1, \dots, N\}, \mathbf{m} \in \mathcal{M}, p \in \mathcal{Q}_\psi\}}.$$

At each time slice $v \in \{0, \mathcal{V} - 1\}$ starting from $v = 0$ we perform the following algorithm:

1. Take the best estimation of market parameters θ_v^m from the prior distribution for the current slice v .
2. Compute the optimal trading strategy on \mathcal{T}_v using the set of parameters θ_v^m .
3. Observe market events during the current slice (executions, changes of the state).
4. At T_{v+1} , update the parameters θ_{v+1}^m following the Bayes rules described in Section 9.3.

To summarize, we use the output of the control model (the optimal volumes and limits in each venue) over a slice of execution and then run the model again with the updated market parameters. This method, which is clearly time inconsistent, is common practice when one applies optimal control with online parameter estimation, see for example Chapter 6 of this thesis.

We now present some possible extensions of the presented model.

9.4 Model extensions

In this section we describe different potential model extensions and their impact on the problem's dimensionality.

9.4.1 Extension 1: incorporation of signals in the price process

9.4.1.1 Short-term price signals

The two main sources of signals at the microstructural level are the imbalance and the bid-ask spread. Therefore, one can assume a parametric dependence $f^{\text{short}}(\psi_t, I_t)$ of the price process on these two sources, such that the price process becomes

$$dS_t = (\mu + f^{\text{short}}(\psi_t, I_t))dt + \sigma dW_t.$$

In a modified stochastic control problem the term μq in the HJBQVI is replaced by $(\mu + f^{\text{short}}(\psi, I))$, which causes no increase in the dimensionality of the state process.

9.4.1.2 Mid/Long term and path-dependent price signals

When trading on longer time horizon, one can incorporate mid- or long-term signals such as Bollinger bands, moving average or cointegration ratio. For example, consider a signal taking into account the moving average and the maximum of the price process S_t , that is

$$\bar{S}_t = \frac{1}{t} \int_0^t S_s dt, \quad S_t^* = \max_{s \leq t} S_s.$$

The triplet (\bar{S}_t, S_t^*, S_t) is Markovian. Therefore, we can add a long term signal $f^{\text{long}}(S_t, \bar{S}_t, S_t^*)$ into the asset's drift:

$$dS_t = (\mu + f^{\text{long}}(S_t, \bar{S}_t, S_t^*))dt + \sigma dW_t.$$

The HJBQVI then becomes:

$$\begin{aligned} 0 = \min \left\{ -\partial_t u(t, q, S, \bar{S}, S^*, \psi, I) + g(q - q_t^*) - (\mu + f^{\text{long}}(S, \bar{S}, S^*))\partial_S u - \frac{S - \bar{S}}{t} \partial_{\bar{S}} u - \frac{1}{2} \sigma^2 \partial_{SS} u \right. \\ \left. - \sum_{\mathbf{k} \in \mathcal{K}} r_{(\psi, I), (\mathbf{k}^\psi, \mathbf{k}^I)} (u(t, q, S, \bar{S}, S^*, \mathbf{k}^\psi, \mathbf{k}^I) - u(t, q, S, \bar{S}, S^*, \psi, I)) \right. \\ \left. - \sup_{p \in Q_\psi, \ell \in \mathcal{A}} \sum_{n=1}^N \lambda^n(\psi, I, p^n, \ell) \mathbb{E} \left[\epsilon^n \ell^n \left(S + \frac{\psi^n}{2} + p^n \delta^n \right) + u(t, q - \ell^n \epsilon^n, S, \bar{S}, S^*, \psi, I) \right. \right. \\ \left. \left. - u(t, q, S, \bar{S}, S^*, \psi, I) \right]; \sum_{n=1}^N u(t, q, S, \bar{S}, S^*, \psi, I) - \sup_{m^n \in [0, \bar{m}]} m^n \left(S - \frac{\psi^n}{2} \right) + u(t, q - m^n, S, \bar{S}, S^*, \psi, I) \right\}, \end{aligned}$$

for $S \leq S^*$, with $\partial_S u = 0$ for $S = S^*$. To obtain this equation we just use the change of variable

$$v(t, x, q, S, \bar{S}, S^*, \psi, I) = x + u(t, q, S, \bar{S}, S^*, \psi, I),$$

linear with respect to the cash process X_t . We end up with a $2N + 4$ dimensional HJBQVI, that we can still solve using our deep reinforcement learning algorithm (but unlikely with finite differences).

More generally, adding a path-dependent state variable that gives information on the price trend adds one dimension to the HJBQVI (in the example above (\bar{S}_t, S_t^*, S_t) add one dimension each).

9.4.2 Extension 2: market impact

So far we assumed no market impact on the price process. It is common knowledge that cost of market impact can cut down a large proportion of the trading strategy's profit. Therefore, we can use a simple permanent-temporary market impact model, inspired by Almgren, Thum, Hauptmann, and Li [14].

The impacted mid-price process can be modeled as follows:

$$dS_t = (\mu + h(\ell_t))dt + \sigma dW_t + \sum_{n=1}^N (\xi^{n,l}(t, \ell_t^n) dN_t^n + \xi^{n,m}(t, m_t^n) dJ_t^n),$$

where the functions $h, \xi^{n,l}, \xi^{n,m}$ are the permanent and temporary market impact functions. Following Gatheral [123], we assume linear permanent market impact, that is

$$h(\ell_t) = \sum_{n=1}^N \kappa^{n,\text{per}} \ell_t^n, \quad \kappa^{n,\text{per}} > 0 \text{ for all } n \in \{1, \dots, N\}.$$

For the temporary market impact, we can follow the well-known ‘‘square-root law’’ and set

$$\xi^{n,l}(t, \ell_t^n) = \kappa^{n,l} (\ell_t^n)^{\gamma^{n,l}}, \quad \xi^{n,m}(t, m_t^n) = \kappa^{n,m} (m_t^n)^{\gamma^{n,m}},$$

where $\kappa^{n,l}, \kappa^{n,m}, \gamma^{n,l}, \gamma^{n,m} > 0$ and $\gamma^{n,l}, \gamma^{n,m} \approx 1/2$. On the other hand, in order to take into account the transient part of the impact, we can set the following form for S_t :

$$S_t = S_0 + \int_0^t \mu + h(\ell_s) ds + \sigma W_t + \sum_{n=1}^N \int_0^t \xi^{n,l}(t-s) \tilde{\xi}^{n,l}(\ell_s^n) dN_s^n + \xi^{n,m}(t-s) \tilde{\xi}^{n,m}(m_s^n) dJ_s^n, \quad (9.6)$$

where $\xi^{n,l}, \xi^{n,m}$ are decreasing kernels, and $\tilde{\xi}^{n,l}, \tilde{\xi}^{n,m}$ are decreasing functions of the posted volume. It is well known that by taking an exponentially decreasing kernel, Equation (9.6) admits a Markovian representation as the couples $(N_t^n, \int_0^t \xi^{n,\{l,m\}}(t-s) dN_s^n)_{t \in [0, T]}$ are Markovian. Practically, this will add $2N$ dimension to the HJBQVI.

Functions $h, \xi^{n,l}, \xi^{n,m}$ could also be approximated by neural networks. Determination of a cross-impact function between liquidity pools can lead to possible arbitrage detection across liquidity venues.

9.4.3 Extension 3: hidden liquidity

Hidden liquidity represents a great proportion of the liquidity especially in the US markets, see for example Jain and Jain [173]. Therefore, if one wants to design trading tactics for assets cross-listed in a European and an American market, taking into account the hidden part of the liquidity is crucial.

Assume that the n -th venue is a US liquidity pool. Borrowing the notations of Avellaneda, Reed, and Stoikov [27], we denote by H^n the hidden liquidity of the n -th venue at the first limit of the order book. Therefore, the corresponding imbalance process represented by the continuous-time Markov chain I^n can be rewritten as $\frac{N_t^{n,a,m} - N_t^{n,b,m}}{N_t^{n,a,m} + N_t^{n,b,m} + 2H^n}$, where $N_t^{n,b,m}, N_t^{n,a,m}$ are the bid and ask market order flow processes on the n -th venue. Empirical estimation of the prior parameters for the transition matrix of I^n have to take into account this additional term in the imbalance processes. Furthermore, incorporating the imbalance process with hidden liquidity into trading signals allows to detect arbitrage opportunities between different venues. This does not increase the dimensionality of Equation (9.3).

9.5 Numerical results

9.5.1 Global parameters

We take the example of a trader acting on a stock cross-listed on 2 different venues ($N = 2$), with the following global parameters:

- $\bar{\psi}^n = \{\delta, 2\delta\}$: the processes $(\psi_t^n)_{t \in [0, T]}$ can take two values, which correspond to a low or high spread regimes, and the tick size is $\delta = 0.05$.

- $\bar{I}^n = \{-0.5, 0, 0.5\}$: the processes $(I_t^n)_{t \in [0, T]}$ can take three values, which correspond to a negative, neutral or positive imbalance regime.
- $R = 2, (\omega^1, \omega^2) = (0.5, 1)$: the processes $(\epsilon_t^n)_{t \in [0, T]}$ can take two values, which correspond to a total or half-execution of the posted volume $(\ell_t^n)_{t \in [0, T]}$.
- $q_0 = 5 \times 10^4$: initial inventory of the trader.
- $\mathcal{T}_v = [v, v + \Delta_v]$, where $\Delta_v = 1$ min, which means that each slice lasts one minute, with $\mathcal{V} = 10$ slices and $T = 10$ min.
- $\Delta_t = 0.1$: we take 10 time steps in each slice, that is the agent takes 10 trading decisions during each slice.

The pre-computed trading curve is borrowed from an implementation shortfall execution using market orders, that is:

$$q_t^* = q_0 \frac{\sinh\left(\sqrt{\frac{\gamma\sigma^2 V}{2\eta}}(T-t)\right)}{\sinh\left(\sqrt{\frac{\gamma\sigma^2 V}{2\eta}}T\right)},$$

with the following set of parameters

- $\eta = 0.1$: coefficient of quadratic costs.
- $V = 1 \times 10^8$: average market volume.
- $\gamma = 1 \times 10^{-6}$: risk aversion of the trader using a CARA utility function.
- $\sigma = 0.05$: volatility of the asset.
- $f^\lambda(\ell_t) = \exp\left(-\kappa \sum_{n=1}^N \ell_t^n\right)$ with $\kappa = 2.5 \times 10^{-5}$: sensitivity of the execution with respect to the total volume posted.
- $f^\rho(\ell_t) = 1$: no sensitivity of the executed proportion with respect to the total volume posted.

For the sake of clarity, in the numerical experiment we consider the trader sending only limit orders.

9.5.2 Numerical methods

9.5.2.1 Finite differences

To find optimal strategy for limit orders we consider the following equation:

$$\begin{aligned} 0 = & -\partial_t v(t, q, \psi, I) + g(q - q_t^*) - \mu q \\ & - \sum_{n=1}^N \sum_{j=1}^J r_{\psi, j\delta}^{n, \psi} (v(t, q, \psi_{j\delta}^{-n}, I) - v(t, q, \psi, I)) - \sum_{n=1}^N \sum_{k=1}^K r_{I, I_k}^{n, I} (v(t, q, \psi, I_{I_k}^{-n}) - v(t, q, \psi, I)) \\ & - \sup_{p \in Q_\psi, \ell \in \mathcal{A}} \sum_{n=1}^N \lambda^n(\psi, I, p^n, \ell) \mathbb{E} \left[\epsilon^n \ell^n \left(\frac{\psi^n}{2} + p^n \delta^n \right) + v(t, q - \ell^n \epsilon^n, \psi, I) - v(t, q, \psi, I) \right], \end{aligned}$$

where

$$\psi_{j\delta}^{-n} = (\psi^1, \dots, \psi^{n-1}, j\delta, \psi^{n+1}, \dots), \quad I_{I_k}^{-n} = (I^1, \dots, I^{n-1}, I_k, I^{n+1}, \dots).$$

In order to apply the finite difference method we introduce the discretization of time and state space. For inventories we have $\mathfrak{Q} = \{q_1 = 0 < \dots < q_{\#\mathfrak{Q}} = q_0\}$. Time discretization in the slice is defined by $\mathfrak{T} = \{t_0 = 0 < t_1 = t_0 + \Delta_t < \dots < t_{\#\mathfrak{T}} = \Delta_v\}$. We also discretize the order volumes the trader can send $\mathfrak{L} = \{l_1 = 0 < \dots < l_{\#\mathfrak{L}} = q_0\}$.

Using the first difference for the value function derivative with respect to time we can rewrite the above equation as $\forall i \in \{0, \dots, \#\mathfrak{T} - 1\}, \forall q \in \mathfrak{Q}, \forall (\psi, I) \in \mathcal{M}$

$$\begin{aligned} v(t_{i+1}, q, \psi, I) &= v(t_i, q, \psi, I) - \Delta_t \left(g(q - q_t^*) - \mu q \right. \\ &\quad - \sum_{n=1}^N \sum_{j=1}^J r_{\psi, j\delta}^{n, \psi} (v(t, q, \psi_{j\delta}^{-n}, I) - v(t, q, \psi, I)) - \sum_{n=1}^N \sum_{k=1}^K r_{I, I_k}^{n, I} (v(t, q, \psi, I_k^{-n}) - v(t, q, \psi, I)) \\ &\quad \left. - \sup_{p \in \{-1, 0, 1\}^N, \ell \in \mathfrak{L}^N} \sum_{n=1}^N \lambda^n(\psi, I, p^n, \ell) \mathbb{E} \left[\epsilon^n \ell^n \left(\frac{\psi^n}{2} + p^n \delta^n \right) + v(t, q - \ell^n \epsilon^n, \psi, I) - v(t, q, \psi, I) \right] \right), \end{aligned}$$

with terminal condition $v(T, q, \psi, I) = 0$.

In terms of calculations the most demanding part is obviously the search of the supremum which is needed to be performed on the dimension $3^N \times \#\mathfrak{Q} \times \#\mathfrak{L}^N \times \#\mathcal{M}$ for each time step. From what follows that finite differences can be applied to solve the problem of optimal orders posting for the stock cross-listed in $N = 2$ venues with reasonable precision and calculation time. However, if we introduce more venues finite differences are not going to be any more efficient because the complexity is growing exponentially.

For our numerical example, we used the discretization with $\#\mathfrak{Q} = 101$ and $\#\mathfrak{L} = 51$ which assures the calculation time (on a simple PC) around 1min for the whole slice.

9.5.2.2 Neural networks

In this section, we briefly introduce the method using neural networks to solve HJB equations. In this chapter, we used a method which can be referred to as Actor-Critic method to approximate optimal controls and corresponding value function for the problem. Applications of this approach have shown to be fruitful, especially when we talk about equations in high dimension, more elaborate description of the method can be found for example in Pham, Warin, and Germain [230], Bachouch, Huré, Langrené, and Pham [29], Guéant and Manziuk [140], Huré, Pham, Bachouch, and Langrené [167], Germain, Pham, and Warin [130, 129] and Chapter 1 of this thesis.

The core of this approach is to represent the strategy of the trader with a neural network as well as the corresponding value function. Then one needs to formalize the target functions for both neural networks and to perform the gradient descent on the parameters (weights) of these networks. This procedure needs to be done for every time step, and so one ends up with $2\#\mathfrak{T}$ networks.

Let us start from the description of the value function approximation. We consider the neural networks taking as an input the spreads and the imbalances in the venues of interest and the inventory of the trader giving as an output the value function at this point. As in the finite difference method we solve our problem backward, starting from $t_{\#\mathfrak{T}-1} = \Delta_v - \Delta_t$, because the value function at the end of the slice is known from the terminal condition. To calculate the value function at time $t_i, \forall i \in \{0, \dots, \#\mathfrak{T} - 1\}$ we use the minimization of the mean-squared error between values given by the neural network and the target values calculated with the use of the value function approximation for time t_{i+1} and the network for the controls at the current step. Let us assume that we have the controls ℓ^*, p^* (obtained via neural networks, for example) for time t_i , then the target for the value function can be found as

$$\begin{aligned} v^{\text{target}}(t_{i-1}, q, \psi, I) &= v[\theta_i^v](t_i, q, \psi, I) + \Delta_t \left(g(q - q_t^*) - \mu q \right. \\ &\quad \left. - \sum_{n=1}^N \sum_{j=1}^J r_{\psi, j\delta}^{n, \psi} (v[\theta_i^v](t, q, \psi_{j\delta}^{-n}, I) - v[\theta_i^v](t, q, \psi, I)) \right) \end{aligned}$$

$$\begin{aligned}
 & - \sum_{n=1}^N \sum_{k=1}^K r_{I, I_k}^{n, I} (v[\theta_i^v](t, q, \psi, I_{I_k}^{-n}) - v[\theta_i^v](t, q, \psi, I)) \\
 & - \sum_{n=1}^N \lambda^n(\psi, I, p^{*n}, \ell^*) \mathbb{E} \left[\epsilon^n \ell^{*n} \left(\frac{\psi^n}{2} + p^{*n} \delta^n \right) + v[\theta_i^v](t, q - \ell^{*n} \epsilon^n, \psi, I) - v[\theta_i^v](t, q, \psi, I) \right],
 \end{aligned}$$

with $v[\theta_{\#I}^v](t_{\#I}, q, \psi, I) = 0$ and where $[\theta_i^v]$ stands for the weights of the neural network for the value function at time t_i .

The trader’s inventory is of continuous nature, however, spread and imbalance are categorical, so we need to verify if we should use some special techniques to ensure better fitting in this case.

Let us see first in Figure 9.1 the example of the target value function of the trader for $q \in [0, q_0]$ at different spreads and imbalances. We see considerable changes in the value function level depending on the market state which we would like to capture by our approximation.




Figure 9.1: Target value function for increasing inventory and random market states.

Now, let us compare the fitting of the value function parametrization taking as inputs raw spread and imbalance values with the parametrization working with encoded values of the spread and the imbalance. Here we are going to use the so-called one-hot encoding for categorical variables, which consists in the representation of different values of the variable by a one-hot vector $e_{\psi}^i \in \{0, 1\}^{\#\Psi}$ for the spread and $e_I^i \in \{0, 1\}^{\#\mathcal{I}}$ for the imbalance. And e^i (both for e_{ψ}^i and e_I^i) are such that that $e_j^i = 0, \forall j \neq i$, and $e_i^i = 1$ otherwise.




Figure 9.2: Comparison of the target value with approximation continuous in spread and imbalance.




Figure 9.3: Comparison of the target value with approximation discrete in spread and imbalance.

In Figures 9.2 and 9.3 we see the comparison between values predicted by two parametrizations with target values for the same number of learning epochs. There is a considerable gain in precision when the parametrization takes into account the categorical nature of market states. Therefore we apply it for both value function network approximation and the strategy neural network approximation.

Now, let us describe the learning procedure for the strategy. First of all, the inputs of the strategy network are the same as for the value function network, *i.e.* the trader's inventory, spreads and imbalances for both venues. As an output, we need to have volumes of the orders and limits on which the trader needs to send his orders. Volumes to send to each venue are bounded by the current inventory because we do not want the trader to execute more shares than he possesses. Limits should equal -1 , 0 , or 1 , but as soon as we want to use the tools of automatic differentiation, we need to represent them by differentiable function. The softmax activation function serves well to this purpose, so we represent the limits for each venue by the probabilities to send an order to each precise limit. In practice, the trader can choose the maximum of the three to perform his action.

The optimization criterium used for the strategy neural network is the function under supremum from the HJBQVI (9.4), with limit probabilities taken into account (let us denote them by $\mathbb{P}(p = a)$, for $a \in \{-1, 0, 1\}$) we need to maximize with respect to $\theta_i^\ell, i \in \{0, \dots, \#\mathfrak{T} - 1\}$:

$$\sum_{n=1}^N \sum_{a \in \{-1, 0, 1\}} \mathbb{P}[\theta_i^\ell](p^n = a) \lambda^n(\psi, I, a, \ell[\theta_i^\ell]) \mathbb{E} \left[\epsilon^n \ell[\theta_i^\ell]^n \left(\frac{\psi^n}{2} + a \delta^n \right) + v[\theta_{i+1}^\ell](t_i, q - \ell[\theta_i^\ell]^n \epsilon^n, \psi, I) - v[\theta_{i+1}^\ell](t_i, q, \psi, I) \right],$$

where θ_i^ℓ stand for the weights of the neural network of controls at time t_i . So we want to maximize this function for all possible values of market states and inventories. To avoid the dimensionality trap we need to optimize this function on some subset of possible values, which we are going to draw randomly.

When optimizing neural networks approximations, it is important to normalize the data, to have if possible a universal set of hyperparameters. First of all, the inventory entering as an input of the value function neural network and of the strategy neural network is normalized by q_0 to always stay in $[0, 1]$. Also, we are going to learn not the target value function itself, but the target value function normalized by q_0 , which sufficiently reduces the order of values. For strategy network, we are going to learn the proportion of the inventory to be sent and not the volume itself. And finally, we can notice that for high inventories the difference between value functions (which are quadratic in the inventory) in the supremum can become much more important than the profit of the trader coming from the tick (which is not more than linear in inventory). This fact can hinder us from finding optimal values for the limit to which the trader should send his order, especially for small inventories. We normalize the values of the optimized function for different inventories to make small inventories more important by multiplying all values by $\frac{1}{q}$. However, this latter normalization is used when we optimize over the part of the strategy responsible for the limits only, leaving volume updates untouched.

To summarize in Figures 9.4 and 9.5, we presented the structures of the neural networks used to represent the approximators for the strategy and the value function. Another feature worth mentioning here is the separation of market state and inventory inputs for some layers, both for the strategy and the value function. This allows capturing features of the market state independently of the inventory. Also, we separated some layers preceding the outputs of the strategy network to be able to perform the learning process with different learning rates for volumes and limits of limit orders.

While the finite difference schemes must complete the entire recalculation of values for the whole grid every time the trader wants to adapt his strategy using the updated market parameters, neural networks can be adapted progressively, starting from some pre-trained strategy, for example, the one corresponding to the previous parameters. In practice, a pre-trained model can be reused for different problem settings due to normalization. Therefore a long and elaborate training procedure should be done only once. The resulting model can be ameliorated by small adjustment trainings which take only 1 minute on the simplest instance of the AWS platform (2CPU, no GPU), and have great speed-up potential when performed on more complex infrastructures.

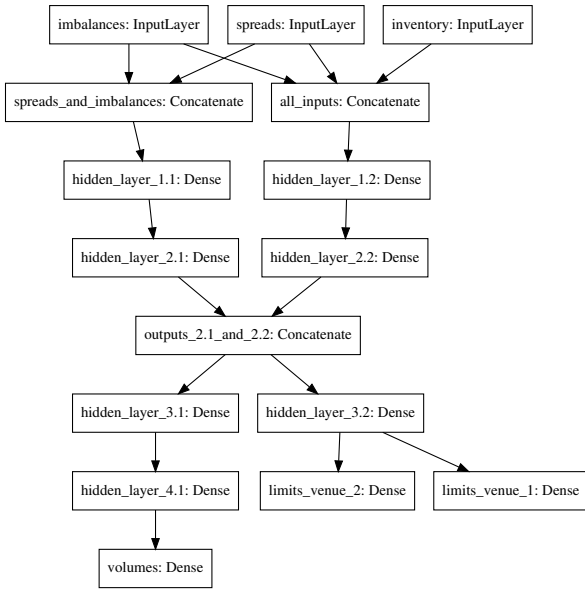


Figure 9.4: Neural network structure for the trader’s strategy.

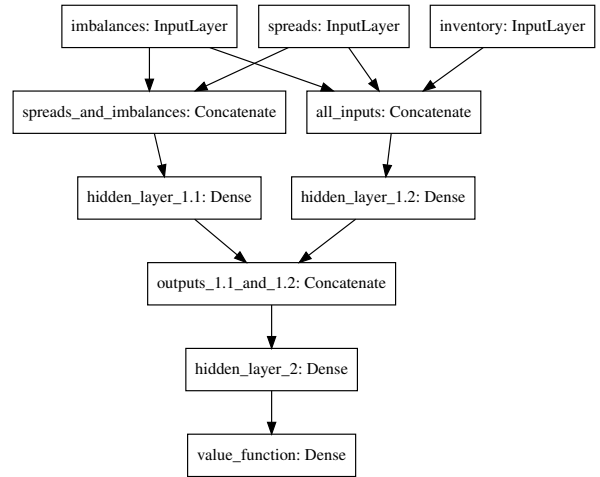


Figure 9.5: Neural network structure for the value function.

9.5.3 Two identical venues

We assume that the trader is confident about his estimation of σ . Therefore he uses Bayesian update only on the drift μ of the asset. The venues share identical parameters, which will be inferred by the trader through time.

9.5.3.1 Value function

We first plot in Figures 9.6 and 9.7 the evolution through time of the value function of the trader in the state $\psi^1 = \psi^2 = 1$ and $I^1 = I^2 = 0$ during a slice of execution, obtained through finite difference method.

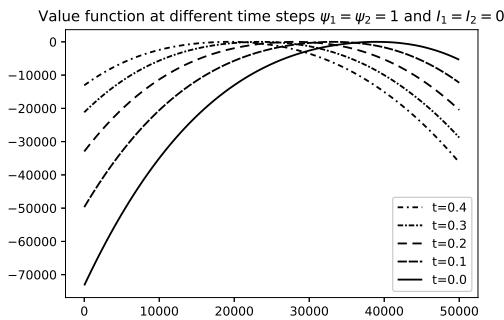


Figure 9.6: Value function with respect to the inventory between $t = 0$ and $t = 0.4$.




Figure 9.7: Evolution of the value function v between $t = 0.5$ and $t = 0.9$.




Figure 9.8: Evolution of the value function v between $t = 0$ and $t = 0.4$ using neural networks.




Figure 9.9: Evolution of the value function v between $t = 0.5$ and $t = 0.9$ using neural networks.

The parabolic form of the value function comes from the term $g(q - q_t^*)$ in (9.3). The maximum value indicates the optimal inventory for the next step in the slice. When t increases, the maximum shifts toward zero, which means that the trader wants to finish the execution at the end of the slice.

We plot in Figures 9.8 and 9.9 the value function of (9.3) obtained using neural networks. We can see that the neural networks approximate accurately the value function.

Next, we plot the strategy (in terms of limits and volumes) of the trader in both venues, using finite difference schemes.

9.5.3.2 Strategy: limit orders and volumes with finite difference schemes

In Figures 9.10 and 9.11, we plot the limits at which the trader posts his limit orders in the two venues, given equal spread and imbalance processes.




Figure 9.10: Limit strategy in the first venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$.




Figure 9.11: Limit strategy in the second venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$.

As the trader has the same prior distribution in the two venues, his strategy is the same in both venues. At the beginning of the slice, *i.e.* at $t = 0$, the maximum of the value function is near $q = 32000$. Therefore, if the trader has a lower inventory, he does not post any orders and wait for the next time step. If he has a higher inventory, he tries to reach $q = 32000$ inventory. For $q \in [32000, 34000]$, being sufficiently close to the next step optimal inventory, he posts limit orders on the second best limit to collect an additional tick. For $q \in [34000, 40000]$, he posts at the first best limit to increase his probability of execution. If he has $q > 40000$, he creates a new best limit and accepts to lose one tick in order to be executed faster and reach the optimal inventory at the following time step. We can see in this behavior the trade-off between the possibility of being executed at a more favorable price and the necessity to complete the execution.

For the sake of homogeneity (for all \mathcal{M} market states, the trader faces similar trade-off), we considered the same set of controls for the limit where the trader can send his order. For this reason, we can see that even for the spread equal to δ the trader can submit an order to the limit $p = -1$, which in practice can obviously be treated as $p = 0$ due to piecewise monotonous nature of the optimal limit strategy (which is, in fact, monotonous, though it cannot be reflected by finite differences when the optimal volume equals to 0).

When the trader is near the end of the slice, he starts posting limit orders earlier (can be seen if both volumes and limits are considered). For example if $t = 0.6$, he begins to trade at the second best limit when $q \in [8000, 11000]$, at the first best limit when $q \in [11000, 19000]$, and creates a new best limit when $q \in [19000, 50000]$. Therefore, if the trader still has a very positive inventory at the end of the slice, he prefers to sacrifice one tick at the first best limit in order to complete his execution at this step.

It is important to highlight the fact that, when $t = 0.9$, the trader does not rush to liquidate his inventory completely. This comes from the absence of a terminal penalty, often used in optimal liquidation problem to guarantee the complete execution of the inventory. It enables in some sense to “relax” the optimal execution framework on a slice, as the part of the inventory that has not been executed during one slice is split between the remaining ones.

We plot in Figures 9.12 and 9.13 the volumes posted in both venues, for the same spread and imbalance. We see that, at the beginning of the slice, the trader begins to post a nonzero volume only when $q > 32000$. Moreover, he posts a higher volume when he is near the end of the slice.

Volumes at different time steps when $\psi_1 = \psi_2 = 1$ and $I_1 = I_2 = 0$ venue :




Figure 9.12: Volume sent to the first venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$.

Volumes at different time steps when $\psi_1 = \psi_2 = 1$ and $I_1 = I_2 = 0$ venue :




Figure 9.13: Volume sent to the second venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$.

When the second venue has a higher spread, we plot the strategy of the trader in both venues in Figures 9.14 and 9.15.

Limits at different time steps when $\psi_1 = 1, \psi_2 = 2$ and $I_1 = I_2 = 0$




Figure 9.14: Limit strategy, $\psi^1 = \delta, \psi^2 = 2\delta, I^1 = I^2 = 0$.

Volumes at different time steps when $\psi_1 = 1, \psi_2 = 2$ and $I_1 = I_2 = 0$




Figure 9.15: Volume strategy, $\psi^1 = \delta, \psi^2 = 2\delta, I^1 = I^2 = 0$.

For $t = 0.5$, we see in Figure 9.14 that the trader starts to post at the second best limit in the second venue when $q = 10000$ and in the first when $q = 11000$. For $q \in [18000, 21000]$, he creates a new best limit in the second venue to execute his inventory faster but keeps posting at the best limit in the first venue in order to collect a higher spread. Finally, for $q \in [30000, 50000]$, he stops posting in the second venue in order to consume more liquidity in the first one where the probability of getting his order filled is higher. Similar interpretations apply for $t = 0.1$.

In Figure 9.15, we see that the trader posts a higher volume in the first venue compared to the second one. For $t = 0.5$, he starts to trade at $q = 10000$ for the second venue and at $q = 11000$ for the first one. The volume posted in the first venue increases almost linearly with respect to the inventory. In contrast, the volume posted in the second venue increases until an inventory of $q = 22000$, then stays constant until $q = 30000$ and decreases to zero afterward. This means that for $q \in [10000, 30000]$, the trader prefers to collect the spread from both venues. When $q > 30000$, he prefers to stop posting in the second venue, the one with a higher spread, in order to maximize his chances of being executed in the first one. Similar interpretations apply for $t = 0.1$.

In Figures 9.16 and 9.17, we show the choice of limits and volumes of the trader if the imbalance is more favorable in the second venue. In Figure 9.16, we observe for $t = 0.5$ that the trader posts in the first venue at the second best limit for $q \in [12000, 15000]$, at the first limit for $q \in [15000, 20000]$ and at a new best limit for $q \in [20000, 32000]$. At the same time, he posts in the first limit of the second venue when $q \in [12000, 22000]$ and at a new best limit for $q \in [22000, 50000]$. We see that the trader prefers to post at a higher limit in the second venue because of the higher probability of execution due to a more favorable imbalance. For large inventories, he stops posting in the first venue in order to increase his probability of execution using limit orders in the second venue at a new best limit. Same results hold for $t = 0.1$.

In Figure 9.17, we see that the trader posts a majority of his volume in the second venue due to a more favorable imbalance. When his inventory is not too high, he collects the spread from both venues. However, when his inventory is relatively high, he sends all the volume to the first venue in order to increase the probability of filling.




Figure 9.16: Limit order strategy, $\psi^1 = \psi^2 = \delta$, $I^1 = -0.5, I^2 = 0.5$.




Figure 9.17: Volume strategy, $\psi^1 = \psi^2 = \delta$, $I^1 = -0.5, I^2 = 0.5$.

We now describe the strategies on the limits and the volumes obtained by a reinforcement learning approach.

9.5.3.3 Strategy: limit orders and volumes with neural networks

We plot in Figures 9.18 and 9.19 the strategies on the limits used by the trader. As soon as limits are represented by probabilities to send an order to each precise limit, for graphical representation, we plot the limit corresponding to the highest of the three probabilities. We see that the choice of the limits is in line with the ones of Figures 9.10 and 9.11 up to states where optimal order volume is at 0 (in this case limit values are indistinguishable for finite differences). When the trader is at the beginning of the slice, for a small inventory, he prefers to collect a higher spread by being executed at the second best limit. When he is near the end of the slice, he prefers to be filled at a less favorable price, at the best or new best limit, in order to lower his execution risk. We can also see that neural networks preserve the monotonicity of the optimal limit function.




Figure 9.18: Limit order strategy in the first venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$ using neural networks.




Figure 9.19: Limit order strategy in the second venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$ using neural networks.

In Figures 9.20 and 9.21, we plot the posted volumes of the trader in both venues for the same spread and imbalance. We see that the strategy is a smoothed approximation of the one obtained using finite differences in Figures 9.12 and 9.13. We see that at the very beginning of the slice, the trader is not going to trade if his inventory is already small enough. The strategy in both venues is the same up to some negligible numerical effects.

Volumes at different time steps when $\psi_1 = \psi_2 = 1$ and $I_1 = I_2 = 0$ venue :




Figure 9.20: Volume posted in the first venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$ using neural networks.

Volumes at different time steps when $\psi_1 = \psi_2 = 1$ and $I_1 = I_2 = 0$ venue :




Figure 9.21: Volume posted in the second venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$ using neural networks.

If the spread of the second venue is higher, we see in Figure 9.22 that the strategy with the limits is the same as in Figure 9.14. It is interesting to note in Figure 9.23 that the trader does not stop posting in the second venue, as in Figure 9.15, again because of the approximation coming from neural networks. However, this behavior enables to perform some exploration of the venue parameters. For example, if the trader follows the strategy given by finite differences in Figure 9.15, he posts a volume equal to 0 in the second venue when $q > 32000$ for $t = 0.5$. However, if the trader underestimates the prior on the filling probability in the second venue $\hat{\lambda}^2$, he will keep sending orders in the first venue, neglecting the possibility of splitting his orders which can potentially improve his execution. Moreover, Figures 9.8 and 9.9 show that this slight difference in the obtained controls does not change drastically the performance of the trader in terms of the value function.

Limits at different time steps when $\psi_1 = 1, \psi_2 = 2$ and $I_1 = I_2 = 0$




Figure 9.22: Limit order strategy, $\psi^1 = \delta, \psi^2 = 2\delta, I^1 = I^2 = 0$ using neural networks.

Volumes at different time steps when $\psi_1 = 1, \psi_2 = 2$ and $I_1 = I_2 = 0$




Figure 9.23: Volume strategy, $\psi^1 = \delta, \psi^2 = 2\delta, I^1 = I^2 = 0$ using neural networks.

The same comments apply to Figures 9.24 and 9.25, where we see that the trader posts a small but nonzero volume in the first venue with a less favorable imbalance which potentially allows to perform exploration in this venue and faster improve parameter estimations.

Limits at different time steps when $\psi_1 = \psi_2 = 1$ and $I_1 = -0.5, I_2 = 0.5$




Figure 9.24: Limit order strategy, $\psi^1 = \psi^2 = \delta, I^1 = -0.5, I^2 = 0.5$ using neural networks.

Volumes at different time steps when $\psi_1 = \psi_2 = 1$ and $I_1 = 0.5, I_2 = -0.5$




Figure 9.25: Volume strategy, $\psi^1 = \psi^2 = \delta, I^1 = -0.5, I^2 = 0.5$ using neural networks.

9.5.4 Two different venues

In this section, we analyze the behavior of the trader believing that the first venue is better than the second venue in terms of filling rate. We compare the solutions obtained via finite difference schemes and neural networks.

9.5.4.1 Value function

We show in Figures 9.26 and 9.27 the evolution of the value function of the trader during a slice of execution, obtained through the finite difference method.




Figure 9.26: Evolution of the value function v between $t = 0$ and $t = 0.4$.




Figure 9.27: Evolution of the value function v between $t = 0.5$ and $t = 0.9$.

One can see that the value function deteriorates compared to the previous example, which is predictable in view of the fact that one of the venues is exactly like in the above example, and another one is worse in terms of filling ratio. For example in Figure 9.27, the minimum of the function v at $t = 0.5$ when $q = 50000$ is -49000 compared to a minimum of -45000 in the example above. This is a natural consequence of a worse prior distribution on the filling ratio of the second venue while keeping the prior on the first venue unchanged.




Figure 9.28: Evolution of the value function v between $t = 0$ and $t = 0.4$ using neural networks.




Figure 9.29: Evolution of the value function v between $t = 0.5$ and $t = 0.9$ using neural networks.

We check in Figures 9.28 and 9.29 that we obtain a similar shape for the value function using neural networks. We now describe the strategy of the trader on the limits and the posted volumes and compare it to the case of two identical venues.

9.5.4.2 Strategy: limit orders and volumes with finite difference schemes

In Figures 9.30 and 9.31, we show the limit order strategy of the trader in the two venues for the same spreads and imbalances. As the second venue is less favorable for execution, the trader prefers to create a new best limit for smaller inventories. For example, when $t = 0.6$, he posts an order on the new best limit starting from $q = 19000$, and in the second venue, he prefers to create a new limit starting from $q = 18000$. Generally, either at the beginning or at the end of the slice, the trader prefers to post at a lower limit in the second venue in order to increase his execution rate there, sacrificing the spread that could have been collected.




Figure 9.30: Limit order strategy in the first venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$.




Figure 9.31: Limit order strategy in the second venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$.

The strategy of the trader differs drastically in terms of order volumes. In Figures 9.32 and 9.33, we see that the trader posts the majority of his volume in the first venue. Especially when at $t = 0.9$ the trader stops posting in the second venue to reduce his liquidity consumption and maximize his probability of execution in the first venue.




Figure 9.32: Volume posted in the first venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$.




Figure 9.33: Volume posted in the second venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$.

In Figures 9.34 and 9.35, we see the limits and the volumes recommended to the trader when the second venue has a higher spread, and the imbalances are equal. The trader posts an even smaller volume in the second venue, compared to Figure 9.15. As the filling rate is lower in the second venue, the trader decreases his liquidity consumption in this venue, because of the smaller probability of collecting a higher spread.

The strategy on the limits in Figure 9.34 is also different from the one in Figure 9.14. When $t = 0.5$ and the two venues are the same, the trader posts at the second best limit in the second venue when $q \in [11000, 13000]$, then at the first best limit when $q \in [13000, 18000]$ and at a new best limit for $q \in [18000, 30000]$. When the venues are different, the trader posts at the second best limit in the second venue for $q \in [10000, 12000]$, at the first best limit for $q \in [12000, 17000]$ and at a new best limit when $q \in [17000, 19000]$. Therefore, when the second venue has a worse filling rate, the trader posts in the second venue earlier (for a higher inventory) and less compared to the case with two equivalent venues.




Figure 9.34: Limit order strategy, $\psi^1 = \delta, \psi^2 = 2\delta, I^1 = I^2 = 0$.




Figure 9.35: Volume strategy, $\psi^1 = \delta, \psi^2 = 2\delta, I^1 = I^2 = 0$.

If the imbalance is more favorable in the second venue, we see in Figures 9.36 and 9.37 that the strategy is very different from the one in Figures 9.16 and 9.17 where the two venues shared the same characteristics. As the second venue has a more favorable imbalance, the trader posts a higher volume in it. However, he posts a nonzero volume in the first venue, because of the overall better filling ratio. This contrasts with Figure 9.17 where at some sufficiently high inventories, the trader stops sending orders to the first venue. Due to the trade-off between an overall higher filling ratio in the first venue and a more favorable imbalance in the second venue, the trader splits his liquidity consumption between the two venues.

The strategy on the limits in Figure 9.36 also differs from the one with two identical venues in Figure 9.16. For $t = 0.5$ in Figure 9.16, the trader posts in the first venue at the second best limit for $q \in [10000, 13000]$, at the first best limit for $q \in [13000, 20000]$ and at a new best limit for $q \in [20000, 32000]$. In Figure 9.36, the trader posts in the first venue at the second best limit for $q \in [10000, 12000]$, at the first best limit for $q \in [12000, 18000]$ and at a new best limit for $q > 18000$. Therefore, he posts at a more favorable limit in terms of filling rate in the first venue in order to compensate for the unfavorable imbalance compared to the second venue.




Figure 9.36: Limit order strategy, $\psi^1 = \psi^2 = \delta$, $I^1 = -0.5, I^2 = 0.5$.




Figure 9.37: Volume strategy, $\psi^1 = \psi^2 = \delta$, $I^1 = -0.5, I^2 = 0.5$.

Before moving to the analysis of the effectiveness of the Bayesian update of market parameters, we conclude with a comparison of the strategies obtained via neural networks optimization.

9.5.4.3 Strategy: limit orders and volumes with neural networks

We observe in Figures 9.38 and 9.39 that the strategy of the trader on the limits is in line with the one in Figures 9.30 and 9.31 up to the states where the optimal volume of the order equals 0.




Figure 9.38: Limit order strategy in the first venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$ using neural networks.




Figure 9.39: Limit order strategy in the second venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$ using neural networks.

In Figures 9.40 and 9.41, we see that the strategy of the trader on the posted volumes is well approximated and smoothed by neural networks.

In Figures 9.42 and 9.43, we see in the case of a higher spread in the second venue that, because of neural network parametrization of the strategy, the trader posts a nonzero volume in the second venue leaving the possibility to better explore the filling ratios. Results are in line with the ones in Figures

9.34 and 9.35: the trader posts the majority of his volume in the first venue because of a lower spread and a more favorable filling ratio.

Volumes at different time steps when $\psi_1 = \psi_2 = 1$ and $I_1 = I_2 = 0$ venue :




Figure 9.40: Volume posted in the first venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$ using neural networks.

Volumes at different time steps when $\psi_1 = \psi_2 = 1$ and $I_1 = I_2 = 0$ venue :




Figure 9.41: Volume posted in the second venue, $\psi^1 = \psi^2 = \delta, I^1 = I^2 = 0$ using neural networks.

Limits at different time steps when $\psi_1 = 1, \psi_2 = 2$ and $I_1 = I_2 = 0$




Figure 9.42: Limit order strategy, $\psi^1 = \delta, \psi^2 = 2\delta, I^1 = I^2 = 0$ using neural networks.

Volumes at different time steps when $\psi_1 = 1, \psi_2 = 2$ and $I_1 = I_2 = 0$




Figure 9.43: Volume strategy, $\psi^1 = \delta, \psi^2 = 2\delta, I^1 = I^2 = 0$ using neural networks.

Finally, we show in Figures 9.44 and 9.45 a similar behavior compared to the finite difference schemes in Figures 9.36 and 9.37: the trader posts a higher volume in the second venue due to a more favorable imbalance, and keeps posting in the first venue due to an overall more favorable filling ratio.

Limits at different time steps when $\psi_1 = \psi_2 = 1$ and $I_1 = -0.5, I_2 = 0.5$




Figure 9.44: Limit order strategy, $\psi^1 = \psi^2 = \delta, I^1 = -0.5, I^2 = 0.5$ using neural networks.

Volumes at different time steps when $\psi_1 = \psi_2 = 1$ and $I_1 = 0.5, I_2 = -0.5$




Figure 9.45: Volume strategy, $\psi^1 = \psi^2 = \delta, I^1 = -0.5, I^2 = 0.5$ using neural networks.

9.5.5 Bayesian update

In this section, we analyze the effectiveness of the Bayesian update framework through several execution slices.

9.5.5.1 Market simulation on a slice

We first show an example of a market simulation of one slice and demonstrate the trading strategy through the slice, which are illustrated in Figure 9.46. At $t = 0.2$, the spread in both venues is equal to δ , with an unfavorable imbalance in both venues. In that case, as two venues share the same

characteristics, and the inventory is sufficiently close to the optimal for the next step, so the trader sends the same quantity to both venues, which is close to zero. When $t = 0.5$, the first venue has an unfavorable imbalance, and the second venue has a higher spread. In this configuration, the trader sends a higher volume in the first venue, in order to get a better filling rate due to a lower spread. Finally, at $t = 0.7$, the first venue has a higher spread and a more favorable imbalance compared to the second venue. This leads to a higher volume in the first venue at the second best limit and a lower volume in the second venue at first best limit. The favorable imbalance in the first venue indicates a higher probability of execution for an order at a higher limit, because the price may move in this direction. Therefore, even if the spread is equal to two ticks, the trader posts in this venue in order to be executed at a more favorable price. As the spread in the second venue is lower, but the imbalance is less favorable, he posts at the first best limit to benefit from the trade-off between execution and profit through collecting the spread.

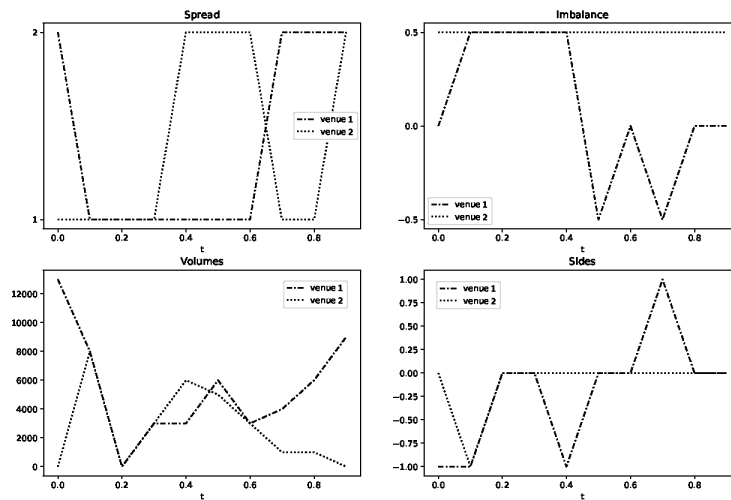


Figure 9.46: Market simulation: spreads (upper left), imbalances (upper right), volumes (lower left) and limits (lower right) in both venues.

The corresponding execution trajectory is shown in Figure 9.47, where we can see the typical Implementation Shortfall execution shape, coming from the pre-computed trading curve q^* .




Figure 9.47: Evolution of the inventory of the trader on a slice of execution.

9.5.5.2 Update of the execution proportion

We show how the trader updates the market parameters through observations and trading. The update of the execution proportion is quite fast, as it can be seen in Figures 9.48 and 9.49 the good estimation can be achieved after completing 1-2 slices. In this example we started from the correct prior for the second venue and the inaccurate one for the first:

$$\rho^1 = [0.1 \quad 0.9], \quad \rho^2 = [0.1 \quad 0.9].$$




Figure 9.48: Bayesian update of the executed proportion in the first venue.




Figure 9.49: Bayesian update of the executed proportion in the second venue.

9.5.5.3 Update of the imbalance and the spread dynamics

We plot the convergence of the estimated transition matrices $r^{1,\psi}, r^{2,\psi}$ in Figures 9.50 and 9.51. We observe fast convergence to a good approximation of the spread dynamics parameters with prior values being respectively:

$$r^{1,\psi} = \begin{bmatrix} -5 & 5 \\ 5 & -5 \end{bmatrix},$$

$$r^{2,\psi} = \begin{bmatrix} -5 & 5 \\ 5 & -5 \end{bmatrix}.$$

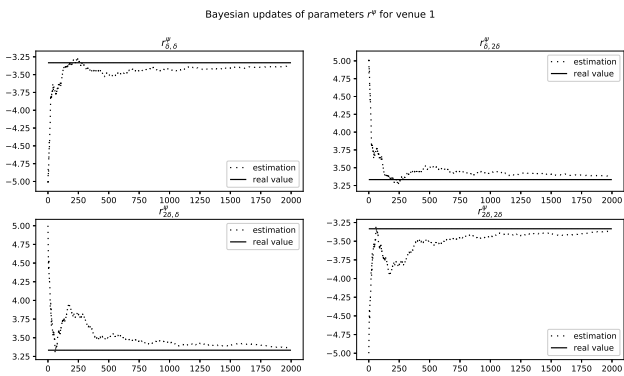


Figure 9.50: Bayesian update of the transition matrix $r^{1,\psi}$.

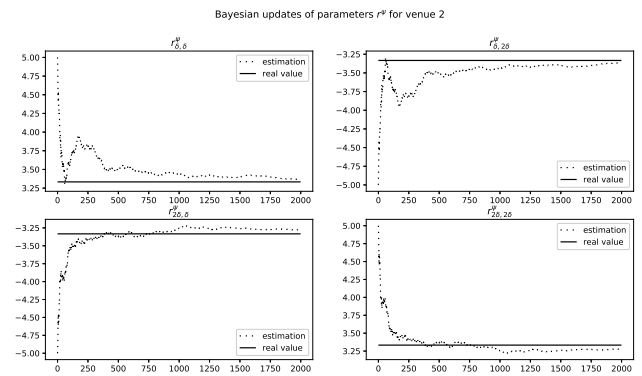


Figure 9.51: Bayesian update of the transition matrix $r^{2,\psi}$.




Figure 9.52: Bayesian update of the transition matrix $r^{1,I}$.

We perform the same study for the transition matrices $r^{1,I}, r^{2,I}$ of the imbalance processes through the observed one in Figures 9.52 and 9.53.

Starting from the following prior parameters:

$$r^{1,I} = \begin{bmatrix} -5 & 2.8 & 2.2 \\ 2.2 & -5 & 2.8 \\ 2.2 & 2.8 & -5 \end{bmatrix}, \quad r^{2,I} = \begin{bmatrix} -5 & 2.8 & 2.2 \\ 2.2 & -5 & 2.8 \\ 2.2 & 2.8 & -5 \end{bmatrix},$$

we see that we need just a couple of slices to have a quite good approximation and only a dozen of slices (less for more granular slices) to achieve the right estimation.

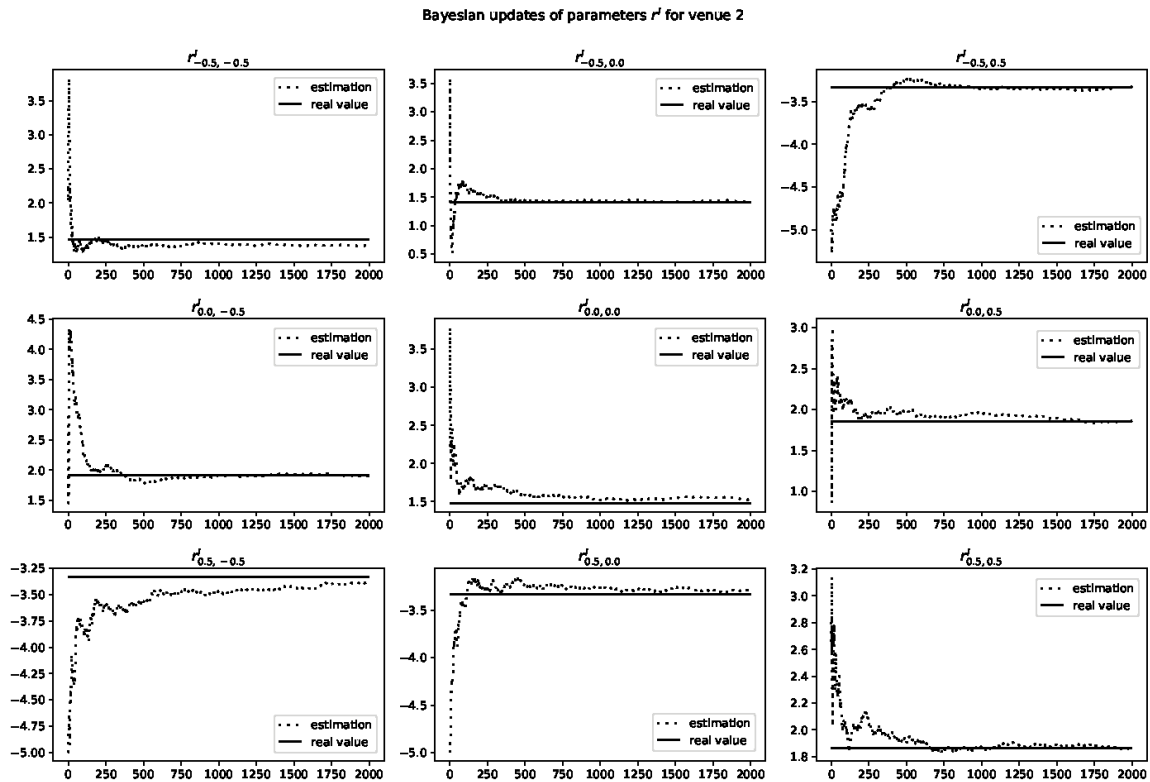


Figure 9.53: Bayesian update of the transition matrix $r^{2,I}$.

9.5.5.4 Update of the long term drift of the asset

As we observe the increments of the price process S_t continuously, it is easy to converge toward a real market drift, the example is in Figure 9.54, we find $\mu^{\text{true}} = -0.5$, starting from a prior of $\mu = 0.1$.




Figure 9.54: Bayesian update of the drift of the asset.

It took 20 slices to find a real value even if in the considered example we supposed to be sure in our prior estimation $\nu = 0.02$, which appeared to be incorrect.

9.5.5.5 Update of the intensity of limit orders

The hardest parameter to update quickly is obviously the intensity of filling which depends on states of both venues. In our numerical setting we have 32 possible states, so during one slice of 10 time steps we have no possibility to even visit all the states.

The results of convergence of the parameters λ can be found in Figures 9.55 and 9.56. We see that full convergence requires a lot of observations, however we should keep in mind that to have a strategy close to the optimal one we do not necessitate an excessive precision.

In this example, we started from the priors same for both venues, whereas the real parameters are different. The priors are:

$$\lambda_{\delta,\delta}^1 = \lambda_{\delta,\delta}^2 = \begin{bmatrix} 5.35 & 6.52 & 7.11 \\ 2.75 & 3.4 & 3.79 \\ 1.5 & 1.86 & 2.1 \end{bmatrix}, \quad \lambda_{\delta,2\delta}^1 = \lambda_{\delta,2\delta}^2 = \begin{bmatrix} 8.28 & 10.03 & 10.9 \\ 4.38 & 5.35 & 5.9 \\ 2.5 & 3.05 & 3.4 \end{bmatrix},$$

$$\lambda_{2\delta,\delta}^1 = \lambda_{2\delta,\delta}^2 = \begin{bmatrix} 1.81 & 2.27 & 2.5 \\ 0.78 & 1.04 & 1.19 \\ 0.29 & 0.43 & 0.53 \end{bmatrix}, \quad \lambda_{2\delta,2\delta}^1 = \lambda_{2\delta,2\delta}^2 = \begin{bmatrix} 2.96 & 3.65 & 4. \\ 1.42 & 1.81 & 2.04 \\ 0.68 & 0.9 & 1.04 \end{bmatrix}.$$




Figure 9.55: Bayesian update of the intensity of limit orders in the first venue.




Figure 9.56: Bayesian update of the intensity of limit orders in the second venue.

9.A Appendix

9.A.1 Proof of Theorem 9.3

It can be show with the dynamic programming principle that the HJBQVI (9.3) does not depend on the cash variable x . We set $(\tilde{q}, \tilde{\psi}, \tilde{I}) \in \mathcal{D} = \mathcal{Q} \times \mathcal{K}$ and $(t_i, S_i) \in [0, T) \times \mathbb{R}$ such that

$$t_i \xrightarrow{i \rightarrow +\infty} \hat{t}, \quad S_i \xrightarrow{i \rightarrow +\infty} \hat{S}, \quad v(t_i, \tilde{q}, S_i, \tilde{\psi}, \tilde{I}) \xrightarrow{i \rightarrow +\infty} v_*(\hat{t}, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I}).$$

We begin with $\hat{t} = T$. By taking $\ell^n = 0$ for all $n \in \{1, \dots, N\}$ we get

$$v(t_i, \tilde{q}, S_i, \tilde{\psi}, \tilde{I}) \geq \mathbb{E}_{t_i, \tilde{q}, S_i, \tilde{\psi}, \tilde{I}} \left[Q_T S_T - \int_0^T g(q_t^* - q_t) dt \right].$$

By dominated convergence, we get $v_*(T, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I}) \geq \tilde{q} \hat{S}$.

Assume now that $\hat{t} < T$ and that the minimum in the HJBQVI is given by the first term. We take $\phi : [0, T) \times \mathbb{R} \times \mathcal{D} \rightarrow \mathbb{R}$ be C^1 in time, C^2 in \hat{S} and such that $0 = \min_{[0, T) \times \mathbb{R} \times \mathcal{D}} (v_* - \phi) = (v_* - \phi)(\hat{t}, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I})$. If there exists $\eta > 0$ such that

$$2\eta < \partial_t \phi(\hat{t}, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I}) - g(q - q_t^*) + \mu \partial_S \phi + \frac{1}{2} \sigma^2 \partial_{SS} \phi + \sum_{\mathbf{k} \in \mathcal{K}} r(\tilde{\psi}, \tilde{I}, (\mathbf{k}^\psi, \mathbf{k}^I)) (\phi(\hat{t}, \tilde{q}, \hat{S}, \mathbf{k}^\psi, \mathbf{k}^I) - \phi(\hat{t}, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I}))$$

$$+ \sup_{p \in Q_\psi, \ell \in \mathcal{A}} \sum_{n=1}^N \lambda^n(\tilde{\psi}, \tilde{I}, p^n, \ell) \mathbb{E} \left[\epsilon^n \ell^n \left(\hat{S} + \frac{\tilde{\psi}^n}{2} + p^n \delta^n \right) + \phi(\hat{t}, \tilde{q} - \ell^n \epsilon^n, \hat{S}, \tilde{\psi}, \tilde{I}) - \phi(\hat{t}, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I}) \right],$$

we should have

$$0 \leq \partial_t \phi(t, \tilde{q}, S, \tilde{\psi}, \tilde{I}) - g(q - q_t^*) + \mu \partial_S \phi + \frac{1}{2} \partial_{SS} \phi + \sum_{\mathbf{k} \in \mathcal{K}} r_{(\tilde{\psi}, \tilde{I}), (\mathbf{k}^\psi, \mathbf{k}^I)}(\phi(t, \tilde{q}, S, \mathbf{k}^\psi, \mathbf{k}^I) - \phi(t, \tilde{q}, S, \tilde{\psi}, \tilde{I})) \\ + \sup_{p \in Q_\psi, \ell \in \mathcal{A}} \sum_{n=1}^N \lambda^n(\tilde{\psi}, \tilde{I}, p^n, \ell) \mathbb{E} \left[\epsilon^n \ell^n \left(S + \frac{\tilde{\psi}^n}{2} + p^n \delta^n \right) + \phi(t, \tilde{q} - \ell^n \epsilon^n, S, \tilde{\psi}, \tilde{I}) - \phi(t, \tilde{q}, S, \tilde{\psi}, \tilde{I}) \right],$$

for all $(t, S) \in B = \left((\hat{t} - r, \hat{t} + r) \cap [0, T] \right) \times \left(\hat{S} - r, \hat{S} + r \right)$ for a given $r \in (0, T - \hat{t})$. We can assume without loss of generality that B contains the sequences $(t_i, S_i)_i$ and, by taking η arbitrarily small

$$\phi(t, \tilde{q}, S, \tilde{\psi}, \tilde{I}) + \eta \leq v_*(t, \tilde{q}, S, \tilde{\psi}, \tilde{I}) \leq v(t, \tilde{q}, S, \tilde{\psi}, \tilde{I})$$

on the boundary of B , denoted by $\partial_p B$. Without loss of generality we can also assume that

$$\phi(t, q, S, \psi, I) + \eta \leq v_*(t, q, S, \psi, I) \leq v(t, q, S, \psi, I),$$

for $(t, q, S, \psi, I) \in \tilde{B}$ where

$$\tilde{B} = \left\{ (t, q, S, \psi, I) : (t, S) \in B, q \in \{\tilde{q} - \min_n \epsilon^n, \tilde{q}, \tilde{q} + \min_n \epsilon^n\}, \right. \\ \left. \psi \in \prod_{n=1}^N \{\tilde{\psi}^n - \delta^n, \tilde{\psi}^n, \tilde{\psi}^n + \delta^n\}, I \in \prod_{n=1}^N \{\tilde{I}_{-1}^n, \tilde{I}^n, \tilde{I}_{+1}^n\}, (q, \psi, I) \neq (\tilde{q}, \tilde{\psi}, \tilde{I}) \right\}.$$

We introduce the set

$$B_{\mathcal{D}} = \left\{ (t, \tilde{q}, S, \tilde{\psi}, \tilde{I}) : (t, S) \in B \right\},$$

and denote by τ_i the first exit time of $(t, q_t, S_t, \psi_t, I_t)_{t \geq t_i}$ from $B_{\mathcal{D}}$, with $q_{t_i} = \tilde{q}, S_{t_i} = \hat{S}, \psi_{t_i} = \tilde{\psi}, I_{t_i} = \tilde{I}$, and the processes are controlled by the optimal controls $(\ell^n, p^n)_{n \in \{1, \dots, N\}} \in \mathcal{A} \times Q_\psi$. By Itô's formula, we get

$$\phi(\tau_i, q_{\tau_i}, S_{\tau_i}, \psi_{\tau_i}, I_{\tau_i}) = \phi(t_i, q_{t_i}, S_{t_i}, \psi_{t_i}, I_{t_i}) + M(\tau_i, t_i) + \int_{t_i}^{\tau_i} \partial_t \phi(s, q_s, S_s, \psi_s, I_s) + \mu \partial_S \phi \\ + \frac{1}{2} \sigma^2 \partial_{SS} \phi + \sum_{\mathbf{k} \in \mathcal{K}} r_{(\psi_s, I_s), (\mathbf{k}^\psi, \mathbf{k}^I)}(\phi(s, q_s, S_s, \mathbf{k}^\psi, \mathbf{k}^I) - \phi(s, q_s, S_s, \psi_s, I_s)) \\ + \sum_{n=1}^N \lambda^n(\psi_s, I_s, p_s^n, \ell_s) \mathbb{E} \left[\phi(s, q_s - \ell_s^n \epsilon_s^n, S_s, \psi_s, I_s) - \phi(s, q_s, S_s, \psi_s, I_s) \right] ds,$$

where $M(\tau_i, t_i)$ is a martingale. This can be rewritten as

$$\phi(\tau_i, q_{\tau_i}, S_{\tau_i}, \psi_{\tau_i}, I_{\tau_i}) = \phi(t_i, q_{t_i}, S_{t_i}, \psi_{t_i}, I_{t_i}) + M(\tau_i, t_i) + \int_{t_i}^{\tau_i} \partial_t \phi(s, q_s, S_s, \psi_s, I_s) + \mu \partial_S \phi - g(q_s - q^*(s)) \\ + \frac{1}{2} \sigma^2 \partial_{SS} \phi + \sum_{\mathbf{k} \in \mathcal{K}} r_{(\psi_s, I_s), (\mathbf{k}^\psi, \mathbf{k}^I)}(\phi(s, q_s, S_s, \mathbf{k}^\psi, \mathbf{k}^I) - \phi(s, q_s, S_s, \psi_s, I_s)) + g(q_s - q^*(s)) \\ + \sum_{n=1}^N \lambda^n(\psi_s, I_s, p_s^n, \ell_s) \mathbb{E} \left[\epsilon_s^n \ell_s^n \left(S_s + \frac{\psi_s^n}{2} + p_s^n \delta^n \right) + \phi(s, q_s - \ell_s^n \epsilon_s^n, S_s, \psi_s, I_s) \right. \\ \left. - \phi(s, q_s, S_s, \psi_s, I_s) \right] ds - \sum_{n=1}^N \int_{t_i}^{\tau_i} \lambda^n(\psi_s, I_s, p_s^n, \ell_s) \mathbb{E} \left[\epsilon_s^n \ell_s^n \left(S_s + \frac{\psi_s^n}{2} + p_s^n \delta^n \right) \right] ds.$$

We derive

$$\begin{aligned} \phi(\tau_i, q_{\tau_i}, S_{\tau_i}, \psi_{\tau_i}, I_{\tau_i}) &\geq \phi(t_i, q_{t_i}, S_{t_i}, \psi_{t_i}, I_{t_i}) \\ &+ M(\tau_i, t_i) - \sum_{n=1}^N \int_{t_i}^{\tau_i} \lambda^n(\psi_s, I_s, p_s^n, \ell_s) \mathbb{E} \left[\epsilon_s^n \ell_s^n \left(S_s + \frac{\psi_s^n}{2} + p_s^n \delta^n \right) \right] + g(q_s - q^*(s)) ds. \end{aligned}$$

As the martingale term vanishes with the expectation, we get

$$\begin{aligned} \phi(t_i, q_{t_i}, S_{t_i}, \psi_{t_i}, I_{t_i}) &\leq \mathbb{E} \left[\phi(\tau_i, q_{\tau_i}, S_{\tau_i}, \psi_{\tau_i}, I_{\tau_i}) \right. \\ &\quad \left. + \sum_{n=1}^N \int_{t_i}^{\tau_i} \lambda^n(\psi_s, I_s, p_s^n, \ell_s) \mathbb{E} \left[\epsilon_s^n \ell_s^n \left(S_s + \frac{\psi_s^n}{2} + p_s^n \delta^n \right) \right] - g(q_s - q^*(s)) ds \right]. \end{aligned}$$

and thus

$$\begin{aligned} \phi(t_i, q_{t_i}, S_{t_i}, \psi_{t_i}, I_{t_i}) &\leq -\eta + \mathbb{E} \left[v(\tau_i, q_{\tau_i}, S_{\tau_i}, \psi_{\tau_i}, I_{\tau_i}) \right. \\ &\quad \left. + \sum_{n=1}^N \int_{t_i}^{\tau_i} \lambda^n(\psi_s, I_s, p_s^n, \ell_s) \mathbb{E} \left[\epsilon_s^n \ell_s^n \left(S_s + \frac{\psi_s^n}{2} + p_s^n \delta^n \right) \right] - g(q_s - q^*(s)) ds \right]. \end{aligned}$$

For i sufficiently large, we deduce

$$\begin{aligned} v(t_i, \tilde{q}, S_{t_i}, \tilde{\psi}, \tilde{I}) &\leq -\frac{\eta}{2} + \mathbb{E} \left[v(\tau_i, q_{\tau_i}, S_{\tau_i}, \psi_{\tau_i}, I_{\tau_i}) \right. \\ &\quad \left. + \sum_{n=1}^N \int_{t_i}^{\tau_i} \lambda^n(\psi_s, I_s, p_s^n, \ell_s) \mathbb{E} \left[\epsilon_s^n \ell_s^n \left(S_s + \frac{\psi_s^n}{2} + p_s^n \delta^n \right) \right] - g(q_s - q^*(s)) ds \right], \end{aligned}$$

which contradicts the dynamic programming principle. In conclusion, we necessarily have

$$\begin{aligned} 0 &\geq \partial_t v(\hat{t}, \hat{q}, \hat{S}, \hat{\psi}, \hat{I}) - g(q - q_t^*) + \mu \partial_S v + \frac{1}{2} \sigma^2 \partial_{SS} v + \sum_{\mathbf{k} \in \mathcal{K}} r_{(\hat{\psi}, \hat{I}), (\mathbf{k}^\psi, \mathbf{k}^I)} (v(\hat{t}, \hat{q}, \hat{S}, \mathbf{k}^\psi, \mathbf{k}^I) - v(\hat{t}, \hat{q}, \hat{S}, \hat{\psi}, \hat{I})) \\ &\quad + \sup_{p \in Q_\psi, \ell \in \mathcal{A}} \sum_{n=1}^N \lambda^n(\tilde{\psi}, \tilde{I}, p^n, \ell) \mathbb{E} \left[\epsilon^n \ell^n \left(\hat{S} + \frac{\tilde{\psi}^n}{2} + p^n \delta^n \right) + v(\hat{t}, \hat{q} - \ell^n \epsilon^n, \hat{S}, \tilde{\psi}, \tilde{I}) - v(\hat{t}, \hat{q}, \hat{S}, \hat{\psi}, \hat{I}) \right]. \end{aligned}$$

The second part of the HJBQVI being straightforward, we prove that v is a viscosity supersolution of the HJBQVI on $[0, T) \times \mathbb{R} \times \mathcal{D}$. The proof for the subsolution is identical, except that we need to prove

$$\sum_{n=1}^N \sup_{m^n \in [0, \bar{m}]} m^n \left(S - \frac{\psi^n}{2} \right) + v(t, q - m^n, S, \psi, I) - v(t, q, S, \psi, I) \geq 0,$$

which is direct by choosing the constant controls $\bar{m}^n = 0$ for all $n \in \{1, \dots, N\}$.

For the proof of the uniqueness, we recall the definition of subjet and superjet.

Definition 9.1. Let $v : [0, T) \times \mathbb{R} \times \mathcal{D} \mapsto \mathbb{R}$ be l.s.c (resp u.s.c) with respect to (\hat{t}, \hat{S}) . For $(\hat{t}, \hat{q}, \hat{S}, \hat{\psi}, \hat{I}) \in [0, T) \times \mathbb{R} \times \mathcal{D}$ we say that $(y, p, A) \in \mathbb{R}^3$ is in the subjet $\mathcal{P}^- v(\hat{t}, \hat{q}, \hat{S}, \hat{\psi}, \hat{I})$ (resp. the superjet $\mathcal{P}^+ v(\hat{t}, \hat{q}, \hat{S}, \hat{\psi}, \hat{I})$) if

$$v(t, \tilde{q}, S, \tilde{\psi}, \tilde{I}) \geq v(\hat{t}, \hat{q}, \hat{S}, \hat{\psi}, \hat{I}) + y(t - \hat{t}) + p(S - \hat{S}) + \frac{1}{2} A(S - \hat{S})^2 + o(|t - \hat{t}| + |S - \hat{S}|^2),$$

(resp. $v(t, \tilde{q}, S, \tilde{\psi}, \tilde{I}) \geq v(\hat{t}, \hat{q}, \hat{S}, \hat{\psi}, \hat{I}) + y(t - \hat{t}) + p(S - \hat{S}) + \frac{1}{2} A(S - \hat{S})^2 + o(|t - \hat{t}| + |S - \hat{S}|^2)$), for all (t, S) such that $(t, \tilde{q}, S, \tilde{\psi}, \tilde{I}) \in [0, T) \times \mathbb{R} \times \mathcal{D}$. We also define $\bar{\mathcal{P}}^- (\hat{t}, \hat{q}, \hat{S}, \hat{\psi}, \hat{I})$ as the set of points $(y, p, A) \in \mathbb{R}^3$ such that there exists a sequence $(t_I, \tilde{q}, S_i, \tilde{\psi}, \tilde{I}, y_i, p_i, A_i)$ satisfying

$$(t_i, \tilde{q}, S_i, \tilde{\psi}, \tilde{I}, y_i, p_i, A_i) \xrightarrow{i \rightarrow +\infty} (\hat{t}, \hat{q}, \hat{S}, \hat{\psi}, \hat{I}, y, p, A).$$

The set $\bar{\mathcal{P}}^+ (\hat{t}, \hat{q}, \hat{S}, \hat{\psi}, \hat{I})$ is defined similarly.

We now introduce an analogous of the Ishii's lemma, whose proof can be found in Bouchard [53].

Lemma 9.1. *A l.s.c (resp u.s.c) function v is a supersolution (resp. subsolution) of the HJBQVI on $[0, T] \times \mathbb{R} \times \mathcal{D}$ if and only if for all $(\hat{t}, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I}) \in [0, T] \times \mathbb{R} \times \mathcal{D}$, and all $(\hat{y}, \hat{p}, \hat{A}) \in \overline{\mathcal{P}}^-(\hat{t}, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I})$ (resp. $\overline{\mathcal{P}}^+(\hat{t}, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I})$), we have*

$$\begin{aligned} 0 \leq \min & \left\{ -\hat{y} + g(\tilde{q} - q^*(\hat{t})) - \mu\hat{p} - \frac{1}{2}\sigma^2\hat{A} - \sum_{\mathbf{k} \in \mathcal{K}} r_{(\tilde{\psi}, \tilde{I}), (\mathbf{k}^\psi, \mathbf{k}^I)}(v(\hat{t}, \tilde{q}, \hat{S}, \mathbf{k}^\psi, \mathbf{k}^I) - v(\hat{t}, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I})) \right. \\ & - \sup_{p \in Q_\psi, \ell \in \mathcal{A}} \sum_{n=1}^N \lambda^n(\tilde{\psi}, \tilde{I}, p^n, \ell) \mathbb{E} \left[\epsilon^n \ell^n \left(\hat{S} + \frac{\tilde{\psi}^n}{2} + p^n \delta^n \right) + v(\hat{t}, \tilde{q} - \ell^n \epsilon^n, \hat{S}, \tilde{\psi}, \tilde{I}) \right. \\ & \left. \left. - v(\hat{t}, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I}) \right]; \sum_{n=1}^N v(\hat{t}, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I}) - \sup_{m^n \in [0, \bar{m}]} m^n \left(\hat{S} - \frac{\tilde{\psi}^n}{2} \right) + v(\hat{t}, \tilde{q} - m^n, \hat{S}, \tilde{\psi}, \tilde{I}) \right\}, \end{aligned}$$

(resp. ≤ 0).

We now prove the following comparison principle:

Proposition 9.1. *Let u (resp. v) be a l.s.c supersolution (resp. u.s.c subsolution) with polynomial growth of the HJBQVI on $[0, T] \times \mathbb{R} \times \mathcal{D}$. If $u \geq v$ on $\{T\} \times \mathbb{R} \times \mathcal{D}$, then $u \geq v$ on $[0, T] \times \mathbb{R} \times \mathcal{D}$.*

Proof. For $\rho > 0$ we introduce the following change of variables:

$$\tilde{u}(t, q, S, \psi, I) = e^{\rho t} u(t, q, S, \psi, I), \quad \tilde{v}(t, q, S, \psi, I) = e^{\rho t} v(t, q, S, \psi, I).$$

Then, \tilde{u} and \tilde{v} are respectively supersolution and subsolution of the following equation:

$$\begin{aligned} 0 = \min & \left\{ -\partial_t w(t, q, S, \psi, I) + \rho w(t, q, S, \psi, I) + g(q - q_t^*) - \mu \partial_S w - \frac{1}{2} \sigma^2 \partial_{SS} w \right. \\ & - \sum_{\mathbf{k} \in \mathcal{K}} r_{(\psi, I), (\mathbf{k}^\psi, \mathbf{k}^I)}(w(t, q, S, \mathbf{k}^\psi, \mathbf{k}^I) - w(t, q, S, \psi, I)) \\ & - \sup_{p \in Q_\psi, \ell \in \mathcal{A}} \sum_{n=1}^N \lambda^n(\psi, I, p^n, \ell) \mathbb{E} \left[\epsilon^n \ell^n e^{\rho t} \left(S + \frac{\psi^n}{2} + p^n \delta^n \right) + w(t, q - \ell^n \epsilon^n, S, \psi, I) \right. \\ & \left. \left. - w(t, q, S, \psi, I) \right]; \sum_{n=1}^N w(t, q, S, \psi, I) - \sup_{m^n \in [0, \bar{m}]} m^n e^{\rho t} \left(S - \frac{\psi^n}{2} \right) + w(t, q - m^n, S, \psi, I) \right\}, \end{aligned}$$

on $[0, T] \times \mathbb{R} \times \mathcal{D}$, with $\tilde{u} \geq \tilde{v}$ on $\{T\} \times \mathbb{R} \times \mathcal{D}$. In order to prove the proposition, we only have to show that $\tilde{u} \geq \tilde{v}$ on $[0, T] \times \mathbb{R} \times \mathcal{D}$. Assume that the minimum is given by the first term and that $\sup_{[0, T] \times \mathbb{R} \times \mathcal{D}} \tilde{v} - \tilde{u} > 0$. We fix $p \in \mathbb{N}^*$ such that

$$\lim_{\|S\|_2 \rightarrow +\infty} \sup_{(t, q, \psi, I) \in [0, T] \times \mathcal{D}} \frac{|\tilde{u}(t, q, S, \psi, I)| + |\tilde{v}(t, q, S, \psi, I)|}{1 + \|S\|_2^{2p}} = 0.$$

Then, there exists $(\hat{t}, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I}) \in [0, T] \times \mathbb{R} \times \mathcal{D}$ such that

$$\begin{aligned} 0 & < \tilde{v}(\hat{t}, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I}) - \tilde{u}(\hat{t}, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I}) - \phi(\hat{t}, \tilde{q}, \hat{S}, \hat{S}, \tilde{\psi}, \tilde{I}) \\ & = \max_{(t, q, S, \psi, I)} \tilde{v}(t, q, S, \psi, I) - \tilde{u}(t, q, S, \psi, I) - \phi(t, q, S, S, \psi, I), \end{aligned}$$

where $\epsilon > 0$ is small enough and

$$\phi(t, S, R) = \epsilon \exp(-\tilde{\kappa} t) (1 + \|S\|_2^{2p} + \|R\|_2^{2p}), \quad \tilde{\kappa} > 0.$$

Since $\tilde{u} \geq \tilde{v}$ on $\{T\} \times \mathbb{R} \times \mathcal{D}$, we directly have $\hat{t} < T$.

For all $i \in \mathbb{N}$, we can find a sequence (t_i, S_i, R_i) such that

$$\begin{aligned} 0 &< \tilde{v}(t_i, \tilde{q}, S_i, \tilde{\psi}, \tilde{I}) - \tilde{u}(t_i, \tilde{q}, R_i, \tilde{\psi}, \tilde{I}) - \phi(t_i, S_i, R_i) - i|S_i - R_i|^2 - (|t_i - \hat{t}|^2 + |S_i - \hat{S}|^4) \\ &= \max_{(t, S, R)} \tilde{v}(t, \tilde{q}, S, \tilde{\psi}, \tilde{I}) - \tilde{u}(t, \tilde{q}, R, \tilde{\psi}, \tilde{I}) - \phi(t, S, R) - i|S - R|^2 - (|t - \hat{t}|^2 + |S - \hat{S}|^4). \end{aligned}$$

Then we have:

$$(t_i, S_i, R_i) \xrightarrow{i \rightarrow +\infty} (\hat{t}, \hat{S}, \hat{S})$$

up to a subsequence, and

$$\begin{aligned} &\tilde{v}(t_i, \tilde{q}, S_i, \tilde{\psi}, \tilde{I}) - \tilde{u}(t_i, \tilde{q}, R_i, \tilde{\psi}, \tilde{I}) - \phi(t_i, S_i, R_i) - i|S_i - R_i|^2 - (|t_i - \hat{t}|^2 + |S_i - \hat{S}|^4) \\ &\xrightarrow{n \rightarrow +\infty} \tilde{v}(\hat{t}, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I}) - \tilde{u}(\hat{t}, \tilde{q}, \hat{S}, \tilde{\psi}, \tilde{I}) - \phi(\hat{t}, \hat{S}, \hat{S}) \end{aligned}$$

Let us then denote for $i \in \mathbb{N}^*$

$$\varphi_i(t, S, R) = \phi(t, S, R) + i|S - R|^2 + |t - \hat{t}|^2 + |S - \hat{S}|^4, \quad \text{for all } (t, S, R) \in [0, T] \times \mathbb{R}^2.$$

Then Ishii's Lemma (see Barles and Imbert [36], Crandall, Ishii, and Lions [83]) guarantees that for all $\eta > 0$, we can find $(y_i^1, p_i^1, A_i^1) \in \bar{\mathcal{P}}^+ \tilde{v}(t_i, \tilde{q}, S_i, \tilde{\psi}, \tilde{I})$ and $(y_i^2, p_i^2, A_i^2) \in \bar{\mathcal{P}}^- \tilde{u}(t_i, \tilde{q}, R_i, \tilde{\psi}, \tilde{I})$ such that:

$$y_i^1 - y_i^2 = \partial_t \varphi_i(t_i, S_i, R_i), \quad (p_i^1, p_i^2) = (\partial_S \varphi_i, -\partial_R \varphi_i)(t_i, S_i, R_i),$$

and

$$\begin{pmatrix} A_i^1 & 0 \\ 0 & -A_i^2 \end{pmatrix} \leq H_{SR} \varphi_i(t_i, S_i, R_i) + \eta (H_{SR} \varphi_n(t_i, S_i, R_i))^2,$$

where $H_{SR} \varphi_i(t_i, \cdot, \cdot)$ denotes the Hessian matrix of $\varphi_i(t_i, \cdot, \cdot)$. Applying Lemma 9.1, we get

$$\begin{aligned} \rho(\tilde{v}(t_i, \tilde{q}, S_i, \tilde{\psi}, \tilde{I}) - \tilde{u}(t_i, \tilde{q}, R_i, \tilde{\psi}, \tilde{I})) &\leq y_i^1 - y_i^2 + \frac{1}{2} \sigma^2 (A_i^1 - A_i^2) + \mu(p_i^1 - p_i^2) \\ &+ \sum_{\mathbf{k} \in \mathcal{K}} r_{(\tilde{\psi}, \tilde{I}), (\mathbf{k}^\psi, \mathbf{k}^I)} (\tilde{v}(t_i, \tilde{q}, S_i, \mathbf{k}^\psi, \mathbf{k}^I) - \tilde{v}(t_i, \tilde{q}, S_i, \tilde{\psi}, \tilde{I})) \\ &+ \sup_{p \in Q_\psi, \ell \in \mathcal{A}} \sum_{n=1}^N \lambda^n(\tilde{\psi}, \tilde{I}, p^n, \ell) \mathbb{E} \left[\epsilon^n \ell^n e^{\rho t_i} \left(S_i + \frac{\tilde{\psi}^n}{2} + p^n \delta^n \right) + \tilde{v}(t_i, \tilde{q} - \ell^n \epsilon^n, S_i, \tilde{\psi}, \tilde{I}) - \tilde{v}(t_i, \tilde{q}, S_i, \tilde{\psi}, \tilde{I}) \right] \\ &- \sum_{\mathbf{k} \in \mathcal{K}} r_{(\tilde{\psi}, \tilde{I}), (\mathbf{k}^\psi, \mathbf{k}^I)} (\tilde{u}(t_i, \tilde{q}, R_i, \mathbf{k}^\psi, \mathbf{k}^I) - \tilde{u}(t_i, \tilde{q}, R_i, \tilde{\psi}, \tilde{I})) \\ &- \sup_{p \in Q_\psi, \ell \in \mathcal{A}} \sum_{n=1}^N \lambda^n(\tilde{\psi}, \tilde{I}, p^n, \ell) \mathbb{E} \left[\epsilon^n \ell^n e^{\rho t_i} \left(R_i + \frac{\tilde{\psi}^n}{2} + p^n \delta^n \right) + \tilde{u}(t_i, \tilde{q} - \ell^n \epsilon^n, R_i, \tilde{\psi}, \tilde{I}) - \tilde{u}(t_i, \tilde{q}, R_i, \tilde{\psi}, \tilde{I}) \right]. \end{aligned}$$

Moreover, we have

$$H_{SR} \varphi_i(t_i, S_i, R_i) = \begin{pmatrix} \partial_{SS}^2 \phi(t_i, S_i, R_i) + 2i + 12(S_i - \hat{S})^2 & \partial_{SR}^2 \phi(t_i, S_i, R_i) - 2i \\ \partial_{SR}^2 \phi(t_i, S_i, R_i) - 2i & \partial_{RR}^2 \phi(t_i, S_i, R_i) + 2i \end{pmatrix},$$

and

$$\begin{aligned} \partial_S \varphi_i(t_i, S_i, R_i) &= \partial_S \phi(t_i, S_i, R_i) + 2i|S_i - R_i| + 4|S_i - \hat{S}|^3, \\ \partial_R \varphi_i(t_i, S_i, R_i) &= \partial_R \phi(t_i, S_i, R_i) - 2i|S_i - R_i|, \end{aligned}$$

so from what precedes we can write

$$\begin{aligned} \rho(\tilde{v}(t_i, \tilde{q}, S_i, \tilde{\psi}, \tilde{I}) - \tilde{u}(t_i, \tilde{q}, R_i, \tilde{\psi}, \tilde{I})) &\leq \partial_t \phi(t_i, S_i, R_i) + 2(t_i - \hat{t}) + \mu(\partial_S \phi(t_i, S_i, R_i) + \partial_R \phi(t_i, S_i, R_i) + 4(S_i - \hat{S})^3) \\ &+ \frac{1}{2} \sigma^2 (\partial_{SS}^2 \phi(t_i, S_i, R_i) + \partial_{RR}^2 \phi(t_i, S_i, R_i) + 2\partial_{SR}^2 \phi(t_i, S_i, R_i) + 12(S_i - \hat{S})) + \eta C_i \end{aligned}$$

$$\begin{aligned}
& + \sum_{\mathbf{k} \in \mathcal{K}} r_{(\tilde{\psi}, \tilde{I}), (\mathbf{k}^\psi, \mathbf{k}^I)} (\tilde{v}(t_i, \tilde{q}, S_i, \mathbf{k}^\psi, \mathbf{k}^I) - \tilde{v}(t_i, \tilde{q}, S_i, \tilde{\psi}, \tilde{I})) \\
& + \sup_{p \in Q_\psi, \ell \in \mathcal{A}} \sum_{n=1}^N \lambda^n(\tilde{\psi}, \tilde{I}, p^n, \ell) \mathbb{E} \left[\epsilon^n \ell^n e^{\rho t_i} \left(S_i + \frac{\tilde{\psi}^n}{2} + p^n \delta^n \right) + \tilde{v}(t_i, \tilde{q} - \ell^n \epsilon^n, S_i, \tilde{\psi}, \tilde{I}) - \tilde{v}(t_i, \tilde{q}, S_i, \tilde{\psi}, \tilde{I}) \right] \\
& - \sum_{\mathbf{k} \in \mathcal{K}} r_{(\tilde{\psi}, \tilde{I}), (\mathbf{k}^\psi, \mathbf{k}^I)} (\tilde{u}(t_i, \tilde{q}, R_i, \mathbf{k}^\psi, \mathbf{k}^I) - \tilde{u}(t_i, \tilde{q}, R_i, \tilde{\psi}, \tilde{I})) \\
& - \sup_{p \in Q_\psi, \ell \in \mathcal{A}} \sum_{n=1}^N \lambda^n(\tilde{\psi}, \tilde{I}, p^n, \ell) \mathbb{E} \left[\epsilon^n \ell^n e^{\rho t_i} \left(R_i + \frac{\tilde{\psi}^n}{2} + p^n \delta^n \right) + \tilde{u}(t_i, \tilde{q} - \ell^n \epsilon^n, R_i, \tilde{\psi}, \tilde{I}) - \tilde{u}(t_i, \tilde{q}, R_i, \tilde{\psi}, \tilde{I}) \right],
\end{aligned}$$

where C_i does not depend on η . As \tilde{v} is u.s.c., \tilde{u} is l.s.c. and $(t_i, S_i, R_i)_i$ is convergent, when $\eta \rightarrow 0$ it is clear, that, when $i \rightarrow +\infty$, for a certain constant M we get

$$\begin{aligned}
\rho(\tilde{v}(\hat{t}, \hat{q}, \hat{S}, \hat{\psi}, \hat{I}) - \tilde{u}(\hat{t}, \hat{q}, \hat{S}, \hat{\psi}, \hat{I})) & \leq \partial_t \phi(\hat{t}, \hat{S}, \hat{S}) + \mu(\partial_S \phi(\hat{t}, \hat{S}, \hat{S}) + \partial_R \phi(\hat{t}, \hat{S}, \hat{S})) \\
& + \frac{1}{2} \sigma^2 (\partial_{SS}^2 \phi(\hat{t}, \hat{S}, \hat{S}) + \partial_{RR}^2 \phi(\hat{t}, \hat{S}, \hat{S}) + 2\partial_{SR}^2 \phi(\hat{t}, \hat{S}, \hat{S})) + M.
\end{aligned}$$

For $\tilde{\kappa} > 0$ large enough, the right-hand side is strictly negative, and as $\rho > 0$ we get

$$\tilde{v}(\hat{t}, \hat{q}, \hat{S}, \hat{\psi}, \hat{I}) - \tilde{u}(\hat{t}, \hat{q}, \hat{S}, \hat{\psi}, \hat{I}) < 0,$$

which yields to a contradiction. The proof for the other part of the HJBQVI is direct.

With the two above propositions, it is easy to conclude the proof of the theorem. Indeed, as v_\star is a supersolution such that $v_\star \geq v$ on $\{T\} \times \mathbb{R} \times \mathcal{D}$, and v^\star is a subsolution such that $v^\star \leq v$ on $\{T\} \times \mathbb{R} \times \mathcal{D}$, we can apply the maximum principle to get $v_\star \geq v^\star$ on $[0, T] \times \mathbb{R} \times \mathcal{D}$. But by definition of v_\star and v^\star , we must have $v_\star \leq v \leq v^\star$ on $[0, T] \times \mathbb{R} \times \mathcal{D}$, which proves that we have $v_\star = v = v^\star$ and v is continuous. The maximum principle implies that if two continuous viscosity solutions of the HJBQVI satisfy the same terminal condition, they are equal on $[0, T] \times \mathbb{R} \times \mathcal{D}$, hence the uniqueness. \square

9.A.2 Application to OTC market-making

9.A.2.1 Framework

The model we present in this chapter is designed for trading in cross-listed stocks in limit order books. However, it can be adapted straightforwardly to handle the problem of an OTC market-maker, who often deals with a large number of assets driven by a few factors. We borrow here the factorial method market-making framework of Bergault and Guéant [41] (we are also going to keep their notation only for this section). We consider a market-maker who is in charge of providing bid and ask quotes on d assets, whose dynamics are

$$dS_t^i = \mu^i dt + \sigma^i dW_t^i, \quad i \in \{1, \dots, d\},$$

where μ^i is the drift of the i -th asset, σ^i is its volatility and (W_t^1, \dots, W_t^d) is a d -dimensional Brownian motion. We consider a non-singular variance-covariance matrix $\Sigma = (\rho^{i,j} \sigma^i \sigma^j)_{i,j \in \{1, \dots, d\}}$ for the vector of assets (S_t^1, \dots, S_t^d) . The market-maker sets bid and ask prices on every asset:

$$S^{i,b}(t, z) = S_t^i - \delta^{i,b}(t, z), \quad S^{i,a}(t, z) = S_t^i + \delta^{i,a}(t, z), \quad z \in \mathbb{R},$$

where $\delta = (\delta^{i,a}, \delta^{i,b})_{i \in \{1, \dots, d\}}$ are the (predictable and uniformly lower bounded) bid and ask quotes around the mid-price of each asset. The volume of transactions on the bid and ask sides are modeled by marked point processes $N^{i,b}(dt, dz), N^{i,a}(dt, dz)$ of intensity $\nu_t^{i,b}(dz), \nu_t^{i,a}(dz)$ defined by

$$\nu_t^{i,j}(dt, dz) = \Lambda^{i,j}(\delta^{i,j}(t, z)) \eta^{i,j}(dz), \quad i \in \{1, \dots, d\},$$

where $\Lambda^{i,j}$ is a sufficiently regular function (exponential, logistic, SU Johnson etc.) modeling the probability to trade on the asset i , on the side j for a given spread $\delta^{i,j}(t, z)$ and a size z . The functions $\eta^{i,j}(dz)$ are probability densities over \mathbb{R}_+ modeling the distribution of a trade size. The market-maker manages his inventory process $q_t = (q_t^1, \dots, q_t^d)$ of dynamics given by

$$dq_t^i = \int_{\mathbb{R}_+} z N^{i,b}(dt, dz) - \int_{\mathbb{R}_+} z N^{i,a}(dt, dz), \quad i \in \{1, \dots, d\}.$$

The market-maker manages his cash process given at time t by

$$dX_t = \sum_{i=1}^d \int_{\mathbb{R}_+} z S^{i,a}(t, z) N^{i,a}(dt, dz) - \int_{\mathbb{R}_+} z S^{i,b}(t, z) N^{i,b}(dt, dz).$$

Its optimization problem is defined as

$$\sup_{\delta} \mathbb{E} \left[X_T + \sum_{i=1}^d q_T^i S_T^i - \int_0^T \phi(q_t) dt \right],$$

where ϕ is a running penalty preventing from too large positions and $\sum_{i=1}^d q_T^i S_T^i$ is the marked-to-market value of the market-maker's portfolio at time t . The corresponding HJB equation is given by

$$\begin{aligned} 0 = & \partial_t v(t, q) + \sum_{i=1}^d q^i \mu^i - \phi(q) + \sum_{i=1}^d \int_{\mathbb{R}_+} z H^{i,b} \left(\frac{v(t, q) - v(t, q + ze^i)}{z} \right) \eta^{i,b}(dz) \\ & + \sum_{i=1}^d \int_{\mathbb{R}_+} z H^{i,a} \left(\frac{v(t, q) - v(t, q - ze^i)}{z} \right) \eta^{i,a}(dz), \end{aligned}$$

with terminal condition $v(T, q) = 0$, $H^{i,j}(p) = \sup_{\delta} \Lambda^{i,j}(\delta)(\delta - p)$, and (e^1, \dots, e^d) is the canonical basis of \mathbb{R}^d .

9.A.3 Bayesian update for OTC market-makers

Usually, the functions $\Lambda^{i,j}$ are of the form

$$\Lambda^{i,j}(\delta^{i,j}(t, z)) = \lambda_{\text{RFQ}}^{i,j} f(\delta^{i,j}(t, z)),$$

where $\lambda_{\text{RFQ}}^{i,j}$ is the constant intensity of arrival of requests for quote, and $f(\delta^{i,j}(t, z))$ gives the probability that a request will result in a transaction given the quote δ proposed by the market-maker. The estimation of the quantity $\lambda_{\text{RFQ}}^{i,j}$ is of particular importance for the market-maker so that he can adjust his quotes depending on his view on the number of request for a certain asset and a certain side. In the same spirit as in Section 9.3.1.1, we assume the following prior distribution:

$$\lambda_{\text{RFQ}}^{i,j} \sim \Gamma(\alpha^{i,j}, \beta^{i,j}), \quad (\alpha^{i,j}, \beta^{i,j}) > 0.$$

For an asset $i \in \{1, \dots, d\}$ on the side $j \in \{a, b\}$, this corresponds to an average intensity of $\frac{\alpha^{i,j}}{\beta^{i,j}}$, with variance equal to $\frac{\alpha^{i,j}}{(\beta^{i,j})^2}$. If the market-maker is confident in his estimation of the intensity $\lambda_{\text{RFQ}}^{i,j}$, he can choose a large $\beta^{i,j}$ so that the variance of his Bayesian estimator is small. Given all the information accumulated up to time t , its best estimation of the quantity $\lambda_{\text{RFQ}}^{i,j}$, is given by

$$\mathbb{E}[\lambda_{\text{RFQ}}^{i,j} | N(t, dz)] = \frac{\alpha^{i,j} + \int_{\mathbb{R}_+} N(t, dz)}{\beta^{i,j} + \int_{\mathbb{R}_+} \int_0^t f(\delta^{i,j}(s, z)) ds \eta^{i,j}(dz)}. \quad (9.7)$$

By the law of large numbers, when the market-maker has accumulated a sufficiently large number of observations, his best estimation of $\lambda_{\text{RFQ}}^{i,j}$ converges to the “real” intensity of the market. As time passes, the prior parameters $(\alpha^{i,j}, \beta^{i,j})$ of the market-maker are less important as the estimation will rely mostly on the observations.

Another important parameter of the model is the size of transactions, which impacts the quotes of the market-maker as well as his inventory risk. In Bergault and Guéant [41], the authors choose in their numerical experiments a $\Gamma(a^{i,j}, b^{i,j})$ distribution for $\eta^{i,j}$. The trader can choose between Bayesian updates (revise only $a^{i,j}$, only $b^{i,j}$, or both), depending on his confidence on parameters’ estimation. If he is confident with respect to the shape parameter $a^{i,j}$, that is he knows approximately the average size of a request but not the standard deviation, he sets $b^{i,j} \sim \Gamma(a_0^{i,j}, b_0^{i,j})$. Given n observations of size z^1, \dots, z^n , the best Bayesian estimate of $b^{i,j}$ (the scale parameter of the size of the request) is

$$\mathbb{E}[b^{i,j} | (z^1, \dots, z^n)] = \frac{a_0^{i,j} + na^{i,j}}{b_0^{i,j} + \sum_{i=1}^n z^i}. \quad (9.8)$$

The use of different prior distribution to take into account the uncertainty on the shape parameter $a^{i,j}$ (if $b^{i,j}$ is known) or on both $(a^{i,j}, b^{i,j})$ can be done in the same way.

Another sensitive parameter, especially for the multi-asset market-making, is the variance-covariance matrix Σ . This quantity is usually estimated on a long run, but parameters are subject to a brutal change. For example, let us assume that the market-maker is in charge of d assets on 2 different sectors (for instance, technology and aerospace). Following the factorial approach, the market-making problem’s dimension will be reduced from d to 3. The three factors mainly correspond to the three highest eigenvalues of the variance-covariance matrix Σ , and will drive the quotes of the market-maker. However, in case of a sectorial tail event, for example the bankruptcy of one of the companies of the tech sector, it is likely that all the correlations between the assets of this sector will rise to one. This will impact the eigenvalue related to the technology sector, and change the quotes of the market-maker as he has to avoid long inventory positions on assets whose values are decreasing. To design adaptive market-making strategy based on Bayesian update of the correlation matrix and the drift of the assets, we define the Normal-Inverse-Wishart prior on $(\mu, \Sigma) \sim \text{NIW}(\mu_0, \kappa_0, \nu_0, \psi)$, where $(\mu_0, \kappa_0, \nu_0, \psi) \in \mathbb{R}^d \times \mathbb{R}_+^* \times (d-1, +\infty) \times \mathcal{M}_d(\mathbb{R})$. This distribution is built as follows:

$$\mu | (\mu_0, \kappa_0, \Sigma) \sim \mathcal{N}\left(\mu_0, \frac{1}{\kappa_0} \Sigma\right), \quad \Sigma | (\psi, \nu_0) \sim \mathcal{W}^{-1}(\psi, \nu_0), \quad \text{then } (\mu, \Sigma) \sim \text{NIW}(\mu_0, \kappa_0, \nu_0, \psi),$$

where \mathcal{W}^{-1} is the standard inverse Wishart distribution. In other words, the drift vector μ of the assets follows a multivariate Gaussian distribution whereas the variance-covariance matrix Σ follows a standard inverse Wishart distribution. At time t , if we note $S_t = (S_t^1, \dots, S_t^d)$ the prices observed up to time t , the Bayesian update of (μ, Σ) is

$$(\mu, \Sigma | S_t - S_0) \sim \text{NIW}\left(\frac{\kappa_0 \mu_0 + (S_t - S_0)}{\kappa_0 + t}, \kappa_0 + t, \nu_0 + t, \psi + \left(S_t - \frac{S_t}{t}\right) \left(S_t - \frac{S_t}{t}\right)^\top + \frac{\kappa_0 t}{\kappa_0 + t} \left(\mu_0 - \frac{S_t}{t}\right) \left(\mu_0 - \frac{S_t}{t}\right)^\top\right).$$

Following the law of large numbers, as $t \rightarrow +\infty$ we have a larger number of information and we converge toward the drift and variance-covariance of the market-maker’s portfolio. Therefore, the market-maker can recompute his factors derived from the updated variance-covariance matrix and adjust his quotes.

This extension deserves several remarks. First, the problems encountered by an OTC market-maker are quite different from a high-mid frequency trader in an order book. The model is more parsimonious,

especially for the intensity functions. Therefore, the convergence toward the “true” market parameters will be faster than in order book model. The objective of the Bayesian update on the quantities $\lambda_{\text{RFQ}}^{i,j}$ is to determine the average behavior or the counterparts of the market-maker. If he observes a large number of requests on the ask (resp. bid) side of the asset i , the Bayesian update (9.7) enables the market-maker to adjust his quotes to set a higher ask (resp. bid) price for this asset. If the market-maker observes a higher discrepancy than expected for the transaction sizes, the Bayesian update (9.8) helps to adjust his quotes. Finally, the Bayesian learning on the drift and covariance of the assets enables to update the factors from which the market-maker chooses his quotes.

Chapter 10

Optimal trading without optimal control

10.1 Introduction

Consider an investor whose preferences are described by a utility function of wealth, $u(w)$, as per Arrow [24] and Pratt [233]. Let w_T denote the investor's wealth at some known final time T . The investor attempts to maximize the expectation of utility of final wealth, $\mathbb{E}[u(w_T)]$, by trading financial assets. The mechanism by which buyers meet sellers and trades occur is known as the market microstructure. In this work, the microstructure is assumed to be a continuous double auction electronic order book with time priority, although our methods could be generalized to include other kinds of market microstructure. In continuous limit-order-book microstructure, trades are effected by submitting limit orders to an exchange's matching engine. For each security being traded, the investor must determine the price levels at which to submit buy and sell orders and the associated share quantities attached to those orders. In real markets, the price levels are discrete; the minimum possible price increment is the quote resolution allowed by the exchange, known as the tick size. Alternatively, the investor may decide to refrain from placing any orders or cancel some existing orders. Other decision variables include order type and venue. Considering all of these details, we see that the instantaneous action space is an inconveniently large discrete space; we discuss ways of simplifying it later on.

Problem 10.1. *The investor seeks the optimal dynamic strategy for choosing an action $a_t \in \mathcal{A}_t$ at each time t , where \mathcal{A}_t is the set of possible actions at time t , optimal in the sense of maximizing the expected utility of final wealth, $\mathbb{E}[u(w_T)]$.*

Problem 10.1 is mathematically deep and perhaps intractable; it is essentially a stochastic optimal control problem over high-dimensional discrete action and state spaces. According to Cont and Kukanov, 'Although simultaneous optimization of order timing, type, and routing decisions is an interesting problem, it also appears to be intractable' [81, p. 4] and even this is a special case of Problem 10.1.

The purpose of the chapter is to give practically implementable methods which execution desks could start using right away to solve Problem 10.1 approximately. Our approximation method breaks the problem up into two parts. The first part is to construct a *smooth relaxation* of the problem, which is essentially the continuous-time and continuous-space limit in which market microstructure is hidden behind smooth cost functions; the second part is to adjust our microstructure decisions to track the smooth relaxation optimally. Smooth relaxation problems provide no guidance on microstructure-level decisions, effectively assuming all executions use market orders. In particular, if we can predict the probability of a passive fill at any given instant (e.g., based on order book imbalance) our model provides a very clear and obvious way for the implementor to take advantage of predictions of passive fill completion.

Let $V(t, x)$ denote the Bellman value function for the smooth relaxation, defined to be the remaining expected-utility gain from time t obtained from following the best policy when the current state at time

t is x . There is also, in principle, a value function for Problem 10.1 defined in the same way, but the latter appears to be intractable. The present chapter's key idea is to exploit the value function of the smooth relaxation to provide effective microstructure alphas that adjust the microstructure decisions toward long-term optimality. The vector

$$p := \nabla V(t, x) \in \mathbb{R}^d \quad (10.1)$$

plays a central role in our approach, where $d \geq 1$ is the number of traded assets. As defined, p is the direction of steepest ascent for the value function. In our heuristics, (10.1) encodes all of the information about the long-term utility function is needed to make the microstructure decision, so it provides the key link between the trading schedule and the order routing problem.

In order to describe our policy for selecting the best microstructure action, we must first introduce some more notation. Write $R_{i,t}(v, a)$ for the (random) profit (or loss, if negative) from an order of quantity v on stock i using action a (aggressive or passive) over a short interval $[t, t + \delta t]$. The expected profit of the trader $\mathbb{E}[R_{i,t}(v, a)]$ is a function of the trading cost associated with the pair (v, a) , and also on his views concerning the short-horizon direction of the mid-price

$$r_i^{\text{mid}} = \text{mid}_i(t + \delta t) / \text{mid}_i(t) - 1.$$

With no subscript, r^{mid} denotes the d -dimensional vector of all midpoint returns for all assets.

Definition 10.1. *The term effective microstructure alphas, as used in this chapter, will denote a set of parameters given to a microstructure trader in the place of $\mathbb{E}[r^{\text{mid}}]$, for the purpose of satisfying either a short-term goal or a long-term goal.*

Microstructure alphas, as defined, could be a simple prediction of r^{mid} , or, more interestingly, they could be purposely skewed to encourage trading to increase expected utility (ie. increase long-term alpha and reduce risk) as we shall suggest in Equation (10.3). Our heuristics is, at time t , for each security indexed by $i \in \{1, \dots, d\}$, choose an instantaneous action a_i^* which solves the following maximization (over the finite set $\mathcal{A}_{t,i}$ of possible actions on the asset i at time t):

$$a_i^* = \operatorname{argmax}_{a \in \mathcal{A}_{t,i}} \mathbb{E}[R_{i,t}(v_a, a)], \quad \text{for all } i \in \{1, \dots, d\}, \quad (10.2)$$

$$\text{where: } \mathbb{E}[r^{\text{mid}}] = p := \nabla V(t, x), \quad (10.3)$$

and follow this action over the interval $[t, t + \delta t]$. Here v_a denotes the quantity associated with action a ; for example “ $a = \text{aggressive buy 100 shares}$ ” means $v_a = 100$. Note also that r^{mid} and p are n -vectors, so the equation $\mathbb{E}[r^{\text{mid}}] = p$ expresses the trader's views in all n assets.

There is a very good intuitive justification for (10.2)-(10.3). We show later that, under certain conditions, the optimal instantaneous trading rate at any time t in the smooth relaxation is given by

$$\operatorname{argmax}_v \{ \langle p, v \rangle - c(v) \}, \quad (10.4)$$

where $c(v)$ is the average cost of trading at rate v . The expression $\langle p, v \rangle - c(v)$ is the instantaneous analogue of expected profit minus cost, if your expected return is p .

The rest of the chapter is organized as follows. In Section 10.2, we present an example of a long-term trading schedule, the Almgren-Chriss case, and how to compute the value function $V(t, x)$ and its gradient p . We further show how p is related to the generalized momenta of the Hamiltonian approach. In Section 10.3, we derive the heuristics (10.2) by analogy with the smooth case and show how the trader can choose its short-term alpha in order to minimize the error with respect to the trading schedule. In Section 10.4, we present a general microstructure trading framework on a portfolio of cross-listed assets, taking into account long and short-term trading signals as well as many

components of market microstructure (spread, imbalance, probability of filling etc). We also show how this heuristics can be applied to the problem of multi-asset market-making. Finally, Section 10.5 shows a detailed numerical example which illustrates the dangers of separating portfolio construction from execution, and in which our method generates an improvement that is both statistically and economically significant.

10.2 The long-term trading curve

In this section, we recall the optimal portfolio liquidation framework of Almgren and Chriss [10] in continuous time and show how to solve the problem using the Hamiltonian method. In the process, we introduce notation that is used in the rest of the chapter.

Variations of the Almgren-Chriss model are used in execution desks all over the world. Indeed it is safe to say that the Almgren-Chriss model (usually with some custom extensions) has been used to execute trillions of dollars' worth of customer orders. We also believe that the Almgren-Chriss model, used in conjunction with the presented heuristics for microstructure decisions, is a reasonable choice if the trader has no alpha forecasts (with the possible exception of microstructure forecasts). Moreover, the computations we do here generalize in a straightforward way to extensions of the Almgren-Chriss model, such as to include long-term alpha. One of the key pieces of intuition that allows us to generalize the model, we feel, comes from Theorem 10.1 which makes the connection between microstructure alpha, Hamilton's generalized momentum, and the steepest-ascent direction of the long-term value function. For all of these reasons, we feel this is a useful example to do in detail.

We consider a trader in charge of a portfolio of $d \geq 1$ assets, of initial positions $q_0 = (q_0^1, \dots, q_0^d)^\top$ where $q_0^i \in \mathbb{R}$ for all $i \in \{1, \dots, d\}$. The trader wants to unwind this portfolio over the time horizon $[0, T]$, where $T > 0$. Given a control process $(v_t^i)_{t \in [0, T]}$ representing the trading rate on asset i , the inventory process of the i -th asset is given by

$$q_t^i = q_0^i - \int_0^t v_s^i ds, \quad i \in \{1, \dots, d\}. \quad (10.5)$$

For each stock, we consider Gaussian price dynamics:

$$dS_t^i = \sigma^i dW_t^i,$$

where the Brownian motions $(W_t^i)_{t \in [0, T]}$ are such that $(S_t^1, \dots, S_t^d)_{t \in [0, T]}$ has a nonsingular covariance matrix Σ .

We treat all temporary impact as instantaneous and permanent impact as linear, hence irrelevant in the continuous-time case. In the single-asset case, if we trade Δq dollars in some small time interval of length Δt , and this costs $\lambda \Delta q / \Delta t$ times traded notional for some $\lambda > 0$, then the total cost in dollars per unit time is

$$\lambda(\delta q / \delta t)^2 \equiv c(\delta q / \delta t),$$

where $c(v) = \lambda v^2$. However, $c(\cdot)$ does not need to be quadratic, merely convex. In the multi-asset case, we simply set $c : \mathbb{R}^d \rightarrow \mathbb{R}_+$. As an aside, we note that these impact assumptions are an approximation that is only ever assumed to be valid within a certain regime. For example, if we repeatedly aggress with medium to large order sizes within a short time-frame, it is unrealistic to assume that the impact will revert instantly.

The trader attempts to maximize the expectation of utility of final wealth. Due to the nonlinear nature of most utility functions, this is inconvenient to work with. If time is discrete and if the multi-period

asset return vector follows an elliptical distribution, then there exists some constant $\kappa > 0$ for which it is equivalent to maximize the mean-variance quadratic form:

$$\mathbb{E}[w_T] - \frac{\kappa}{2}\mathbb{V}[w_T],$$

where w_T is the trader's wealth at terminal time. This is essentially the Markowitz prescription. Following Almgren and Chriss [10], most authors and practitioners replace the variance of final wealth $\mathbb{V}[w_T]$ with the integrated instantaneous variance, leading to the standard continuous-time approximation of the mean-variance form,

$$\int_0^T \left(-\frac{\kappa}{2} q_t^T \Sigma q_t - c(\dot{q}_t) \right) dt.$$

We shall now recast the maximization of the mean-variance form of the utility function as a problem in the calculus of variations. As in classical mechanics, it has both a Lagrangian and a Hamiltonian formulation, which are convex duals to each other. The Hamiltonian is related to the Lagrangian by the Legendre-Fenchel transform. In the following, we define what we mean by the term "smooth relaxation" which is not standard terminology.

Definition 10.2. Let $q_T \in \mathbb{R}^d$ be a desired final portfolio to be achieved at time T . The smooth relaxation problem associated to $c(\cdot)$, Σ , q_T is defined to be:

$$V(0, q_0) = \min_{q \in C^2([0, T], \mathbb{R}^d)} \int_0^T L(q_s, \dot{q}_s) ds \text{ subject to } q_0 = q_0, q_T = q_T, \quad (10.6)$$

where the (autonomous) Lagrangian is given by:

$$L(q, v) = c(v) + \frac{1}{2} \kappa (q - q_T)^\top \Sigma (q - q_T), \quad (10.7)$$

and $\kappa > 0$ is the risk-aversion constant.

The terminology of Definition 10.2 arises because Problem 10.1 is non-smooth and perhaps intractable, given that the action space of Problem 10.1 is discrete and quite large. In this sense (10.6) is a relaxation of the intractable problem to the space of twice-differentiable paths $C^2([0, T], \mathbb{R}^d)$. In this chapter, most of our results assume autonomous Lagrangians for simplicity. However, non-autonomous Lagrangians also arise in trading problems. For example, an alpha forecast which attenuates for large t entails a time-dependent linear term in (10.6).

An application of the Euler-Lagrange formula to (10.6) leads to a system of second-order equations. By a standard trick, a first-order system can be obtained if we introduce the so-called generalized momenta p , defined as

$$p := \partial_v L(q, \dot{q}). \quad (10.8)$$

If the conditions of the implicit function theorem are satisfied, we could solve (10.8) for \dot{q} , obtaining

$$\dot{q} = \phi(q, p),$$

for some function ϕ defined implicitly by (10.8). The Euler equation then takes the form

$$\dot{p} = \partial_q L(q, \dot{q}) = \partial_q L(q, \phi(q, p)) \equiv \psi(q, p),$$

where this defines ψ .

As the functions ϕ, ψ are algebraic (not involving derivatives), we have a system of $2d$ first-order ODEs given by

$$\dot{q} = \phi(q, p), \quad \dot{p} = \psi(q, p) \quad (10.9)$$

These equations can be expressed more symmetrically by introducing the Hamiltonian

$$H(q, p) := p\phi(q, p) - L(q, \phi(q, p)).$$

Equations (10.9) are equivalently written in a form known as *Hamilton's equations*:

$$\dot{q} = H_p(q, p), \quad \dot{p} = -H_q(q, p). \quad (10.10)$$

Suppose $c(v) = \frac{1}{2}v^\top \Lambda v$ where $\Lambda = \text{diag}(\lambda^1, \dots, \lambda^d)$ is a diagonal matrix. We assume that no trading is free of cost, so $\lambda^i > 0$ for all i . From (10.7) and (10.8), we see that the generalized momenta are $p = \Lambda \cdot q$ and hence algebraically solving, one has $\phi(q, p) = \Lambda^{-1}p$. The Hamiltonian is then

$$H(q, p) = \frac{1}{2}p\Lambda^{-1}p - \frac{1}{2}\kappa q^\top \Sigma q.$$

Hamilton's equations then become:

$$\dot{q} = \Lambda^{-1}p, \quad \dot{p} = \kappa \Sigma q, \quad (10.11)$$

and some technical computations lead to the following proposition.

Lemma 10.1. *The solution to (10.6) is given by*

$$q_t^* = (C^\top)^{-1} \Omega \left(e^{D^{\frac{1}{2}}(T-t)} - e^{-D^{\frac{1}{2}}(T-t)} \right) \left(e^{D^{\frac{1}{2}}T} - e^{-D^{\frac{1}{2}}T} \right)^{-1} \Omega^\top C^\top q_0, \quad (10.12)$$

where $\Sigma = CC^\top$ is the Cholesky decomposition of C and $\Omega D \Omega^\top$ is a spectral decomposition of the positive definite matrix $\frac{\kappa}{2}C^\top \Lambda^{-1}C$.

The trading curve (10.12) can be computed in advance, and corresponds to the order scheduling decision. We end this section by showing that the gradient of the value function (10.6) is equal to the generalized momenta p .

Theorem 10.1. *Let $V : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ be continuously differentiable in time and space such that:*

$$V(t, q_t) = - \min_{q \in C^2([t, T], \mathbb{R}^d)} \int_t^T L(q_s, v_s) ds \text{ subject to } q_0 = q_0, q_T = q_T, \quad (10.13)$$

where for all $s \in [0, T]$, $\dot{q}_s = v_s$ and $L(q, v)$ is separable.¹ The function V defined in (10.13) satisfies the Hamilton-Jacobi differential equation:

$$\partial_t V(t, q) + H(q, \nabla V) = 0, \quad (10.14)$$

where $H(q, p) = \sup_v \{ \langle p, v \rangle - L(q, v) \}$, with the singular final condition:

$$V(T, q) = \begin{cases} 0, & \text{if } q = q_T, \\ \infty, & \text{if } q \neq q_T. \end{cases}$$

Proof. Let q^* be the path that solves (10.13) on $[t, T]$ with initial condition $q_t^* = q$. By the dynamic programming principle, we have for $h > 0$

$$V(t, q) = - \int_t^{t+h} L(q_s^*, \dot{q}_s^*) ds + V(t+h, q_{t+h}^*). \quad (10.15)$$

As $V(t+h, q_{t+h}^*) = V(t, q) + \int_t^{t+h} (\partial_t V(s, q_s) + \langle \nabla V(s, q_s) \cdot \dot{q}_s^* \rangle) ds$, Equation (10.15) can be rewritten as

$$0 = \int_t^{t+h} (\partial_t V(s, q_s) + \langle \nabla V(s, q_s) \cdot \dot{q}_s^* - L(q_s^*, \dot{q}_s^*) \rangle) ds.$$

The conclusion follows from an application of the Bellman's optimality principle, see Dreyfus [101], which gives the desired Hamilton-Jacobi equation. \square

¹The value function (10.6) satisfies these hypothesis, as the associated optimal control (10.12) is continuously differentiable with respect to time and the Lagrangian of the problem is quadratic with respect to both of its variables.

The above theorem gives the desired interpretation of the generalized momenta p in terms of the value function. Indeed, along an optimal trajectory q^* , we have

$$\nabla V|_{x=q_t^*} = \partial_v L(q_t^*, \dot{q}_t^*) = p_t \quad \text{for all } t \in [0, T],$$

where the second equality is just the definition of p from (10.8).

10.3 From smooth relaxation to microstructure decision

In this section, we prove our main theorem, Theorem 10.2, which shows how a risk-neutral instantaneous-profit maximizer (or “myopic agent”) can achieve long-term optimality given a judicious choice of microstructure alpha model. In other words, there is a specific microstructure alpha model related to Hamilton’s generalized momenta, which, if used by a microstructure trader, encourages the trader to take positions that are optimal at a much longer horizon.

Suppose now that L is coercive of degree $r > 1$. One may prove that H coincides with the Fenchel conjugate of L :

$$H(q, p) = \sup_{v \in \mathbb{R}^d} \{ \langle p, v \rangle - L(q, v) \}. \quad (10.16)$$

It follows that H is convex in the p variable. We now restrict attention to autonomous and separable Lagrangians that take the form

$$L(q, v) = c(v) + f(q), \quad (10.17)$$

which includes the mean-variance example discussed before. Under the assumption (10.17), duality between the Lagrangian and Hamiltonian implies that the optimal instantaneous trade $\dot{q}^*(t)$ at each time t is the argument v which solves the maximization problem in (10.16). That is,

$$\dot{q}^* = \operatorname{argmax}_v \{ \langle p, v \rangle - c(v) \}. \quad (10.18)$$

where (10.17) implies we can safely omit the term $f(q)$.

One can interpret (10.18) as optimization in a risk-neutral world. Indeed, suppose a risk-neutral agent had a vector of expected returns, which happened to exactly equal the vector of generalized momenta, p , and sought only to maximize net profit, irrespective of risk. In that case, (10.18) is the problem faced by this agent. This agent can be considered “myopic” because any information concerning more than one period ahead is available to the agent only indirectly, insofar as p depends on the rest of the trading path. This intuition is related to the interpretation of the generalized momenta as the gradient of the value function. Indeed, for a myopic risk-neutral investor who does not face transaction costs, the “value function” of a position of size x is simply the expected profit over the next period, i.e., $V(t, q) = r \cdot q$ where r is the vector of expected returns, but then $\nabla V = r$. The following definition characterizes what a myopic agent is. Moreover, we emphasize that, throughout the chapter, $V(\cdot, \cdot)$ corresponds to the value function of a long-term optimization problem (whose gradient provide the effective microstructure alphas), whereas $W(\cdot, \cdot)$ denotes the value function of the myopic agent.

Definition 10.3. *A myopic agent with microstructure alphas p is defined to be a risk-neutral trader seeking to maximize instantaneous net profit by choosing trading rate given by*

$$\dot{q}^* = \operatorname{argmax}_{v \in \mathbb{R}^d} \{ \langle p, v \rangle - c(v) \}, \quad (10.19)$$

where p denotes a set of microstructure alphas, as in Definition 10.1. The value function of the myopic trader at time $t \in [0, T]$ for an inventory $q \in \mathbb{R}^d$ is defined as

$$W(t, q) = \int_t^T (\langle p_s, v_s^* \rangle - c(v_s^*)) ds, \quad (10.20)$$

where v^* is defined as the solution of (10.19).

In other words, the value function of a myopic trader defined in (10.20) is simply the sum of his instantaneous trading gains over time. The following proposition shows that a myopic trader sending market orders only has to choose $p = \nabla V$ in order to minimize the error between his value function and the long-term objective function V .

Theorem 10.2. *Assume that L takes the separable form (10.17). A myopic agent with instantaneous cost function $c(\cdot)$ must choose microstructure alphas $p = \nabla V$ in order to minimize the absolute error between his value function and the long-term objective function V defining the trading schedule. More precisely, for all $(t, q) \in [0, T] \times \prod_{i=1}^d [\min(q_0^i, q_T^i), \max(q_0^i, q_T^i)]$,*

$$|W(t, q) - V(t, q)| \leq \kappa(T - t) \frac{|q - q_T|^\top \Sigma |q - q_T|}{2},$$

and we have the uniform bound

$$\sup_{(t, q) \in [0, T] \times \prod_{i=1}^d [\min(q_0^i, q_T^i), \max(q_0^i, q_T^i)]} |W(t, q) - V(t, q)| \leq \kappa T \frac{|q_0 - q_T|^\top \Sigma |q_0 - q_T|}{2},$$

where $|q_0 - q_T| = (|q_0^1 - q_T^1|, \dots, |q_0^d - q_T^d|)$.

Proof. The myopic trader aims at minimizing $c(v) - pv$, where for the moment p remains undetermined, at each trading time. Over $[t, T]$ the trader's problem can be written:

$$W(t, q) = \max_v \int_t^T [\langle p_s, v_s \rangle - c(v_s)] ds.$$

We choose quadratic costs $c(v) = \frac{\eta}{2} \|v\|_2^2$, where $\|\cdot\|_2$ is the Euclidian norm and the first order condition with respect to v gives $W(t, q) = \int_t^T \frac{\|p_s\|_2^2}{2\eta} ds$. On the other hand, the value function $V(t, q)$ becomes

$$\int_t^T \left(\frac{1}{2\eta} \|\partial_v L(q_s^*, v_s^*)\|_2^2 + \frac{1}{2} \kappa (q_s^* - q_T)^\top \Sigma (q_s^* - q_T) \right) ds,$$

where q^* is defined by (10.12) and v^* is its derivative with respect to time. Therefore,

$$|V(t, q) - W(t, q)| = \left| \int_t^T \left(\frac{1}{2\eta} \|\partial_v L(q_s^*, v_s^*)\|_2^2 + \frac{1}{2} \kappa (q_s^* - q_T)^\top \Sigma (q_s^* - q_T) - \frac{\|p_s\|_2^2}{2\eta} \right) ds \right|,$$

and the minimum with respect to p is attained at $p_s = \partial_v L(q_s^*, v_s^*) = \nabla V(s, q_s^*)$, because p must depend only of the instantaneous trading rate and L is separable and additive. The bounds are obtained easily by definition of the space of inventories. \square

This simple result has several important consequences. Suppose one wants to avoid the use of optimal control and still wants to follow the Almgren-Chriss trading curve. In that case, one can simply solve the static optimization problem (10.18) at discrete times (the times of trading), using $p = \nabla V$. Equation (10.18) does not give a full set of instructions for the trader with a long-term trading schedule who has to interact with a continuous limit order book market, but it can serve as a guide. Indeed, the order routing problem, treated notably in Cont and Kukanov [81], takes into account the possibility to send limit, market, or cancel orders to several liquidity venues, depending on their spread and imbalance. Stochastic control appears to be inefficient for this problem, as one needs to solve a high-dimensional Hamilton-Jacobi-Bellman equation. Methods involving deep reinforcement learning have been developed for optimal trading, see for example Chapter 9 of this thesis, but they lead to high computation time, especially if one wants to deal with a portfolio of assets traded on several venues. The advantage of the methodology presented in this chapter is that one can avoid optimal control and solve a simple static optimization problem to determine the optimal action at each discrete trading time.

Remark 46. *The bounds on the absolute error between the value function of the myopic agent and the long-term objective function enable to compute the accuracy of the myopic trader. For example, take the liquidation over $T = 1$ day of $q_0^1 = q_0^2 = 2000$ shares of 2 assets, with correlation $\rho = 0.6$ and daily volatilities $\sigma^1 = 0.015$ and $\sigma^2 = 0.02$. The absolute error between the two value functions is uniformly bounded in time and inventories by 2×10^{-2} .*

In Theorem 10.2, the myopic trader does not consider the properties of an order book, such as the possibility to submit limit and market orders or to wait. In this case, the optimal effective microstructure alpha p (in the sense of minimization of the error with respect to the Almgren-Chriss value function V) should not be equal to the generalized momenta because we add microstructure effects for the myopic trader that are not present in the trading schedule represented by ∇V .

With a sufficiently simple fill model, the myopic trader's problem dealing with microstructure effects can be solved in closed form. This closed-form expression (see Example 10.1) illustrates the contrast between the two possible decisions the trader must face, as mentioned above.

Example 10.1. *For the sake of readability, we assume $d = 1$. Suppose that the myopic trader can choose between submitting a limit order (with fill probability $0 < f < 1$) or a market order (with the cost of crossing the spread equal to $s > 0$). The myopic trader's optimization problem is:*

$$\sup_{v \in \mathbb{R}} \left\{ \left(pv - \frac{\eta}{2} v^2 - sv \right) \mathbf{1}_{\{pv - \frac{\eta}{2} v^2 - sv > f(pv - \frac{\eta}{2} v^2)\}} + f \left(pv - \frac{\eta}{2} v^2 \right) \mathbf{1}_{\{pv - \frac{\eta}{2} v^2 - sv < f(pv - \frac{\eta}{2} v^2)\}} \right\}.$$

Computations lead to the following decisions:

$$\text{passive order if } p < \frac{2s}{1-f}, \text{ aggressive otherwise,}$$

and the optimal microstructure alpha is given by:

$$p_t^* = \frac{1}{\sqrt{f}} |\nabla V(t, q_t^*)| \operatorname{sgn}(\nabla V(t, q_t^*)) \text{ if } p_t^* < \frac{2s}{1-f}, \quad p_t^* = |\nabla V(t, q_t^*)| \operatorname{sgn}(\nabla V(t, q_t^*)) + s \text{ otherwise.}$$

The use of the generalized momenta as effective microstructure alpha has a wide range of practical implications. First, it offers a way to bridge the gap between order placement decisions and scheduling decisions, usually decoupled in practice. Second, the microstructure formulation helps to tackle classic optimal control on limit order books. For example, a realistic optimal trading framework dealing with a portfolio of assets on several liquidity venues is in practice intractable due to the dimensionality of the problem. In the method presented in this chapter, the optimal controls of the trader (that is, the volume sent on each venue for each asset by the mean of limit and market orders) are derived through a simple static optimization problem, which can be solved for a large number of assets on a large number of venues. The convergence through the trading schedule is guaranteed by choice of the effective microstructure alpha p .

Remark 47. *Note that, in the framework of Example (10.1), if the spread s tends to zero and the filling probability f tends to one, we recover the framework of a myopic sending market orders only, and the optimal effective microstructure alpha is given by ∇V .*

This method can easily handle the increasing complexity coming from the microstructure effects (short term alpha, imbalance, and spread of each venue, etc.). In the next section, we present a general microstructure trading model taking into account the main stylized facts combining order placement and order routing of a portfolio of assets. We show that using the method proposed in this chapter can be applied to solve in practice two important problems in systematic trading: the multi-asset, multi-venue optimal trading problem and the multi-asset, multi-venue optimal market-making problem.

10.4 A general microstructure trading model with long-term trading schedule

10.4.1 Optimal trading

We first provide a definition of the problem.

Definition 10.4. *Consider an agent trading a portfolio of correlated assets, where each asset is listed on one or more liquidity venues. The multi-asset, multi-venue optimal trading problem consists in determining at a given time and for each asset, the optimal quantity to buy or sell on each venue, for given market conditions and a pre-computed trading schedule, as well as the optimal limit at which such quantity should be posted.*

The framework described here is inspired by Chapter 9. Consider a trading schedule for $d \geq 1$ assets $q^* \in \mathbb{R}^d$ (the Almgren-Chriss trading schedule described in Equation (10.12), for example) with associated value function $V(t, q_t^*)$. For each asset $i \in \{1, \dots, d\}$, the trader splits his limit and market orders between $N^i \geq 1$ liquidity venues. We assume that he wants to unwind the portfolio so that $q_T = 0^d$. For all $i \in \{1, \dots, d\}, n \in \{1, \dots, N^i\}$, the order book of the asset i on the venue n is characterized by the following quantities:

- the bid-ask spread process $(\psi_t^{i,n})_{t \in [0, T]}$ taking values in the state space $\bar{\psi}^{i,n} = \{\delta^{i,n}, \dots, J\delta^{i,n}\}$,
- the imbalance process $(I_t^{i,n})_{t \in [0, T]}$ taking values in the state space $\bar{I}^{i,n} = \{I_1^{i,n}, \dots, I_K^{i,n}\}$,

where $J, K \in \mathbb{N}$ denote the number of possible spreads and imbalances respectively and $\delta^{i,n}$ stands for the tick size of i -th asset on the n -th venue. Note that the dynamics are unspecified, meaning that any continuous-time stochastic process with discrete values can be considered for the purpose of simulation.

Definition 10.5. *A market regime is, for asset $i \in \{1, \dots, d\}$, a set of spread and imbalance values on the different venues $n \in \{1, \dots, N^i\}$.*

We define the sets $\Psi = \{\Psi_1, \dots, \Psi_{\#\Psi}\}, \mathcal{I} = \{\mathcal{I}_1, \dots, \mathcal{I}_{\#\mathcal{I}}\}$ of disjoint intervals, representing different market regimes of interest in terms of spreads and imbalances.

Example 10.2. *Assume $d = 1$ and for all $n \in \{1, \dots, N\}$, $\delta^n = \delta$. The set $\Psi = \{\delta, \{2\delta, 3\delta\}, \{4\delta, 5\delta\}\}$ denotes three spread regimes: low (one tick), medium (two or three ticks), and high (four or five ticks).*

Example 10.3. *Assume $d = 1$ and for all $n \in \{1, \dots, N\}$ and $k \in \{1, \dots, K\}$ that $I_k^n = I_k$. In this case the set $\mathcal{I} = \{[-1, -0.66], (-0.66, -0.33], (-0.33, 0.33], (0.33, 0.66], (0.66, 1]\}$ denotes five regimes of imbalance: low (-33% to 33%), medium on the ask (resp. bid) from 33% to 66% (resp. from -66% to -33%) and high on the ask (resp. bid) from 66% to 100% (resp. from -100% to -66%).*

The number of, possibly partially, filled ask orders on the asset i in the venue n is modeled by a Cox process denoted by $N^{i,n}$ with intensities $\lambda^{i,n}(\psi_t^i, I_t^i, p_t^{i,n}, \ell_t^i)$ where $p_t^{i,n} \in Q_\psi^{i,n}$ represent the limit at which the trader sends a limit order of size $\ell_t^{i,n}$, and

$$Q_\psi^{i,n} = \{0, 1\} \text{ if } \psi^{i,n} = \delta^{i,n}, \text{ and } \{-1, 0, 1\} \text{ otherwise.}$$

Practically, on asset i , for $n \in \{1, \dots, N^i\}$, when the spread is equal to the tick size, the trader can post at the first best limit ($p^{i,n} = 0$) or the second best limit (if $p^{i,n} = 1$). When the spread is equal to two ticks or more, the trader can either create a new best limit ($p^{i,n} = -1$) or post at the best or the second best limit as previously. The arrival intensity of a buy market order at time t on the venue $n \in \{1, \dots, N^i\}$ for asset i at the limit $p \in Q_\psi^{i,n}$, given a couple $(\psi_t^i, I_t^i) = \mathbf{m}$ of spread and imbalance on each venue, is equal to $\lambda^{i,n, \mathbf{m}, p} > 0$. When the trader posts limit orders of volume $\ell_t^{i,n}$ on the n -th venue for $n \in \{1, \dots, N^i\}$, the probability that it is executed is equal to $f^\lambda(\ell_t^i)$, where $f^\lambda(\cdot) \in [0, 1]$ is a continuously differentiable function, decreasing with respect to each of its coordinate. Therefore,

the arrival intensity of an ask market order filling the buy limit order of the trader for asset i on the n -th venue at the limit $p_t^{i,n}$, given spread and imbalance (ψ_t^i, I_t^i) is a multi-regime function defined by

$$\lambda^{i,n}(\psi_t^i, I_t^i, p_t^{i,n}, \ell_t^i) = f^\lambda(\ell_t^i) \sum_{\mathbf{m} \in \mathcal{M}^i, p \in Q_\psi^{i,n}} \lambda^{i,n,\mathbf{m},p} \mathbf{1}_{\{(\psi_t^i, I_t^i) \in \mathbf{m}, p_t^{i,n} = p\}}, \quad (10.21)$$

where $\mathcal{M}^i = \Psi^{N^i} \times \mathcal{I}^{N^i}$. Moreover, we allow for partial execution, the fact of which we represent by random variables $\epsilon_t^{i,n} \in [0, 1]$. The proportion of executed volume for limit orders in each venue depends on the spread and the imbalance in all N^i venues for asset i , as well as the volume and the limit of the order chosen by the trader. We assume a categorical distribution with $R > 0$ different execution proportions $\omega^r, r \in \{1, \dots, R\}$ for each venue with $\mathbb{P}(\epsilon_t^{i,n} = \omega^r) = \rho^{i,n,r}(\psi_t^i, I_t^i, p_t^{i,n}, \ell_t^i)$, where

$$\rho^{i,n,r}(\psi_t^i, I_t^i, p_t^{i,n}, \ell_t^i) = f^\rho(\ell_t^i) \sum_{\mathbf{m} \in \mathcal{M}^i, p \in Q_\psi^{i,n}} \rho^{i,n,\mathbf{m},p,r} \mathbf{1}_{\{(\psi_t^i, I_t^i) \in \mathbf{m}, p_t^{i,n} = p\}}, \quad (10.22)$$

where $f^\rho(\cdot)$ is a continuously differentiable function, decreasing with respect to each of its coordinate.

We allow for the execution of market orders (denoted by a point process $(J_t^{i,n})_{t \in [0, T]}$) on each venue of size $(v_t^{i,n})_{t \in [0, T]} \in [0, \bar{v}]$ where $\bar{v} > 0$ and $J_t^{i,n} = J_{t^-}^{i,n} + 1$. We assume that market orders are always fully executed but this assumption can be relaxed easily. As each asset must be bought or sold, we define $\Delta = (\Delta^1, \dots, \Delta^d)$ where for $i \in \{1, \dots, d\}$, $\Delta^i = 1$ if $q_0^i > 0$, -1 otherwise. The inventory process on each asset is defined by

$$q_t^i = q_0^i - \Delta^i \sum_{n=1}^{N^i} \left(\int_0^t \ell_s^{i,n} \epsilon_s^{i,n} dN_s^{i,n} + \int_0^t v_s^{i,n} dJ_s^{i,n} \right).$$

The myopic trader has an effective microstructure $(p_t^{\text{eff},i})_{t \in [0, T]}$ in order to follow the pre-computed execution curve q_t^{*i} on each asset, but also a short-term alpha $(p_t^{\text{short},i})_{t \in [0, T]}$ which is a function of the current spread and imbalance \mathbf{m}^i of all the venues where asset i is listed.

Remark 48. *The microstructure alpha considered for each asset is the sum of a direct microstructure alpha depending on the market regimes and an effective microstructure alpha that gives a signal to follow the long-term objective function. The sum of these two terms gives the magnitude of the buy or sell signal. For example, suppose $p^{\text{eff},i}$ is small and positive, indicating that filling a buy order would be a slight improvement to the value function. Suppose with $p^{\text{eff},i}$ alone, the system would have recommended a passive buy order. Now suppose a strongly-positive microstructure alpha, denoted $p^{\text{short},i}$, is also present; then the combination $p^{\text{eff},i} + p^{\text{short},i}$ in place of $p^{\text{eff},i}$ should recommend a more aggressive action, such as a spread-crossing buy order.*

Finally, the cost function of a limit order of size ℓ at limit p on venue n for asset i is defined as $c^{i,n,\mathcal{L}}(\ell, p)$ and $c^{i,n,\mathcal{M}}(v)$ for the cost function of a market order of size v on venue n for asset i .

The myopic trader acts at discrete times and at time $t \in [0, T]$ for $(\psi_t^i, I_t^i) \in \mathbf{m}^i, i \in \{1, \dots, d\}$, his optimization problem is

$$\begin{aligned} \max \left\{ \sup_{\ell, p} \left\{ \sum_{i=1}^d \sum_{n=1}^{N^i} \lambda^{i,n}(\mathbf{m}^i, p^{i,n}, \ell^i) \mathbb{E} \left[(p^{\text{eff},i} + p^{\text{short},i,n}(\mathbf{m}^i)) \ell^{i,n} \epsilon^{i,n} - c^{i,n,\mathcal{L}}(\ell^{i,n}, p^{i,n}) \right] \right\}, \right. \\ \left. \sup_v \left\{ \sum_{i=1}^d \sum_{n=1}^{N^i} (p^{\text{eff},i} + p^{\text{short},i,n}(\mathbf{m}^i)) v^{i,n} - c^{i,n,\mathcal{M}}(v^{i,n}) \right\} \right\}, \quad (10.23) \end{aligned}$$

where the expectation is taken with respect to the variables $\epsilon^{i,n}$ for all $i \in \{1, \dots, d\}, n \in \{1, \dots, N^i\}$. This is a simple static optimization which can be solved for a large number of assets and venues using a multidimensional root-finding method. The output is, for each state \mathbf{m}^i , the optimal volumes and limits $\ell^{*,i,n}(\mathbf{m}^i), p^{*,i,n}(\mathbf{m}^i)$ for each asset on each liquidity venue. We define the value function of the myopic trader at time t as

$$W(t, \psi, I, q) = \int_t^T \max \left\{ \sum_{i=1}^d \sum_{n=1}^{N^i} \lambda^{i,n}(\psi_s^i, I_s^i, p_s^{*,i,n}, \ell_s^{*,i,n}) \mathbb{E} \left[(p_s^{\text{long},i} + p_s^{\text{short},i,n}(\psi_s^i, I_s^i)) \ell_s^{*,i,n} \epsilon_s^{i,n} - c^{i,n,\mathcal{L}}(\epsilon_s^{i,n} \ell_s^{*,i,n}, p_s^{*,i,n}) \right], \sum_{i=1}^d \sum_{n=1}^{N^i} (p_s^{\text{long},i} + p_s^{\text{short},i,n}(\psi_s^i, I_s^i)) v_s^{*,i,n} - c^{i,n,\mathcal{M}}(v_s^{*,i,n}) \right\} ds$$

where $\psi_t = \psi, I_t = I, q_t = q$. As in Theorem 10.2, the myopic trader has now to choose the long-term alpha p^{eff} to match the trading schedule q^* . This leads to the following optimization setting for all $\mathbf{m}^i \in \mathcal{M}^i$:

$$\begin{aligned} & \max \left\{ \sup_{\ell, p} \left\{ \sum_{i=1}^d \sum_{n=1}^{N^i} \lambda^{i,n}(\mathbf{m}^i, p^{i,n}, \ell^i) \mathbb{E} \left[(p^{\text{eff},i} + p^{\text{short},i,n}(\mathbf{m}^i)) \ell^{i,n} \epsilon^{i,n} - c^{i,n,\mathcal{L}}(\epsilon^{i,n} \ell^{i,n}, p^{i,n}) \right] \right\}, \right. \\ & \quad \left. \sup_v \left\{ \sum_{i=1}^d \sum_{n=1}^{N^i} (p^{\text{eff},i} + p^{\text{short},i,n}(\mathbf{m}^i)) v^{i,n} - c^{i,n,\mathcal{M}}(v^{i,n}) \right\} \right\}, \\ & p^{\text{eff}} = \operatorname{argmin}_{p^{\text{eff}}} |V(\cdot, q^*) - W(\cdot, \psi, I, q)|. \end{aligned} \tag{Opt-Trd}$$

In this general framework, order scheduling with a long-term target is easily tractable even for a large portfolio of assets, as the trader has to solve a static optimization problem at each trading time. For a parsimonious model of filling probabilities, the effective microstructure alpha can be computed in closed form. Note that each time a fill is received that changes the portfolio holdings, and/or each time a significant amount of time passes, the effective microstructure alpha p^{eff} must be recomputed.

Remark 49. *The methodology presented in this chapter leads to entirely tractable optimization problems, even for a large number of assets. This is the case when we have a closed-form solution for the long-term value function V , which can be computed quickly. It also suggests an approximation of the effective microstructure alpha, that is to take $p_t^{\text{eff},i} = \nabla_i V(t, q^*)$. This heuristics will be used in the next section to solve a different control problem.*

10.4.2 Market-making

The great advantage of the framework presented in this chapter is that it avoids the use of optimal control to tackle optimal trading problems. The trader solves a simple static optimization problem, and the use of the generalized momenta as a long-term alpha plays the role of the trading schedule. Similar ideas can be applied to the market-making problem, with some minor changes.

Definition 10.6. *Consider an agent trading on a portfolio of correlated assets, where each of them is listed on one or several liquidity venues. His goal is to earn the difference between the bid and ask prices (the bid-ask spread) while keeping his inventory close to zero to avoid an unwanted large exposure and be forced to buy at a higher price or sell at a lower price in order to unwind this position. The multi-asset, multi-venue optimal market-making problem consists in deriving at a given time, for each asset, the optimal quantity to buy or sell in each venue, for given market conditions, as well as the optimal limit at which such quantity should be posted, with an inventory vector mean-reverting around zero or some predetermined target.*

The market-making problem has been introduced in the financial literature by Ho and Stoll [156], Glosten and Milgrom [131]. Ho and Stoll presented a framework to tackle inventory management,

while Grossman and Miller proposed a 3 periods model that encompassed both market-makers and final customers, enabled them to understand what happens at equilibrium, and contributed to the important literature on the price formation process. The seminal reference of the recent literature on market-making is the work of Avellaneda and Stoikov [26], who proposed a stochastic control framework to tackle the quoting and inventory management problems. Since then, a vast literature on optimal market-making has emerged, basically adding many features to the Avellaneda and Stoikov [26] framework, see for example Cartea, Jaimungal, and Ricci [69], Guéant, Lehalle, and Fernandez-Tapia [142] and the two textbooks Cartea, Jaimungal, and Penalva [70], Guéant [137]. These works deal with single asset market-making, and the considered framework is more suitable for OTC markets rather than order-driven markets. The problem of multi-asset market-making, dealing with the curse of dimensionality, has been addressed via deep reinforcement learning methods, see for example Guéant and Manziuk [140]. Models for optimal market-making in limit order books have been developed for the single asset case, see Guilbaud and Pham [143], for example. All these models suffer from the same problem when dealing with a portfolio of assets: solving a high-dimensional Hamilton-Jacobi-Bellman equation makes the problem almost intractable in practice. In this section, we propose an adaptation of the previously described heuristics to tackle the multi-asset market-making problem in limit order books.

10.4.2.1 The long-term objective function

Our methodology to solve optimal control problems in high dimension relies on the fact that the effective microstructure alphas come from a long-term objective function computed analytically. This is the case of the Almgren-Chriss trading curve, which hides the microstructure effects that are incorporated in the myopic optimization problem. However, as stated previously, the main constraint of the optimal market-making problem is that, even for market-making on OTC markets, the value function's computation is very time-consuming. We propose to use the gradient of an approximation of the value function of the optimal market-making problem on OTC markets as the effective microstructure alphas for the optimal market-making problem in order books. To this end, we borrow the OTC framework of Bergault, Evangelista, Guéant, and Vieira [42] and recall their modeling assumptions briefly.

For $i \in \{1, \dots, d\}$, the reference price of asset i is modeled by a process S_t^i with dynamics

$$dS_t^i = \sigma^i dW_t^i,$$

where (W_t^1, \dots, W_t^d) is a d -dimensional Brownian motion with variance-covariance matrix Σ . At each $t \in [0, T]$, the market-maker chooses the prices $P_t^{i,b}, P_t^{i,a}$ at which she is ready to buy/sell each asset i . These prices are given by

$$P_t^{i,b} = S_t^i - \delta_t^{i,b}, \quad P_t^{i,a} = S_t^i + \delta_t^{i,a},$$

where $\delta_t = (\delta_t^{1,b}, \delta_t^{1,a}, \dots, \delta_t^{d,b}, \delta_t^{d,a})$ are the control processes of the market-maker corresponding to the bid and ask spreads set on each asset i . For $i \in \{1, \dots, d\}$, the point processes $N_t^{i,b}, N_t^{i,a}$ denote the total number of bid and ask transactions between 0 and t on asset i . Their intensities are given by $\Lambda^{i,b}(\delta_t^{i,b}), \Lambda^{i,a}(\delta_t^{i,a})$ where the functions $\Lambda^{i,b}, \Lambda^{i,a}$ satisfy some technical conditions, see Bergault, Evangelista, Guéant, and Vieira [42] for details. These conditions are sufficiently general to allow the use of several form of intensity such as exponential, logistic, SU Johnson etc.

The transaction size for asset i is constant and denoted by z^i , and the inventory process of the market-maker for asset i is

$$dq_t^i = z^i (dN_t^{i,b} - dN_t^{i,a}), \quad q_t = (q_t^1, \dots, q_t^d)^\top.$$

The cash process of the market-maker has the following dynamics:

$$dX_t = \sum_{i=1}^d (P_t^{i,a} dN_t^{i,a} - P_t^{i,b} dN_t^{i,b}).$$

The optimization problem of the market-maker is defined by

$$\sup_{\delta} \mathbb{E} \left[X_T + \sum_{i=1}^d q_T^i S_T^i - \frac{\gamma}{2} \int_0^T q_s^\top \Sigma q_s dt \right],$$

and simple computations² give the associated value function at time t for a given inventory vector $q_t = q$:

$$V(t, q) = \sup_{\delta} \mathbb{E}_t \left[\sum_{i=1}^d \int_t^T \left(\delta_s^{i,a} \Lambda^{i,a}(\delta_s^{i,a}) + \delta_s^{i,b} \Lambda^{i,b}(\delta_s^{i,b}) - \frac{\gamma}{2} q_s^\top \Sigma q_s \right) dt \right], \quad (10.24)$$

where \mathbb{E}_t denote the conditional expectation with respect to the canonical filtration at time t and $\gamma > 0$ is the risk-aversion of the market-maker. He wishes to maximize the sum of his cash process and the mark-to-market value of his inventory. The running penalty forces him to mean-revert his inventories to zero. We now state the main proposition of Bergault, Evangelista, Guéant, and Vieira [42] that provides a closed form approximation of $V(t, q)$, and refer to this article for the proof.

Proposition 10.1. *Define the functions*

$$H^{i,b}(p) = \sup_{\delta} \{ \Lambda^{i,b}(\delta)(\delta - p) \}, \quad H^{i,a}(p) = \sup_{\delta} \{ \Lambda^{i,a}(\delta)(\delta - p) \},$$

and the constants $\alpha_j^{i,b} = (H^{i,b})^j(0)$, $\alpha_j^{i,a} = (H^{i,a})^j(0)$, where the superscript $j \in \{0, 1, 2\}$ denote the derivative of order j . Define also for $k \in \mathbb{N}$

$$\begin{aligned} \Delta_{j,k}^{i,b} &= \alpha_j^{i,b} (z^i)^k, & \Delta_{j,k}^{i,a} &= \alpha_j^{i,a} (z^i)^k, \\ V_{j,k}^b &= \left(\Delta_{j,k}^{1,b}, \dots, \Delta_{j,k}^{d,b} \right), & V_{j,k}^a &= \left(\Delta_{j,k}^{1,a}, \dots, \Delta_{j,k}^{d,a} \right), \\ D_{j,k}^b &= \text{diag} \left(\Delta_{j,k}^{1,b}, \dots, \Delta_{j,k}^{d,b} \right), & D_{j,k}^a &= \text{diag} \left(\Delta_{j,k}^{1,a}, \dots, \Delta_{j,k}^{d,a} \right). \end{aligned}$$

Then if $\alpha_2^{i,b} + \alpha_2^{i,a} > 0$, the value function of the optimal control problem (10.24) can be approximated by the function

$$\tilde{V}(t, q) = -q^\top A(t)q - q^\top B(t) - C(t), \quad (10.25)$$

where $A : [0, T] \rightarrow \mathcal{S}_d^{++}$, $B : [0, T] \rightarrow \mathbb{R}^d$ and $C : [0, T] \rightarrow \mathbb{R}$ are deterministic functions given by

$$\begin{aligned} A(t) &= \frac{1}{2} D_+^{-\frac{1}{2}} \hat{A} \left(e^{\hat{A}(T-t)} - e^{-\hat{A}(T-t)} \right) \left(e^{\hat{A}(T-t)} + e^{-\hat{A}(T-t)} \right)^{-1} D_+^{-\frac{1}{2}}, \\ B(t) &= -2e^{-2 \int_t^T A(u) D_+ du} \int_t^T e^{2 \int_s^T A(u) D_+ du} A(s) \left(V_- + D_- \mathcal{D}(A(s)) \right) ds, \\ C(t) &= -\text{Tr}(D_{0,1}^b + D_{0,1}^a)(T-t) - \text{Tr} \left((D_{1,2}^b + D_{1,2}^a) \int_t^T A(s) ds \right) - V_-^\top \int_t^T B(s) ds \\ &\quad - \frac{1}{2} \int_t^T \mathcal{D}(A(s))^\top (D_{2,3}^b + D_{2,3}^a) \mathcal{D}(A(s)) ds - \frac{1}{2} \int_t^T B(s)^\top D_+ B(s) ds \\ &\quad - \int_t^T B(s)^\top D_- B(s) ds, \end{aligned}$$

²It can be shown by simple change of variables that the value function of this control problem is only a functional of the time and the inventories.

with

$$D_+ = D_{2,1}^b + D_{2,1}^a, D_- = D_{2,2}^b - D_{2,2}^a, V_- = V_{1,1}^b - V_{1,1}^a, \hat{A} = \sqrt{\gamma} (D_+^{\frac{1}{2}} \Sigma D_+^{\frac{1}{2}})^{\frac{1}{2}},$$

and \mathcal{D} is the linear operator mapping a matrix onto the vector of its diagonal and \mathcal{S}_d^{++} is the set of $d \times d$ definite positive matrix.

The approximated value function (10.25) is quadratic, therefore, sub-differentiable with respect to the vector of inventories q and the deterministic functions $A(t), B(t), C(t)$ can be computed in closed form. It takes into account the main property of a “high-level” multi-asset market-making problem, that is, the correlation structure between the assets. By analogy with Section 10.4.1, its sub-gradient can be chosen as an effective microstructure alpha to mean revert toward a flat inventory.

Remark 50. *In order to compute efficiently the value function in (10.25), note that the matrix A can be diagonalized and therefore approximated with a principal component analysis. The expressions of A, B, C do not provide intuition about the long-term behavior of the market-maker. However, in an asymptotic framework, that is when $T \rightarrow +\infty$, we obtain*

$$\begin{aligned} A &\xrightarrow{T \rightarrow +\infty} \frac{1}{2} \sqrt{\gamma} \Gamma, \\ B &\xrightarrow{T \rightarrow +\infty} -D_+^{-\frac{1}{2}} \hat{A} \hat{A}^+ D_+^{-\frac{1}{2}} (V_- + \frac{1}{2} \sqrt{\gamma} D_- \mathcal{D}(\Gamma)), \end{aligned}$$

where $\Gamma = D_+^{-\frac{1}{2}} (D_+^{\frac{1}{2}} \Sigma D_+^{\frac{1}{2}})^{\frac{1}{2}} D_+^{-\frac{1}{2}}$ and \hat{A}^+ is the Moore-Penrose generalized inverse of \hat{A} . If we perform a principal component analysis on the variance-covariance matrix Σ , we observe that the buy or sell signal (depending on the sign of the inventories) coming from the sub-gradient of $V(t, q)$ is an increasing function of the eigenvalues of Σ and the risk-aversion parameter γ . Thus, choosing the sub-gradient of $V(t, q)$ as an effective microstructure alpha should provide a mean-reverting signal for a myopic agent, taking into account the correlation between the assets.

10.4.2.2 Multi-asset multi-venue optimal market-making in limit order book

We now derive the solution to the multi-asset multi-venue optimal market-making problem using our heuristics. We take the same modeling notation as in 10.4.1. Assuming bid and ask symmetry for sake of simplicity, we introduce the processes $N^{i,n,b}, N^{i,n,a}$ to model the number of (possibly partially-filled) bid and ask orders on the asset i in the venue n of intensity $\lambda^{i,n}(\psi^i, I^i, p^{i,n,b}, \ell^{i,b})$ for the bid side and $\lambda^{i,n}(\psi^i, I^i, p^{i,n,a}, \ell^{i,a})$ for the ask side, where the function λ is defined by (10.21). The quantity $\ell^{i,n,b}$ (resp. $\ell^{i,n,a}$) is the volume sent on the n -th venue of the i -th asset on the bid (resp. ask) side. The quantity $p^{i,n,b}$ (resp. $p^{i,n,a}$) is the limit chosen by the market-maker on the n -th venue of the i -th asset on the bid (resp. ask) side to send a limit order. The distribution of the random variables $\epsilon_t^{i,n,b}$ and $\epsilon_t^{i,n,a}$ are defined as in Equation (10.22). The market-maker can also send market orders on the bid and ask sides (denoted by point processes $J_t^{i,n,b}, J_t^{i,n,a}$) on each venue of size $v_t^{i,n,b}, v_t^{i,n,a}$. Its inventory process on each asset is defined by

$$q_t^i = \left(\int_0^t \ell_s^{i,n,b} \epsilon_s^{i,n,b} dN_s^{i,n,b} + \int_0^t v_s^{i,n,b} dJ_s^{i,n,b} \right) - \left(\int_0^t \ell_s^{i,n,a} \epsilon_s^{i,n,a} dN_s^{i,n,a} + \int_0^t v_s^{i,n,a} dJ_s^{i,n,a} \right) \in \mathbb{Z}.$$

The problem faced by a market-maker is slightly different compared to a classic trader. While the trader must follow a predetermined target, the market-maker’s inventory must revert toward zero. Therefore, we seek a long-term alpha that gives a signal to our myopic market-maker of the form “sell for high inventory, buy for low inventory” with different type of aggressiveness (limit or market order) depending on the level of inventory. Contrary to optimal execution, there is no optimal inventory in market-making problems at a given time $t \in [0, T]$, which explains the dependence of the long-term alpha on the current inventory.

As stated previously, the effective microstructure alphas should be the gradient of a value function corresponding to a “high-level” multi-asset market-making problem (which hides the microstructure effects). This value function should be in closed form to recompute the gradient quickly when the market-maker trades and too much time passed. Thus, we propose to use the sub-gradient of the value function (10.25) corresponding to an approximation of the multi-asset market-making value function in OTC markets as a proxy for effective microstructure alphas used by a market-maker acting on a portfolio of assets listed on several order book platforms. By analogy with Section 10.4.1, at each time step t , for an inventory vector $q \in \mathbb{R}^d$, the market-maker solves the following optimization problem:

$$\begin{aligned} \max \left\{ \sup_{\ell,p} \left\{ \sum_{i=1}^d \sum_{n=1}^{N^i} \sum_{j \in \{b,a\}} \lambda^{i,n}(\mathbf{m}^i, p^{i,n,j}, \ell^{i,j}) \mathbb{E} \left[(p^{\text{eff},i} + p^{\text{short},i,n}(\mathbf{m}^i)) \ell^{i,n,j} \epsilon^{i,n,j} - c^{i,n,\mathcal{L}}(\epsilon^{i,n,j}, \ell^{i,n,j}, p^{i,n,j}) \right] \right\}, \right. \\ \sup_{v^b} \left\{ \sum_{i=1}^d \sum_{n=1}^{N^i} (p^{\text{eff},i} + p^{\text{short},i,n}(\mathbf{m}^i)) v^{i,n,b} - c^{i,n,\mathcal{M}}(v^{i,n,b}) \right\}, \\ \left. \sup_{v^a} \left\{ \sum_{i=1}^d \sum_{n=1}^{N^i} (p^{\text{eff},i} + p^{\text{short},i,n}(\mathbf{m}^i)) v^{i,n,a} - c^{i,n,\mathcal{M}}(v^{i,n,a}) \right\} \right\}, \\ p^{\text{eff},i} = \nabla^i \tilde{V}(t, q). \end{aligned} \quad (\text{Opt-MM})$$

The control problem is essentially a choice between sending limit orders or market orders in each venue for each asset. The effective microstructure alpha helps the market-maker to mean revert his inventory toward zero. For example, assume that the market-maker received a large buy passive filling in asset i . The effective microstructure alpha, that is the i -th component of the gradient of the long-term utility function \tilde{V} , will point down which is a strong sell signal. Therefore, the market-maker will send a sell market order to reduce his long position. Note that the effective microstructure alpha takes into account the correlation structure between the assets, meaning that the market-maker can hedge a long position in an asset with a short position in another positively correlated asset.

10.5 Numerical results

Mathematical elegance and simplicity are to be prized, of course, but an execution model cannot pass the test of practicality until it helps us execute portfolio transitions.

One of the most important features of our framework, as compared with a plain-vanilla, Almgren-Chriss executor, it allows the executor to consider market microstructure and use passive orders, hence avoiding certain types of market impact and spread costs. The main point we wish to make in this example is that our method potentially avoids the pitfalls of a purely-passive execution model because it can consider the utility gradient (and its multi-period analog, the gradient of the value function) in the formation of aggression levels. With this specific aim in mind, we consider the liquidation of a market-neutral portfolio with our method and contrast this with comparable results for a purely-passive method.

The specific example we choose is the liquidation of a market-neutral portfolio on October 15, 2008. The portfolio to be liquidated is long IBM and short AAPL. We choose the long position in IBM arbitrarily to be 1000 shares. We estimate the CAPM beta of each security, denoted $\hat{\beta}_i$ where $i \in \{1, 2\}$, using three years of daily data, and size the short position so that the beta exposure of the portfolio $\sum_i h_i \hat{\beta}_i$ is near zero.

10.5.1 Transaction cost model and microstructure simulation

We take $\kappa = 10^{-3}$ and we assume $c(v) = v^\top \Lambda v$ for some diagonal matrix $\Lambda = \text{diag}(\lambda^1, \dots, \lambda^d)$. This reduces the transaction cost modeling problem to one of estimating appropriate values for each λ_i . Let

advp_i denote our prediction of the daily dollar volume in the i -th security. The notation “advp” comes from the fact that it is computed as the average daily volume “adv” in shares, times the price “p”. For simplicity we assume trading one percent of advp_i will cause 20 basis points of market impact, with extension by linearity, meaning that

$$\lambda^i = 20 \times 10^{-4} \times \frac{1}{0.01 \text{advp}_i}. \quad (10.26)$$

For very large trades (say, more than 0.05advp_i), simple models such as (10.26) break down. For this reason, we restrict our attention in this example to trades that are relatively small with respect to the anticipated volume.

One of the most challenging aspects of this study is simulating passive execution, which we defined previously as a process of continually joining the queue on the near side of the limit order book until the order is filled, but never crossing the spread.

We are limited to the academic data sets available via the Wharton Research Data Services (WRDS). For this exercise, we used the New York Stock Exchange Trade and Quote (TAQ) database, which contains intraday transactions data on trades and quotes for all securities listed on the New York Stock Exchange (NYSE) and American Stock Exchange (AMEX), as well as Nasdaq National Market System (NMS) and SmallCap issues.

The TAQ database represents the aggregate inside quote for each exchange. Therefore, it includes both specialists and the public limit order book. Only having access to the consolidated feed, we construct a conservative simulation of when passive fills occur. Specifically, if we have a “buy” limit order (the entire process is similar for limit “sell” orders with “bid” replaced by “ask”) which is simulated as existing in the queue on the bid side of the order book, when can we assume such an order was filled? Conservatively, if the order book changes and the new ask price is less or equal to the existing limit order price, we assume that markets would have cleared in the process of this change, and our limit order would have been filled, at least partially. We limit the amount of fill to the posted quantity at the new ask price. If this quantity is simulated to have been taken out, then no further fills are allowed to occur in the simulation until the price level changes. We assume that when the price level of the NBBO has changed, the liquidity is also replenished to the reported value at the new price level. This is a fairly conservative set of conventions; in reality, a larger number of passive fills could occur than merely the ones we simulate. This is because if there are multiple limit orders in the queue, one limit order can, of course, be filled without either bid or ask price levels changing.

Predicting the probability of a passive fill, denoted f_i above, is equivalent to predicting the next transition of the limit order book and hence requires a model of limit order book dynamics. Indeed, such fill probabilities are one of the possible outputs of the very detailed model of Cont, Stoikov, and Talreja [82] or the microstructure trading model presented in the previous section. As our data set is only the consolidated feed, we simply take $f_i = 0.1$ as the passive fill probability.

10.5.2 Results

As indicated above, we construct a market-neutral portfolio of $d = 2$ securities in which the long side is initially 1000 shares of IBM. Security $i = 1$ is IBM and $i = 2$ is AAPL. We estimate the security betas to the S&P 500 (via regression on several years of daily data) as

$$\hat{\beta}_1 = 0.705, \quad \hat{\beta}_2 = 1.276. \quad (10.27)$$

We begin the simulation at 10:00 am on October 15, 2008, rather than immediately at the open since there are often outlier quotes, wide spreads and other effects around the open. The most recent midpoint price of IBM at 10:00 am was $p_1 = 93.06$ and for AAPL, $p_2 = 105.985$.

For convenience, we keep track of a cash balance for each position. The $n_1 = 1000$ shares of IBM are financed by borrowing $n_1 p_1 = \text{USD } 93,060$ in cash and purchasing a position initially worth USD 93,060, so the net value (cash plus stock) of that position is initially zero. Similarly, the short position in AAPL is obtained by borrowing $n_2 = -485$ shares and immediately selling them for USD 51,403, and this position also initially has a net (cash + stock) value of zero. Note that with these holdings, (10.27) implies that the portfolio's beta is

$$n_1 p_1 \hat{\beta}_1 + n_2 p_2 \hat{\beta}_2 \approx 0.$$

Any cash generated from further stock sales or cash used for further purchases of the same security is considered part of the separate cash balance allocated to that position. As prices change and as orders are filled, the values of each position will fluctuate.

Let $n_{i,t}$ denote the number of shares held in the i -th security at time t , and $p_{i,t}$ the latest midpoint price as of time t . Also, let $c_{i,t}$ denote the amount of cash (which can be positive or negative) attributed to the i -th security at time t , according to the accounting conventions outlined above. These variables change throughout the lifetime of the execution.

The value of a position is the number of shares held times the most recent midpoint price, plus the total amount of cash associated to the position, i.e. $n_{i,t} p_{i,t} + c_{i,t}$. The value of a portfolio is the sum of the values of all its positions, i.e.

$$\text{value}_t := \sum_{i=1}^d (n_{i,t} p_{i,t} + c_{i,t}). \quad (10.28)$$

The value process (10.28), and especially its drift, is one measure of the execution's quality. If the value tends to drift downward, as in the AlwaysPassive model detailed below, then the execution desk is losing money due to slippage. This is perhaps the typical situation – one expects execution to have associated costs. A particularly pleasant situation arises when the drift of the portfolio value process (10.28) is zero, as in Figure 10.2, and it is possible that with very good microstructure alphas added to the generalized momenta, the drift could even become positive. All monetary values are reported in USD. The predicted daily volumes are estimated to

$$\text{advp}_1 \approx 1.16 \times 10^9, \quad \text{advp}_2 \approx 6.13 \times 10^9.$$

The covariance matrix is

$$\Sigma = 10^{-4} \times \begin{pmatrix} 15.5728 & 17.7558 \\ 17.7558 & 28.6519 \end{pmatrix},$$

which implies a correlation of 0.84 among the two assets and daily volatilities of approximately 3.9% and 5.4%.

The output of our algorithm is the instantaneous aggression level: aggressive, passive, or wait. There is not a unique benchmark to gauge such an algorithm's performance, but it is sensible to compare a complicated method of choosing the aggression level to a simple method for choosing the aggression level, to see if the additional complexity is justified. Hence one could compare it to a constant aggression level – always passive.

Figure 10.1 reveals that, as the market was falling, the passive “buy” orders in AAPL were all filled very quickly, while unsurprisingly the “sell” orders in IBM were filled very slowly, and indeed were not even finished by the end of the trading day. This drove the Gross Market Value (gmv) down while pushing the net and beta higher, where we define

$$\beta_t := \sum_i n_{i,t} p_{i,t} \hat{\beta}_i, \quad \text{net}_t := \sum_i n_{i,t} p_{i,t}, \quad \text{gmv}_t := \sum_i |n_{i,t} p_{i,t}|, \quad (10.29)$$

with $\hat{\beta}_i$ given by (10.27). Thus the portfolio had $\beta_t > 0$ in a falling market. Note that the losses incurred in this manner do not become gains if the sign of the market move is reversed; they remain losses irrespective of the market's direction. In a rising market, the “always passive” model would have the same problem: the “sell” orders would be filled quickly, the “buy” orders would linger, and the portfolio would build up negative beta in a rising market.




Figure 10.1: Portfolio holdings in the “always passive” model, and portfolio characteristics: gross market value (gmv), net, and β_t given by (10.29).

We now show the analogous graphs for the simplest version of our execution model developed in the previous section. Note that the model retains a fairly small beta exposure throughout the lifetime of the execution. This is because CAPM beta is also a factor in the APT risk model, and the generalized momenta point along the gradient of the Hamilton-Jacobi-Bellman value function and hence drive trading towards the optimal value of multiperiod utility (including the risk term). This is the key advantage of our model over simpler execution algorithms.




Figure 10.2: Portfolio holdings in the sophisticated model, and portfolio characteristics: gross market value (gmv_t), net_t , and β_t given by (10.29).

Finally, we consider the portfolio value over the lifetime of the execution. Note that in our model, the value process (10.28) is approximately driftless, which as explained above is a desirable property, and outperforms the “always passive” value process realization. In particular, in our model value_t is able to avoid negative drift in a falling market precisely because the portfolio remains approximately beta-neutral. In a portfolio with many assets (large n), our method would allow it to remain approximately neutral to all factors in the APT model.




Figure 10.3: Portfolio value (10.28) over the lifetime of the execution, for both execution methods.

The difference in Figure 10.3 is both statistically and economically significant. The t-statistic for the difference is about 78, hence significant at the 99.999% level. Moreover, the dollar value of the difference between the two methods is about 1.5% of the initial gross market value to be liquidated.

10.6 Conclusion

In this chapter, we present a framework to perform optimal trading, taking into account market microstructure and a long-term trading schedule without the use of optimal control. This approach relies on the use of the generalized momenta $p = \nabla V(t, q)$ as the effective microstructure alpha. We show that a myopic agent sending only market orders with such alpha will minimize the error with respect to the long-term trading schedule. Moreover, when we add the possibility of passive execution, the long-term alpha can be chosen as a transformation of the generalized momenta p . We also present a general microstructure trading framework for the multi-asset multi-venue optimal trading problem. For a parsimonious model of fill probabilities, the effective microstructure alpha can be computed in closed form. We apply the same heuristics to derive an optimal market-making model that is tractable for a large number of assets and venues.

Based on the dual formulation of the classic Almgren-Chriss optimization problem, this simple heuristics has wide-ranging practical implications. In addition to bridging the gap between order placement and scheduling, it simplifies optimal trading problems that are usually intractable using optimal control due to the high-dimensional Hamilton-Jacobi-Bellman equation resulting from the control problem. This is of particular importance for a quantitative execution desk wishing to trade a high number of cross-listed assets. It opens up many avenues for future exploration. One set of projects is to consider trading problems beyond the typical buy-side utility-maximization, which can still be viewed within the unifying framework of a myopic risk-neutral wealth-maximizer, whose microstructure alphas are aligned with the value function gradient.

Chapter 11

A note on Almgren-Chriss optimal execution problem with geometric Brownian motion

11.1 Introduction

Optimal liquidation is a problem faced by a trader when he needs to liquidate a large number of shares. The trader faces a tradeoff between fast execution, reducing a risk related to price changes and slow execution, allowing to avoid high trading costs. Since the seminal paper Almgren and Chriss [10], various extensions of optimal liquidation problems have been studied, see for example Alfonsi, Fruth, and Schied [5], Cartea and Jaimungal [66], Guéant, Lehalle, and Fernandez-Tapia [141]. The common framework to address this issue introduced in Almgren and Chriss [10] assumes the following:

- the efficient price process follows an arithmetic Brownian motion (ABM),
- permanent market impact is linear,
- transaction costs are a linear function of the trading rate.

The execution of a large order is then formulated in discrete time as a tradeoff between expected costs and risk of the trading strategy, with variance as a risk measure. Under this framework, there exists a unique optimal liquidation strategy, which is a deterministic function of time and the initial position of the trader.

Continuous versions of this problem have been considered, notably in Forsyth [119], Forsyth, Kennedy, Tse, and Windcliff [120], where the author shows the ill-posedness of the mean-variance framework leading to time-inconsistent solutions. To overcome this issue, the authors suggest using alternative objective functions, particularly mean-quadratic variation. Under this choice, the authors solve a two-dimensional Hamilton-Jacobi-Bellman equation numerically. Moreover, in Guéant, Lehalle, and Fernandez-Tapia [141], the authors consider the optimal execution problem with CARA utility objective function. To the best of our knowledge, there is no closed-form solution to the continuous version of the Almgren-Chriss framework with quadratic variation as a risk measure and geometric Brownian motion (GBM) assumption for the efficient price process. In Gatheral and Schied [126], the authors solve a modified version of the problem with GBM, accounting for the risk with a linear function of the trading rate.

In this chapter, we solve the optimal liquidation problem under the Almgren-Chriss framework in continuous time, in the case where the efficient price process follows a GBM, and the risk measure is quadratic variation. Motivated by Collin-Dufresne and Fos [77], we assume that trading costs are

a quadratic function of the amount of cash, instead of shares. Thus, we reformulate the problem in terms of cash traded and derive in closed-form the optimal control of the trader liquidating his position. As the method is based on the resolution of a system of ODEs, it does not suffer from the curse of dimensionality. In particular, we show how to extend this framework to the case of the liquidation of a portfolio of N , possibly correlated, assets. It also enables us to treat the case where the return's drift is a stochastic process without any BSDE methods with a singular condition. The specifications of the optimal execution models previously described are summarized in the following table¹:

Model	Asset dynamics	Penalty	Closed-form solution	Extension to multiple assets
Almgren and Chriss [10]	ABM	Quadratic	Yes	Yes
Gatheral and Schied [126]	GBM	Linear (VaR)	Yes	No
Forsyth et al. [120]	GBM	Quadratic	No	No
Baldacci and Benveniste	GBM	Quadratic	Yes	Yes

The chapter is organized as follows. In Section 11.2, we describe the Almgren-Chriss framework in continuous time and reformulate the optimization problem in terms of cash. In Section 11.3, we obtain a closed-form solution of the Almgren-Chriss framework with GBM for the efficient price process. Finally in Section 11.4, we present numerical applications under different market conditions.

11.2 The model

We define $(\Omega, \mathcal{F}_{t \in [0, T]}, \mathbb{P})$ a filtered probability space, on which all stochastic process are defined, a trading horizon is $T > 0$ and $\mathbb{F} = (\mathcal{F}_t)_{t \in [0, T]}$.

11.2.1 Almgren-Chriss framework in continuous time

We rapidly recall the well-known Almgren-Chriss problem in continuous time. We consider the issue of the liquidation of $q_0 \in \mathbb{R}$ shares of a stock whose price at time t is defined by S_t . The number of shares hold by the trader is defined by an absolutely continuous measurable process $q_t := q_0 - \int_0^t \dot{q}_s ds$ where $(\dot{q}_s)_{s \in [0, T]}$ is the trading rate, controlled by the trader. In the case where $q_0 > 0$, the trading rate \dot{q} is a positive process, and conversely for $q_0 < 0$.

The transaction price is

$$\tilde{S}_t := S_t + \frac{\lambda}{2} \dot{q}_t + \gamma(q_t - q_0),$$

where $\lambda, \gamma \in \mathbb{R}^+$ are constants related respectively to temporary and permanent price impact. Indeed, the term $\frac{\lambda}{2} \dot{q}_t$ is the impact of trading \dot{q}_t shares at time t , whereas the term $\gamma(q_t - q_0)$ is the impact generated by the flow of transactions up to time t . In the original framework in discrete time, see Almgren and Chriss [10], and in most of the extensions in continuous time, see Guéant, Lehalle, and Fernandez-Tapia [141] for example, the price process follows an ABM. The number of shares hold by the trader satisfies the boundary condition $q_T = 0$. Therefore, the cost of this strategy during the trading period is

$$\mathcal{C}(\dot{q}) := \int_0^T \tilde{S}_t \dot{q}_t dt.$$

Aiming at remedying the time inconsistency of the optimal strategies in the pre-commitment mean-variance framework, inspired by Forsyth, Kennedy, Tse, and Windcliff [120], we replace the variance

¹A similar comparison of optimal execution models with respect to price dynamics and risk measures is presented in Brigo and Di Graziano [56].

by the quadratic variation in the penalty. The optimal execution problem consists in the optimization of a mean-quadratic variation objective function over the strategies $(\dot{q}_t)_{t \in [0, T]} \in \mathcal{A}$ where

$$\mathcal{A} := \left\{ (\dot{q}_t)_{t \in [0, T]} : \mathbb{F}\text{-predictable, } \int_0^T \dot{q}_s ds = q_0 \right\}.$$

The problem can be written as follows:

$$\sup_{v \in \mathcal{A}} \mathbb{E} \left[-\mathcal{C}(\dot{q}) - \frac{\kappa}{2} \langle \mathcal{C} \rangle_T \right],$$

with $\kappa > 0$ and

$$\langle \mathcal{C} \rangle_T := \int_0^T \dot{q}_t^2 d\langle S \rangle_t.$$

The use of quadratic variation leads to time-consistent strategies. Moreover, in contrast to the variance, quadratic variation takes into account the trajectory of liquidation. A direct integration by parts on $\mathcal{C}(q)$ gives

$$\mathcal{C}(q) = -q_0 S_0 - \int_0^T q_t dS_t + \frac{\lambda}{2} \int_0^T \dot{q}_t^2 dt + \frac{\gamma}{2} q_0^2.$$

Therefore, the problem is

$$\sup_{\dot{q} \in \mathcal{A}} \mathbb{E} \left[\int_0^T q_t dS_t - \frac{\lambda}{2} \int_0^T \dot{q}_t^2 dt - \frac{\kappa}{2} \int_0^T \dot{q}_t^2 d\langle S \rangle_t \right]. \quad (11.1)$$

When the price process follows an ABM, Problem (11.1) boils down to a simple calculus of variations problem, which has been solved, for example, in Guéant, Lehalle, and Fernandez-Tapia [141]. The case where the dynamics are given by a GBM is more intricate. In Gatheral and Schied [126], the authors consider it analytically intractable when a quadratic variation penalty is used. Moreover, in Forsyth, Kennedy, Tse, and Windcliff [120], the authors derive a numerical solution of (11.1) by solving the corresponding Hamilton-Jacobi-Bellman equation. Note that strategies under ABM assumption are good proxies of the ones under GBM assumption in period of low volatility.

11.2.2 Reformulation in terms of cash

We now reformulate the optimal execution problem in terms of cash. We emphasize that we treat the very same problem as in (11.1), except that we modify the transaction costs such that the penalty for \dot{q}_t becomes $\dot{q}_t S_t$.

We assume that the price process follows a GBM:

$$dS_t = \sigma S_t dW_t.$$

Multiplying above and below by S_t , we obtain that

$$\int_0^T q_t dS_t = \int_0^T \theta_t dy_t,$$

where $dy_t := \sigma dW_t$ is the return of the price process, and $\theta_t := q_t S_t$ is the trader's position expressed in dollars. Moreover, the quadratic variation penalty has the form

$$\frac{\kappa \sigma^2}{2} \int_0^T \theta_t^2 dt.$$

Applying Itô's formula, we derive that the cash position $\theta_t := \theta_t^u$ has the following dynamics:²

$$d\theta_t^u = u_t dt + \theta_t^u dy_t = u_t dt + \sigma \theta_t^u dW_t, \quad (11.2)$$

where $u_t = \dot{q}_t S_t$ is the trading's rate in dollar at time t .

Recall that in the classical Almgren-Chriss framework (11.1), trading costs are a quadratic function of the number of shares traded at time t defined by \dot{x}_t (the second term in (11.1)). The only modification we make here is to assume that instantaneous costs are a quadratic function of the amount of cash. According to Collin-Dufresne and Fos [77], working with dollar holdings and returns is more consistent with common practice. We define the set of admissible control processes $(u_t)_{t \in [0, T]}$ as

$$\mathcal{A} := \left\{ (u_t)_{t \in [0, T]} : \mathbb{F}\text{-predictable, } \int_0^T |u_t| dt < +\infty \right\}.$$

We must ensure the complete liquidation of the trader's position at terminal time T , meaning $\theta_T^u = 0$. Following the problem formulation in terms of cash instead of shares, we consider the following mean-quadratic variation optimization problem with terminal penalty:

$$\lim_{a \rightarrow +\infty} \sup_{u \in \mathcal{A}} \mathbb{E} \left[\int_0^T - \left(\frac{\lambda}{2} u_t^2 + \frac{\kappa \sigma^2}{2} (\theta_t^u)^2 \right) dt - \frac{a}{2} \theta_T^u \int_0^T u_t dt \right], \quad (11.3)$$

where the limit over $a > 0$ aims at representing the singular condition $\theta_T^u = 0$ ³. Equation (11.3) can be seen as a classical linear-quadratic optimization problem, which is reduced to the resolution of a Riccati equation in dimension one. However such equations are not well suited for multidimensional extensions of this problem, that is to say the liquidation of a portfolio of N assets. Furthermore when adding a possibly non-Markovian drift $(\alpha_t)_{t \in [0, T]}$ to the price process, one has to rely on BSDE methods to compute the optimal control.

Our method, developed in the next section, has several advantages. First, it enables us to solve the original Almgren-Chriss problem explicitly, under the GBM assumption, only by assuming that instantaneous costs are a function of the amount of cash. In addition to this, it applies to the case of a stochastic drift $(\alpha_t)_{t \in [0, T]}$ without using the BSDE framework. Finally, an explicit solution can be obtained in the case of the liquidation of a portfolio of N possibly correlated assets.

We solve in the next section Problem (11.3) under the dynamics (11.2) for the trader's position. We treat the non-zero drift case in Section 11.5.1.

11.3 Solving explicitly the Almgren-Chriss problem with GBM

Throughout this section, we work on the following functional space:

$$\mathbb{H}^2 := \left\{ (v_t)_{t \in [0, T]} : \mathbb{E} \left[\int_0^T v_t^2 dt \right] < +\infty \right\},$$

with its associated inner product and norm

$$\langle u, v \rangle_t := \mathbb{E} \left[\int_0^t u_s v_s ds \right], \quad \|u\|_t^2 := \mathbb{E} \left[\int_0^t u_s^2 ds \right].$$

We also define for all $t \in [0, T]$ the exponential martingale $M_t := \exp \left(\sigma W_t - \frac{\sigma^2}{2} t \right)$ and the associated change of measure $\frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_{\mathcal{F}_T} = M_T$. We begin with a lemma characterizing the trader's position.

²We write the superscript u since $(u_t)_{t \in [0, T]}$ is the control process.

³Note that the form of the terminal penalty does not impact the nature of our results. In particular, the trader can consider $a = 0$ if he does not necessarily want a complete liquidation of the portfolio at time T , and still obtain a closed form solution for the optimal trading rate.

Lemma 11.1. *The unique solution of (11.2) is given by*

$$\int_0^t M_t M_s^{-1} u_s ds.$$

For all $v \in \mathbb{H}^2$, we define the operator

$$(Kv)_t := \int_0^t M_t M_s^{-1} v_s ds.$$

The adjoint process (K^*v) is equal for all $s \in [0, T]$ to

$$(K^*v)_s := \int_s^T \mathbb{E}^{\mathbb{Q}}[v_t | \mathcal{F}_s] dt.$$

The proof is given in Appendix 11.A.1 and relies on a straightforward application of Itô's formula. Therefore the optimization problem (11.3) can be rewritten, with a fixed $a > 0$, as

$$\sup_{u \in \mathcal{A}} -\frac{\lambda}{2} \|u\|_T^2 - \frac{\kappa\sigma^2}{2} \|Ku\|_T^2 - \frac{a}{2} (Ku)_T \int_0^T u_s ds. \quad (11.4)$$

The problem is a supremum over a concave function of u , which is Gateaux-differentiable on \mathbb{H}^2 . Thus first order condition gives:⁴

$$\frac{\kappa\sigma^2}{\lambda} K^*Ku + u + \frac{a}{\lambda} (Ku)_T = 0, \quad (11.5)$$

or equivalently

$$\frac{\kappa\sigma^2}{\lambda} \int_s^T \int_0^t \mathbb{E}^{\mathbb{Q}}[M_t M_\tau^{-1} u_\tau | \mathcal{F}_s] d\tau dt + u_s + \frac{a}{\lambda} M_T \int_0^T M_\tau^{-1} u_\tau d\tau = 0. \quad (11.6)$$

For all $(s, s_0) \in [0, T]^2$ such that $s \geq s_0$, we apply $\mathbb{E}^{\mathbb{Q}}[\cdot | \mathcal{F}_{s_0}]$ on both sides of (11.6). This leads to the following technical lemma.

Lemma 11.2. *We define $v(s) := \mathbb{E}^{\mathbb{Q}}[u_s | \mathcal{F}_{s_0}]$ such that $v(s_0) = u_{s_0}$, and assume that it is differentiable with respect to s .⁵ We also set*

$$z(t) = e^{\sigma^2(t-s_0)} \theta_{s_0} + \int_{s_0}^t e^{\sigma^2(t-\tau)} v(\tau) d\tau,$$

where we recall that $\theta_{s_0} := (Ku)_{s_0}$.

(i) Equation (11.5) can be rewritten

$$v(s) + \frac{\kappa\sigma^2}{\lambda} \int_s^T z(t) dt + \frac{a}{\lambda} z(T) = 0.$$

(ii) The couple (v, z) satisfies the following system of differential equations

$$\begin{cases} v'(s) = \frac{\kappa\sigma^2}{\lambda} z(s), \\ z'(s) = \sigma^2 z(s) + v(s), \end{cases}$$

with boundary conditions

$$\begin{cases} v(T) = -\frac{a}{\lambda} z(T), \\ z(s_0) = \theta_{s_0}. \end{cases}$$

⁴See Appendix 11.A.2 for well-definedness of the first order condition.

⁵It will be shown ex-post, by a direct verification argument, that $v(\cdot)$ is differentiable.

Thus, the control problem (11.3) is reduced to the resolution of a linear system of ODEs with constant coefficients. We can now state our main theorem.

Theorem 11.1. *Consider the problem (11.3), the optimal control is given explicitly for all time $t \in [0, T]$ by*

$$u_t^* = \theta_t^{u^*} \Gamma(t),$$

where $\Gamma(\cdot)$ is a deterministic function of time defined in (11.11) and the optimal trader's position satisfies

$$\theta_t^{u^*} = \theta_0^{u^*} \exp\left(\int_0^t \left(\Gamma(s) - \frac{\sigma^2}{2}\right) ds + \sigma W_t\right).$$

The proof is included in the one of Theorem 11.2, where we prove a similar result in a more general framework by allowing a stochastic drift in the dynamics of the price process, and is reported in Appendix 11.A.4. The theorem shows that the optimal control is a linear function of the trader's position. Therefore, we find an *aggressive in-the-money* selling strategy, similar to Gatheral and Schied [126], in the sense that the trader liquidates faster when the stock price increases and conversely. This is illustrated in the following section. Moreover, the trader's position is a geometric Brownian motion, so that it always stays positive, in contrast to Gatheral and Schied [126]. As the function $\Gamma(t) \xrightarrow{t \rightarrow T} -\infty$ superlinearly, we have $\theta_t^u \xrightarrow{t \rightarrow T} 0$.

11.4 Numerical results

We simulate one Brownian motion trajectory, and plot the corresponding stock price process, as well as trading strategy $(u_t^*)_{t \in [0, T]}$ and trader's cash position $(\theta_t^*)_{t \in [0, T]}$ and in shares $(\theta_t^*/S_t)_{t \in [0, T]}$ for different values of σ . We take a stock with initial price $S_0 = 100\$$ following a GBM without drift (whose trajectories for different values of σ are in Figure 11.5), a portfolio of 10^3 shares to liquidate over $T = 20$ days, with $\lambda = \kappa = 0.2$. In Figure 11.1, we see an increase of the cash position at the beginning, which can be misleading but is only due to the initial increase of the stock price process. This is also represented in the trading strategy of Figure 11.2, where we see that the trader liquidates his position faster when the stock process has a higher volatility. Figures 11.3 and 11.4 show the position and the trading strategy in terms of shares. We also compare in Figure 11.6 our trading strategy in shares to the one in Gatheral and Schied [126], which is defined as

$$q_t^* := \frac{T-t}{T} \left(q_0 - \frac{\kappa T}{4} \int_0^t S_u du \right), \quad (11.7)$$

where $q_0 = \frac{\theta_0}{S_0}$ is the initial number of shares hold by the trader, and $dS_t = \sigma S_t dW_t$. The trader still liquidates faster with a high volatility but his trading strategy, in this rather extreme regime, can go negative.




Figure 11.1: Evolution of the cash position with respect to time.




Figure 11.2: Trading strategy in cash with respect to time.




Figure 11.3: Evolution of the share's position with respect to time.




Figure 11.4: Trading strategy in shares with respect to time.




Figure 11.5: Evolution of the stock price with respect to time.




Figure 11.6: Evolution of the share's position with respect to time using (11.7).

We now fix $\sigma = 0.1$ and $\kappa = 0.2$. The various cases of the impact of the transaction costs λ on the trader's behavior are represented in Figures 11.7, 11.8, 11.9 and 11.10. Obviously, the price process is insensitive to a variation of λ . Moreover, the trading strategies in Figures 11.8 and 11.10 are decreasing functions of λ meaning that the trader liquidates his position using smaller sell orders when transactions costs are higher. This is also shown in the trader's position in Figures 11.7 and 11.9.




Figure 11.7: Evolution of the cash position with respect to time.




Figure 11.8: Trading strategy in cash with respect to time.




Figure 11.9: Evolution of the share's position with respect to time.




Figure 11.10: Trading strategy in shares with respect to time.

Finally, we set $\sigma = 0.1, \lambda = 0.2$ and study the influence of the risk aversion parameter κ .




Figure 11.11: Evolution of the cash position with respect to time.




Figure 11.12: Trading strategy in cash with respect to time.

In Figures 11.12 and 11.14, we see that a highly risk averse trader will liquidate faster than a low risk averse trader. This is shown in terms of his position in Figures 11.11 and 11.13.




Figure 11.13: Evolution of the share's position with respect to time.




Figure 11.14: Trading strategy in shares with respect to time.

We now show how to extend our framework to the case of a stochastic drift for the price process and the liquidation of a portfolio of N assets.

11.5 Extensions of the model

11.5.1 Stochastic drift

We now consider the case of a stochastic drift, that is we solve

$$\lim_{a \rightarrow +\infty} \sup_{u \in \mathcal{A}} \mathbb{E} \left[\int_0^T \alpha_t \theta_t - \left(\frac{\lambda}{2} u_t^2 + \frac{\kappa \sigma^2}{2} (\theta_t^u)^2 \right) dt - \frac{a}{2} \theta_T^u \int_0^T u_t dt \right],$$

where $(\alpha_t)_{t \in [0, T]}$ is a stochastic drift of the price process. We consider a slight modification of the problem where we neglect the part $\alpha_t S_t$ of the price's drift.⁶ Therefore, we simplify the dynamics of the price process, and assume

$$d\theta_t^u = u_t dt + \sigma \theta_t^u dW_t. \tag{11.8}$$

The first-order condition associated to this optimization problem is

$$\frac{\kappa \sigma^2}{\lambda} \int_s^T \int_0^t \mathbb{E}^{\mathbb{Q}} [M_t M_\tau^{-1} u_\tau | \mathcal{F}_s] d\tau dt + u_s + \frac{a}{\lambda} M_T \int_0^T M_\tau^{-1} u_\tau d\tau = \frac{1}{\lambda} \int_s^T \mathbb{E}^{\mathbb{Q}} [\alpha_t | \mathcal{F}_s] dt.$$

⁶It can be shown that, if there exists $\eta > 0$ such that $\sup_t |\alpha_t| < \eta$, then the trading strategy derived in this section is arbitrary close (as a function of η, T, σ) to the optimal strategy without simplification of the drift.

Then, the analogous of Equation (11.5) can be rewritten

$$v(s) + \frac{\kappa\sigma^2}{\lambda} \int_s^T z(t)dt + \frac{a}{\lambda} z(T) = \frac{1}{\lambda} \int_s^T \mathbb{E}^{\mathbb{Q}}[\alpha_t | \mathcal{F}_{s_0}] dt.$$

where the couple (v, z) satisfies the following system of differential equations

$$\begin{cases} v'(s) = \frac{\kappa\sigma^2}{\lambda} z(s) - \frac{1}{\lambda} \mathbb{E}^{\mathbb{Q}}[\alpha_s | \mathcal{F}_{s_0}], \\ z'(s) = \sigma^2 z(s) + v(s), \end{cases}$$

with conditions

$$\begin{cases} v(T) = -\frac{a}{\lambda} z(T), \\ z(s_0) = \theta_{s_0}. \end{cases}$$

We finally obtain the following theorem:

Theorem 11.2. *The optimal control at any time $t \in [0, T]$ is given by*

$$u_t^* = \theta_t^{u^*} \Gamma(t) + \nu(t),$$

where the optimal trader's position is defined as

$$\theta_t^{u^*} = H_t \int_0^t H_s^{-1} \nu(s) ds,$$

with $dH_t = \Gamma(t)H_t dt + \sigma H_t dW_t$, and $\Gamma(\cdot), \nu(\cdot)$ are deterministic functions defined in (11.11).

The term $\nu(\cdot)$ is a linear function of both α_t and $\mathbb{E}^{\mathbb{Q}}[\alpha_T | \mathcal{F}_t]$, representing the influence of the drift on the optimal strategy. It is an increasing function of the drift α_t meaning that we aim at liquidating faster our position when the stock price increases. Moreover, it is a decreasing function of $\mathbb{E}^{\mathbb{Q}}[\alpha_T | \mathcal{F}_t]$: when the expected drift at the terminal time is high, the trader prefers to liquidate slower, waiting for a future stock price increase. As in the zero-drift case, we observe an *aggressive in-the-money* selling strategy.

11.5.2 Multi-dimensional case

This model extends directly to the problem of optimal execution of a portfolio of N assets. We define the return of the i -th asset as

$$dy_t^i = \sigma_i dW_t^i,$$

where (W^1, \dots, W^N) are Brownian motions with non singular covariance matrix $\Sigma = (\sigma_i \sigma_j \rho^{i,j})_{1 \leq i, j \leq N}$, $\sigma_i > 0$ is the volatility of the i -th asset and $\rho^{i,j}$ is the correlation between the i -th and the j -th Brownian motion. The cash position of the trader with respect to the i -th asset is defined by

$$d\theta_t^{u,i} = u_t^i dt + \theta_t^{u,i} dy_t^i = u_t^i dt + \sigma_i \theta_t^{u,i} dW_t^i, \quad (11.9)$$

where $(u_t^i)_{t \in [0, T]}$ is the trading rate on the i -th asset. Therefore the optimization problem (11.3) rewrites as

$$\lim_{a \rightarrow +\infty} \sup_{u \in \mathcal{A}} \mathbb{E} \left[\int_0^T \sum_{i=1}^N -\frac{\lambda}{2} (u_t^i)^2 - \frac{\kappa}{2} \left(\sum_{i=1}^N \sigma_i^2 (\theta_t^{u,i})^2 dt + \sum_{\substack{i,j=1 \\ i \neq j}}^N \rho^{i,j} \sigma_i \sigma_j \theta_t^{u,i} \theta_t^{u,j} dt \right) - \frac{a}{2} \sum_{i=1}^N \theta_T^{u,i} \int_0^T u_t^i dt \right],$$

where

$$\mathcal{A} := \left\{ (u_t^i)_{t \in [0, T]} : i \in \{1, \dots, N\} \text{ } \mathcal{F}_t\text{-measurable, } \int_0^T |u_t^i| dt < +\infty \right\},$$

and the terminal condition ensures that $\theta_T^i = 0$ for $i \in \{1, \dots, N\}$. We define $(K^i u_t^i)_{t \in [0, T], i \in \{1, \dots, N\}}$ as the solution of the SDE (11.9):

$$(K^i u^i)_t = M_t^i \int_0^t (M_s^i)^{-1} u_s^i ds,$$

where $\frac{d\mathbb{Q}^i}{d\mathbb{P}} \Big|_{\mathcal{F}_T} = M_T^i := \exp(\sigma_i W_T^i - \frac{(\sigma_i)^2 T}{2})$. The adjoint operator is defined as

$$(K^{*i} u^i)_s = \int_s^T \mathbb{E}^{\mathbb{Q}^i} [u_t^i | \mathcal{F}_s] dt.$$

For a fixed $a > 0$, the first order condition gives the following system

$$u^i + \frac{\kappa}{2\lambda} \left(2\sigma_i^2 K^{*i} K^i u^i + \sum_{j \neq i}^N \rho^{i,j} \sigma_i \sigma_j K^{*i} K^j u^j \right) + \frac{a}{\lambda} (K^i u^i)_T = 0, \quad i \in \{1, \dots, N\}, \quad (11.10)$$

or equivalently for all $i \in \{1, \dots, N\}$,

$$u_s^i + \frac{\kappa}{2\lambda} \left(2\sigma_i^2 \int_s^T \int_0^t \left(\mathbb{E}^{\mathbb{Q}^i} [M_t^i (M_\tau^i)^{-1} u_\tau^i | \mathcal{F}_s] d\tau dt + \sum_{j \neq i}^N \rho^{i,j} \sigma_i \sigma_j \mathbb{E}^{\mathbb{Q}^i} [M_t^j (M_\tau^j)^{-1} u_\tau^j | \mathcal{F}_s] d\tau dt \right) \right) + \frac{a}{\lambda} (K^i u^i)_T = 0.$$

For any $s \geq s_0$, apply $\mathbb{E}^{\mathbb{Q}^i} [\cdot | \mathcal{F}_{s_0}]$ on both sides of the equations. Simple but tedious computations lead to

$$\begin{aligned} 0 &= \mathbb{E}^{\mathbb{Q}^i} [u_s^i | \mathcal{F}_{s_0}] + \frac{\kappa \sigma_i^2}{\lambda} \left(\int_s^T e^{\sigma_i^2(t-s_0)} (K^i u^i)_{s_0} dt + \int_s^T \int_{s_0}^t e^{\sigma_i^2(t-\tau)} \mathbb{E}^{\mathbb{Q}^i} [u_\tau^i | \mathcal{F}_{s_0}] d\tau dt \right) \\ &+ \frac{\kappa}{\lambda} \sum_{j \neq i}^N \rho^{i,j} \sigma_i \sigma_j \left(\int_s^T e^{\sigma_i \sigma_j \rho^{i,j}(t-s_0)} (K^j u^j)_{s_0} dt + \int_s^T \int_{s_0}^t e^{\sigma_i \sigma_j \rho^{i,j}(t-\tau)} \mathbb{E}^{\mathbb{Q}^j} [u_\tau^j | \mathcal{F}_{s_0}] d\tau dt \right) \\ &+ \frac{a}{\lambda} \left(e^{\sigma_i^2(T-s_0)} (K^i u^i)_{s_0} + \int_{s_0}^T e^{\sigma_i^2(T-\tau)} \mathbb{E}^{\mathbb{Q}^i} [u_\tau^i | \mathcal{F}_{s_0}] d\tau \right). \end{aligned}$$

By denoting for all $i \in \{1, \dots, N\}$, $v^i(s) = \mathbb{E}^{\mathbb{Q}^i} [u_s^i | \mathcal{F}_{s_0}]$, and $\theta_{s_0}^i = (K^i u^i)_{s_0}$ the system becomes

$$\begin{aligned} 0 &= v^i(s) + \frac{\kappa \sigma_i^2}{\lambda} \left(\int_s^T e^{\sigma_i^2(t-s_0)} \theta_{s_0}^i dt + \int_s^T \int_{s_0}^t e^{\sigma_i^2(t-\tau)} v^i(\tau) d\tau dt \right) \\ &+ \frac{\kappa}{\lambda} \sum_{j \neq i}^N \rho^{i,j} \sigma_i \sigma_j \left(\int_s^T e^{\sigma_i \sigma_j \rho^{i,j}(t-s_0)} \theta_{s_0}^j dt + \int_s^T \int_{s_0}^t e^{\sigma_i \sigma_j \rho^{i,j}(t-\tau)} v^j(\tau) d\tau dt \right) \\ &+ \frac{a}{\lambda} \left(e^{\sigma_i^2(T-s_0)} \theta_{s_0}^i + \int_{s_0}^T e^{\sigma_i^2(T-\tau)} v^i(\tau) d\tau \right). \end{aligned}$$

We define

$$\begin{aligned} z^i(t) &:= e^{\sigma_i^2(t-s_0)} \theta_{s_0}^i + \int_{s_0}^t e^{\sigma_i^2(t-\tau)} v^i(\tau) d\tau, \\ z^{i,j}(t) &:= e^{\sigma_i \sigma_j \rho^{i,j}(t-s_0)} \theta_{s_0}^j + \int_{s_0}^t e^{\sigma_i \sigma_j \rho^{i,j}(t-\tau)} v^j(\tau) d\tau, \end{aligned}$$

and obtain for all $i \in \{1, \dots, N\}$:

$$v^i(s) + \frac{\kappa\sigma_i^2}{\lambda} \left(\int_s^T z^i(t) dt \right) + \frac{\kappa}{\lambda} \sum_{j \neq i}^N \rho^{i,j} \sigma_i \sigma_j \int_s^T z^{i,j}(t) dt + \frac{a}{\lambda} z^i(T) = 0.$$

Therefore first-order condition (11.10) is equivalent to the system of differential equations

$$\begin{cases} v'^i(s) - \frac{\kappa\sigma_i^2}{\lambda} z^i(s) - \frac{\kappa}{\lambda} \sum_{j \neq i}^N \rho^{i,j} \sigma_i \sigma_j z^{i,j}(s) = 0, \\ z'^i(s) = \sigma_i^2 z^i(s) + v^i(s), \\ z'^{i,j}(s) = \sigma_i \sigma_j \rho^{i,j} z^{i,j}(s) + v^j(s), \end{cases}$$

with initial conditions

$$\begin{cases} v^i(T) = -\frac{a}{\lambda} z^i(T), \\ z^i(s_0) = \theta_{s_0}^i, \\ z^{i,j}(s_0) = \theta_{s_0}^j. \end{cases}$$

We obtain a system of linear differential equations with constant coefficients. Thus, by noting that for all $i \in \{1, \dots, N\}$ and $s_0 \in [0, T]$, $v^i(s_0) = u_{s_0}^i$, we obtain the controls u_t^i for all $t \in [0, T]$ and $i \in \{1, \dots, N\}$ by solving this system of ODEs.

11.6 Conclusion

In this chapter, we present a way to solve the traditional Almgren-Chriss liquidation problem when the underlying asset is driven by a GBM. By working in terms of cash and using functional analysis tools, we can provide the optimal control of the problem explicitly. We provide an extension to the case of a GBM with stochastic drift and the liquidation of a portfolio of correlated assets. In particular, our method does not suffer from the curse of dimensionality.

11.A Appendix

11.A.1 Proof of Lemma 11.1

An application of Itô's formula gives

$$d(Ku)_t = u_t dt + \sigma(Ku)_t dW_t,$$

hence solving (11.2). The adjoint of K is the operator K^* such that for all $(u, v) \in \mathcal{A}$,

$$\langle Ku, v \rangle = \langle u, K^*v \rangle.$$

Using Bayes formula, we have

$$\begin{aligned} \langle Ku, v \rangle &= \mathbb{E} \left[\int_0^T (Ku)_t v_t dt \right] = \mathbb{E} \left[\int_0^T \int_0^t M_t M_s^{-1} u_s v_t ds dt \right] = \mathbb{E} \left[\int_0^T \int_s^T M_t M_s^{-1} u_s v_t dt ds \right] \\ &= \mathbb{E} \left[\int_0^T u_s (K^*v)_s ds \right] = \langle u, K^*v \rangle, \end{aligned}$$

where $(K^*v)_s = \int_s^T M_t M_s^{-1} v_t dt = \int_s^T \mathbb{E}^{\mathbb{Q}}[v_t | \mathcal{F}_s] dt$.

11.A.2 Gâteaux differentiability in (11.4)

We define the map $\Xi_T : \mathbb{H}^2 \longrightarrow \mathbb{R}$ by

$$\Xi_T(u) = -\frac{\lambda}{2}\|u\|_T^2 - \frac{\kappa\sigma^2}{2}\|Ku\|_T^2 - \frac{a}{2}(Ku)_T \int_0^T u_s ds.$$

As K is a linear operator of $u \in \mathbb{H}^2$, and $\lambda, \kappa > 0$, we deduce that Ξ is continuous, strictly concave and Gateaux differentiable; with Gateaux derivative given, for any $h \in \mathbb{H}^2$, by

$$\Xi_T(u)[h] = -\lambda\langle u, h \rangle_T - \kappa\sigma^2\langle Ku, Kh \rangle_T - a(Kh)_T.$$

By setting $\Xi(u)[h] = 0$, and using the definition of an adjoint operator, we obtain Equation (11.5).

11.A.3 Proof of Lemma 11.2

By the fact that

$$\begin{aligned} \mathbb{E}^{\mathbb{Q}} \left[\int_s^T \int_0^t \mathbb{E}^{\mathbb{Q}}[M_t M_\tau^{-1} u_\tau | \mathcal{F}_s] d\tau dt | \mathcal{F}_{s_0} \right] &= \int_s^T \int_0^t \mathbb{E}^{\mathbb{Q}}[M_t M_\tau^{-1} u_\tau | \mathcal{F}_{s_0}] d\tau dt \\ &= \int_s^T \left(\int_0^{s_0} M_\tau^{-1} u_\tau \mathbb{E}^{\mathbb{Q}}[M_t | \mathcal{F}_{s_0}] d\tau + \int_{s_0}^t \mathbb{E}^{\mathbb{Q}}[M_t M_\tau^{-1} u_\tau | \mathcal{F}_{s_0}] d\tau \right) dt \\ &= \int_s^T \left(e^{\sigma^2(t-s_0)} (Ku)_{s_0} dt + \int_{s_0}^t \mathbb{E}^{\mathbb{Q}}[\mathbb{E}^{\mathbb{Q}}[M_t M_\tau^{-1} u_\tau | \mathcal{F}_{s_0}] | \mathcal{F}_\tau] d\tau \right) dt \\ &= \int_s^T \left(e^{\sigma^2(t-s_0)} (Ku)_{s_0} dt + \int_{s_0}^t e^{\sigma^2(t-\tau)} \mathbb{E}^{\mathbb{Q}}[u_\tau | \mathcal{F}_{s_0}] d\tau \right) dt. \end{aligned}$$

Condition (11.6) can be rewritten

$$\begin{aligned} 0 &= \frac{\kappa\sigma^2}{\lambda} \int_s^T e^{\sigma^2(t-s_0)} (Ku)_{s_0} dt + \frac{\kappa\sigma^2}{\lambda} \int_s^T \int_{s_0}^t e^{\sigma^2(t-\tau)} \mathbb{E}^{\mathbb{Q}}[u_\tau | \mathcal{F}_{s_0}] d\tau dt + \frac{a}{\lambda} e^{\sigma^2(T-s_0)} \theta_{s_0} \\ &\quad + \frac{a}{\lambda} \int_{s_0}^T e^{\sigma^2(T-\tau)} \mathbb{E}^{\mathbb{Q}}[u_\tau | \mathcal{F}_{s_0}] + \mathbb{E}^{\mathbb{Q}}[u_s | \mathcal{F}_{s_0}], \end{aligned}$$

which proves the first statement of the theorem. We obtain the second point by a straightforward derivation of the functions z and v .

11.A.4 Proof of Theorem 11.2

The solution of the ODE for z is given by

$$z(s) = \theta_{s_0} e^{\sigma^2(s-s_0)} + \int_{s_0}^s e^{\sigma^2(s-u)} v(u) du,$$

and we can rewrite

$$v'(s) = \frac{\kappa\sigma^2}{\lambda} (\theta_{s_0} e^{\sigma^2(s-s_0)} + \int_{s_0}^s e^{\sigma^2(s-u)} v(u) du) - \frac{1}{\lambda} \mathbb{E}^{\mathbb{Q}}[\alpha_s | \mathcal{F}_{s_0}].$$

Multiplying by $e^{-\sigma^2 s}$ on both sides we have

$$e^{-\sigma^2 s} v'(s) = \frac{\kappa\sigma^2}{\lambda} \left(e^{-\sigma^2 s_0} \theta_{s_0} + \int_{s_0}^s e^{-\sigma^2 u} v(u) du \right) - \frac{e^{-\sigma^2 s}}{\lambda} \mathbb{E}^{\mathbb{Q}}[\alpha_s | \mathcal{F}_{s_0}],$$

and defining $w(s) = e^{-\sigma^2 s} v(s)$ we obtain

$$w'(s) = \frac{\kappa\sigma^2}{\lambda} \left(e^{-\sigma^2 s_0} \theta_{s_0} + \int_{s_0}^s w(u) du \right) - \sigma^2 w(s) - \frac{e^{-\sigma^2 s}}{\lambda} \mathbb{E}^{\mathbb{Q}}[\alpha_s | \mathcal{F}_{s_0}].$$

We note $y(s) = \int_{s_0}^s w(u) du$, satisfying the following differential equation

$$y''(s) = \frac{\kappa\sigma^2}{\lambda} y(s) - \sigma^2 y'(s) + \frac{\kappa\sigma^2}{\lambda} e^{-\sigma^2 s_0} \theta_{s_0} - \frac{e^{-\sigma^2 s}}{\lambda} \mathbb{E}^{\mathbb{Q}}[\alpha_s | \mathcal{F}_{s_0}].$$

Solving this ODE without second member, we have

$$y(s) = C_1 e^{\gamma_1 s} + C_2 e^{\gamma_2 s},$$

where $C_1, C_2 \in \mathbb{R}$, $\gamma_1 = \frac{-\sigma^2 - \sqrt{\Delta}}{2}$, $\gamma_2 = \frac{-\sigma^2 + \sqrt{\Delta}}{2}$, $\Delta = \sigma^2(\sigma^2 + 4\frac{\kappa}{\lambda}) > 0$. A particular solution is given by the function $y(s) = -\theta_{s_0} e^{-\sigma^2 s_0} + \frac{e^{-\sigma^2 s}}{\kappa\sigma^2} \mathbb{E}^{\mathbb{Q}}[\alpha_s | \mathcal{F}_{s_0}]$. The general solution is therefore given by:

$$y(s) = C_1 e^{\gamma_1 s} + C_2 e^{\gamma_2 s} - \theta_{s_0} e^{-\sigma^2 s_0} + \frac{e^{-\sigma^2 s}}{\kappa\sigma^2} \mathbb{E}^{\mathbb{Q}}[\alpha_s | \mathcal{F}_{s_0}].$$

To find C_1, C_2 we use the fact that $y(s_0) = 0$ and $y'(T) = w(T) = e^{-\sigma^2 T} v(T) = -\frac{e^{-\sigma^2 T} a}{\lambda} z(T)$. Substituting the previous expression of y , and making $a \rightarrow +\infty$ to ensure liquidation at terminal time, we obtain

$$\begin{aligned} C_1^\infty(s_0) &= \beta^\infty(s_0) \left(e^{\gamma_2 T - \sigma^2 s_0} \left(\theta_{s_0} - \frac{\alpha_{s_0}}{\kappa\sigma^2} \right) + e^{\gamma_2 s_0 - \sigma^2 T} \frac{\mathbb{E}^{\mathbb{Q}}[\alpha_T | \mathcal{F}_{s_0}]}{\kappa\sigma^2} \right), \\ C_2^\infty(s_0) &= \beta^\infty(s_0) \left(-e^{\gamma_1 T - \sigma^2 s_0} \left(\theta_{s_0} - \frac{\alpha_{s_0}}{\kappa\sigma^2} \right) - e^{\gamma_1 s_0 - \sigma^2 T} \frac{\mathbb{E}^{\mathbb{Q}}[\alpha_T | \mathcal{F}_{s_0}]}{\kappa\sigma^2} \right), \end{aligned}$$

where $\beta^\infty(s_0) = \frac{1}{e^{\gamma_1 s_0 + \gamma_2 T} - e^{\gamma_1 T + \gamma_2 s_0}} > 0$. Note that $v(s_0) = u_{s_0} = e^{\sigma^2 s_0} y'(s_0)$, which gives

$$u_{s_0} = \theta_{s_0} \Gamma(s_0) + \nu(s_0),$$

where

$$\begin{aligned} \Gamma(s_0) &:= \beta^\infty(s_0) \left(e^{\gamma_1 s_0 + \gamma_2 T} \gamma_1 - e^{\gamma_1 T + \gamma_2 s_0} \gamma_2 \right), \\ \nu(s_0) &= \beta^\infty(s_0) \left(\gamma_1 e^{(\gamma_1 + \sigma^2) s_0} \left(-\frac{\alpha_{s_0}}{\kappa\sigma^2} e^{\gamma_2 T - \sigma^2 s_0} + e^{\gamma_2 s_0 - \sigma^2 T} \frac{\mathbb{E}^{\mathbb{Q}}[\alpha_T | \mathcal{F}_{s_0}]}{\kappa\sigma^2} \right) \right. \\ &\quad \left. + \gamma_2 e^{(\gamma_2 + \sigma^2) s_0} \left(\frac{\alpha_{s_0}}{\kappa\sigma^2} e^{\gamma_1 T - \sigma^2 s_0} - e^{\gamma_1 s_0 - \sigma^2 T} \frac{\mathbb{E}^{\mathbb{Q}}[\alpha_T | \mathcal{F}_{s_0}]}{\kappa\sigma^2} \right) \right) - \frac{\alpha_{s_0}}{\kappa}. \end{aligned} \quad (11.11)$$

Substituting this expression in (11.8), the trader's position becomes

$$d\theta_t^u = (\nu(t) + \Gamma(t)\theta_t^u) dt + \sigma\theta_t^u dW_t.$$

Therefore, we have the optimal position defined by

$$\theta_t^{u^*} = H_t \int_0^t H_s^{-1} \nu(s) ds,$$

where $dH_t = \Gamma(t)H_t dt + \sigma H_t dW_t$. The optimal control is finally given explicitly at any time $t \in [0, T]$ by

$$u_t^* = \theta_t^{u^*} \Gamma(t) + \nu(t).$$

Chapter 12

Liquidity stress testing using optimal portfolio liquidation

12.1 Introduction

The European Securities and Markets Authority (ESMA) has set out guidance on liquidity stress testing supplementary to the existing requirements enshrined in the AIFMD and UCITS directives, with the ESMA guidelines coming into force on 30th September 2020. The core of the liquidity stress testing framework is a model that can be used to estimate liquidation times and costs in a reasonably realistic way for a portfolio of investments, including funds that can take short positions such as hedge funds. The main components of a liquidity stress testing framework are predefined stress tests, a market liquidity model that estimates liquidation cost and time, and a governance framework. This chapter focuses on the model of market liquidity applied to optimal portfolio liquidation for corporate bonds. The model needs to produce liquidation times and costs depending on the market volatility, market daily volumes, and bid-ask spreads. The trader faces a trade-off between liquidating quickly, resulting in unfavourable price changes or liquidating slowly, thus incurring Profit and Loss (hereafter PnL) volatility.

Since the seminal work of Almgren and Chriss in Almgren and Chriss [10], a vast literature on optimal execution has emerged. In the initial Almgren-Chriss framework, a trader, allowed to send only market orders, aims to liquidate a large position in one or several assets and minimize the direct costs. Direct costs include both transaction costs and market impact, where the latter refers to the fact that, on average, a large order moves the price in the sense of the order's direction (price goes up for a buy order and conversely for a sell order). Their framework has numerous extensions, such as models incorporating an order-flow, introducing a price limiter, considering both limit and market orders, and taking into account an adverse selection mechanism (see Cardaliaguet and Lehalle [63], Cartea and Jaimungal [66, 68], Huang, Jaimungal, and Nourian [166], Jaimungal and Kinzebulatov [172]).

The vast majority of the optimal liquidation models are designed for electronic markets with a central limit order book where the trader or broker can send limit and market orders. Surprisingly, the issue of optimal liquidation on OTC markets, especially for fixed income products such as corporate bonds or credit default swaps (CDS), has been the subject of little interest in academic research. In these markets, the notion of market microstructure is radically different as there is no market impact in the usual sense because of liquidity fragmentation. There is no unique source of liquidity but several dealers who receive requests for quotes (hereafter RFQ) for a given size of security. Given the price proposed by the dealer, the client can accept the transaction or suggest a better price. This particular mechanism of OTC markets is prone to more opacity, which has significant consequences for the nature of the price impact of a trade. For example, in Hendershott and Madhavan [151], the authors show a

notable price impact due to information leakage, when the client requests prices from several dealers. In Shachar [249], the authors provide evidence of “hot potato” trading: an initial client’s request changes the hands of dealers several times, while it is gradually incorporated into the price. The very notion of liquidity is ambiguous in such markets, and so is the price impact of a trade.

One of the works studying the price impact of individual transactions in an OTC market is Eisler and Bouchaud [104], where the authors focus on credit indices. They find only a quantitative (but not qualitative) difference in terms of order flow and price impact compared to electronic markets. Moreover, the study argues in favor of the idea of latent liquidity in OTC markets. This concept of liquidity modeling, introduced in Tóth, Lempriere, Deremble, De Lataillade, Kockelkoren, and Bouchaud [264], suggests that there is a latent volume which is not revealed in the observable order book because of the trading strategies of the market participants.

In this work, we adapt the concept of latent liquidity in order to capture portfolio liquidation costs as a function of liquidation time based on a relatively small number of market data inputs such as estimated average daily volume (ADV), volatility and bid-ask spread. In Section 12.2.1, we recall the framework of the latent order book model of Tóth, Lempriere, Deremble, De Lataillade, Kockelkoren, and Bouchaud [264] and derive the price impact equation, which allows one to estimate the volume available in a price range $[p, p + \Delta p]$, where p is a certain price level and $\Delta p > 0$ is arbitrarily small. In Section 12.2.2, we study the optimal liquidation for the one-asset case and due to the simplicity of the formulae, in the small size limit, we derive the optimal liquidation time explicitly. Using the same equation for the available volume we study a linearly liquidated portfolio of assets in Section 12.2.3. Finally, Section 12.3 is devoted to the numerical results where we applied our algorithm to a test portfolio of corporate bonds, the parameters of which are computed using real market data.

12.2 Framework

In this section, we first recall the framework of the Locally Linear Order Book (LLOB for short) model introduced in Tóth, Lempriere, Deremble, De Lataillade, Kockelkoren, and Bouchaud [264]. We then show how to use it to assess the costs of liquidation, in particular on OTC markets, for example, the corporate bonds market. Even though there is no order book for the corporate bond market, electronic trading platforms form a rough approximation of it. It is as if there exists an unobservable order book, hereafter called a latent order book, and that one can observe block prices as a function of the price-volume dynamics of this order book.

12.2.1 The locally linear order book model

Initially the LLOB model emerged from an empirical fact that in limit order books the latent volumes around the best ticks are linear in price deviation from the best price, even if it is not directly reflected in the order book. In this section we describe the LLOB model in the initial context of order driven markets. The general idea of the LLOB is that there exists a latent order book which, at any time t , aggregates the total intended volume to be potentially sold at price $p > 0$ or above $\mathcal{V}_+(t, p)$ and the total intended volume to be potentially bought at price p or below $\mathcal{V}_-(t, p)$. The latent volumes $\mathcal{V}_+(t, p)$ and $\mathcal{V}_-(t, p)$ are not the volumes revealed in the observable order book but the volumes that would be revealed as limit or market orders if the price comes closer to p at some point (in short, as stated in Tóth, Lempriere, Deremble, De Lataillade, Kockelkoren, and Bouchaud [264], the latent volumes reflect intentions that do not necessarily materialize).

Between t and $t + dt$, new buy and sell orders of unit volume may arrive at levels $p_t \mp u$ where $u > 0$, with corresponding intensity rates $\lambda(u)$. At the same time the buyers and sellers who have already sent orders at $p_t \mp u$ might want to change the price to $p_t \mp u'$, for $u' > 0$, at rate $\nu(u, u')$, or even cancel an order temporarily in the case $u' = +\infty$.

Let us assume that the price process p_t is a Brownian motion, which may not be well suited to order books due to microstructural effects, but is suitable to approximate the price process on OTC markets. We assume that either $u' = +\infty$ with rate $\nu_\infty(u)$ or that the change of price is a Brownian motion. We define $D(u) = \int_0^{+\infty} (u - u')^2 \nu(u, u') du'$ interpretable as the squared volatility of intentions. Let us denote by $\rho_\pm(t, u)$ a latent volume averaged over price paths, the equation for which is

$$\begin{cases} \partial_t \rho_\pm(t, u) = \frac{1}{2} \partial_{uu}^2 (\mathcal{D}(u) \rho_\pm(t, u)) - \nu_\infty(u) \rho_\pm(t, u) + \lambda(u), \\ \rho_\pm(t, u) = 0, \end{cases}$$

for $(t, u) \in \mathbb{R}_+ \times \mathbb{R}_-$, where $\mathcal{D}(u) = D(u) + \sigma^2$ and $\sigma > 0$ is the daily price volatility.

For arbitrary functions $D(u), \lambda(u), \nu_\infty(u)$ the explicit form of the stationary solution of the above PDE is not known. However, in the case where new orders appear uniformly, i.e. $\lambda(u) = \lambda$ and $\mathcal{D}(u) = \mathcal{D}$ independent of u , the exact stationary solution is

$$\rho(u) = \rho_\infty \left(1 - \exp \left(- \frac{u}{u^*} \right) \right), \quad (\rho)$$

where $\rho_\infty = \frac{\lambda}{\nu_\infty}$, $u^* = \sqrt{\frac{\mathcal{D}}{2\nu_\infty}}$. The function $\rho(u)$ is the density of order book trades, or more precisely for all $u' > 0$ $\rho(u') = \frac{dV}{du}(u')$, where V is the order book volume as a function of the distance from the mid-price. The meaning of u^* is the width of the linear price change zone: for $u \ll u^*$, that is for small deviations from the mid-price, the price depends linearly on the volume, whereas it stays constant for $u \gg u^*$. The constant ρ_∞ is understood to be the density of an order book trade far away from the mid-price p_t . Precisely, it is the inverse of the asymptotic large size market elasticity.¹ We define the asymptotic market elasticity as $\epsilon_{\text{asympt.}} = \frac{1}{\rho_\infty}$, and the “naive” market elasticity as $\epsilon_{\text{naive}} = \frac{\sigma}{\text{ADV}}$.

Let us consider a buy order. We integrate Equation (ρ) over u from mid-price to Δp . The resulting equation gives the order book volume as a function of price change:

$$V(\Delta p) = \rho_\infty \left(\Delta p - u^* \left(1 - e^{-\frac{\Delta p}{u^*}} \right) \right).$$

In the case $\Delta p \ll u^*$, the price impact varies as the square root of the trade size. For large trade sizes, the price impact is a linear function of the trade size, imitating the increasing cost of trading when the traded volume is bigger than one the market can digest.

Note that the above model corresponds to trades that can be done in a single day, and is considered as a one day model. We consider a “linear” liquidation in which the block of assets is unwound in equal parts over a number of days T which needs to be determined. If we denote the total block size as N and a daily trade size of $\frac{N}{T}$, the cost of trading each block is given by

$$C_{\text{block}}(T) = \frac{N}{T} \Delta p \left(\frac{N}{T} \right),$$

where $\Delta p(v)$ is defined as the solution of $V(w) = v$ for $w \in \mathbb{R}, v \in \mathbb{R}_+$. The total direct costs are given by the sum of C_{block} over the number of days, which is

$$\text{DC}(T) = C_{\text{block}}(T)T = N \Delta p \left(\frac{N}{T} \right).$$

We introduce the following parametrization for the quantity ρ_∞ :

$$\rho_\infty = \alpha_\infty \frac{\text{ADV}}{\sigma},$$

¹The asymptotic large size market elasticity is the incremental price needed to trade an incremental volume when trading volume is large relative to ADV.

to be compared with $\rho_\infty = \frac{\lambda}{\nu_\infty}$ in Tóth, Lemperiere, Deremble, De Lataillade, Kockelkoren, and Bouchaud [264]. In other words, we take the number of daily orders λ (of unit volume) equal to the average daily number of unit volumes (equal to ADV), and the rate at which buy or sell orders are canceled equal to $\frac{\sigma}{\alpha_\infty}$ where $\alpha_\infty > 0$ is a free parameter. We can rewrite the dimensionless parameter α_∞ in terms of market elasticity so that

$$\alpha_\infty = \frac{\frac{\sigma}{\text{ADV}}}{\frac{1}{\rho_\infty}} = \frac{\epsilon_{\text{naive}}}{\epsilon_{\text{asympt.}}},$$

So the physical meaning of and intuition behind this is the ratio of the “naive” market elasticity for trades not large compared with ADV to the value of market elasticity for trades materially larger than ADV. We also assume that the width of the linear region is of the order of one day’s price move, so that $u^* = \sigma$.

12.2.2 Single asset liquidation

To account for a trading firm’s risk aversion, we consider a running penalty proportional to the standard deviation of the PnL of the entire block liquidation, which is a measure of risk usually used in practice. As we assume a linear liquidation schedule, the penalty can be written

$$\phi(T) = \gamma \sqrt{\mathbb{V}(\text{PnL})} = \frac{\gamma}{\sqrt{3}} P_0 N \sigma \sqrt{T},$$

where γ is a number of standard deviations of the PnL representing the risk tolerance of the firm and P_0 is taken to be the bond price (in units in which par is 1) and N is the face amount. The effect of the volatility penalty is to incentivize the optimizer to not take too much time with liquidating the position.

Let us consider the case $\Delta p \ll u^*$. By taking a Taylor expansion of Equation (ρ) at $\frac{\Delta p}{u^*} = 0$, we have

$$\Delta p(V) = \sqrt{\frac{2Vu^*}{\rho_\infty}}.$$

This is the commonly assumed square root law for price penalty as a function of volume, see Bouchaud [54], for example. Given an assumed linear liquidation schedule the cost for each block is therefore given by

$$C_{\text{block}} = \frac{N}{T} \Delta p_{\text{block}} = \sqrt{\frac{2u^*}{\rho_\infty}} \left(\frac{N}{T}\right)^{3/2}.$$

The total cost is a sum of costs of all blocks and the volatility penalty

$$\text{TC}(T) = \text{DC}(T) + \phi(T) = \sqrt{\frac{2u^*}{\rho_\infty}} \frac{N^{3/2}}{\sqrt{T}} + \frac{\gamma P_0 N \sigma \sqrt{T}}{\sqrt{3}}.$$

Let us express this money amount in terms of a cost per bond:

$$c(T) = \frac{\text{TC}(T)}{N} = \sqrt{\frac{2Nu^*}{T\rho_\infty}} + \frac{\gamma P_0 \sigma \sqrt{T}}{\sqrt{3}}.$$

We now solve the optimal liquidation problem by setting the first derivative of c with respect to T equal to zero. Computations lead to the following optimal liquidation time:

$$T^* = \frac{\sqrt{3N}}{\gamma P_0 \sigma} \sqrt{\frac{2u^*}{\rho_\infty}},$$

which provides the cost per bond:

$$c(T^*) = 2 \left(\frac{2u^*}{\rho_\infty}\right)^{1/4} \frac{\sqrt{\gamma P_0 \sigma} N^{1/4}}{3^{1/4}}.$$

Remark 51. By setting $u^* = \sigma, \alpha_\infty = 1$ we obtain

$$\tilde{T}^* = \frac{\sqrt{6}}{\gamma P_0} \sqrt{\frac{N}{ADV}}, \quad \tilde{c}(T^*) = \frac{2^{5/4} \sqrt{\gamma P_0} \sigma}{3^{1/4}} \left(\frac{N}{ADV} \right)^{1/4}.$$

A similar analysis to the above may be done with the large trade size limit:

$$\lim_{N \rightarrow +\infty} c^* = 3^{-1/3} (2^{-2/3} + 2^{1/3}) (\gamma P_0)^{1/3} \left(\frac{N}{ADV} \right)^{1/3},$$

So in between the small and large limits the cost dependence on trade size changes from $N^{1/4}$ to $N^{1/3}$.

In the following section, we show how to extend this framework to the multi-asset case.

12.2.3 Portfolio liquidation

When moving to the multi-asset version of the optimization, one needs to create multi-asset versions of both the direct cost and the volatility penalty. The problem of cross-impact emerging when, for example, trades of a certain amount of one asset influence the price of another asset, is not treated in our model. Optimal liquidation models taking into account cross-impact (see, for example, Mastromatteo, Benzaquen, Eisler, and Bouchaud [207]) exist, however it is hard to estimate cross-impact matrices in OTC markets, notably due to fragmentation. For the sake of simplicity, we assume that the total direct cost is the sum of the individual direct costs.

The multi-asset volatility penalty for the portfolio is a straightforward extension of the single asset version. It is an integral over time of the covariances of the remaining positions. The position function $N^i(t)$ for the bond i that is linearly liquidated over time is given by

$$N^i(t) = N_0^i \left(1 - \frac{t}{T^{*i}} \right)_+,$$

where $N_0^i \in \mathbb{R}$ is the initial position in the bond i , and T^{*i} refers to the final liquidation time, such that $N^i(T^{*i}) = 0$. The total variance of the PnL can be expressed as a sum over covariance terms

$$\sum_{i,j=1}^d \sigma^i \sigma^j \rho^{i,j} N_0^i N_0^j \int_0^{\min(T^{*i}, T^{*j})} \left(1 - \frac{t}{T^{*i}} \right) \left(1 - \frac{t}{T^{*j}} \right) dt = \sum_{i,j=1}^d \frac{\sigma^i \sigma^j \rho^{i,j} N_0^i N_0^j}{2} \min(T^{*i}, T^{*j}) \left(1 - \frac{\min(T^{*i}, T^{*j})}{3 \max(T^{*i}, T^{*j})} \right).$$

Remark 52 (Calibration of α_∞). A simple and intuitive approach for fixing the value of α_∞ can be found by looking at the small size asymptotic limit formula for the optimal liquidation time:

$$T_{asympt}^* = \frac{\sqrt{6}}{\gamma P_0 \sqrt{\alpha_\infty}} \sqrt{\frac{N}{ADV}},$$

where α_∞ has been reintroduced. Let us assume a bond priced at par, and impose the condition $T_{asympt}^*(N = ADV) = 1$, implying that it is reasonable to trade the ADV in one day, therefore

$$\alpha_\infty = \frac{6}{\gamma^2}.$$

12.3 Numerical results

In this section we present numerical examples of optimal liquidation using our methodology. We first show an application to a long-short portfolio of two correlated bonds sharing same characteristics except that one is much more liquid than the other. Then, we present the results obtained on a long-short portfolio of 20 bonds. In all the numerical results, we choose a risk aversion parameter $\gamma = 0.5$.

12.3.1 Long-short portfolio with two correlated bonds

This test case demonstrates the disadvantages of a line by line liquidation of a long/short portfolio, typically used in vendor liquidity stress testing offerings. The test portfolio consists of two bond positions of the same size (27 and -27 respectively), where the bonds have same price of 141.49\$ and 7% annualized volatility, but different ADVs: 30 for the first bond and 3 for the second, so that the first one is more liquid.

A liquidation strategy based on individual liquidation would result in the more liquid bond being unwound rapidly and the less liquid one slower. But clearly the optimal way to liquidate this portfolio is to unwind these positions with the same liquidation strategy, especially the same timescale. This would minimize total PnL variance thereby allowing for a longer liquidation time and less costs.




Figure 12.1: Optimal portfolio liquidation costs with respect to correlation for different strategies.




Figure 12.2: Liquidation times for 2-bond case.

In Figure 12.1, we show the liquidation costs as a function of the correlation between the two bond price returns. We refer to individual optimization to be the line by line liquidation of the positions, each bond is liquidated independently of the others. Such individual liquidation implies that the liquidation cost for the portfolio is the sum of the liquidation for each bond, including the standard deviation penalty. By a naive strategy we refer to a strategy with $T_{naive}^{*i} = \frac{N^i}{ADV^i}$.

The liquidation costs are decreasing functions of the correlation for all strategies. The difference between portfolio optimization costs and individual optimization costs even for small correlations is the consequence of the choice of the penalty function, and we have no intention to compare the values directly. We are mostly interested to compare the costs dependence on the correlation level. Notably, in Figure 12.1, we see that the costs of portfolio liquidation are decreasing more steeply when the correlation level increases compared to the individual optimization. Before the 60% correlation level the decrease in costs is mostly linear for portfolio optimization, and for the correlation levels above 60% it becomes more concave. In this specific case, the line by line liquidation provides lower costs than the naive liquidation. However in general, this has no reason to be true.

In Figure 12.2, we show the liquidation times for 2 bonds as a function of correlation in the case of portfolio optimization. Below 45% correlation, the optimal liquidation times appear to be almost independent of correlation (3 days for the first bond and 6 days for the second bond). Then the liquidation time of the less liquid bond decreases so as to approach the liquidation time of the liquid bond. As correlation increases, both times increase, converging to the same optimal time of 10 days.

12.3.2 Long-short portfolio of 20 bonds

We have chosen a set of 20 random bonds from the USD Investment Grade and High Yield universes, with somewhat random position sizes assigned. In Table 12.1 we show the main characteristics of the portfolio and in Table 12.4 its correlation matrix.

We can summarize the bonds' parameters as (from 3rd April 2020 unless otherwise noted):

- The gross value is about \$40M: \$25M long and \$15M short.
- ADV is estimated available daily volume calibrated on TRACE volume data and varies from \$2M to \$21M per day across the 20 bond in this portfolio.
- Volatility is 22 business days (one month) historical volatility and varies from 3.8% to 63% annual.
- Bid-ask is set to 20bp to provide a minimum level for all bonds.

For the particular portfolio chosen and for the time period chosen (three months to early June 2020) the average correlation between bonds was about 20%. However, this may not be representative of typical bonds in typical time periods. It is therefore worth looking at the behavior of optimal liquidation cost and time versus correlation.

Bond	ADV \$M/day	min bid-ask	Annual. vol	Bond face amount \$M
1	3.0	0.20 %	7.0 %	27
2	3.0	0.20 %	8.8 %	-33
3	8.0	0.20 %	12.5 %	-24
4	2.5	0.20 %	4.9 %	-31.5
5	3.5	0.20 %	13.0 %	27
6	6.0	0.20 %	7.1 %	-2
7	4.5	0.20 %	21.5 %	-1.5
8	2.0	0.20 %	18.4 %	-1
9	5.0	0.20 %	3.8 %	-1
10	2.5	0.20 %	11.1 %	-0.71
11	5.0	0.20 %	32.9 %	42
12	3.0	0.20 %	13.0 %	-42
13	4.5	0.20 %	11.3 %	40
14	17.5	0.20 %	11.8 %	-40
15	21.5	0.20 %	10.8 %	37.5
16	20.5	0.20 %	63.4 %	2
17	2.5	0.20 %	60.7 %	1.5
18	9.5	0.20 %	11.7 %	1
19	3.0	0.20 %	26.4 %	-1
20	2.0	0.20 %	12.9 %	-0.77

Table 12.1: Test portfolio characteristics.

In Figure 12.3, we show the optimal portfolio liquidation cost - both the direct cost and the full cost including volatility penalty - versus pairwise correlation, with all off-diagonal correlations set to the same value.




Figure 12.3: Optimal liquidation costs with portfolio optimizer.




Figure 12.4: Optimal liquidation times with portfolio optimizer.

As expected, this long-short portfolio has a decreasing liquidation cost as correlation increases. As none of the other models are sensitive to correlation, their direct cost values are constant, though their volatility penalty naturally decreases with increasing correlation when evaluated using the portfolio cost function. As in the 2-bond example we can notice that the decrease of the costs for the portfolio optimization is steeper for all correlation levels compared to other strategies considered. We can also notice that the costs of portfolio optimization becomes concave for the correlation levels above 50%. Note that, in this example, the line by line optimization provides higher liquidation costs compared to a simple naive liquidation strategy.

It is also interesting to look at liquidation times for the optimal portfolio liquidation as correlation increases, shown in Figure 12.4. The optimizer is taking advantage of the higher correlation causing a reduced volatility penalty for slower liquidation. Both the median and maximum time are increasing monotonously with respect to the correlation level.

In all of these optimizations, there have been no constraints, apart from a very high time constraint of 100 days which was never effective. However, this model can be used to calculate an optimal liquidation strategy under time constraint, which is very useful for liquidity stress testing. Taking our example case portfolio, we can see that with a maximum liquidation deadline of 100 days, the median and maximum liquidation times are about 9 and 32 days, with a direct cost of \$0.805m. It is interesting to see the impact on the cost function as we decrease the deadline.

In Table 12.2, we compare the short deadline results where a time upper bound was used to constrain the optimizer. In Figure 12.5, we show the excess cost above the optimal liquidation cost due to deadline full liquidation shortening. Even though the median time to liquidate the portfolio underlyings was about 9 days in the optimal case, shortening the liquidation time cutoff to a maximum of 10 days only causes a minor increase in cost, but as the deadline becomes shorter the costs increase drastically.

Deadline	Portfolio liq.cost	Portfolio direct cost	Portfolio T-median	Portfolio T-max
100	1.605	0.805	9.4	32.5
20	1.607	0.812	9.4	20
15	1.612	0.829	8.9	15
10	1.635	0.903	8.2	10
7.5	1.677	1.004	7.5	7.5
5	1.780	1.176	5	5
3	1.999	1.503	3	3
2	2.282	1.875	2	2
1	3.061	2.772	1	1

Table 12.2: Short deadline costs and times comparison.




Figure 12.5: Short deadline premium.

In Table 12.3, we compare liquidation costs for naive, individual and portfolio optimization strategies and present the median and the maximum liquidation times for the portfolio optimization across different portfolios. The first portfolio corresponds to the test portfolio considered above with the correlation matrix presented in Table 12.4, and other portfolios are the ones with all correlations set to a certain level.

For every level of correlation and for the example of 20 bonds described in Table 12.1, we present the liquidation costs in the naive, individual and portfolio optimization case. For the last case, we also present the direct costs, the median and maximum liquidation time.

The methodology presented in this chapter allows one to obtain the optimal liquidation strategy for a portfolio of bonds in time proportional to $O(d^2)$ where d is the number of bonds. For the test portfolio

example, the method works in 5 seconds and for a portfolio of 1000 bonds, it takes less than 6 hours, which is reasonable in the context of liquidity stress testing.

Correlation	Naive liq. cost	Individual liq. cost	Portfolio liq. cost	Portfolio direct cost	Portfolio T-median	Portfolio T-max
Test	1.82	1.99	1.61	0.81	9.43	32.51
0%	1.90	2.05	1.73	0.86	7.19	23.45
10%	1.87	2.03	1.69	0.84	7.66	24.97
20%	1.84	2.01	1.65	0.82	8.22	25.67
30%	1.81	1.99	1.61	0.80	8.06	27.88
40%	1.78	1.97	1.57	0.78	8.36	29.22
50%	1.74	1.94	1.52	0.76	7.99	27.37
60%	1.71	1.92	1.46	0.74	8.73	29.86
70%	1.67	1.90	1.39	0.70	9.28	30.77
80%	1.63	1.87	1.30	0.67	10.45	36.20
90%	1.59	1.84	1.20	0.64	11.76	40.62

Table 12.3: Comparison between three type of liquidation.

	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16	17	18	19	20
1	1.00	0.10	0.12	0.14	0.14	0.16	0.42	0.30	0.30	0.10	0.16	-0.11	0.32	0.35	0.48	0.24	-0.01	0.52	0.24	0.43
2	0.10	1.00	0.30	-0.23	0.16	0.26	-0.03	0.23	0.04	0.26	0.30	0.24	0.29	0.04	-0.09	0.22	0.19	0.11	0.36	0.25
3	0.12	0.30	1.00	-0.05	-0.17	0.59	0.28	0.29	0.22	-0.02	0.26	0.11	0.40	0.33	0.29	0.39	0.14	-0.24	0.05	0.05
4	0.14	-0.23	-0.05	1.00	0.23	0.06	0.14	0.11	0.37	-0.13	0.05	0.31	0.22	0.30	0.31	0.25	0.36	0.25	0.25	0.19
5	0.14	0.16	-0.17	0.23	1.00	0.15	-0.31	0.34	0.09	0.12	-0.00	0.24	0.36	-0.03	0.07	0.13	-0.05	-0.02	0.23	-0.01
6	0.16	0.26	0.59	0.06	0.15	1.00	-0.00	0.40	0.46	0.22	0.15	0.13	0.35	0.25	0.13	0.23	-0.00	0.23	0.13	0.14
7	0.42	-0.03	0.28	0.14	-0.31	-0.00	1.00	0.08	0.39	0.13	0.03	0.01	0.17	0.08	0.08	0.09	0.09	0.27	0.03	0.28
8	0.30	0.23	0.29	0.11	0.34	0.40	0.08	1.00	0.48	0.36	0.10	-0.08	0.13	0.06	0.15	0.20	0.02	0.25	0.35	0.12
9	0.30	0.04	0.22	0.37	0.09	0.46	0.39	0.48	1.00	0.20	0.07	0.24	0.28	0.23	0.12	0.05	0.14	0.34	0.06	0.34
10	0.10	0.26	-0.02	-0.13	0.12	0.22	0.13	0.36	0.20	1.00	-0.05	0.19	0.05	-0.04	-0.12	-0.08	0.05	0.13	0.14	0.32
11	0.16	0.30	0.26	0.05	-0.00	0.15	0.03	0.10	0.07	-0.05	1.00	0.12	0.22	0.29	0.29	0.12	0.27	0.19	0.23	0.23
12	-0.11	0.24	0.11	0.31	0.24	0.13	0.01	-0.08	0.24	0.19	0.12	1.00	0.39	0.18	0.05	0.11	0.52	0.03	0.08	0.25
13	0.32	0.29	0.40	0.22	0.36	0.35	0.17	0.13	0.28	0.05	0.22	0.39	1.00	0.40	0.40	0.49	0.12	0.39	0.13	0.38
14	0.35	0.04	0.33	0.30	-0.03	0.25	0.08	0.06	0.23	-0.04	0.29	0.18	0.40	1.00	0.82	0.48	0.25	0.69	0.38	0.57
15	0.48	-0.09	0.29	0.31	0.07	0.13	0.08	0.15	0.12	-0.12	0.29	0.05	0.40	0.82	1.00	0.58	0.20	0.64	0.35	0.42
16	0.24	0.22	0.39	0.25	0.13	0.23	0.09	0.20	0.05	-0.08	0.12	0.11	0.49	0.48	0.58	1.00	0.18	0.50	0.42	0.34
17	-0.01	0.19	0.14	0.36	-0.05	-0.00	0.09	0.02	0.14	0.05	0.27	0.52	0.12	0.25	0.20	0.18	1.00	0.15	0.26	0.06
18	0.52	0.11	0.24	0.25	-0.02	0.23	0.27	0.25	0.34	0.13	0.19	0.03	0.39	0.69	0.64	0.50	0.15	1.00	0.28	0.61
19	0.24	0.36	0.05	0.25	0.23	0.13	0.03	0.35	0.06	0.14	0.23	0.08	0.13	0.38	0.35	0.42	0.26	0.28	1.00	0.27
20	0.43	0.25	0.05	0.19	-0.01	0.14	0.28	0.12	0.34	0.32	0.23	0.25	0.38	0.57	0.42	0.34	0.06	0.61	0.27	1.00

Table 12.4: Correlation matrix for the example set of bonds.

12.4 Conclusion

In this chapter, we presented an optimal portfolio liquidation model based on the Locally Linear Order Book framework with an application to liquidity stress testing on OTC markets. The model has only one free parameter to be calibrated. When the traded volume is small, the optimal liquidation time in the single asset case is obtained analytically and is proportional to the square root of the ratio between the volume being liquidated and the average daily volume. In the case of portfolio liquidation, our simple and reasonably fast optimization procedure established in this chapter can be applied.

Bibliography

- [1] F. Abergel, C.-A. Lehalle, and M. Rosenbaum. Understanding the stakes of high-frequency trading. *The Journal of Trading*, 9(4):49–73, 2014.
- [2] T. Adrian, A. Capponi, M. Fleming, E. Vogt, and H. Zhang. Intraday market making with overnight inventory costs. *Journal of Financial Markets*, 50:100564, 2020.
- [3] E. Agliardi and R. Agliardi. Financing environmentally–sustainable projects with green bonds. *Environment and Development Economics*, 24(6):608–623, 2019.
- [4] R. Aïd, D. Possamaï, and N. Touzi. Optimal electricity demand response contracting with responsiveness incentives. *ArXiv preprint arXiv:1810.09063*, 2018.
- [5] A. Alfonsi, A. Fruth, and A. Schied. Optimal execution strategies in limit order books with general shape functions. *Quantitative Finance*, 10(2):143–157, 2010.
- [6] Global Sustainable Investment Alliance. 2016 global sustainable investment review. Technical report, GSIA, 2017.
- [7] R.F. Almgren. Optimal execution with nonlinear impact functions and trading-enhanced risk. *Applied Mathematical Finance*, 10(1):1–18, 2003.
- [8] R.F. Almgren. Optimal trading in a dynamic market. Mathematics in finance working paper series 2009-2, Courant Institute of Mathematical Sciences, New York University, 2009.
- [9] R.F. Almgren and N. Chriss. Value under liquidation. *Journal of Risk*, 12:61–63, 1999.
- [10] R.F. Almgren and N. Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3(2): 5–40, 2001.
- [11] R.F. Almgren and B. Harts. A dynamic algorithm for smart order routing. Technical report, Princeton University and Modern Markets Initiative, 2008.
- [12] R.F. Almgren and J. Lorenz. Bayesian adaptive trading with a daily cycle. *The Journal of Trading*, 1(4):38–46, 2006.
- [13] R.F. Almgren and J. Lorenz. Adaptive arrival price. *Trading*, 2007(1):59–66, 2007.
- [14] R.F. Almgren, C. Thum, E. Hauptmann, and H. Li. Equity market impact. *Risk*, 7:58–62, 2005.
- [15] E. Alòs, J.A. León, and J. Vives. On the short-time behavior of the implied volatility for jump-diffusion models with stochastic volatility. *Finance and Stochastics*, 11(4):571–589, 2007.
- [16] H. Alsayed and F. McGroarty. Arbitrage and the law of one price in the market for American depository receipts. *Journal of International Financial Markets, Institutions and Money*, 22(5): 1258–1276, 2012.
- [17] O. Alvarez and A. Tourin. Viscosity solutions of nonlinear integro-differential equations. *Annales de l’Institut Henri Poincaré (C) Non Linear Analysis*, 13(3):293–317, 1996.

- [18] E. Alvi. First-order approach to principal-agent problems: a generalization. *The Geneva Papers on Risk and Insurance Theory*, 22(1):59–65, 1997.
- [19] S. Ambec and P. Lanoie. Does it pay to be green? A systematic overview. *The Academy of Management Perspectives*, 22(4):45–62, 2008.
- [20] A. Ang, V. Bhansali, and Y. Xing. Build America bonds. *The Journal of Fixed Income*, 20(1):67–73, 2010.
- [21] J.J. Angel, L.E. Harris, and C.S. Spatt. Equity trading in the 21st century. *The Quarterly Journal of Finance*, 1(01):1–53, 2011.
- [22] A. Araujo and H. Moreira. A general Lagrangian approach for non-concave moral hazard problems. *Journal of Mathematical Economics*, 35(1):17–39, 2001.
- [23] S. Arnuk and J. Saluzzi. Concept release on equity market structure. Technical report, The Federal Register, 2010.
- [24] K.J. Arrow. Liquidity preference. Lecture VI in *Lecture Notes for Economics 285, The Economics of Uncertainty*, pages 33–53, 1963.
- [25] International Capital Market Association. Green bond principles, 2016. Technical report, ICMA, 2016.
- [26] M. Avellaneda and S. Stoikov. High-frequency trading in a limit order book. *Quantitative Finance*, 8(3):217–224, 2008.
- [27] M. Avellaneda, J. Reed, and S. Stoikov. Forecasting prices from level-I quotes in the presence of hidden liquidity. *Algorithmic Finance*, 1(1):35–43, 2011.
- [28] M.J. Bachelet, L. Becchetti, and S. Manfredonia. The green bonds premium puzzle: the role of issuer characteristics and third-party verification. *Sustainability*, 11(4):1098, 2019.
- [29] A. Bachouch, C. Huré, N. Langrené, and H. Pham. Deep neural networks algorithms for stochastic control problems on finite horizon, part 2: numerical applications. *ArXiv preprint arXiv:1812.05916*, 2018.
- [30] E. Bacry and J.-F. Muzy. Hawkes model for price and trades high-frequency dynamics. *Quantitative Finance*, 14(7):1147–1166, 2014.
- [31] E. Bacry, S. Delattre, M. Hoffmann, and J.-F. Muzy. Modelling microstructure noise with mutually exciting point processes. *Quantitative Finance*, 13(1):65–77, 2013.
- [32] E. Bacry, S. Delattre, M. Hoffmann, and J.-F. Muzy. Some limit theorems for Hawkes processes and application to financial statistics. *Stochastic Processes and their Applications*, 123(7):2475–2499, 2013.
- [33] E. Bacry, I. Mastromatteo, and J.-F. Muzy. Hawkes processes in finance. *Market Microstructure and Liquidity*, 1(01):1550005, 2015.
- [34] M. Baker, D. Bergstresser, G. Serafeim, and J. Wurgler. Financing the response to climate change: the pricing and ownership of US green bonds. Technical Report w25194, National Bureau of Economic Research, 2018.
- [35] G. Bakshi, C. Cao, and Z. Chen. Do call prices and the underlying stock always move in the same direction? *The Review of Financial Studies*, 13(3):549–584, 2000.
- [36] G. Barles and C. Imbert. Second-order elliptic integro-differential equations: viscosity solutions’ theory revisited. *Annales de l’Institut Henri Poincaré (C) Non-Linear Analysis*, 25(3):567–585, 2008.

- [37] R. Bauer and P. Smeets. Social identification and investment decisions. *Journal of Economic Behavior & Organization*, 117:121–134, 2015.
- [38] C. Bayer, P. Friz, and J. Gatheral. Pricing under rough volatility. *Quantitative Finance*, 16(6): 887–904, 2016.
- [39] E. Bayraktar and M. Ludkovski. Liquidation in limit order books with controlled intensity. *Mathematical Finance*, 24(4):627–650, 2014.
- [40] M. Bellia. High-frequency market making: liquidity provision, adverse selection, and competition. *SSRN preprint 3074313*, 2017.
- [41] P. Bergault and O. Guéant. Size matters for OTC market makers: general results and dimensionality reduction techniques. *Mathematical Finance*, 31:279–322, 2021.
- [42] P. Bergault, D. Evangelista, O. Guéant, and D. Vieira. Closed-form approximations in multi-asset market making. *ArXiv preprint arXiv:1810.04383*, 2020.
- [43] L. Bergomi. Smile dynamics II. *SSRN preprint 1493302*, 2005.
- [44] L. Bergomi. Smile dynamics III. *SSRN preprint 1493308*, 2008.
- [45] L. Bergomi. *Stochastic volatility modeling*. Financial Mathematics. CRC press, 2015.
- [46] K. Bichteler. Stochastic integration and L^p -theory of semimartingales. *The Annals of Probability*, 9(1):49–89, 1981.
- [47] J. Bieligk, U. Horst, and S. Moreno-Bromberg. Trading under market impact: crossing networks interacting with dealer markets. *Journal of Economic Dynamics and Control*, 100:131–151, 2019.
- [48] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–654, 1973.
- [49] Financial Stability Board. Global shadow banking monitoring report 2015. Technical report, FSB, 2015.
- [50] N. Bollen. A note on the impact of options on stock return volatility. *Journal of Banking & Finance*, 22(9):1181–1191, 1998.
- [51] K. Borch. Equilibrium in a reinsurance market. *Econometrica*, 30(3):424–444, 1962.
- [52] A. Borodin and P. Salminen. *Handbook of Brownian motion - facts and formulae*. Birkhäuser, Basel, 2002.
- [53] B. Bouchard. Introduction to stochastic control of mixed diffusion processes, viscosity solutions and applications in finance and insurance. Lecture notes, Université Paris-IX Dauphine, 2007.
- [54] J.-P. Bouchaud. Price impact. In *Encyclopedia of Quantitative Finance*, ed. R. Cont, volume I-IV, pages 1688–1693. John Wiley & Sons, 2010.
- [55] D.T. Breeden and R.H. Litzenberger. Prices of state-contingent claims implicit in option prices. *Journal of Business*, 51(4):621–651, 1978.
- [56] D. Brigo and G. Di Graziano. Optimal execution comparison across risks and dynamics, with solutions for displaced diffusions. *SSRN preprint 2247951*, 2014.
- [57] M. Brolley and K. Malinova. Informed trading and maker-taker fees in a low-latency limit order market. *SSRN preprint 2178102*, 2013.
- [58] C. Brooks and I. Oikonomou. The effects of environmental, social and governance disclosures and performance on firm value: a review of the literature in accounting and finance. *The British Accounting Review*, 50(1):1–15, 2018.

- [59] E. Budish, P. Cramton, and J. Shim. The high-frequency trading arms race: frequent batch auctions as a market design response. *The Quarterly Journal of Economics*, 130(4):1547–1621, 2015.
- [60] S. Buti, B. Rindi, and I.M. Werner. Dark pool trading strategies, market quality and welfare. *Journal of Financial Economics*, 124(2):244–265, 2017.
- [61] A. Cadenillas, J. Cvitanić, and F. Zapatero. Optimal risk-sharing with effort and project choice. *Journal of Economic Theory*, 133(1):403–440, 2007.
- [62] T.D. Calabrese and T.L. Ely. Borrowing for the public good: the growing importance of tax-exempt bonds for public charities. *Nonprofit and Voluntary Sector Quarterly*, 45(3):458–477, 2016.
- [63] P. Cardaliaguet and C.-A. Lehalle. Mean field game of controls and an application to trade crowding. *Mathematics and Financial Economics*, 12(3):335–363, 2018.
- [64] P. Carr and R. Lee. Robust replication of volatility derivatives. Mathematics in finance working paper series 2008–3, Courant Institute of Mathematical Sciences, New York University, 2008.
- [65] P. Carr and D. Madan. Option valuation using the fast Fourier transform. *Journal of Computational Finance*, 2(4):61–73, 1999.
- [66] Á. Cartea and S. Jaimungal. Optimal execution with limit and market orders. *Quantitative Finance*, 15(8):1279–1291, 2015.
- [67] Á. Cartea and S. Jaimungal. A closed-form execution strategy to target volume weighted average price. *SIAM Journal on Financial Mathematics*, 7(1):760–785, 2016.
- [68] Á. Cartea and S. Jaimungal. Incorporating order-flow into optimal execution. *Mathematics and Financial Economics*, 10(3):339–364, 2016.
- [69] Á. Cartea, S. Jaimungal, and J. Ricci. Buy low, sell high: a high frequency trading perspective. *SIAM Journal on Financial Mathematics*, 5(1):415–444, 2014.
- [70] Á. Cartea, S. Jaimungal, and J. Penalva. *Algorithmic and high-frequency trading*. Cambridge University Press, 2015.
- [71] Á. Cartea, R. Donnelly, and S. Jaimungal. Algorithmic trading with model uncertainty. *SIAM Journal on Financial Mathematics*, 8(1):635–671, 2017.
- [72] Á. Cartea, R. Donnelly, and S. Jaimungal. Enhancing trading strategies with order book signals. *Applied Mathematical Finance*, 25(1):1–35, 2018.
- [73] H. Chade and J. Swinkels. The no-upward-crossing condition, comparative statics, and the moral-hazard problem. *Theoretical Economics*, 15(2):445–476, 2020.
- [74] Q. Chan-Wai-Nam, J. Mikael, and X. Warin. Machine learning for semi-linear PDEs. *Journal of Scientific Computing*, 79(3):1667–1712, 2019.
- [75] X. Chen and J. Sung. Managerial Compensation and Outcome Volatility. *SSRN preprint 3140205*, 2018.
- [76] J.-E. Colliard and T. Foucault. Trading fees and efficiency in limit order markets. *The Review of Financial Studies*, 25(11):3389–3421, 2012.
- [77] P. Collin-Dufresne and V. Fos. Do prices reveal the presence of informed trading? *The Journal of Finance*, 70(4):1555–1582, 2015.
- [78] J.R. Conlon. Two new conditions supporting the first-order approach to multisignal principal-agent problems. *Econometrica*, 77(1):249–278, 2009.

-
- [79] J. Conrad. The price effect of option introduction. *The Journal of Finance*, 44(2):487–498, 1989.
- [80] R. Cont and J. Da Fonseca. Dynamics of implied volatility surfaces. *Quantitative Finance*, 2(1): 45–60, 2002.
- [81] R. Cont and A. Kukanov. Optimal order placement in limit order markets. *Quantitative Finance*, 17(1):21–39, 2017.
- [82] R. Cont, S. Stoikov, and R. Talreja. A stochastic model for order book dynamics. *Operations Research*, 58(3):549–563, 2010.
- [83] M.G. Crandall, H. Ishii, and P.-L. Lions. User’s guide to viscosity solutions of second order partial differential equations. *Bulletin of the American Mathematical Society*, 27(1):1–67, 1992.
- [84] T. Cuk and A. Van Waeyenberge. European legal framework for algorithmic and high frequency trading (MIFID II and MAR): a global approach to managing the risks of the modern trading paradigm. *SSRN preprint 3622307*, 2020.
- [85] J. Cvitanić, X. Wan, and J. Zhang. Optimal contracts in continuous-time models. *International Journal of Stochastic Analysis*, 2006(095203), 2006.
- [86] J. Cvitanić, X. Wan, and J. Zhang. Principal-agent problems with exit options. *The B.E. Journal of Theoretical Economics*, 8(1):1–43, 2008.
- [87] J. Cvitanić, X. Wan, and J. Zhang. Optimal compensation with hidden action and lump-sum payment in a continuous-time model. *Applied Mathematics and Optimization*, 59(1):99–146, 2009.
- [88] J. Cvitanić, D. Possamaï, and N. Touzi. Moral hazard in dynamic risk management. *Management Science*, 63(10):3328–3346, 2017.
- [89] J. Cvitanić, D. Possamaï, and N. Touzi. Dynamic programming approach to principal-agent problems. *Finance and Stochastics*, 22(1):1–37, 2018.
- [90] A. Damodaran and M.G. Subrahmanyam. The effects of derivative securities on the markets for the underlying assets in the united states: a survey. Technical report, New York University Salomon Center, 1992.
- [91] K. Dayri and M. Rosenbaum. Large tick assets: implicit spread and optimal tick size. *Market Microstructure and Liquidity*, 1(01):1550003, 2015.
- [92] T. de Angelis, P. Tankov, and O.D. Zerbib. Environmental impact investing. *SSRN preprint 3562534*, 2020.
- [93] S. Delattre, C.Y. Robert, and M. Rosenbaum. Estimating the efficient price from the order flow: a Brownian Cox process approach. *Stochastic Processes and their Applications*, 123(7): 2603–2619, 2013.
- [94] R. Della Croce, C. Kaminker, and F. Stewart. The role of pension funds in financing green growth initiatives. Working papers on finance, insurance and private pensions 10, OECD, 2011.
- [95] P.M. DeMarzo and Y. Sannikov. Optimal security design and dynamic capital structure in a continuous-time agency model. *The Journal of Finance*, 61(6):2681–2724, 2006.
- [96] P.M. DeMarzo and Y. Sannikov. Learning, termination, and payout policy in dynamic incentive contracts. *The Review of Economic Studies*, 84(1):182–236, 2016.
- [97] J.S. Demski and D. Sappington. Optimal incentive contracts with multiple agents. *Journal of Economic Theory*, 33(1):152–171, 1984.
- [98] E. Derman. Regimes of volatility. *Risk*, 12(4):55–59, 1999.

- [99] E. Derman, I. Kani, and J.Z. Zou. The local volatility surface: unlocking the information in index option prices. *Financial Analysts Journal*, 52(4):25–36, 1996.
- [100] J. Detemple and P. Jorion. Option listing and stock returns: an empirical analysis. *Journal of Banking & Finance*, 14(4):781–801, 1990.
- [101] S.E. Dreyfus. Dynamic programming and the calculus of variations. *Journal of Mathematical Analysis and Applications*, 1(2):228–239, 1960.
- [102] C. Dunis, N. M Kellard, and S. Snaith. Forecasting EUR–USD implied volatility: the case of intraday data. *Journal of Banking & Finance*, 37(12):4943–4957, 2013.
- [103] W. E, J. Han, and A. Jentzen. Deep learning-based numerical methods for high-dimensional parabolic partial differential equations and backward stochastic differential equations. *Communications in Mathematics and Statistics*, 5(4):349–380, 2017.
- [104] Z. Eisler and J.-P. Bouchaud. Price impact without order book: a study of the OTC credit index market. *SSRN preprint 2840166*, 2016.
- [105] I. Ekeland and J. Lefournier. L’obligation verte : homéopathie ou incantation ? Technical report, Université Paris–Dauphine, 2019.
- [106] S. El Aoud and F. Abergel. A stochastic control approach to option market making. *Market Microstructure and Liquidity*, 1(01):1550006, 2015.
- [107] O. El Euch, M. Fukasawa, and M. Rosenbaum. The microstructural foundations of leverage effect and rough volatility. *Finance and Stochastics*, 22(2):241–280, 2018.
- [108] O. El Euch, J. Gatheral, and M. Rosenbaum. Roughening Heston. *Risk*, 5:84–89, 2019.
- [109] O. El Euch, T. Mastrolia, M. Rosenbaum, and N. Touzi. Optimal make-take fees for market making regulation. *Mathematical Finance*, 31:109–148, 2021.
- [110] R. Élie and D. Possamaï. Contracting theory with competitive interacting agents. *SIAM Journal on Control and Optimization*, 57(2):1157–1188, 2019.
- [111] R. Élie, E. Hubert, T. Mastrolia, and D. Possamaï. Mean-field moral hazard for optimal energy demand response management. *Mathematical Finance*, 31:399–473, 2021.
- [112] R.F. Engle and J.V. Rosenberg. Testing the volatility term structure using option hedging criteria. *The Journal of Derivatives*, 8(1):10–28, 2000.
- [113] G.-É. Espinosa and N. Touzi. Optimal investment under relative performance concerns. *Mathematical Finance*, 25(2):221–257, 2015.
- [114] M.-C. Fagart and C. Fluet. The first-order approach when the cost of effort is money. *Journal of Mathematical Economics*, 49(1):7–16, 2013.
- [115] W. Febi, D. Schäfer, A. Stephan, and C. Sun. The impact of liquidity risk on the yield spread of green bonds. *Finance Research Letters*, 27:53–59, 2018.
- [116] M.R. Fengler, W.K. Härdle, and E. Mammen. A semiparametric factor model for implied volatility surface dynamics. *Journal of Financial Econometrics*, 5(2):189–218, 2007.
- [117] C. Flammer. Green bonds: effectiveness and implications for public policy. *Environmental and Energy Policy and the Economy*, 1(1):95–128, 2020.
- [118] C. Flammer. Corporate green bonds. *Journal of Financial Economics*, to appear, 2021.
- [119] P.A. Forsyth. A Hamilton-Jacobi-Bellman approach to optimal trade execution. *Applied Numerical Mathematics*, 61(2):241–265, 2011.

-
- [120] P.A. Forsyth, J.S. Kennedy, S.T. Tse, and H. Windcliff. Optimal trade execution: a mean quadratic variation approach. *Journal of Economic Dynamics and Control*, 36(12):1971–1991, 2012.
- [121] T. Foucault, O. Kadan, and E. Kandel. Liquidity cycles and make/take fees in electronic markets. *The Journal of Finance*, 68(1):299–341, 2013.
- [122] A. Frino, V. Mollica, and S. Zhang. The impact of tick size on high frequency trading: the case for splits. *SSRN preprint 2607391*, 2015.
- [123] J. Gatheral. No-dynamic-arbitrage and market impact. *Quantitative Finance*, 10(7):749–759, 2010.
- [124] J. Gatheral. *The volatility surface: a practitioner’s guide*. Wiley Finance. John Wiley & Sons, 2011.
- [125] J. Gatheral and A. Jacquier. Arbitrage-free SVI volatility surfaces. *Quantitative Finance*, 14(1): 59–71, 2014.
- [126] J. Gatheral and A. Schied. Optimal trade execution under geometric Brownian motion in the Almgren and Chriss framework. *International Journal of Theoretical and Applied Finance*, 14 (03):353–368, 2011.
- [127] J. Gatheral, A. Schied, and A. Slynko. Transient linear price impact and Fredholm integral equations. *Mathematical Finance*, 22(3):445–474, 2012.
- [128] J. Gatheral, T. Jaisson, and M. Rosenbaum. Volatility is rough. *Quantitative Finance*, 18(6): 933–949, 2018.
- [129] M. Germain, H. Pham, and X. Warin. Neural networks-based algorithms for stochastic control and PDEs in finance. *ArXiv preprint arXiv:2101.08068*, 2019.
- [130] M. Germain, H. Pham, and X. Warin. Deep backward multistep schemes for nonlinear PDEs and approximation error analysis. *ArXiv preprint arXiv:2006.01496*, 2020.
- [131] L.R. Glosten and P.R. Milgrom. Bid, ask and transaction prices in a specialist market with heterogeneously informed traders. *Journal of Financial Economics*, 14(1):71–100, 1985.
- [132] S. Graf and H. Luschgy. *Foundations of quantization for probability distributions*, volume 1730 of *Lecture Notes in Mathematics*. Springer-Verlag Berlin Heidelberg, 2000.
- [133] J.R. Green and N.L. Stokey. A comparison of tournaments and contracts. *The Journal of Political Economy*, 91(3):349–364, 1983.
- [134] S.J. Grossman and O.D. Hart. An analysis of the principal-agent problem. *Econometrica*, 51(1): 7–45, 1983.
- [135] S.J. Grossman and M.H. Miller. Liquidity and market structure. *The Journal of Finance*, 43(3): 617–633, 1988.
- [136] O. Guéant. Execution and block trade pricing with optimal constant rate of participation. *Journal of Mathematical Finance*, 4:255–264, 2014.
- [137] O. Guéant. *The financial mathematics of market liquidity: from optimal execution to market making*. Financial Mathematics. Chapman and Hall/CRC, 2016.
- [138] O. Guéant. Optimal market making. *Applied Mathematical Finance*, 24(2):112–154, 2017.
- [139] O. Guéant and C.-A. Lehalle. General intensity shapes in optimal liquidation. *Mathematical Finance*, 25(3):457–495, 2015.

- [140] O. Guéant and I. Manziuk. Deep reinforcement learning for market making in corporate bonds: beating the curse of dimensionality. *Applied Mathematical Finance*, 26(5):387–452, 2019.
- [141] O. Guéant, C.-A. Lehalle, and J. Fernandez-Tapia. Optimal portfolio liquidation with limit orders. *SIAM Journal on Financial Mathematics*, 3(1):740–764, 2012.
- [142] O. Guéant, C.-A. Lehalle, and J. Fernandez-Tapia. Dealing with the inventory risk: a solution to the market making problem. *Mathematics and Financial Economics*, 7(4):477–507, 2013.
- [143] F. Guilbaud and H. Pham. Optimal high-frequency trading with limit and market orders. *Quantitative Finance*, 13(1):79–94, 2013.
- [144] F. Guilbaud and H. Pham. Optimal high-frequency trading in a pro rata microstructure with predictive information. *Mathematical Finance*, 25(3):545–575, 2015.
- [145] B. Hachenberg and D. Schiereck. Are green bonds priced differently from conventional bonds? *Journal of Asset Management*, 19(6):371–383, 2018.
- [146] S. Hamadène and M.-A. Morlais. Viscosity solutions for second order integro-differential equations without monotonicity condition: the probabilistic approach. *Stochastics*, 88(4):632–649, 2016.
- [147] J. Han and W. E. Deep learning approximation for stochastic control problems. *ArXiv preprint arXiv:1611.07422*, 2016.
- [148] J. Han, A. Jentzen, and W. E. Solving high-dimensional partial differential equations using deep learning. *Proceedings of the National Academy of Sciences*, 115(34):8505–8510, 2018.
- [149] L. Harris. Maker-taker pricing effects on market quotations. Technical report, USC Marshall School of Business, 2013.
- [150] M. Harris, C.H. Kriebel, and A. Raviv. Asymmetric information, incentives and intrafirm resource allocation. *Management Science*, 28(6):604–620, 1982.
- [151] T. Hendershott and A. Madhavan. Click or call? Auction versus search in the over-the-counter market. *The Journal of Finance*, 70(1):419–447, 2015.
- [152] P. Henry-Labordère. Deep primal-dual algorithm for BSDEs: applications of machine learning to CVA and IM. *SSRN preprint 3071506*, 2017.
- [153] N. Hernández Santibáñez and T. Mastrolia. Contract theory in a VUCA world. *SIAM Journal on Control and Optimization*, 57(4):3072–3100, 2019.
- [154] S.L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *The Review of Financial Studies*, 6(2):327–343, 1993.
- [155] A. Hirsra and W. Fu. An unsupervised deep learning approach in solving partial-integro differential equations. *ArXiv preprint arXiv:2006.15012*, 2020.
- [156] T. Ho and H.R. Stoll. Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial Economics*, 9(1):47–73, 1981.
- [157] B. Holmström. Moral hazard and observability. *The Bell Journal of Economics*, 10(1):74–91, 1979.
- [158] B. Holmström. Moral hazard in teams. *The Bell Journal of Economics*, 13(2):324–340, 1982.
- [159] B. Holmström and P. Milgrom. Aggregation and linearity in the provision of intertemporal incentives. *Econometrica*, 55(2):303–328, 1987.

- [160] H. Hong and M. Kacperczyk. The price of sin: the effects of social norms on markets. *Journal of Financial Economics*, 93(1):15–36, 2009.
- [161] B. Horvath, A. Jacquier, and P. Tankov. Volatility options in rough volatility models. *SIAM Journal on Financial Mathematics*, 11(2):437–469, 2020.
- [162] K. Hu, Z. Ren, and J. Yang. Principal-agent problem with multiple principals. *ArXiv preprint arXiv:1904.01413*, 2019.
- [163] W. Huang, C.-A. Lehalle, and M. Rosenbaum. Simulating and analyzing order book data: the queue-reactive model. *Journal of the American Statistical Association*, 110(509):107–122, 2015.
- [164] W. Huang, C.-A. Lehalle, and M. Rosenbaum. How to predict the consequences of a tick value change? Evidence from the Tokyo stock exchange pilot program. *Market Microstructure and Liquidity*, 2(03n04):1750001, 2016.
- [165] W. Huang, M. Rosenbaum, and P. Saliba. From Glosten-Milgrom to the whole limit order book and applications to financial regulation. *SSRN preprint 3343779*, 2019.
- [166] X. Huang, S. Jaimungal, and M. Nourian. Mean-field game strategies for optimal execution. *Applied Mathematical Finance*, 26(2):153–185, 2019.
- [167] C. Huré, H. Pham, A. Bachouch, and N. Langrené. Deep neural networks algorithms for stochastic control problems on finite horizon, part I: convergence analysis. *ArXiv preprint arXiv:1812.04300*, 2018.
- [168] C. Huré, H. Pham, and X. Warin. Some machine learning schemes for high-dimensional nonlinear PDEs. *ArXiv preprint arXiv:1902.01599*, 2019.
- [169] C. Huré, H. Pham, and X. Warin. Deep backward schemes for high-dimensional nonlinear PDEs. *Mathematics of Computation*, 89(324):1547–1579, 2020.
- [170] J. Jacod and A.N. Shiryaev. *Limit theorems for stochastic processes*, volume 288 of *Grundlehren der mathematischen Wissenschaften*. Springer-Verlag Berlin Heidelberg, 2003.
- [171] A. Jacquier, C. Martini, and A. Muguruza. On VIX futures in the rough bergomi model. *Quantitative Finance*, 18(1):45–61, 2018.
- [172] S. Jaimungal and D. Kinzebulatov. Optimal execution with a price limiter. *Risk*, 7:49–54, 2014.
- [173] A. Jain and C. Jain. Hidden liquidity on the US stock exchanges. *The Journal of Trading*, 12(3):30–36, 2017.
- [174] T. Jaisson and M. Rosenbaum. Limit theorems for nearly unstable Hawkes processes. *The Annals of Applied Probability*, 25(2):600–631, 2015.
- [175] T. Jaisson and M. Rosenbaum. Rough fractional diffusions as scaling limits of nearly unstable heavy tailed Hawkes processes. *The Annals of Applied Probability*, 26(5):2860–2882, 2016.
- [176] I. Jewitt. Justifying the first-order approach to principal-agent problems. *Econometrica*, 56(5):1177–1190, 1988.
- [177] P. Jusselin and M. Rosenbaum. No-arbitrage implies power-law market impact and rough volatility. *SSRN preprint 3180582*, 2018.
- [178] P. Jusselin, T. Mastrolia, and M. Rosenbaum. Optimal auction duration: a price formation viewpoint. *SSRN preprint 3399151*, 2019.
- [179] O. Kadan and J.M. Swinkels. On the moral hazard problem without the first-order approach. *Journal of Economic Theory*, 148(6):2313–2343, 2013.

- [180] J. Kalagnanam and D.C. Parkes. Auctions, bidding and exchange design. In *Handbook of Quantitative Supply Chain Analysis*, ed. D. Simchi-Levi, D. Wu and Z.-J. Shen, pages 143–212. Springer, 2004.
- [181] M. Kamal and J. Gatheral. Implied volatility surface. In *Encyclopedia of Quantitative Finance*, ed. R. Cont, volume I-IV, pages 926–931. John Wiley & Sons, 2010.
- [182] R. Ke and C.T. Ryan. A general solution method for moral hazard problems. *Theoretical Economics*, 13(3):1425–1481, 2018.
- [183] R. Kirkegaard. Moral hazard and the spanning condition without the first-order approach. *Games and Economic Behavior*, 102:373–387, 2017.
- [184] R. Kirkegaard. A unifying approach to incentive compatibility in moral hazard problems. *Theoretical Economics*, 12(1):25–51, 2017.
- [185] M. Klein. Tax credit bonds. *CitiBank Investment Management Review*, 11:27–31, 2009.
- [186] H. Koo, G. Shim, and J. Sung. Optimal multi-agent performance measures for team contracts. *Mathematical Finance*, 18(4):649–667, 2008.
- [187] S.J. Koopman, B. Jungbacker, and E. Hol. Forecasting daily variability of the S&P 100 stock index using historical, realised and implied volatility measurements. *Journal of Empirical Finance*, 12(3):445–475, 2005.
- [188] P. Kratz. *Optimal liquidation in dark pools in discrete and continuous time*. PhD thesis, Humboldt-Universität zu Berlin, 2011.
- [189] P. Kratz and T. Schöneborn. Optimal liquidation in dark pools. *Quantitative Finance*, 14(9):1519–1539, 2014.
- [190] P. Kratz and T. Schöneborn. Portfolio liquidation in dark pools in continuous time. *Mathematical Finance*, 25(3):496–544, 2015.
- [191] P. Krüger. Corporate goodness and shareholder wealth. *Journal of Financial Economics*, 115(2):304–329, 2015.
- [192] S. Laruelle and C.-A. Lehalle. *Market microstructure in practice*. World Scientific, 2nd edition, 2018.
- [193] S. Laruelle, C.-A. Lehalle, and G. Pagès. Optimal split of orders across liquidity pools: a stochastic algorithm approach. *SIAM Journal on Financial Mathematics*, 2(1):1042–1076, 2011.
- [194] S. Laruelle, C.-A. Lehalle, and G. Pagès. Optimal posting price of limit orders: learning by trading. *Mathematics and Financial Economics*, 7(3):359–403, 2013.
- [195] S. Laruelle, M. Rosenbaum, and E. Savku. Assessing MiFID 2 regulation on tick sizes: a transaction costs analysis viewpoint. *SSRN preprint 3256453*, 2018.
- [196] J. Lee and D. Ryu. The impacts of public news announcements on intraday implied volatility dynamics. *Journal of Futures Markets*, 39(6):656–685, 2019.
- [197] R.W. Lee. The moment formula for implied volatility at extreme strikes. *Mathematical Finance*, 14(3):469–480, 2004.
- [198] C.-A. Lehalle. Rigorous strategic trading: balanced portfolio and mean-reversion. *The Journal of Trading*, 4(3):40–46, 2009.
- [199] R.C.W. Leung. Continuous-time principal-agent problem with drift and stochastic volatility control: with applications to delegated portfolio management. Technical report, Haas School of Business, University of California Berkeley, 2014.

-
- [200] J. Li, R. Rao, and J. Shi. Learning to trade with deep actor critic methods. In *11th International Symposium on Computational Intelligence and Design (ISCID)*, volume 2, pages 66–71. IEEE, 2018.
- [201] P.-L. Lions. Hamilton-Jacobi-Bellman equations and the optimal control of stochastic systems. In Z. Ciesielski and C. Olech, editors, *Proceedings of the International Congress of Mathematicians, 1983, Warsaw*, volume 2, pages 1403—1417, 1983.
- [202] A. Lioui and P. Poncet. Optimal benchmarking for active portfolio managers. *European Journal of Operational Research*, 226(2):268–276, 2013.
- [203] A. Madhavan, M. Richardson, and M. Roomans. Why do security prices change? A transaction-level analysis of NYSE stocks. *The Review of Financial Studies*, 10(4):1035–1064, 1997.
- [204] I. Manziuk. *Optimal control and machine learning in finance: contributions to the literature on optimal execution, market making, and exotic options*. PhD thesis, Université Paris 1 Panthéon-Sorbonne, 2019.
- [205] T. Mastrolia and D. Possamaï. Moral hazard under ambiguity. *Journal of Optimization Theory and Applications*, 179(2):452–500, 2018.
- [206] T. Mastrolia and Z. Ren. Principal-agent problem with common agency without communication. *SIAM Journal on Financial Mathematics*, 9(2):775–799, 2018.
- [207] I. Mastromatteo, M. Benzaquen, Z. Eisler, and J.-P. Bouchaud. Trading lightly: cross-impact and optimal portfolio execution. *SSRN preprint 2949748*, 2017.
- [208] S. Mayhew. Implied volatility. *Financial Analysts Journal*, 51(4):8–20, 1995.
- [209] S. Mayhew and V. Mihov. How do exchanges select stocks for option listing? *The Journal of Finance*, 59(1):447–471, 2004.
- [210] N. Megarbane, P. Saliba, C.-A. Lehalle, and M. Rosenbaum. The behavior of high-frequency traders under different market stress scenarios. *Market Microstructure and Liquidity*, 3(03n04):1850005, 2017.
- [211] A.J. Menkveld. High frequency trading and the new market makers. *Journal of Financial Markets*, 16(4):712–740, 2013.
- [212] J.A. Mirrlees. Population policy and the taxation of family size. *Journal of Public Economics*, 1(2):169–198, 1972.
- [213] J.A. Mirrlees. Notes on welfare economics, information and uncertainty. In M.S. Balch, D.L. McFadden, and S.Y. Wu, editors, *Essays on economic behavior under uncertainty*, pages 243–261. Amsterdam: North Holland, 1974.
- [214] J.A. Mirrlees. The optimal structure of incentives and authority within an organization. *The Bell Journal of Economics*, 7(1):105–131, 1976.
- [215] J.A. Mirrlees. The theory of moral hazard and unobservable behaviour: part I. *The Review of Economic Studies*, 66(1):3–21, 1999.
- [216] C.C. Moallemi and M. Sağlam. OR Forum—the cost of latency in high-frequency trading. *Operations Research*, 61(5):1070–1086, 2013.
- [217] C.C. Moallemi and K. Yuan. A model for queue position valuation in a limit order book. Research paper No. 17–70, Columbia Business School, 2016.
- [218] R. Morel and C. Bordier. Financing the transition to a green economy: their word is their (green) bond. *Climate Brief*, 14, 2012.

- [219] O. Mounjid, M. Rosenbaum, and P. Saliba. From asymptotic properties of general point processes to the ranking of financial agents. *ArXiv preprint arXiv:1906.05420*, 2019.
- [220] H.M. Müller. Asymptotic efficiency in dynamic principal-agent problems. *Journal of Economic Theory*, 91(2):292–301, 2000.
- [221] B.J. Nalebuff and J.E. Stiglitz. Prizes and incentives: towards a general theory of compensation and competition. *The Bell Journal of Economics*, 14(1):21–43, 1983.
- [222] J. Neveu. *Discrete-parameter martingales*. North-Holland, Amsterdam, 1975.
- [223] J. Nilsson. Investment with a conscience: examining the impact of pro-social attitudes and perceived financial performance on socially responsible investment behavior. *Journal of Business Ethics*, 83(2):307–325, 2008.
- [224] A.A. Obizhaeva and J. Wang. Optimal trading strategy and supply/demand dynamics. *Journal of Financial Markets*, 16(1):1–32, 2013.
- [225] OECD. *Mobilising the debt capital markets for a low carbon transition*. Green finance and investment. Éditions OCDE, Paris, 2017.
- [226] OECD. *Investing in climate, investing in growth*. Éditions OCDE, Paris, 2017.
- [227] H. Ou-Yang. Optimal contracts in a continuous-time delegated portfolio management problem. *Review of Financial Studies*, 16(1):173–208, 2003.
- [228] G. Pagès, H. Pham, and J. Printems. Optimal quantization methods and applications to numerical problems in finance. In *Handbook of computational and numerical methods in finance*, ed. S.T. Rachev and G.A. Anastassiou, pages 253–297. Birkhäuser, Boston, MA, 2004.
- [229] S.K. Park. Investors as regulators: green bonds and the governance challenges of the sustainable finance revolution. *Stanford Journal of International Law*, 54:1–47, 2018.
- [230] H. Pham, X. Warin, and M. Germain. Neural networks-based backward scheme for fully nonlinear PDEs. *ArXiv preprint arXiv:1908.00412*, 2019.
- [231] M.E. Porter and C. Van der Linde. Toward a new conception of the environment-competitiveness relationship. *Journal of Economic Perspectives*, 9(4):97–118, 1995.
- [232] D. Possamaï, X. Tan, and C. Zhou. Stochastic control for a class of nonlinear kernels and applications. *The Annals of Probability*, 46(1):551–603, 2018.
- [233] J.W. Pratt. Risk aversion in the small and in the large. *Econometrica*, 32(1-2):122–136, 1964.
- [234] S. Predoiu, G. Shaikhet, and S. Shreve. Optimal execution in a general one-sided limit-order book. *SIAM Journal on Financial Mathematics*, 2(1):183–212, 2011.
- [235] R. Rabinovitch, A.C. Silva, and R. Susmel. Returns on ADRs and arbitrage in emerging markets. *Emerging Markets Review*, 4(3):225–247, 2003.
- [236] P. Renner and K. Schmedders. A polynomial optimization approach to principal-agent problems. *Econometrica*, 83(2):729–769, 2015.
- [237] P. Renner and K. Schmedders. Discrete-time dynamic principal-agent models: contraction mapping theorem and computational treatment. *Quantitative Economics*, 11(4):1215–1251, 2020.
- [238] C.Y. Robert and M. Rosenbaum. A new approach for the dynamics of ultra-high-frequency data: the model with uncertainty zones. *Journal of Financial Econometrics*, 9(2):344–366, 2010.
- [239] C.Y. Robert and M. Rosenbaum. Volatility and covariation estimation when microstructure noise and trading times are endogenous. *Mathematical Finance*, 22(1):133–164, 2012.

-
- [240] W.P. Rogerson. The first-order approach to principal-agent problems. *Econometrica*, 53(6):1357–1368, 1985.
- [241] P. Saliba. *High-frequency trading: statistical analysis, modelling and regulation*. PhD thesis, École Polytechnique, 2019.
- [242] P. Saliba. The information content of high-frequency traders aggressive orders: recent evidence. *Quantitative Finance*, 20:1779–1794, 2020.
- [243] Y. Sannikov. A continuous-time version of the principal-agent problem. *The Review of Economic Studies*, 75(3):957–984, 2008.
- [244] Y. Sannikov. Contracts: the theory of dynamic principal-agent relationships and the continuous-time approach. In D. Acemoglu, M. Arellano, and E. Dekel, editors, *Advances in economics and econometrics, 10th world congress of the Econometric Society, volume 1, economic theory*, number 49 in Econometric Society Monographs, pages 89–124. Cambridge University Press, 2013.
- [245] M. Schäl. A selection theorem for optimization problems. *Archiv der Mathematik*, 25(1):219–224, 1974.
- [246] H. Schättler and J. Sung. The first-order approach to the continuous-time principal-agent problem with exponential utility. *Journal of Economic Theory*, 61(2):331–371, 1993.
- [247] H. Schättler and J. Sung. On optimal sharing rules in discrete-and continuous-time principal-agent problems with exponential utility. *Journal of Economic Dynamics and Control*, 21(2):551–574, 1997.
- [248] A. Schied and T. Schöneborn. Risk aversion and the dynamics of optimal liquidation strategies in illiquid markets. *Finance and Stochastics*, 13(2):181–204, 2009.
- [249] O. Shachar. Exposing the exposed: intermediation capacity in the credit default swap market. *Federal Reserve Bank of New York working paper*, 2012.
- [250] B. Sinclair-Desgagné. The first-order approach to multi-signal principal-agent problems. *Econometrica*, 62(2):459–465, 1994.
- [251] J. Sirignano and K. Spiliopoulos. DGM: a deep learning algorithm for solving partial differential equations. *Journal of Computational Physics*, 375:1339–1364, 2018.
- [252] G. Skiadopoulos, S. Hodges, and L. Clewlow. The dynamics of the S&P 500 implied volatility surface. *Review of Derivatives Research*, 3(3):263–282, 2000.
- [253] A. Sokol. Optimal Novikov-type criteria for local martingales with jumps. *Electronic Communications in Probability*, 18(39):1–8, 2013.
- [254] H.M. Soner, N. Touzi, and J. Zhang. Wellposedness of second order backward SDEs. *Probability Theory and Related Fields*, 153(1-2):149–190, 2012.
- [255] H.M. Soner, N. Touzi, and J. Zhang. Dual formulation of second order target problems. *The Annals of Applied Probability*, 23(1):308–347, 2013.
- [256] M.A. Spence and R. Zeckhauser. Insurance, information, and individual action. *The American Economic Review*, 61(2):380–387, 1971.
- [257] S. Stoikov. The micro-price: a high-frequency estimator of future prices. *Quantitative Finance*, 18(12):1959–1966, 2018.
- [258] S. Stoikov and M. Sağlam. Option market making under inventory risk. *Review of Derivatives Research*, 12(1):55–79, 2009.

- [259] J. Sung. Linearity with project selection and controllable diffusion rate in continuous-time principal-agent problems. *The RAND Journal of Economics*, 26(4):720–743, 1995.
- [260] J. Sung. Corporate insurance and managerial incentives. *Journal of Economic Theory*, 74(2):297–332, 1997.
- [261] J. Sung. Optimal contracting under mean-volatility ambiguity uncertainties. *SSRN preprint 2601174*, 2015.
- [262] D.Y. Tang and Y. Zhang. Do shareholders benefit from green bonds? *Journal of Corporate Finance*, 61:101427, 2020.
- [263] M. Tomas and M. Rosenbaum. From microscopic price dynamics to multidimensional rough volatility models. *ArXiv preprint arXiv:1910.13338*, 2019.
- [264] B. Tóth, Y. Lempriere, C. Deremble, J. De Lataillade, J. Kockelkoren, and J.-P. Bouchaud. Anomalous price impact and the critical nature of liquidity in financial markets. *Physical Review X*, 1(2):021006, 2011.
- [265] B. Toth, Z. Eisler, and J.-P. Bouchaud. The short-term price impact of trades is universal. *Market Microstructure and Liquidity*, 3(02):1850002, 2017.
- [266] Y.-H. Wang and Y.-Y. Wang. The information content of intraday implied volatility for volatility forecasting. *Journal of Forecasting*, 35(2):167–178, 2016.
- [267] I.M. Werner and A.W. Kleidon. UK and US trading of British cross-listed stocks: an intraday analysis of market integration. *The Review of Financial Studies*, 9(2):619–664, 1996.
- [268] N. Williams. On dynamic principal-agent problems in continuous time. *University of Wisconsin, Madison*, 2009.
- [269] N. Williams. Persistent private information. *Econometrica*, 79(4):1233–1275, 2011.
- [270] N. Williams. A solvable continuous time dynamic principal-agent model. *Journal of Economic Theory*, 159(part B):989–1015, 2015.
- [271] M. Wyart, J.-P. Bouchaud, J. Kockelkoren, M. Potters, and M. Vettorazzo. Relation between bid-ask spread, impact and volatility in order-driven markets. *Quantitative Finance*, 8(1):41–57, 2008.
- [272] R. Zeckhauser. Medical insurance: a case study of the tradeoff between risk spreading and appropriate incentives. *Journal of Economic Theory*, 2(1):10–26, 1970.
- [273] O.D. Zerbib. The green bond premium. *SSRN preprint 2890316*, 2017.
- [274] O.D. Zerbib. The effect of pro-environmental preferences on bond prices: evidence from green bonds. *Journal of Banking & Finance*, 98:39–60, 2019.
- [275] O.D. Zerbib. A sustainable capital asset pricing model (S-CAPM): evidence from green investing and sin stock exclusion. *SSRN preprint 3455090*, 2019.

Titre : Finance quantitative à l'échelle de la microstructure : trading algorithmique et régulation.

Mots clés : market-making, trading optimal, microstructure des marchés, régulation, trading d'options, make-take fees

Résumé : Cette thèse est divisée en trois parties. Dans la première partie, nous appliquons la théorie Principal-Agent à certains problèmes de microstructure du marché. Premièrement, nous développons une politique d'incitation afin d'améliorer la qualité de la liquidité de marché dans le cadre d'une activité de market-making dans un lit et un dark pool gérés par la même bourse d'échange. Puis, nous adaptons ce design d'incitations à la régulation de l'activité de market-making lorsque plusieurs market-makers sont en concurrence sur une plateforme. Nous proposons également une forme d'incitation basée sur le choix des tailles de ticks asymétriques à l'achat et à la vente sur un actif. Nous abordons ensuite la question de la conception d'un marché de produits dérivés, en utilisant une méthode de quantization pour sélectionner les options listées sur la plateforme, et la théorie Principal-Agent pour créer des incitations pour un market-maker d'options. Enfin, nous développons un mécanisme d'incitations robuste à la spécification du modèle pour augmenter l'investissement dans les obligations vertes.

La deuxième partie de cette thèse est consacrée au market-making d'options en grande dimension. En faisant l'hypothèse de grecques constants nous proposons dans un premier temps un modèle pour traiter les options de longue maturité. Puis nous proposons une approximation de la fonction valeur permettant de traiter les grecques non-constants et les options de courte maturité. Enfin, nous développons

un modèle pour la dynamique haute fréquence de la surface de volatilité implicite. En utilisant des processus Hawkes multidimensionnels, nous montrons comment ce modèle peut reproduire de nombreux faits stylisés tels que le skew, le smile et la structure par termes de la surface.

La dernière partie de cette thèse est consacrée aux problèmes de trading optimal en grande dimension. Dans un premier temps, nous développons un modèle pour le trading optimal d'actions listées sur plusieurs plateformes. Pour un grand nombre de plateformes, nous utilisons une méthode d'apprentissage par renforcement profond pour calculer les contrôles optimaux du trader. Puis, nous proposons une méthodologie pour résoudre des problèmes de trading de façon approximativement optimale sans utiliser la théorie du contrôle stochastique. Nous présentons un modèle dans lequel un agent exhibe un comportement approximativement optimal s'il utilise le gradient de la trajectoire macroscopique comme signal de court terme. Enfin, nous présentons deux nouveaux développements sur la littérature d'exécution optimale. Tout d'abord, nous montrons que nous pouvons obtenir une solution analytique au problème d'exécution d'Almgren-Chriss avec mouvement Brownien géométrique et pénalité quadratique. Deuxièmement, nous proposons une application du modèle de carnet d'ordres latent au problème d'exécution optimale d'un portefeuille d'actifs, dans le cadre de stress tests de liquidité.

Title : Quantitative finance at the microstructure scale: algorithmic trading and regulation.

Keywords : market-making, optimal trading, market microstructure, regulation, options trading, make-take fees

Abstract : This thesis is split into three parts. In the first part, we apply the Principal-Agent theory to some problems of market microstructure. First, we build an incentives mechanism to improve the market quality in the context of market-making activity in a lit and a dark pool managed by the same exchange. Then, we adapt the incentives design to the regulation of market-making activity when several market-makers compete in a liquidity platform. We also propose a form of incentives based on the choice of tick sizes on the bid and ask sides of a single asset. Next, we tackle the issue of designing a derivatives market, using a quantization method to select the options listed on the exchange and the Principal-Agent framework to create incentives for an option market-maker. Finally, we develop an incentives mechanism to increase the investment in green bonds, robust to model specification, and outperforming current tax-incentives policies of the governments.

The second part of this thesis is dedicated to option market-making in high dimension. We first propose a framework a constant Greek assumption to deal with long-dated options. Then, we propose an approximation of the value function enabling to deal with time-varying Greeks and short-dated options. Finally, we develop a framework for the high-frequency

dynamics of the implied volatility surface. Using multidimensional Hawkes processes, we show how this setting can reproduce easily well-known stylized facts such as the skew, smile and term structure of the surface.

The last part of this thesis is devoted to optimal trading problems in high dimension. First, we develop a framework to tackle the smart order routing (SOR) problem taking into account non-stationarity of markets. For a large number of venues, we use a deep reinforcement learning approach to compute the optimal controls of the trader. Then, we present a methodology to solve approximately optimal trading problems without using stochastic control theory. We propose a framework in which a myopic agent exhibits approximately an optimal behavior if he uses the gradient of the high-level trajectory as short-term alpha. Finally, we present two new developments on the optimal execution literature. First, we show that we can obtain a closed-form solution for the Almgren-Chriss execution problem with geometric Brownian motion and quadratic penalty. Second, we propose an application of the latent order book model to the problem of optimal execution of a portfolio of assets, in the context of liquidity stress testing.