Optimal control in limit order books

Fabien Guilbaud

THÈSE DE DOCTORAT DE
L'UNIVERSITÉ PARIS VII DIDEROT

Spécialité

Mathématiques Appliquées

École doctorale de Paris Centre

Présentée par

Fabien Guilbaud

Pour obtenir le grade de Docteur de l'Université Paris VII Diderot

CONTROLE OPTIMAL DANS DES CARNETS D'ORDRES LIMITES

OPTIMAL CONTROL IN LIMIT ORDER BOOKS

soutenue publiquement le 1er février 2013 devant le jury composé de :

Frédéric Abergel,	Ecole Centrale de Paris	(Rapporteur)
Charles-Albert Lehalle,	Crédit Agricole Cheuvreux	(Examinateur)
Gilles Pagès,	Université Pierre et Marie Curie	(Examinateur)
Huyên Pham,	Université Paris Diderot	(Directeur)
Mathieu Rosenbaum,	Université Pierre et Marie Curie	(Examinateur)
Agnès Sulem,	INRIA Rocquencourt	(Examinateur)

Incertitude, ô mes délices,
Vous et moi nous nous en allons
Comme s'en vont les écrevisses,
A reculons, à reculons

Apollinaire, *in* Le Bestiaire

Contents

1 Introduction générale **15**
- 1.1 Objectifs et motivations . . . 15
- 1.2 Observations qualitatives et contexte . . . 17
  - 1.2.1 Présentation générale . . . 17
  - 1.2.2 Les différents types de carnets d'ordres limites . . . 18
  - 1.2.3 Les enjeux rencontrés dans l'industrie du trading haute fréquence . . . 20
- 1.3 Synthèse des principaux résultats . . . 24
  - 1.3.1 Le problème de l'exécution optimale . . . 24
  - 1.3.2 Trading haute fréquence optimal avec des ordres limites et au marché 28
  - 1.3.3 Trading haute fréquence optimal dans une microstructure au prorata avec information prédictive . . . 37

2 Introduction **45**
- 2.1 General objectives and motivations . . . 45
- 2.2 Qualitative observations and context . . . 47
  - 2.2.1 General presentation . . . 47
  - 2.2.2 The different types of limit order books . . . 47
  - 2.2.3 Issues faced in high-frequency trading industry . . . 50
- 2.3 Thesis outline and main results . . . 54
  - 2.3.1 Optimal execution problem . . . 54
  - 2.3.2 Optimal high-frequency trading with limit and market orders . . . 57
  - 2.3.3 Optimal high-frequency trading in a pro-rata microstructure with predictive information . . . 65

3 Literature survey: quantitative methods in high-frequency trading **76**
- 3.1 Introduction . . . 77
- 3.2 Costs optimization strategies . . . 78
- 3.3 Market-making and mixed strategies . . . 85

4 Numerical methods for an optimal order execution problem 95
4.1 Introduction . . . 96
4.2 Problem formulation . . . 98
4.2.1 The model of portfolio liquidation . . . 98
4.2.2 PDE characterization . . . 102
4.3 Time discretization and convergence analysis . . . 103
4.4 Numerical Algorithm . . . 108
4.5 Numerical Results . . . 116
4.5.1 Procedure . . . 116
4.5.2 Test 0: Convergence of the numerical scheme . . . 118
4.5.3 Test 1: A toy example . . . 120
4.5.4 Test 2: Short term liquidation . . . 123
4.5.5 Test 3: Sensitivity to Bid/Ask spread . . . 127

5 Optimal high frequency trading with limit and market orders 135
5.1 Introduction . . . 136
5.2 A market-making model . . . 139
5.2.1 Mid price and spread process . . . 139
5.2.2 Trading strategies in the limit order book . . . 139
5.2.3 Market making problem . . . 142
5.2.4 Parameters estimation . . . 142
5.3 Optimal limit/market order strategies . . . 147
5.3.1 Value function . . . 147
5.3.2 Dynamic programming equation . . . 149
5.4 Numerical scheme . . . 151
5.4.1 Mean criterion with penalty on inventory . . . 153
5.4.2 Exponential utility criterion . . . 154
5.5 Computational results . . . 156

6 Optimal HF trading in a pro-rata microstructure with predictive information 165
6.1 Introduction . . . 166
6.2 Market model . . . 168
6.3 Market making optimization procedure . . . 173
6.3.1 Control problem formulation . . . 173
6.3.2 Dynamic programming equation . . . 175
6.3.3 Dimension reduction in the Lévy case . . . 176
6.4 Numerical resolution . . . 179
6.4.1 Numerical scheme . . . 179

6.4.2 Convergence of the numerical scheme 181
6.4.3 Numerical tests 185
6.5 Best execution problem and overtrading risk 193
6.6 Conclusion 196

Remerciements

Je tiens en premier lieu à exprimer ma gratitude envers mon directeur de thèse Huyên Pham, qui m'a soutenu et accompagné pendant ces trois dernières années. J'ai eu la chance de bénéficier de son expertise, de son expérience et de ses précieux conseils tout au long de mon parcours de doctorant. Je lui suis également reconnaissant pour les multiples opportunités qu'il a pu m'offrir, et je ne doute pas que son apport à été déterminant non seulement pour mon éducation, mais aussi pour mon avenir professionnel. Enfin, je le remercie pour son sens de l'écoute et pour la confiance qu'il m'a accordée.

Je remercie également Fréderic Abergel et Sebastian Jaimungal d'avoir accepté d'être les rapporteurs de ma thèse. Je suis aussi reconnaissant envers le jury, Charles-Albert Lehalle, Gilles Pagès, Mathieu Rosenbaum et Agnès Sulem pour avoir examiné ma thèse et participé à ma soutenance. Je remercie la Fondation des Sciences Mathématiques de Paris pour son soutien financier pendant mon projet de recherche à l'University of California at Berkeley. Merci à Xin Guo et au laboratoire d'Industrial Engineering and Operational Research pour leur accueil durant ce projet, ainsi que pour les nombreux enseignements qu'ils m'ont apportés.

Je suis aussi reconnaissant envers les membres du Laboratoire de Probabilités et de Modéles Aléatoires des Universités Paris 6 et Paris 7, en particulier envers Jean Jacod et Lorenzo Zambotti pour m'avoir initié aux probabilités et aux mathématiques appliquées à la finance. Je remercie mes collègues et amis du laboratoire, notamment Clément Foucart, Adrien de Larrard, Thomas Lim, Simone Scotti, Paul Gassiat et Sophie Laruelle.

Pour leur soutien, leurs conseils, leur présence et leur amitié, je remercie chaleureusement Thomas Marta, Julien Krantz, Laure Peyrin, Emilie Zinsou, Anastasia Podzorova, Samuel de Bernard, Jonathan Lagier et Arnaud Filliol.

Pour conclure ces remerciements, je voudrais exprimer ma profonde gratitude envers ma famille, mon frère, mes parents et mes grands-parents pour leur soutien sans faille pendant ces années d'études, pour leurs encouragements et leur confiance. Enfin, j'ai une pensée affectueuse pour Lisa et Mathieu, à qui je dédie ce travail.

Résumé

On propose un traitement quantitatif de différentes problématiques du trading haute fréquence. On s'intéresse à plusieurs aspects de cette pratique, allant de la minimisation des frais indirects de trading, jusqu'à la tenue de marché, et plus généralement des stratégies de maximisation du profit sur un horizon de temps fini. On établit un cadre de travail original qui permet de refléter les spécificités du trading haute fréquence, notamment la distinction entre le trading passif et le trading actif, à l'aide de méthodes de contrôle stochastique mixte. On porte un soin particulier à la modélisation des phénomènes de marché en haute fréquence, et on propose pour chacun des méthodes de calibration compatibles avec les contraintes pratiques du trading algorithmique.

Dans le chapitre 3, on passe en revue la littérature sur les méthodes quantitatives appliquées au trading haute fréquence. En particulier, on s'intéresse en premier lieu aux travaux de modélisation des problèmes d'exécution optimale d'ordres. Dans un deuxième temps, on fait un tour d'horizon des stratégies déjà documentés de trading haute fréquence.

Dans le chapitre 4, on propose un modèle simple d'impact de marché non linéaire, qui permet de refléter les caractéristiques des frais indirects de trading. On étudie la situation d'un investisseur qui souhaite vendre son portefeuille, et on propose une stratégie qui optimise le revenu de cette vente. On résout numériquement le problème de contrôle impusionnel correspondant, à l'aide d'un schéma numérique dont on montre la convergence, et on étudie la performance de cette stratégie.

Dans le chapitre 5, on construit un modèle pour une microstructure de marché standard, correspondant aux actions européennes, et on développe les outils statistiques qui permettent sa calibration. Dans ce cadre, on s'intéresse à une stratégie de tenue de marché mixte ordres limites/ordres au marché. On formalise le problème de maximisation du profit sous la forme d'un problème de contrôle stochastique mixte, que l'on résout numériquement, à l'aide d'un schéma numérique dont on montre la convergence, et l'on procède à une analyse de performance détaillée.

Dans le chapitre 6, on construit un modèle pour une microstructure de marché exotique, correspondant par exemple aux futures de taux d'intérets, et on montre comment calibrer ce modèle. Ici encore, on étudie une stratégie de tenue de marché mixte, à l'aide de méthodes de contrôle mixte régulier/impulsionnel, que l'on résout numériquement. Dans ce cadre, on introduit un outil important en pratique qui nous permet d'utiliser une information prédictive sur l'évolution à court terme du prix dans la stratégie. Enfin, on étudie la performance de cette stratégie, et on développe d'autres exemples d'application de cette stratégie.

Abstract

We propose a quantitative approach to some high frequency trading problematics. We are interested in several aspects of this field, from minimizing indirect trading costs to market making, and more generally in profit maximization strategies over a finite time horizon. We build an original framework that reflects specificities of high frequency trading, and especially the distinction between passive and active trading, thanks to mixed stochastic control methods. We carefully model high fequency market phenomena, and for each of them we propose calibration methods that are compatible with practical constraints of algorithmic trading.

In chapter 3, we review the litterature on quantitative methods for high frequency trading. Firstly, we are interested in market impact and best execution problems modelling. Secondly, we provide an overview of profit-seeking high frequency trading strategies documented in academic litterature.

In chapter 4, we propose a simple model for non-linear market impact, which reflects general properties of indirect trading costs. We study the case of an investor that wants to unwind their portfolio, and provide a strategy that maximizes the revenue of this sale. We numerically solve the corresponding impulse control problem, using a numerical whose convergence is proven, and we study the behaviour and performance of the strategy.

In chapter 5, we build a model for a standard market microstructure, as encountered on european stocks for example, and we develop the statistical methods to calibrate the model. In this context, we consider a mixed market making strategy, where the investor can both provide and take liquidity in the market. We solve the problem of maximizing the profit using mixed stochastic control methods, that we solve numerically with a numerical scheme whose convergence is proven. We also provide a detailled performance analysis.

In chapter 6, we build a model for an exotic market microstructure, as encountered on interest rates for example, and we show how this model can be calibrated. In this case again, we consider a mixed market making strategy, that we study by means of mixed stochastic control, and that we solve numerically. In this context, we introduce a central tool for industrial application of high frequency trading that allows us to use predictive information on short term evolution of price. Finally, we provide a performance analysis for this strategy, and we show how to modify the model to cover the costs optimization problem as well.

Chapter 1

Introduction générale

1.1 Objectifs et motivations

L'objectif de cette thèse est de présenter une approche mathématique rigoureuse des aspects les plus communs du trading haute fréquence.

Du point de vue de l'ingénierie financière, nous contribuons au traitement quantitatif des enjeux suivants: minimisation des frais indirects de trading, tenue de marché, stratégie mixte de trading haute fréquence. Nous modélisons et étudions à la fois la microstructure standard à priorité prix/date et la microstructure exotique à priorité pro-rata. Nous cherchons à fournir un traitement complet de chaque situation, depuis la modélisation des phénomènes de marché, la résolution mathématique, jusqu'à la calibration et aux expériences numériques *a posteriori*, qui contiennent des résultats sur données réelles lorsque celles-ci sont disponibles.

Du point de vue mathématique, nous étendons des cadres de travail qui ont été récemment développés en mathématiques financières. Nous proposons la résolution détaillée de problèmes de contrôle stochastique mixte réguliers/impulsionnels. Nous construisons des schémas numériques originaux pour la résolution d'inégalités variationnelles de la programmation dynamique, qui correspondent aux contraintes pratiques du trading haute fréquence: nous mettons en place des méthodes de réduction de la dimension, ainsi que des algorithmes de calcul explicites, qui nous permettent d'accélerer la résolution de tels problèmes, donc d'utiliser ces algorithmes en temps réel. Nous montrons la convergence de chacun de ces schémas et disposons d'exemples numériques.

Dans le chapitre 4, nous examinons la situation d'un investisseur souhaitant vendre un grand portefeuille au meilleur prix. Nous fournissons la solution numérique au problème de contrôle impulsionnel correspondant à la modélisation de cette situation, et nous prenons notamment en compte les effets des frais indirects de trading et de l'impact de marché qui pénalisent les transactions trop rapides ou trop volumineuses. L'équation de la program-

mation dynamique correspondante est une quasi-inégalité aux variations avec une fonction de valeur satisfaisant une contrainte de solvalibilité au sens des solutions de viscosité contraintes. Nous proposons une solution numérique fondé sur un schéma numérique explicite rétrograde pour l'équation. La convergence de ce schéma est montrée par des arguments de solutions de viscosité. Nous proposons ensuite des résultats numériques illustrant le comportement du schéma numérique, la forme de la stratégie optimale, et une analyse de performance comparée.

Dans le chapitre 5, nous proposons un cadre de travail pour étudier les stratégies optimales de trading haute fréquence dans la microstructure standard à priorité prix/date, et nous proposons une application au trading d'actions européennes. Nous examinons la situation d'un investisseur dont l'objectif est de maximiser l'utilité espérée de son revenu sur un horizon de temps court, sachant que celui-ci est capable de traiter avec des ordres limites ou des ordres de marché, tout en controllant son niveau d'inventaire. Ceci est formulé comme un problème de contrôle mixte régulier/impulsionnel à changement de régime que nous caractérisons en termes de système quasi-variationnel par des méthodes de la programmation dynamique. Les procédures de calibration sont explicitées, de même que les exemples pratiques d'ajustement sur des données réelles. Nous construisons un schéma numérique explicite rétrograde par séparation pour résoudre ce problème, et montrons comment réduire le nombre des variables d'état jusqu'à un système où n'interviennent que les niveau de fourchette bid/ask et d'inventaire. Nous procédons ensuite à des expériences numériques sur données simulées et réelles, et nous fournissons une analyse de performance comparée pour la stratégie qui en résulte.

Dans le chapitre 6, nous proposons un cadre de travail pour étudier les strategies optimales de trading haute fréquence dans une microstructure exotique, la microstructure au pro-rata, avec application au trading de futures de taux courts. Ici encore, le trader haute fréquence est capable de traiter avec des ordres au marché ou avec des ordres limites, qui sont respectivement modélisés par des contrôles impulsionnels et réguliers. Nous modélisons et discutons les principaux risques caractéristiques de cette microstructure, qui sont liés aux fait que la taille des transactions n'est pas controllée. Nous nous intéressons à leur conséquences dans le cadre de l'exécution optimale. Le problème de trading optimal est étudié par des méthodes de contrôle stochastique et de programmation dynamique, ce qui conduit à la caractérisation de la fonction de valeur dans les termes d'une quasi-inégalité integro-variationnelle. Nous fournissons la procédure de résolution numérique qui y est associée, et sa convergence est prouvée. Nous proposons aussi des simplications de cet algorithme dans des cas particuliers à l'intérêt pratique: nous faisons en particulier la démonstration d'une stratégie de trading haute fréquence dans le cas où un indicateur prédictif sur le prix est disponible au trader. Chacun des exemples est illustré par des mesures de performance empirique.

Introduisons maintenant le contexte financier qui a conduit à ces modèles, ainsi que

les objets et mécanismes financiers qui sont au coeur du trading haute fréquence. Dans 1.2.1, nous proposons une vue d'ensemble du contexte du trading haute fréquence. Dans 1.2.2, nous présentons le vocabulaire général que nous allons utiliser au cours de cette thèse, et nous rappelons aussi des résultats généraux et des observations qualitatives sur la microstructure des marchés. Dans 1.2.3, nous résumons les principaux enjeux de l'industrie financière où une solution impliquant le trading haute fréquence est disponible. Enfin, nous ferons une synthèse des principaux résultats de cette thèse à la partie .

1.2 Observations qualitatives et contexte

1.2.1 Présentation générale

Le trading haute fréquence (HFT) est l'utilisation de stratégies automatisées pour réaliser des transactions sur des instruments financiers tels que les actions au comptant, les devises ou les produits dérivés, avec la caractéristique que les positions ne sont maintenues que pour une très courte période, allant de quelques secondes à quelques heures. Le terme recouvre plusieurs techniques de trading distinctes, qui sont souvent associées à l'usage de méthodes de décision purement quantitatives ou faisant un usage intense de données de marché, une infrastructure technologique lourde, ainsi qu'une exposition nulle en fin de journée.

Pourtant, grâce à l'essor des technologies de trading électronique, et à des évolutions règlementaires, de nombreux types d'investisseurs sont à présents capables de mettre en place des stratégies de trading haute fréquence. Les principales règlementations concernant le trading haute fréquence sont la MiFID en Europe (*Markets in Financials Instruments Directive*, directive sur les marchés d'instruments financiers, 1er novembre 2007) et la Reg-NMS aux Etats-Unis (*Regulation National Market System*, régulation du système du marché national, 2007). Toutes les deux visent à favoriser la compétition entre les places de marché, et promouvoir un mécanisme de formation des prix non-biaisé. Les résultats pratiques de ces cadres règlementaires sont le développement de places de marché alternatives, comme BATS ou Chi-X par exemple, et coïncidemment de nouveaux besoins en terme de fourniture de liquidité, acheminement des ordres de transactions et d'arbitrage.

Dans la plupart des marchés d'instruments financiers, le processus de formation des prix, ou de découverte du prix, résulte de la concurrence entre différents agents de marché qui prennent part à une enchère publique. En particulier, les sessions journalières de trading, ou phases de trading continu, consistent en des enchères doubles continues. Le trading haute fréquence est installé dans cette phase de trading continu, et ainsi l'étude des mécanismes précis qui réalisent cette double enchère continue est d'une importante centrale lors de la construction d'une stratégie. Ceci est précisement le sujet de la théorie de la microstructure des marchés : dans [56], la théorie de la microstructure des marchés est "l'étude des processus et des résultats de l'échange de biens sous un certain ensemble de règles. Alors que

la plupart de la théorie économique s'abstrait de la mécanique du trading, la théorie de la microstructure des marchés se concentre sur la question de savoir les effets de mécanismes spécifiques de trading sur le processus de formation des prix." Dans la section suivante, nous présentons les principaux mécanismes impliqués dans la formation des prix.

1.2.2 Les différents types de carnets d'ordres limites

Dans cette sous-section, nous présentons les mécanismes d'appariemment des ordres de transaction lors de la phase de trading continu, avec le vocabulaire que nous allons utiliser dans le reste de cette thèse.

La phase de trading continu est réalisée par un dispositif d'enchères doubles continues. Ceci signifie que la place de marché (par exemple la bourse Euronext Paris) diffuse publiquement de l'information au moins partielle sur les prix offerts pour la vente ou pour l'achat de produits financiers. Les fournisseurs de liquidité sont des agents de marché qui offrent ces prix, attendant qu'une contrepartie saisisse leur offre et crée ainsi une transaction. Ces fournisseurs de liquidité sont en concurrence dans une enchère à la fois l'achat (appelé le côté bid) et à la vente (appelé le côté ask). Les praticiens distinguent parfois entre les marchés conduits par les prix, et les marchés conduits par les ordres. Bien que la définition de ces termes puissent varier selon l'auteur, la distinction générale est la suivante:

- Les marchés conduits par les prix sont des marchés où les fournisseurs de liquidité (i.e. teneurs de marché) proposent un prix pour n'importe quel volume de transaction. En général, il y a dans ce cas un nombre réduit de teneurs de marché dédiés qui ont l'exclusivité de la fourniture de liquidité. Dans les marchés réels pourtant, par exemple les marchés de devises, le prix proposé par les teneurs de marché va dépendre du volume de transaction demandé par la contrepartie. Cette microstructure peut aussi se rencontrer sur des marchés plus rudimentaires, tels que les marchés de paris en ligne, où un teneur de marché monopolistique maintient des quotations de telle sorte que les parieurs soient approximativement au même nombre sur toutes les pattes du jeu.

- Les marchés conduits par les ordres sont des marchés où les fournisseurs de liquidité offrent une quantité donnée à un prix donné, soit à la vente, soit à l'achat, d'un actif financier. Contrairement à l'organisation précédente, tout participant du marché peut agir en tant que fournisseur de liquidité, grâce à l'usage d'ordres limites (voir ci-après). Ce mécanisme est la microstructure la plus commune sur les marchés financiers électroniques, elle peut être rencontrée par exemple sur les actions au comptant européennes, les matières premières, les taux d'intérets et les produits dérivés. Cette microstructure est réalisé par l'opération d'un carnet d'ordres limites, un objet que nous décrirons dans les paragraphes qui suivent.

Dans cette thèse, nous nous concentrons sur les marchés conduits par les ordres, puisque c'est l'organisation la plus courante. Définissons à présent ce qu'est un carnet d'ordres limites, et examinons deux règles différentes d'appariemment.

Nous mentionnons la revue de littérature complète [33] à propos du carnet d'ordres limites, de laquelle nous avons adapté les définitions qui suivent. Le rôle d'une place de marché est de rassembler et d'apparier les ordres de transaction, qui proviennent des participants de marché et peuvent êtres soumis à n'importe quel instant de la journée pendant la phase de trading continu. Ils sont de deux types:

Définition. Un ordre au marché de taille m est un ordre d'acheter (de vendre) m unités du bien que l'on traite au plus bas (haut) prix disponible sur le marché.

Définition. Un ordre limite de taille ℓ au prix p est un ordre d'acheter (de vendre) ℓ unités du bien que l'on traite au prix spécifié p.

Un ordre limite peut être soumis au marché, mis à jour en prix ou en quantité, ou encore annulé à tout instant, et ainsi on appelle:

Définition. Un ordre limite actif à l'instant t est un ordre limite qui a été soumis à $t_0 < t$, mais qui n'a pas été encore rempli ou annulé à la date t.

C'est précisemment l'ensemble des ordres limites actifs sur un marché qui font le carnet d'ordres limites:

Définition. Le carnet d'ordres limites (LOB), pour un actif financier spécifié, est l'ensemble des ordres limites actifs sur le marché à la date t pour cet actif.

Par ailleurs, chaque place de marché suit sa propre politique de diffusion publique de l'information sur les carnets d'ordres. Par exemple, les places de marchés les plus transparentes (*lit microstructure*) diffusent les volumes aggrégés offerts pour chaque prix, et usuellement ces donnés se présentent comme ceci 1.1:

	Ask		Bid	
	Prix	Quantité	Prix	Quantité
Level 1	50.01	80	49.98	120
Level 2	50.02	53	49.97	89
Level 3	50.03	81	49.96	64
Level 4	50.04	112	49.95	163
Level 5	50.05	44	49.94	101

Table 1.1: Représentation schématique à une certaine date t d'un carnet d'ordres limite. Dans cet exemple, à la date t, la taille du tick est 0.01 le prix ask est 50.01, le prix bid est 49.98, et la fourchette bid/ask est 0.03

Les marchés autorisent seulement un ensemble discret de prix possibles, et alors l'incrément

minimum entre deux prix possibles est appelé la taille du tick. Nous pouvons donc introduire les définitions suivantes:

Définition. Le prix bid à la date t est le prix le plus haut parmi tous les ordres limites d'achat actifs à la date t.

Définition. Le prix ask à la date t est le prix le plus bas parmi tous les ordres limites de vente actifs à la date t.

Définition. La fourchette bid/ask t est la différence entre le prix ask à la date t et le prix bid à la date t.

De plus, des notions plus complexes sont parfois associés au comportement dynamique du LOB. Parmi eux, le concept d'impact de marché fait référence au fait que le prix ask (resp. bid) soit détérioré, i.e. soit plus haut (resp. bas), après qu'un ordre au marché d'achat (resp. de vente) ait consommé plusieurs niveaux du LOB d'un coup. Le concept opposé à celui-ci est la résilience du carnet d'ordres, qui représente le fait que les niveaux du LOB ont tendance à se repeupler après avoir subi un impact de marché.

Finalement, donnons deux exemples pratiques de tels carnets d'ordres. D'abord, l'implémentation la plus commune du LOB est la microstructure prix/date. Cette microstructure est la plus fréquente sur les marchés modernes. Elle peut être rencontrée par exemple sur les actions européennes, avec différents niveaux de confidentialité des données selon le marché. Son principe est très simple: un ordre au marché entrant est apparié avec l'ordre limite actif au meilleur prix, le plus ancien dans le LOB. Une description détaillée de cette microstructure est fournie dans [21] et [37].

Une autre microstructure importante, quoique plus exotique, est la microstructure au pro rata. La microstructure au pro rata (voir [43] pour une présentation complète) peut être schématiquement décrite comme suit: lorsqu'un ordre au marché arrive dans le LOB au pro rata, son volume est réparti sur tous les ordres limites actifs au meilleur prix, proportionnellement au volume de chaque ordre, et ainsi est à l'origine de plusieurs transactions (voir figure 1.1

Cette microstructure est utilisée sur certains marchés de produits dérivés (par exemple le London International Financials Futures and options Exchange, ou Chicago Mercantile Exchange), et sera le sujet d'un chapitre entier de cette thèse.

1.2.3 Les enjeux rencontrés dans l'industrie du trading haute fréquence

Dans cette sous-partie, nous résumons les principaux enjeux industriels où l'on connait une solution impliquant le trading haute fréquence. Nous nous concentrons sur les enjeux stratégiques, et mettons de côté les enjeux technologiques, tels que les accès directs aux marchés ou l'optimisation de la vitesse du matériel de trading, quoiqu'ils soient pourtant des aspects cruciaux de cette pratique.

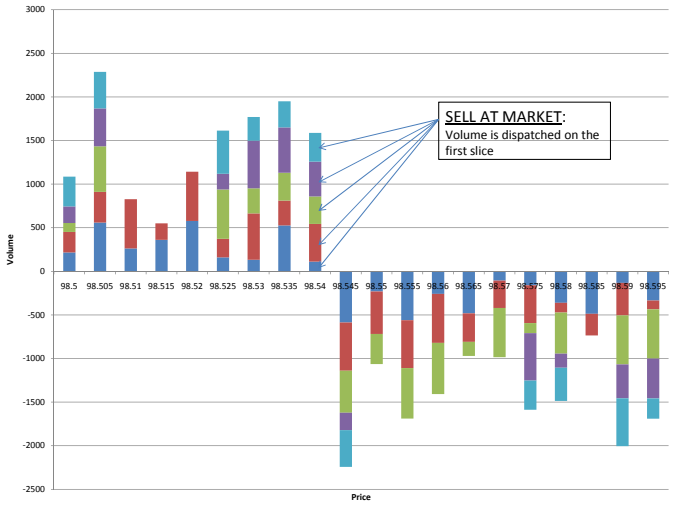

Figure 1.1: Schematic view of the pro-rata Limit Order Book.

Minimisation des coûts indirects de trading

La minimisation des coûts indirect de trading consiste à obtenir le prix le plus élevé possible pour une vente, ou obtenir le prix le moins élevé pour un achat.

Ce problème se présente naturellement lorsque le volume traité est grand, en raison des quantités finies de liquidité offerte dans le LOB (voir la section ci-dessus): en effet, une unique transaction de grand volume peut déséquilibrer le LOB en consommant plusieurs de ses niveaux d'un coup. Par exemple, si un investisseurs envoie un ordre au marché pour acheter 200 unités dans le carnet représenté à la table 1.1, le résultat de cette transaction est:

- 80 unités à 50.01
- 53 unités à 50.02
- 67 unités à 50.03

donc le prix ask à la fin de cette transaction est 50.03 avec un volume offert de 14. Ainsi, le prix moyen pondéré par les volumes de cette opération est: $(80 \times 50.01 + 53 \times 50.02 + 67 \times 50.03)/200 = 50.0193$, ce qui est à peu près un tick au-dessus du meilleur ask avant la transaction, ce qui fait une perte de 2bp. Cet effet s'appelle l'impact de marché. Afin de donner un point de comparaison, une stratégie traitant une fois par jour, et donc la performance attendu sur un an est de 5%, a un rendement journalier moyen d'environ 2bp,

lequel est balayé par l'impact de marché. Ainsi on voit que cet impact a une importance cruciale pour les gestionnaires de portefeuilles.

Les acteurs impliqués dans l'optimisation des coûts de trading sont à la fois des investisseurs comme les hedge funds ou les banques d'investissement, qui développent leur solutions propriétaires à ce problème, et les brokers, qui typiquement exécutent de grands volumes journaliers pour le compte de leurs clients. Les brokers sont, de plus, tenus par les règles MiFID et RegNMS d'utiliser des algorithmes d'exécution optimale. Certaines estimations annoncent que $70\% - 80\%$ des transactions sur les actions comptant européennes sont réalisés par de tels algorithmes [34].

Les solutions classiques à ce problème se rangent autour de deux idées centrales: les méthodes d'optimisation temporelles et les méthodes d'optimisation spatiales.

Les procédures d'optimisation spatiales ont été relativement peu documentées dans la littérature académique, quoique certains travaux soit disponible à ce sujet, par exemple [48]. L'idée soutenant cette méthode est de profiter du fait qu'un actif peut être traité sur plusieurs places de marché différentes. Ainsi, en découpant un ordre parent de grand volume en plusieurs petits ordres enfants, et en répartissant ceux-ci sur plusieurs places de marché différentes, l'investisseurs est en mesure de prendre plus de liquidité au même instant. Cette technique s'appelle Smart Order Routing (SOR), et est très populaire chez les brokers. La procédure d'optimisation dans de tels outils fait intervenir des outils de HFT de sorte à pouvoir rapidement mettre à jour l'agenda de trading.

Au contraire, les procédures d'optimisation temporelles ont reçu un traitement académique extensif, par exemple [3], [31] ou [35]. L'idée soutenant cette méthode est de découper un ordre parent de grand volume en plusieurs ordres enfants de plus petit volume, et de passer ceux-ci sur une periode de temps plus étendue. On peut voir la procédure d'optimisation ici comme un un équilibre à rechercher entre traiter rapidement, et alors être exposé à des risques d'impacter le marché, mais avoir moins de risque que le prix fluctue; ou traiter lentement, et avoir peu d'impact sur le marché, mais risquer que le prix fluctue pendant l'opération. Plusieurs solutions à ce problème ont été proposées, avec différents modèles de marché, mais la technique générale est de suivre un agenda de trading prédéfini (schéma de trading optimal) qui provient de cette procédure d'optimisation avec des hypothèses simplificatrices. Nous donnerons plus de détails sur ce sujet dans les sections suivantes.

Finalement, d'un point de vue industriel, plusieurs enjeux restent ouverts autour de ce sujet. D'abord, la détectabilité des algorithmes d'exécution est un enjeu central pour les brokers et les gestionnaires de portefeuille. En effet, l'utilisation massive d'algorithmes d'exécution est prise pour être à la source de l'autocorrélation dans les signes des transactions (voir [18]) et aussi de correlation retardée entre les données de transactions sur le même actif sur des marchés différents. Ainsi, de tels algorithmes sont très sensibles à la réponse du marché sur lesquels ils sont déployés, et moins efficaces lorsqu'ils sont facilement détectables par les concurrents. Ensuite, les stratégies mixtes ordres limites/marchés ont

jusqu'à présent reçu un traitement de moindre ampleur dans la littératures académique (voir [67] ou [37]), bien que les ordres limites soient beaucoup moins chers et dès lors massivement utilisés dans l'industrie.

Les stratégies "pure alpha"

Concentrons-nous maintenant sur les strategies "pure alpha", un terme de jargon qui fait référence aux stratégies de maximisation du profit qui sont largement indépendantes des conditions de marché. Cette catégorie inclus les stratégies suivantes:

- Les stratégies de tenue de marché. Cette classe de stratégie est fondée sur l'idée qu'en utilisant des ordres limites, on peut acheter au prix bid, revendre au prix ask, et ainsi gagner la fourchette bid/ask dans l'opération. De telles stratégies impliquent en général de fournir continuellement des cotations au bid et à l'ask, et de choisir de manière optimale les prix et les quantités de ces cotations. Le teneur de marché cherche alors à équilibrer son inventaire, c'est à dire à conserver une position sur l'actif risqué proche de zéro à toute date, et ainsi réduire son risque de marché.

- Les stratégies d'arbitrage statistique. Cette classe de stratégies est fondée sur l'idée que l'on peut exploiter les relations statistiques entre les prix des actifs (par exemple la structure de cointegration d'un certain secteur de marché, ou alors la relation entre un indice et ses composants) pour profiter d'inefficiences transitoires sur le marché. Ces stratégies font typiquement un usage intensif de données, elles sont directionnelles sur un horizon de temps court et répètent un grand nombre de fois le même pari afin de réduire la variance du résultat (d'où le nom d'arbitrage statistique). Souvent, ces stratégies sont agressives, dans le sens où elles prennent de la liquidité dans le LOB (*hit orders*). Elles sont aussi très dépendantes de la vitesse de l'infrastructure de trading, à cause de la concurrence d'autres acteurs utilisant une stratégie similaire.

- Les stratégies mixtes, qui sont une combinaison des deux classes ci-dessus.

Les acteurs impliqués dans de telles stratégies sont les banques d'investissement, les hedge funds, les firmes de trading propriétaire et les teneurs de marché spécialistes. Les avantages à conduire des stratégies pure alpha est que leur performance est très stable quelles que soient les conditions de marché, et ainsi l'investisseur est peu exposé au risque de marché au sens large. Au contraire, les inconvénients à conduire ces stratégies sont de deux types: d'abord, la performance totale est bornée la plupart du temps, dû au fait que les opportunités d'arbitrage sont rares, et ensuite, le risque opérationnel est élevé, puisque la performance technologique est d'une importance cruciale dans cette activité.

Cette classe de stratégie a été étudiée dans la littérature académique, avec une attention particulière pour les stratégies de tenue de marché.

D'abord, les stratégies de tenue de marché ont été présentées comme un problème de gestion d'inventaire depuis le travail pionnier d'Amihud et Mendelsohn en 1980 [5], et cette

approche a été modernisé par le travail d'Avellaneda et Stoikov en 2008 [7]. L'idée sous-jacente ici est d'adopter une approche profit/risque: l'objectif du teneur de marché est de "faire le spread", ce qui signifie acheter au bid et revendre à l'ask, et ainsi gagner la fourchette bid/ask. Alors qu'il fait cette opération, le teneur de marché est exposé au risque de marché, c'est à dire le risque de détenir une position non nulle en actif, sujet à des changements de prix. Ainsi, le teneur de marché a deux objectifs divergents: d'un côté, il souhaite participer au plus grand nombre de transactions possibles, afin de tirer profit de la fourchette bid/ask, et d'un autre côté, il doit maintenir son inventaire proche de zéro à toute date, afin de maintenir le risque de marché bas. Ceci à été l'objet de nombreux travaux récemment [16], [35] et [37].

Ensuite, les stratégies d'arbitrage statistique ont reçu moins de couverture de la littérature académique, malgré leur grande popularité parmi les traders haute fréquence. L'idée générale de telles stratégies est de construire un indicateur prédictif sur le prix de l'actif fondé sur l'observation de certains phénomènes de marché, et de traiter selon lui. Citons trois exemples pour illustrer ce principe. Dans l'article [6], l'auteur développe un approche généralisée du trading de paires: il pratique une analyse en composantes principales des rendements de plusieurs actifs d'un marché, et ainsi obtient un *portefeuille de marché* qui explique les rendement des actifs. A partir de là, l'hypothèse de trading est que le résidu entre chaque actif et le portefeuille de marché doit osciller autour de sa moyenne, et donc on traite en fonction de ce principe. Un autre exemple est disponible dans l'article [21], où les auteurs proposent une stratégie d'arbitrage très simple pour illustrer la pertinence d'un indicateur prédictif de prix fondé sur un modèle de la dynamique du LOB. Conditionnellement à l'état actuel du LOB, les auteurs sont capables de calculer la probabilité pour que le prix monte ou baisse dans les prochaines millisecondes, et proposent une stratégie qui exploite cette information. Enfin, dans le chapitre 6, nous proposons une manière d'inclure un tel indicateur prédictif de prix dans une stratégie mixte.

La prochaine section est consacrée à la synthèse des principaux résultats de cette thèse.

1.3 Synthèse des principaux résultats

1.3.1 Le problème de l'exécution optimale

Dans le chapitre 4, nous examinons le problème d'un investisseur qui souhaite fermer une grande position sur un actif risqué. Cette situation est présentée comme un équilibre à trouver entre le risque de marché et le risque d'impact. En effet, exécuter l'ordre lentement conduit à avoir un faible impact, mais beaucoup de risque de marché, et inversement, exécuter l'ordre rapidement conduit à avoir beaucoup d'impact, mais un faible risque de marché.

Plus précisemment, on vise à contrôler la différence entre la valeur marked to market

(ou valeur faciale) du portefeuille et le revenu réellement retiré de la vente de celui-ci. Cette différence, en défaveur de l'investisseur, est dûe à des effets d'illiquidité de l'actif qui incluent la fourchette bid/ask, les frais de broker et l'impact de marché. Nous discutons de la notion d'impact de marché, que l'on présente comme une réaction adverse du marché résultant de la finitude de la liquidité offerte sur le marché. Ce modèle a été inspiré des travaux précurseurs [10] et [3] qui ont les premiers introduit la notion d'impact dans un modèle à temps discret. Appliquer une approche de contrôle optimal au problème de l'exécution optimale d'ordre a été déjà documenté dans [63] et [28] avec des contrôles continus (l'approximation du trading continu), et dans [44] avec une approche de contrôle optimal impusionnel. Nous utilisons cette dernière approche car elle permet un modèle plus réaliste. Notre objectif est de trouver un agenda de trading optimal.

On propose ici une vue d'ensemble du modèle et de nos contributions.

Modèle de marché et stratégies de trading

On considère un marché où un investisseur veut vendre $y > 0$ unité d'un actif risqué avant la date T. On introduit les processus suivants:

- $(P_t)_{t\in[0,T]}$ le prix de l'actif
- $(X_t)_{t\in[0,T]}$ le montant de cash dans le portefeuille
- $(Y_t)_{t\in[0,T]}$ le nombre d'unités d'actif risqué dans le portefeuille ou inventaire
- $(\Theta_t)_{t\in[0,T]}$ le temps écoulé entre t et le dernier trade avant t

Les stratégies de trading sont une suite de contrôles impulsionnels:

$$\alpha = (\tau_n, \xi_n)_{n\in\mathbb{N}}$$

où (τ_n), représentant les dates de trading, sont des $\mathbb{F}$-temps d'arrêt et (ξ_n), représentant les quantités traitées, sont des variables $\mathcal{F}_{\tau_n}$-mesurable à valeurs dans $\mathbb{R}$. La dynamique de l'inventaire et du cash sont, sous α:

$$\begin{aligned}
\Theta_t &= t - \tau_n,\ \tau_n \le t < \tau_{n+1} \\
\Theta_{\tau_{n+1}} &= 0,\ n \ge 0. \\
Y_s &= Y_{\tau_n},\quad \tau_n \le s < \tau_{n+1} \\
Y_{\tau_{n+1}} &= Y_{\tau_n} + \xi_{n+1} \quad n \ge 0.
\end{aligned}$$

On suppose que le prix est un Brownien géométrique:

$$dP_t = P_t(bdt + \sigma dW_t)$$

Supposons que l'investisseur veuille traiter e. Si le prix actuel de l'actif est p, que le délai depuis le dernier trade est θ, alors le prix qu'il obtient réellement pour e est:

$$Q(e,p,\theta) = pf(e,\theta)$$

Nous pouvons utiliser de nombreuses fonctions f, mais nous prenons l'exemple suivant:

$$f(e,\theta) \quad = \quad \exp\big(\lambda|\frac{e}{\theta}|^{\beta}\mathrm{sgn}(e)\big).\big(\kappa_a\mathbf{1}_{e>0} + \mathbf{1}_{e=0} + \kappa_b\mathbf{1}_{e<0}\big),$$

Dans cette expression, $\kappa_a > 1$ et $\kappa_b < 1$ et ainsi $\big(\kappa_a\mathbf{1}_{e>0} + \mathbf{1}_{e=0} + \kappa_b\mathbf{1}_{e<0}\big)$ représente l'effet de traverser le spread. La partie exponentielle $\exp\big(\lambda|\frac{e}{\theta}|^{\beta}\mathrm{sgn}(e)\big)$ représente l'impact de marché, i.e. l'effet de la finitude de l'offre de liquidité. Des discussions sur le choix de f sont disponibles dans [50].

Le cash a la dynamique suivante:

$$\begin{aligned} X_t &= X_{\tau_n}, \quad \tau_n \le t < \tau_{n+1}, \ \ n \ge 0. \\ X_{\tau_{n+1}} &= X_{\tau_{n+1}^-} - \xi_{n+1}P_{\tau_{n+1}}f(\xi_{n+1}, \Theta_{\tau_{n+1}^-}) - \epsilon, \quad n \ge 0. \end{aligned}$$

Caractérisation par PDE

Nous choisissons une fonction d'utilité CRRA $U(x) = x^{\gamma}$ avec $\gamma \in (0,1)$ et notons $U_L(.) = U(L(.))$, où $L(.)$ est la fonction de liquidation, c'est à dire le revenu obtenu pour la liquidation du portefeuille. La fonction de valeur est définie par (on écrit $z = (x,y,p)$):

$$v(t,z,\theta) = \sup_{\alpha\in\mathcal{A}(t,z,\theta)} \mathbb{E}\big[U_L(Z_T)\big], \quad (t,z,\theta) \in [0,T]\times\mathcal{S}$$

où $\mathcal{A}(t,z,\theta)$ est un ensemble de contrôles approprié et $\mathcal{S} \subset \mathbb{R}^3$ la zone de solvabilité où vivent les variables d'état.

Selon [44] v est la solution de viscosité à l'HJBQVI:

$$\begin{aligned} \min\left[-\frac{\partial}{\partial t}v - \mathcal{L}v \ , \ v - \mathcal{H}v\right] &= 0, \quad \text{on} \quad [0,T)\times\mathcal{S}, \\ \min\left[v - U_L, v - \mathcal{H}v\right] &= 0, \quad \text{on} \quad \{T\}\times\mathcal{S}. \end{aligned}$$

où $\mathcal{L}$ est le générateur infinitésimal du processus (X,Y,P,Θ) dans une période sans trading:

$$\mathcal{L}\varphi = \frac{\partial}{\partial\theta}\varphi + bp\frac{\partial}{\partial p}\varphi + \frac{1}{2}\sigma^2p^2\frac{\partial^2}{\partial p^2}\varphi$$

et $\mathcal{H}$ est l'opérateur impulsionnel:

$$\mathcal{H}\varphi(t,z,\theta) = \sup_{e\in\mathcal{C}(t,z,\theta)} \varphi(t,\Gamma(z,\theta,e),0)$$

avec

$$\Gamma(z,\theta,e) = (x - epf(e,\theta) - \epsilon, y + e, p), \;\; z = (x,y,p) \in \mathcal{S}, \; e \in \mathbb{R}$$

en effet, pendant une période où on ne traite pas, le processus d'état évolue avec le prix P et la variable délai Θ d'une manière diffusive. Lorsqu'un contrôle impulsionnel apparaît, les variables sautent sous l'effet d'une transaction, avec une perte stricte de valeur faciale, ceci à cause des frais indirects de trading.

A présent, nous résolvons cette HJBQVI numériquement.

Schéma numérique explicite

Le choix du schéma numérique est important puisqu'il impacte le temps de calcul de la stratégie optimale. Nous choisissons un schéma rétrograde explicite en utilisant une propriété spécifique de notre problème. On considère le schéma de discrétisation en temps standard:

$$S^h(t,z,\theta,v^h(t,z,\theta),v^h) \;\; = \;\; 0, \quad (t,z,\theta) \in [0,T] \times \bar{\mathcal{S}},$$

avec

$$\begin{array}{rl} & S^h(t,z,\theta,r,\varphi) \\ := & \left\{ \begin{array}{ll} \min\Big[r - \mathbb{E}\big[\varphi(t+h, Z^{0,t,z}_{t+h}, \Theta^{0,t,\theta}_{t+h})\big] \;,\; r - \mathcal{H}\varphi(t,z,\theta)\Big] & \text{if } t \in [0, T-h] \\ \min\Big[r - \mathbb{E}\big[\varphi(T, Z^{0,t,z}_{T}, \Theta^{0,t,\theta}_{T})\big] \;,\; r - \mathcal{H}\varphi(t,z,\theta)\Big] & \text{if } t \in (T-h, T) \\ \min\Big[r - U_L(z,\theta) \;,\; r - \mathcal{H}\varphi(t,z,\theta)\Big] & \text{if } t = T. \end{array} \right. \end{array}$$

qui peut se formuler de manière équivalente sous forme de schéma implicite rétrograde:

$$\begin{array}{rcl} v^h(T,z,\theta) & = & \max\big[U_L(z,\theta) \;,\; \mathcal{H}v^h(T,z,\theta)\big], \\ v^h(t,z,\theta) & = & \max\Big[\mathbb{E}\big[v^h(t+h, Z^{0,t,z}_{t+h}, \theta+h)\big] \;, \mathcal{H}v^h(t,z,\theta)\Big], \quad 0 \le t \le T-h, \end{array}$$

et $v^h(t,z,\theta) = v^h(T-h,z,\theta)$ for $T-h < t < T$.

La manière usuelle de traiter de tels schéma est d'utiliser la récurrence:

$$\begin{array}{rcl} v^{h,n+1}(T,z,\theta) & = & \max\big[U_L(z,\theta) \;,\; \mathcal{H}v^{h,n}(T,z,\theta)\big], \\ v^{h,n+1}(t,z,\theta) & = & \max\Big[\mathbb{E}\big[v^{h,n+1}(t+h, Z^{0,t,z}_{t+h}, \theta+h)\big] \;, \mathcal{H}v^{h,n}(t,z,\theta)\Big], \end{array}$$

depuis $v^{h,0} = \mathbb{E}[U_L(Z^{0,t,z}_T, \Theta^{0,t,\theta}_T)]$. Grâce à l'effet de la variable délai Θ_t dans l'impact de marché, ce n'est pas optimal de traiter immédiatement après une transaction. Dès lors, on se ramène au schéma explicite rétrograde suivant:

$$\begin{array}{rcl} v^h(T,z,\theta) & = & \max\big[U_L(z,\theta) \;,\; \mathcal{H}U_L(z,\theta)\big], \\ v^h(t,z,\theta) & = & \max\Big[\mathbb{E}\big[v^h(t+h, Z^{0,t,z}_{t+h}, \theta+h)\big] \;,\; \sup_{e \in \mathcal{C}_\varepsilon(z,\theta)} \mathbb{E}\big[v^h(t+h, Z^{0,t,z^e_\theta}_{t+h}, h)\big]\Big], \end{array}$$

où $z_\theta^e = \Gamma(z, \theta, e)$

Convergence du schéma numérique

Nous prouvons la stabilité, la monotonicité et la cohérence du schéma numérique, et ainsi qu'il est convergent, suivant un argument adapté de Barles-Souganidis [8].

Analyse de performance

Nous fournisssons des résultats numériques obtenus avec notre implémentation. Nous proposons une analyse de performance comparée. Par exemple, nous illustrons le rendement de cette stratégie par sa performance finale 2.2, avec le détail de la méthodologie et des commentaires.

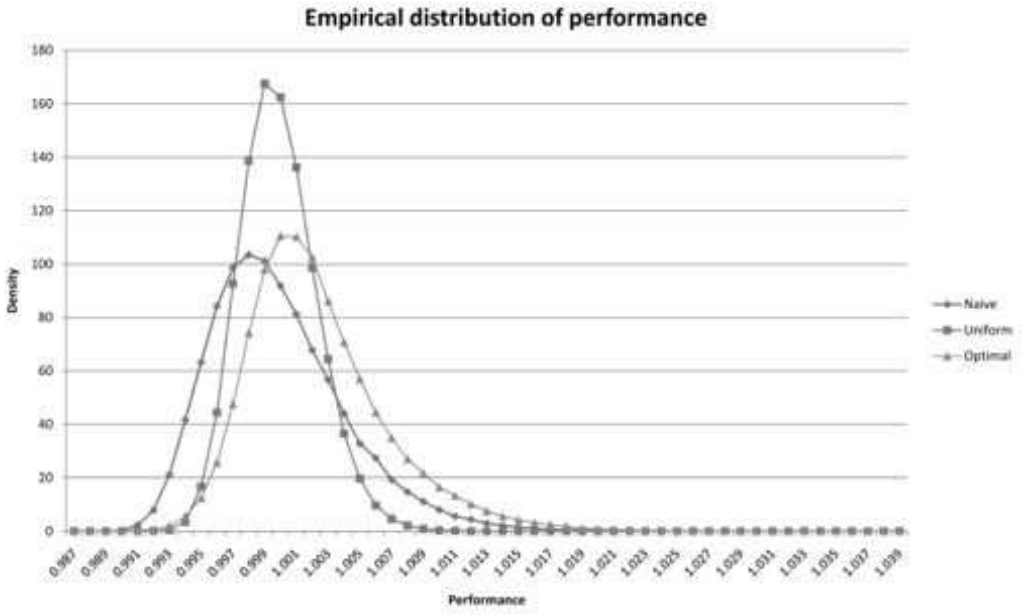

Figure 1.2: Optimal strategy perfrmance empirical distribution.

1.3.2 Trading haute fréquence optimal avec des ordres limites et au marché

Dans le chapitre 5, nous passons à un autre aspect important du trading haute fréquence, les stratégies de tenue de marché. La tenue de marché est l'action de fournir en permanence de la liquidité sur le marché en traiter avec des ordres limites. Dans ce travail, nous

examinons la situation d'un investisseur qui est capable de traiter avec des ordres limites, mais aussi avec des ordres au marché, et ainsi nous considérons une stratégie de type mixte. L'objectif de l'investisseur est de maximiser l'utilité de son profit sur un horizon de temps fini. Notre but est d'obtenir un modèle simple et facile à manier, toute en gardant une modélisation précise de la microstructure sous-jacente. Nous choisissons le contexte de la microstructure à priorité prix/date, qui est la plus standard, et peut être rencontrée par exemple sur les actions au comptant. Nous proposons un modèle facile à calibrer qui reflète les éléments centraux de la microstructure prix/date: en particulier, le modèle permet de reproduire des comportements divers pour le spread, et nous prenons notamment en compte le fait que le marché peut réagir aux actions de l'investisseur. Nous représentons cette situation comme un problème de contrôle stochastique mixte, que l'on étudie par des méthodes de programmation dynamique, et nous fournissons un schéma numérique rapide, grâce une méthode de réduction de la dimension des variables d'état. Nous prouvons que ce schéma est convergent, et proposons des illustrations numériques ainsi qu'une analyse de performance comparée.

Nous proposons d'examiner les risques suivants

- Risque d'inventaire: risque de détenir une position non nulle d'un actif dont le prix fluctue
- Risque d'exécution: incertitude que les ordres limites seront exécutés
- Risque de sélection adverse: le marché réagit de manière adverse aux action de l'investisseur

Notre objectif est de prendre en compte ces trois risques dans notre stratégie de tenue de marché. Nous adoptons l'approche de gestion d'inventaire qui a été développé par le travail Avellaneda et Stoikov [7]: le teneur de marché peut soumettre des cotations au bid et à l'ask avec une taille unitaire, à n'importe quel prix autour d'un prix mid, et l'arrivé d'ordre au marché de contrepartie est modélisée par un processus de Poisson donc l'intensité dépend de la distance avec le prix price. Ce modèle conduit à conserver l'inventaire proche de zéro à toute date. D'autres articles récents proposent des approches suivant cette même ligne [35] et [16].

Modèle de marché et stratégies de trading

On suppose que le prix mid est un processus de Markov P avec générateur $\mathcal{P}$ à valeurs dans $\mathbb{P}$. Le nombre de mise à jour du prix, l'horloge du tick-time est un processus ponctuel $(N_t)_t$ avec une intensité déterministe $\lambda(t)$. Sous l'horloge tick-time, la fourchette bid/ask est supposée être une chaîne de Markov stationnaire $(\hat{S}_n)_{n\in\mathbb{N}}$ à valeurs dans $\mathbb{S} = \delta\mathbb{I}_m$, $\mathbb{I}_m = \{1, \ldots, m\}$, où δ est la taille du tick. On défini aussi sa matrice de transition $(\rho_{ij})_{ij}$: $\rho_{ij} = \mathbf{P}[\hat{S}_{n+1} = j\delta | \hat{S}_n = i\delta]$, $i, j \in \mathbb{I}_m$, $\rho_{ii} = 0$. En temps calendaire, le spread est donc: $S_t =$

$\hat{S}_{N_t}$, supposé indépendant de P. Puis les prix bid et ask sont définis par:

$$P_t^b \;=\; P_t - \frac{S_t}{2}, \qquad P_t^a \;=\; P_t + \frac{S_t}{2}.$$

Décrivons maintenant les stratégies de trading. D'abord les ordres limites (make strategy) sont modélisés comme des contrôles continus:

$$\alpha_t^{make} \;=\; \{(Q_t^b, L_t^b), (Q_t^a, L_t^a)\}$$

où Q_t^b représente la cotation au bid $\mathcal{Q}^b = \{Bb, Bb_+\}$, ce qui signifie:

- Bb: meilleur prix bid, et Bb_+: meilleur prix bid + un tick (pour gagner la priorité d'exécution)
- L^b: taille de l'ordre limite d'achat dans $[0, \bar{\ell}]$

et Q_t^a représente la cotation à l'ask $\mathcal{Q}^a = \{Ba, Ba_-\}$, ce qui signifie:

- Ba: meilleur prix ask, et Ba_-: meilleur prix ask − un tick (pour gagner la priorité d'exécution)
- L^a: taille de l'ordre limite de vente dans $[0, \bar{\ell}]$

Dans ce contexte, on décrit la dynamique des variables d'états qui représentent le portefeuille. Lorsque l'on conduit une stratégie aux ordres limites $\alpha_t^{make} = \{(Q_t^b, L_t^b), (Q_t^a, L_t^a)\}$, l'inventaire Y et le cash X évoluent selon:

$$\begin{aligned} dY_t &= L_t^b dN_t^b - L_t^a dN_t^a, \\ dX_t &= \pi^a(Q_t^a, P_{t^-}, S_{t^-}) L_t^a dN_t^a - \pi^b(Q_t^b, P_{t^-}, S_{t^-}) L_t^b dN_t^b. \end{aligned}$$

où

$$\begin{aligned} \pi^a(q^a, p, s) &= p + \frac{s}{2} - \delta 1_{q^b = Ba_-} \\ \pi^b(q^b, p, s) &= p - \frac{s}{2} + \delta 1_{q^b = Bb_+}, \end{aligned}$$

et où nous avons introduit les processus de trade N^a et N^b, qui comptent les transaction apparaissant à l'ask et au bid, qui sont, plus précisément:

- N_t^a: arrivée d'un ordre au marché d'achat rencontrant un ordre limite de vente $\sim$ $\mathrm{Cox}(\lambda^a(Q_t^a, S_t))$: $\lambda^a(Ba, s) < \lambda^a(Ba_-, s)$
- N_t^b: arrivée d'un ordre au marché de vente rencontrant un ordre limite d'achat $\sim$ $\mathrm{Cox}(\lambda^b(Q_t^b, S_t))$: $\lambda^b(Bb, s) < \lambda^b(Bb_+, s)$

Notons que l'intensité des processus de trade dépend des ordres limites de l'investisseur (Q_t^a, Q_t^b), ce qui est pertinent pour modéliser une réaction adverse du marché, ou comme ici une dépendance au spread actuel.

La stratégie d'ordres au marché est modélisée par les contrôles impulsionnels $\alpha^{take} = (\tau_n, \zeta_n)_{n\geq 0}$ où $(\tau_n)_n$ est une suite croissante de temps d'arrêts représentants les temps de décisions et ζ_n sont $\mathcal{F}_{\tau_n}$-mesurables, représentant les quantités achetées à l'ask (si $\zeta_n \geq 0$), ou vendues au bid (si $\zeta_n < 0$). Ces ordres au marché sont exécutés immédiatement, conduisant aux sauts suivants:

$$\begin{aligned} Y_{\tau_n} &= Y_{\tau_n^-} + \zeta_n \\ X_{\tau_n} &= X_{\tau_n^-} - c(\zeta_n, P_{\tau_n}, S_{\tau_n}), \end{aligned}$$

où

$$c(e,p,s) = ep + |e|\frac{s}{2} + \varepsilon,$$

avec $\varepsilon > 0$ représentant un frais fixe.

Estimation

La section suivante est consacrée à l'estimation des paramètres du modèle. Nous nous intéressons d'abord à calibrer le modèle de spread. Nous supposons que (S_t) est observable. Et nous recontruisons:

- Les ticks times $(\theta_n)_n$ définis par:

$$\theta_{n+1} = \inf\left\{t > \theta_n : S_t \neq S_{t^-}\right\}, \qquad \theta_0 = 0.$$

- Le processus ponctuel qui lui est associé:

$$N_t = \#\left\{\theta_j > 0 : \theta_j \leq t\right\}, \ t \geq 0,$$

- Le spread selon l'horloge tick-time:

$$\hat{S}_n = S_{\theta_n}, \ n \geq 0.$$

Puis, la probabilité de transition $\rho_{ij} = \mathbf{P}[\hat{S}_{n+1} = j\delta | \hat{S}_n = i\delta]$ de la chaîne de Markov stationnaire $(\hat{S}_n)$ est estimée à partir de K échantillons $\hat{S}_n$, $n = 1, \ldots, K$ suivant un estimateur standard. L'intensité de l'horloge tick-time est elle aussi estimée avec un estimateur standard, sous des hypothèses simplificatrices valides en haute fréquence.

Nous présentons ensuite une méthode pour estimer les intensités de N^a et N^b. Concentrons nous sur N^b par exemple, ce processus représentant l'arrivée de transactions au

bid. En supposant que nous puissions observer (Q_t^b, N_t^b, S_t), $t \geq 0$, nous voulons estimer les intensités suivantes pour N^b:

$$\lambda_i^b(q^b) \quad := \quad \lambda^b(q^b, s), \;\; q^b \in \{Bb, Bb_+\}, s = i\delta, i = 1, \ldots, m.$$

L'estimation de cette intensité revient à estimer $2m$ scalaires, ce qui apporte de la flexibilité au modèle, mais qui requiert une méthode spécifique, on définit:

$$\begin{aligned} N_t^{b,q^b,i} &= \int_0^t 1_{\{Q_u^b = q^b, S_{u-} = i\delta\}} dN_u^b, \\ \mathcal{T}_t^{b,q^b,i} &= \int_0^t 1_{\{Q_u^b = q^b, S_{u-} = i\delta\}} du. \end{aligned}$$

et on propose l'estimateur suivant pour $\lambda_i^b(q^b)$:

$$\hat{\lambda}_i^b(q^b) \quad = \quad \frac{N_T^{b,q^b,i}}{\mathcal{T}_T^{b,q^b,i}}$$

qui est consistent lorsque $\mathcal{T}_T^{b,q^b,i} >> 1/\lambda_i^b(q^b)$. En effet, $N_t^{b,q^b,i}$ a pour intensité $\lambda_i^b(q^b) 1_{\{Q_t^b = q^b, S_{t^-} = i\delta\}}$ et on applique la loi des grands nombres pour sa martingale compensée. La figure 1.3 illustre cette procédure sur données réelles.

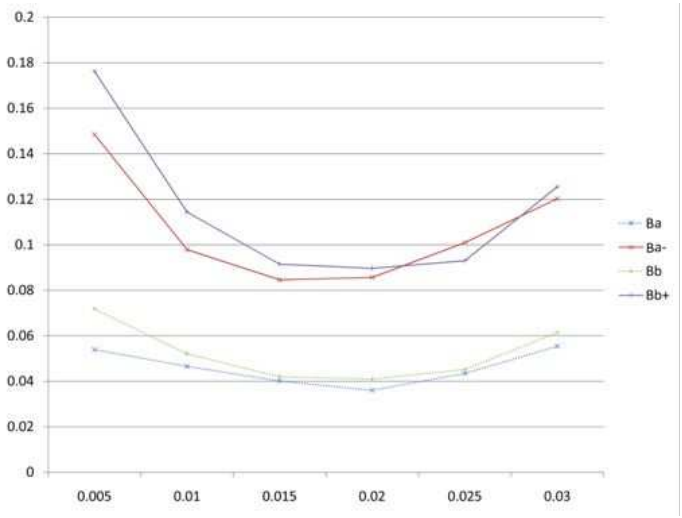

Figure 1.3: Intensités d'exécution sur SOGN.PA le 18 avril, 2011,en s^{-1} (interpolation affine) comme fonction du spread.

Optimisation

On se propose d'optimiser l'utilité terminale du profit du teneur de marché, sur un horizon de temps fini, avec deux exmaple de fonction d'utilité: la fonction d'utilité exponentielle

et la fonction d'utilité moyenne-variance. Dans cette synthèse, pour être plus concis, on se focalise sur le critère moyenne-variance.

$$\text{maximiser} \quad \mathbf{E}\big[X_T - \gamma \int_0^T Y_t^2 d < P >_t \big]$$

sur toutes les stratégies $\alpha = (\alpha^{make}, \alpha^{take}) \in \mathcal{A}$ telles que $Y_T = 0$. Notre objectif est donc de maximiser le cash terminal, sachant que l'on ne détient aucune position sur l'actif risqué à la date T, et l'on pénalise la détention d'un inventaire non nul pendant $[0;T]$ à l'aide de la variance intégrée du portefeuille. $\gamma > 0$ est l'aversion au risque quadratique lié à la détention de Y unités de l'actif P. On peut aisément retirer la contrainte $Y_T = 0$ en introduisant:

$$L(x,y,p,s) \;=\; x - c(-y,p,s) \;=\; x + yp - |y|\frac{s}{2} - \varepsilon.$$

et nous définissons la fonction de valeur:

$$v(t,x,y,p,s) \;=\; \sup_{\alpha\in\mathcal{A}} \mathbf{E}_{t,x,y,p,s}\Big[L(X_T,Y_T,P_T,S_T) - \gamma\int_t^T Y_u^2 \varrho(Y_u)du\Big],$$

où l'on a supposé $d < P >_t = \varrho(P_t)dt$.

Comme le spread prend des valeurs discrètes, $s = i\delta$, $i \in \mathbb{I}_m$, nous notons

$$v_i(t,x,y,p) \;=\; v(t,x,y,p,i\delta)$$

et identifions v avec $(v_i)_{i=1,\dots,m}$: une fonction à valeur vecteur dans $\mathbb{R}^m$ de $[0,T]\times\mathbb{R}\times\mathbb{R}\times\mathbb{P}$. On utilise des notations similaires pour L_i, c_i, π_i^a, π_i^b, λ_i^a, λ_i^b.

Et nous caractérisons v_i comme unique solution d'un QVI tridimensionnelle, que l'on va simplifier.

Réduction de la dimension

Pour améliorer la vitesse de résolution numérique de la HJBQVI, nous nous intéressons à réduire la dimension de l'espace d'état. Si l'on suppose que P est un processus de Lévy, nous avons:

$$\mathcal{P} I_{\mathbb{P}} \;=\; c_P, \qquad d < P >_t = \; \varrho dt,$$

où $I_{\mathbb{P}}$ est l'identité, pour des constantes c_P, ϱ. On a alors la réduction de $v = (v_i)_{i=1,\dots,m}$ sous la forme:

$$v_i(t,x,y,p) \;=\; L_i(x,y,p) + \phi_i(t,y).$$

De plus, il existe une constante κ t.q.

$$0 \;\le\; \phi_i(t,y) \;\le\; (T-t)\kappa,$$

pour tous $(t,y,i) \in [0,T] \times \mathbb{R} \times \mathbb{I}_m$.

Au final, le problème simplifié devient une QVI unidimensionnelle:

$$\begin{aligned}
&\min\Big[-\frac{\partial \phi_i}{\partial t} - yc_P + \gamma \varrho y^2 - \lambda(t)\sum_{j=1}^{m} \rho_{ij}\big[\phi_j - \phi_i + |y|(j-i)\frac{\delta}{2}\big] \\
&- \sup_{(q^b,\ell^b)\in \mathcal{Q}_i^b \times [0,\bar{\ell}]} \lambda_i^b(q^b)\big[\phi_i(t,y+\ell^b) - \phi_i(t,y) + \frac{i\delta}{2}(|y| + \ell^b - |y+\ell^b|) - \delta \ell^b 1_{q^b = Bb_+}\big] \\
&- \sup_{(q^a,\ell^a)\in \mathcal{Q}_i^a \times [0,\bar{\ell}]} \lambda_i^a(q^a)\big[\phi_i(t,y-\ell^a) - \phi_i(t,y) + \frac{i\delta}{2}(|y| + \ell^a - |y-\ell^a|) - \delta \ell^a 1_{q^a = Ba_-}\big] \; ; \\
&\phi_i(t,y) - \sup_{e\in\mathbb{R}}\Big[\phi_i(t,y+e) - \frac{i\delta}{2}(|y+e| + |e| - |y|) - \varepsilon\Big] \;=\; 0,
\end{aligned}$$

pour $(t,y,i) \in [0,T) \times \mathbb{R} \times \{1,\dots,m\}$, avec la condition terminale:

$$\phi_i(T,y) \;=\; 0, \qquad \forall y \in \mathbb{R},\ i = 1,\dots,m.$$

Schéma numérique et résultats

Nous résolvons la QVI numériquement en fournissant un schéma numérique explicite rétrograde. Nous discrétisons d'abord le temps sur une grille régulière $[0,T]$: $\mathbb{T}_n = \{t_k = kh, k = 0,\dots,n\}$, $h = T/n$. Puis nous discrétisons et localisons les variables d'espace: $\mathbb{Y}_{R,M} = \{\ell \frac{R}{M}, \ell = -M,\dots,M\}$.

$(\phi_i)_{i=1,\dots,m}$ approchée par $(\phi_i^{h,R,M})_{i=1,\dots,m}$, avec la condition terminale: $\phi_i^{h,R,M}(t_n,y) = 0$, et nous obtenons le schéma numérique $\mathcal{S}^{h,R,M}$ en remplaçant les quantités suivantes dans la QVI:

$$\frac{\partial \phi_i}{\partial t}(t_k,y) \;\sim\; \frac{\phi_i^{h,R,M}(t_k+h,y) - \phi_i^{h,R,M}(t_k,y)}{h}$$

les termes non-locaux (t_k,z,i) calculés à l'instant $t_k + h$ avec:

$$\phi_i(t_k,z) \;\sim\; \phi_i^{h,R,M}(t_k+h, \mathrm{Proj}_{[-R,R]}(z))$$

Et nous écrivons le schéma numérique:

$$\begin{aligned}
&\phi_i^{h,R,M}(t_k,y) \\
=\;& \mathcal{S}^{h,R,M}\Big(t_k, y, \phi_i^{h,R,M}(t_k+h,.), \big(\phi_j^{h,R,M}(t_k+h,y)\big)_{j=1,\dots,m}\Big),
\end{aligned}$$

nous prouvons que $\mathcal{S}^{h,R,M}$ est stable et monotone dès lors que:

$$\Big[\max_{i\in\mathbb{I}_m, q^b\in\mathcal{Q}_i^b} \lambda_i^b(q^b) + \max_{i\in\mathbb{I}_m, q^a\in\mathcal{Q}_i^a} \lambda_i^a(q^a) + \sup_{t\in[0,T]} \lambda(t)\Big]h < 1,$$

De plus $\mathcal{S}^{h,R,M}$ est cohérent (lorsque $h \to 0$, $M, N \to \infty$), et donc convergent en utilisant un argument de Barles-Souganidis [8].

Enfin, nous proposons des tests numériques détaillés, assortis d'une analyse de performance comparée, dont nous reproduisons ici les figures principales: 1.4, 1.2, 1.3.2 et 1.6.

(a) near date 0 (b) near date T

Figure 1.4: Forme stylisée de la politique optimal dans le plan YS.

		optimal $\alpha^\star$	WoMO α^w	constant α^c	random α^r
Terminal wealth	$m(X_T)/\sigma(X_T)$	2.117	1.999	0.472	0.376
	$m(X_T)$	26.759	25.19	24.314	24.022
	$\sigma(X_T)$	12.634	12.599	51.482	63.849
Num. of exec. at bid	$m(N_T^b)$	18.770	18.766	13.758	21.545
	$\sigma(N_T^b)$	3.660	3.581	3.682	4.591
Num. of exec. at ask	$m(N_T^a)$	18.770	18.769	13.76	21.543
	$\sigma(N_T^a)$	3.666	3.573	3.692	4.602
Num. of exec. at market	$m(N_T^{market})$	6.336	0	0	0
	$\sigma(N_T^{market})$	2.457	0	0	0
Maximum Inventory	$m(\sup_{s\in[0;T]} \lvert Y_s\rvert)$	241.019	176.204	607.913	772.361
	$\sigma(\sup_{s\in[0;T]} \lvert Y_s\rvert)$	53.452	23.675	272.631	337.403

Table 1.2: Synthèse de l'analyse de performance (5.10^5 simulations).

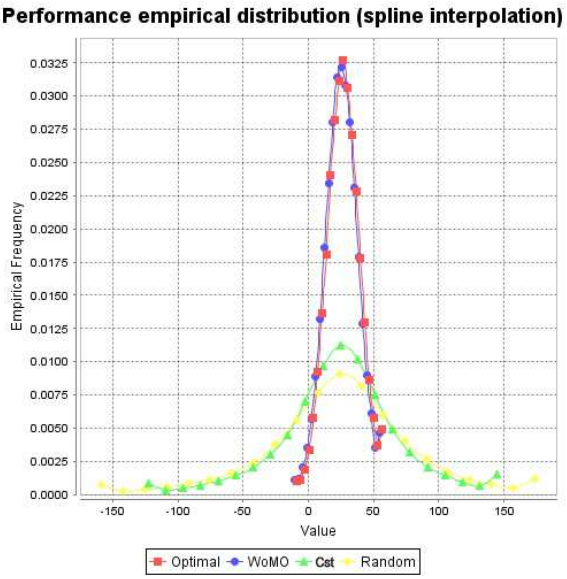

Figure 1.5: Distribution empirique de la performance terminale X_T (spline interpolation).

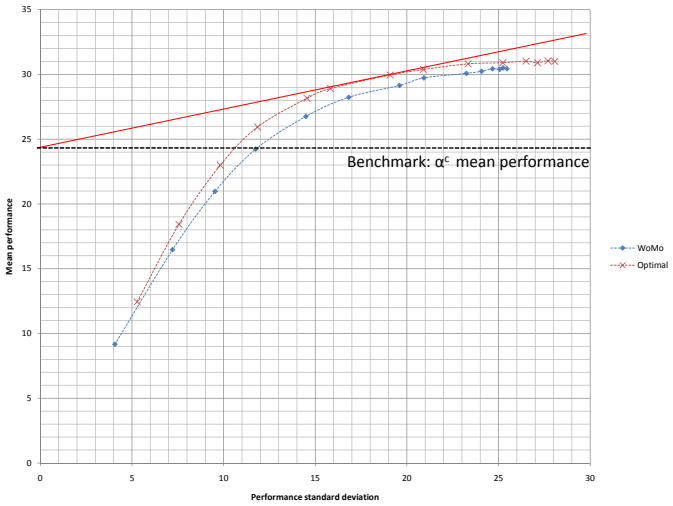

Figure 1.6: Efficient frontier plot

1.3.3 Trading haute fréquence optimal dans une microstructure au pro-rata avec information prédictive

Dans le chapitre 6 nous étudions une stratégie mixte de tenue de marché dans une microstructure exotique appelée la microstructure pro rata. Cette microstructure peut être rencontrée par example sur les futures de taux d'intérets courts. Ici encore, on examine la situation d'un investisseur qui souhaite maximiser son profit sur un horizon de temps fini, et qui est capable de traiter avec des ordres limites et des ordres de marché. Nous prenons la perspective de la gestion d'inventaire, ce qui signifie que l'objectif premier du teneur de marché est de maintenir une position sur l'actif risqué proche de zéro à toute date, de sorte à éviter d'avoir une exposition au risque de marché. Dans cette microstructure particulière, on peut introduire et prendre en compte dans notre stratégie deux sorte de risques supplémentaires: le risque d'overtrading, qui est le risque de variation brutale de l'inventaire du teneur de marché, dû au fait qu'il ne contrôle pas la quantité traitée avec des ordres limites; et le risque de sélection adverse, qui est le risque que le marché réagisse de manière adverse aux actions de l'investisseur. Suivant cette dernière problématique, on introduit une variable d'état supplémentaire, que l'on interprète comme une information prédictive sur le prix, qui nous permet d'équilibrer notre inventaire lorsque qu'on prévoit un changement de prix. Cette propriété nous permet également d'avoir une performance supérieure dans nos backtests.

Nous nous intéressons à la microstructure appelé "vanilla pro rata", qui peut être décrite succintement de la manière suivante: chaque ordre de marché est réparti sur tous les ordres limites actifs dans le LOB au meilleur prix, proportionnellement au volume de chaque ordre limite. La figure 1.7 décrit l'appariemment d'un ordre au marché avec des ordres limites actifs du LOB.

Ce type de microstructure, avec la taille du tick caractéristique, amène deux particularités, ainsi qu'illustré par la figure 1.8 (instruments en haut à gauche), reproduite de [25].

- D'abord la fourchette bid/ask est égale la plupart du temps à 1 tick
- Ensuite, la liquidité offerte aux meilleurs prix est largement surdimensionnée par rapport à la taille moyenne d'une transaction.

Notre travail est fondé sur l'approche de gestion d'inventaire ainsi que présenté dans Avelaneda et Stoikov (2008) [7]. Nous utilisons aussi des méthodes développées dans [35], [37] ou [68]. De plus, l'idée d'utiliser un indicateur prédictif de prix vient de [21]. Enfin, nous avons comparé nos résultats empiriques à ceux de [25].

Modèle de marché

Soit un espace de probabilités $(\Omega, \mathcal{F}, \mathbf{P})$ équipé d'une filtration $\mathbb{F} = (\mathcal{F}_t)_{t\geq 0}$, satisfaisant les conditions usuelles. Nous utilisons le modèle simple de prix:

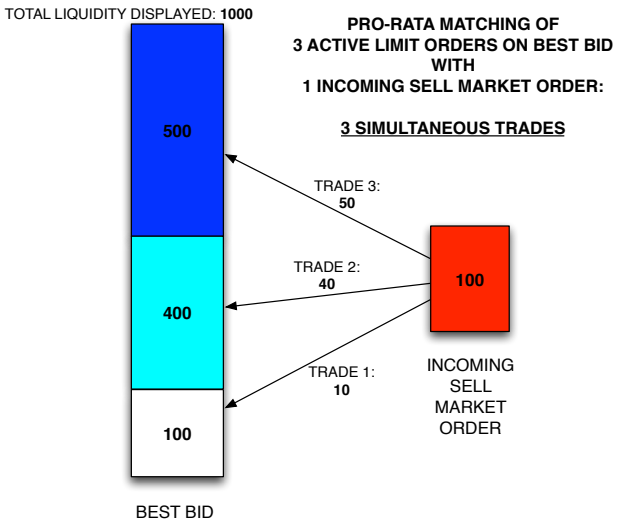

Figure 1.7: Transactions simultanées déclenchées par un ordre au marché

- P le prix mid (observable: *lit* microstructure): un processus de Markov de générateur $\mathcal{P}$ à valeurs dans $\mathbb{P}$. On suppose que c'est une semi-martingale spéciale.
- δ la taille du tick, en général sur les STIR 12.5 EUR par contrat
- P^a (resp. P^b) le prix ask (resp. bid) (*one-tick* microstructure):

$$P^a := P + \delta/2 \ , \ P^b := P - \delta/2$$

Maintenant, nous considérons des stratégies de trading mixtes, c'est à dire faites d'ordres limites et d'ordres au marché, qui sont modélisés respectivement comme des contrôles continus et des contrôles impulsionnels. En effet, la soumission d'ordres limites, leur mise à jour ou leur annulation est gratuite, ainsi il est judicieux de considérer comme des contrôles continus. Au contraire, les exécutions sont coûteuses, et ainsi les ordres au marché, conduisant à une exécution immédiate, sont modélisés comme des contrôles impulsionnels. Plus précisemment, une stratégie de trading est une paire $\alpha := (\alpha^{make}, \alpha^{take})$ de contrôles réguliers/impulsionnels:

$$\alpha^{make} := (L^a_t, L^b_t)_{t\geq 0} \ , \ \alpha^{take} := (\tau_n, \xi_n)_{n\in\mathbb{N}}$$

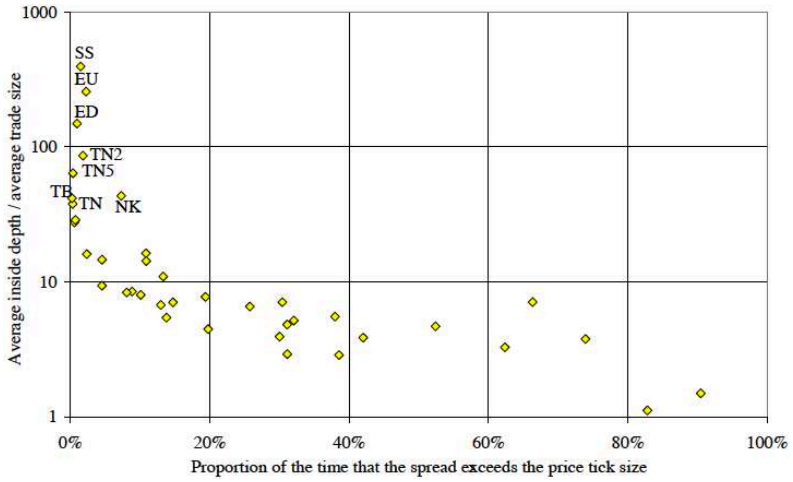

Figure 1.8: Exploration du marché

où L^a et L^b sont des processus prévisibles, à valeurs dans $\{0,1\}$, représentant les régimes *make*. $L^a = 1$ (resp. $L^b = 1$) signifie que l'investisseur a des ordres limite actifs à l'ask (resp. bid). Aussi, (τ_n) est une suite croissante de temps d'arrêts, et ξ_n est une variable $\mathcal{F}_{\tau_n}$-mesurable à valeurs dans $[-\bar{e}, \bar{e}]$, representant la quantité achetée (si $\xi_n > 0$) ou vendue (si $\xi_n < 0$) par l'investisseur. L'ensemble de telles stratégies est noté $\mathcal{A}$.

Décrivons maintenant notre modèle de processus de trade. En raison de la règle du pro rata, les fournisseurs de liquidité doivent exagérer le volume de leurs ordres limites: ils postent des ordres avec un volume bien plus élevé que ce qu'ils n'entendent réellement traiter. Donc, ils ne contrôlent pas la quantité des transactions auxquelles ils participent, et ceci est le risque d'overtrading. Cela diffère de la microstructure prix/date. Le volume entrant à l'ask (resp. bid), reçu par l'investisseur, est modélisé par une mesure de Poisson aléatoire ν^a (resp. ν^b) d'intensité $\lambda dt \times \mu(dz)$ sur $\mathbb{R}^+ \times \mathbb{R}^+$. $\lambda > 0$ représente l'intensité de trading et μ est la distribution du volume d'une transaction. On définit:

- Le nombre de transactions auxquelles à participé le HFT:

$$N_t^a := \int_0^t \int_{z\geq 0} \nu^a(dt, dz) \ , \ N_t^b := \int_0^t \int_{z\geq 0} \nu^b(dt, dz)$$

- Le volume cumulé exécuté par le HFT:

$$\vartheta_t^a := \int_0^t \int_{z\geq 0} z\nu^a(dt,dz) \ , \ \vartheta_t^b := \int_0^t \int_{z\geq 0} z\nu^b(dt,dz)$$

Dans cette sitation, décrivons l'évolution des variables du portefeuille. L'inventaire Y et le cash X ont la dynamique suivante sous le contrôle α:

$$\begin{aligned} dY_t &= L_t^b d\vartheta_t^b - L_t^a d\vartheta_t^a \ , \ \tau_n \leq t < \tau_{n+1} \\ dX_t &= L_t^a(P_t + \frac{\delta}{2})d\vartheta_t^a - L_t^b(P_t - \frac{\delta}{2})d\vartheta_t^b \ , \ \tau_n \leq t < \tau_{n+1} \\ Y_{\tau_n} - Y_{\tau_n -} &= \xi_n \\ X_{\tau_n} - X_{\tau_n -} &= -\xi_n P_{\tau_n} - |\xi_n|(\frac{\delta}{2} + \epsilon) - \varepsilon_0 \end{aligned}$$

où $\epsilon > 0$ est un frais par unité et $\varepsilon_0 > 0$ est un frais fixe. Remarquons que la valeur liquidative du portefeuille, $V := X + YP$ a la dynamique suivante:

$$\begin{aligned} dV_t &= \frac{\delta}{2}(L_t^a d\vartheta_t^a + L_t^b d\vartheta_t^b) + Y_{t-} dP_t \\ V_{\tau_n} - V_{\tau_n -} &= -(\frac{\delta}{2} + \epsilon)|\xi_n| - \varepsilon_0 \end{aligned}$$

Optimisation

Le système est entièrement déterminé par les variables d'état (X, Y, P) controllées par $\alpha \in \mathcal{A}$. Soit $T > 0$ un horizon de temps fini. Nous voulons :

$$\text{maximiser } \mathbf{E}\left[X_T - \gamma \int_0^T Y_t^2 d < P >_t\right] \text{ over all } \alpha \in \mathcal{A} \text{ s.t. } Y_T = 0$$

où $\gamma > 0$ est un paramètre de pénalisation. Ceci est équivalent à:

$$\text{maximiser } \mathbf{E}\left[L(X_T, Y_T, P_T) - \gamma \int_0^T Y_t^2 \varrho(P_t) dt\right] \text{ over all } \alpha \in \mathcal{A}$$

où l'on suppose $d < P >_t = \varrho(P_t)dt$, avec ϱ positif, continu sur $\mathbb{R}$. La fonction de liquidation L est:

$$L(x,y,p) = x + yp - |y|(\frac{\delta}{2} + \epsilon) - \varepsilon_0$$

Définissons maintenant la fonction de valeur:

$$v(t,x,y,p) := \sup_{\alpha \in \mathcal{A}} \mathbf{E}_{t,x,y,p}\left[L(X_T, Y_T, P_T) - \gamma \int_0^T Y_t^2 \varrho(P_t) dt\right]$$

et nous avons des bornes sur cette fonction de valeur (Proposition 6.3.1): il existe une constante $K_P \in \mathbb{R}$ t.q. $L(x,y,p) \leq v(t,x,y,p) \leq x + yp + \delta\lambda\bar{\mu}(T-t) + K_P$ où $\bar{\mu}$ est la moyenne de μ.

Introduisons les opérateurs apparaissant dans le DPP. Pour tous $(\ell^a, \ell^b) \in \{0,1\}^2$ nous définissons l'opérateur non-local associé avec le contrôle de l'ordre limite:

$$\mathcal{L}^{\ell^a,\ell^b} \;:=\; \mathcal{P} + \ell^a \Gamma^a + \ell^b \Gamma^b$$

où

$$\begin{aligned}
\Gamma^a \phi(t,x,y,p) &:= \lambda \int_0^\infty [\phi(t, x + z(p+\delta/2), y - z, p) - \phi(t,x,y,p)]\, \mu(dz) \\
\Gamma^b \phi(t,x,y,p) &:= \lambda \int_0^\infty [\phi(t, x - z(p-\delta/2), y + z, p) - \phi(t,x,y,p)]\, \mu(dz)
\end{aligned}$$

Nous définissons aussi l'opérateur non-local associé au contrôle d'ordre au marché:

$$\mathcal{M}\phi(t,x,y,p) \;:=\; \sup_{e \in [-\bar{e};\bar{e}]} \phi(t, x - ep - |e|(\delta/2 + \epsilon) - \varepsilon_0, y + e, p)$$

L'équation de la programmation dynamique associée à ce problème est une QVI:

$$\min\left\{ -\frac{\partial v}{\partial t} - \sup_{(\ell^a,\ell^b)\in\{0,1\}^2} \mathcal{L}^{\ell^a,\ell^b} v + \gamma g \;;\; v - \mathcal{M} v \right\} = 0 \;\;,\;\; \text{on } [0,T) \times \mathbb{R}^2 \times \mathbb{P}$$

avec la condition terminale:

$$v(T,.) = L \;\;,\;\; \text{on } \mathbb{R}^2 \times \mathbb{P}$$

où nous notons $g(y,p) = y^2 \varrho(p)$. Cette expression peut être écrite explicitement (voir chapitre 5).

Réduction de la dimension

Nous pouvons simplifier cette QVI dans le cas où le prix mid est un processus de Lévy:

$$\mathcal{P}\mathbb{I}_{\mathbb{P}} = c_P \text{ et } \varrho \text{ est constant.}$$

où $\mathbb{I}_{\mathbb{P}}$ est l'identité sur $\mathbb{P}$ i.e. $\mathbb{I}_{\mathbb{P}}(p) = p$ et c_P est une constante dépendant du triplet caractéristique de P.

Dans ce contexte, v se décompose en:

$$v(t,x,y,p) = L(x,y,p) + w(t,y)$$

nous voyons dans cette décomposition la fonction de liquidation. Avec cette simplification, w est solution de l'inégalité variationnelle unidimensionnelle:

$$\min\Big[-\frac{\partial w}{\partial t} - y c_{_P} + \gamma \varrho y^2 - \mathcal{I}^a w - \mathcal{I}^b w \;,\; w - \tilde{\mathcal{M}} w \Big] \;=\; 0, \;\; \text{on } \; [0,T) \times \mathbb{R},$$

avec condition terminale:

$$w(T,y) \;=\; 0, \quad \forall y \in \mathbb{R},$$

où $\mathcal{I}^a$ et $\mathcal{I}^b$ sont des opérateurs intégraux non-locaux:

$$\begin{aligned}\mathcal{I}^a w(t,y) &= \lambda^a\Big(\int_0^\infty \Big[w(t,y-z)-w(t,y)+z\frac{\delta}{2}+(\frac{\delta}{2}+\varepsilon)(|y|-|y-z|)\Big]\mu^a(dz)\Big)_+\\ \mathcal{I}^b w(t,y) &= \lambda^b\Big(\int_0^\infty \Big[w(t,y+z)-w(t,y)+z\frac{\delta}{2}+(\frac{\delta}{2}+\varepsilon)(|y|-|y+z|)\Big]\mu^b(dz)\Big)_+,\end{aligned}$$

et $\tilde{\mathcal{M}}$ est l'opérateur non-local:

$$\tilde{\mathcal{M}}w(t,y) = \sup_{e\in[-|y|,|y|]}\Big[w(t,y+e)-(\frac{\delta}{2}+\varepsilon)(|y+e|+|e|-|y|)-\varepsilon_0\Big].$$

Enfin, nous avons des bornes et une propriété de symétrie pour w.

- Les bornes suivantes sont vérifiées (principe de comparaison):

$$0 \le w(t,y) \le (T-t)\Big[\frac{c_P^2}{4\gamma\rho}+\lambda^a(\delta+\epsilon)\bar{\mu}^a+\lambda^b(\delta+\epsilon)\bar{\mu}^b\Big],$$

- Explicitant la dépendance en c_P, on a que:

$$w(t,y,c_P)=w(t,-y,-c_P)$$

Schéma numérique

Nous étudions un schéma numérique explicite rétrograde pour cette QVI. On définit une grille de temps régulière:

$$\mathbb{T}_N := \{t_k = kh\,,\ k=0,\dots,N\}$$

et une discrétisation/troncation usuelle pour l'espace d'état:

$$\mathbb{Y}_M = \big\{y_i = i\Delta_Y,\ i=-N_Y,\dots,N_Y\big\}.$$

Enfin, on écrit $\mathrm{Proj}_M(y) := -M\vee(y\wedge M)$, et on considère les approximations de μ^a et μ^b, définies par:

$$\hat{\mu}^a = \sum_{i\in\mathbb{Z}^+}\mu^a([i\Delta_Y;(i+1)\Delta_Y))\delta_{i\Delta_Y}\ , \qquad \hat{\mu}^b = \sum_{i\in\mathbb{Z}^+}\mu^b([i\Delta_Y;(i+1)\Delta_Y))\delta_{i\Delta_Y},$$

avec δ_x le Dirac en x. Pour toute φ sur $[0,T]\times\mathbb{R}$, $t\in[0,T]$, et $y\in\mathbb{R}$, nous définissons:

$$\mathcal{S}^{h,\Delta_Y,M}(t,y,\varphi) = \max\Big[\mathcal{T}^{h,\Delta_Y,M}(t,y,\varphi)\ ;\ \tilde{\mathcal{M}}^{h,\Delta_Y,M}(t,y,\varphi)\Big],$$

où

$$
\begin{aligned}
&\mathcal{T}^{h,\Delta_Y,M}(t,y,\varphi) = \\
&\varphi(t,y) - h\gamma\varrho y^2 + hyc_P \\
&+ \lambda^a h\Big(\int_0^\infty \big[\varphi(t,\mathrm{Proj}_M(y-z)) - \varphi(t,y)\big]\hat{\mu}^a(dz) + \big[\frac{\delta}{2}z + (\frac{\delta}{2}+\varepsilon)(|y|-|y-z|)\big]\mu^a(dz)\Big)_+ \\
&+ \lambda^b h\Big(\int_0^\infty \big[\varphi(t,\mathrm{Proj}_M(y+z)) - \varphi(t,y)\big]\hat{\mu}^b(dz) + \big[\frac{\delta}{2}z + (\frac{\delta}{2}+\varepsilon)(|y|-|y+z|)\big]\mu^b(dz)\Big)_+,
\end{aligned}
$$

et

$$
\begin{aligned}
&\tilde{\mathcal{M}}^{h,\Delta_Y,M}(t,y,\varphi) \\
= &\sup_{e\in\mathbb{Y}_M\cap[-|y|,|y|]} \big[\varphi(t,\mathrm{Proj}_M(y+e)) - (\frac{\delta}{2}+\varepsilon)(|y+e|+|e|-|y|) - \varepsilon_0\big].
\end{aligned}
$$

Enfin, nous montrons que ce schéma est monotone, stable et cohérent (Proposition 6.4.1-6.4.2-6.4.3) et donc sa solution $w^{h,\Delta_Y,M}$ converge localement uniformément vers w sur $[0,T)\times\mathbb{R}$, lorsque (h,Δ_Y,M) va en $(0,0,\infty)$ (Theorème 6.4.1).

Application: HFT avec information prédictive sur le prix

Enfin, nous avons conduits des tests numériques en supposant que le prix mid est un processus de Lévy, sur lequel nous avons une information prédictive. Plus précisement, on suppose que:

- Le prix mid P est un processus de saut pur $\delta\mathbb{Z}$.
- On a:

$$
\begin{aligned}
\mathbf{P}\left(P_{t+h}-P_t=\delta\,|\mathcal{F}_t\right) &= \pi^+h+o(h)\\
\mathbf{P}\left(P_{t+h}-P_t=-\delta\,|\mathcal{F}_t\right) &= \pi^-h+o(h)\\
\mathbf{P}\left(|P_{t+h}-P_t|>\delta\,|\mathcal{F}_t\right) &= o(h)
\end{aligned}
$$

  avec $\pi^+,\pi^->0$ et on note $\varpi := \pi^+-\pi^-$

Donc $\mathcal{PI}_{\mathbb{P}} = c_P = \varpi\delta$ et $\varrho(.)\equiv(\pi^++\pi^-)\delta^2$.

Dans ce contexte on peut calculer la politique optimale (figure 1.9).

Comme illustration numérique, nous avons réalisé une analyse de performance comparée détaillée sur des données simulées, et nous reproduisons les résultats principaux à la table 1.10.

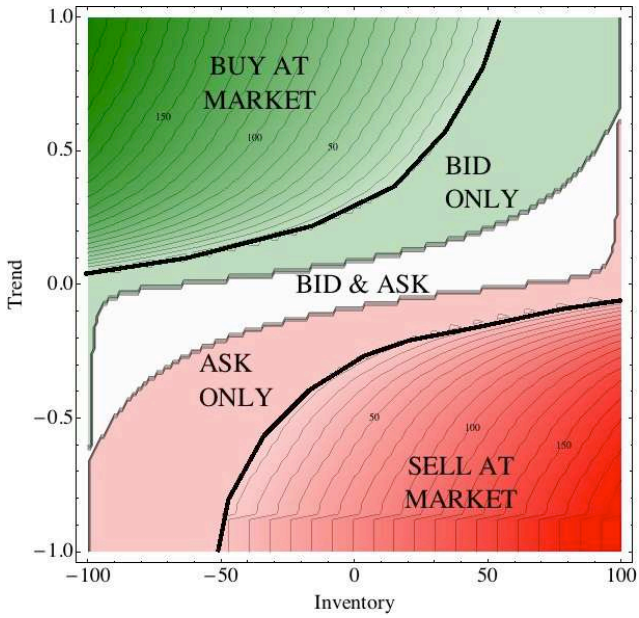

Figure 1.9: Politique optimale $\alpha^\star$ à la date $t = 0$.

Quantity	Definition	$\alpha^\star$	α^{WoMO}	α^{cst}
Info ratio over T	$m(\hat{V}_{\dot{T}})/\sigma(\hat{V}_{\dot{T}})$	3.67	0.89	0.18
Profit per trade	$m(\hat{V}_{\dot{T}})/m(\hat{Q}^{total,\cdot})$	8.06	16.31	5.57
Risk per trade	$\sigma(\hat{V}_{\dot{T}})/m(\hat{Q}^{total,\cdot})$	2.19	18.31	29.56
Mean performance	$m(\hat{V}_{\dot{T}})$	31446.4	28246.3	21737.2
Standard deviation of perf	$\sigma(\hat{V}_{\dot{T}})$	8555.46	31701.2	115312
Skew of perf	skew$(\hat{V}_{\dot{T}})$	0.64	0.16	-0.007
Kurtosis of perf	kurt$(\hat{V}_{\dot{T}})$	3.82	3.31	7.02
Mean total executed volume	$m(\hat{Q}^{total,\cdot})$	3900.68	1730.82	3900.61
Mean at market volume	$m(\hat{Q}^{market,\cdot})$	1932.29	0	0
Ratio market over total exec	$m(\hat{Q}^{market,\cdot})/m(\hat{Q}^{total,\cdot})$	0.495	0	0

Figure 1.10: Synthèse des résultats de backtest.

Chapter 2

Introduction

2.1 General objectives and motivations

This thesis' objective is to provide a mathematically rigourous approach to some of the most common aspects of high frequency trading.

From a financial point of view, we contribute to the quantitative coverage of the following issues: indirect trading costs minimization, market-making, and mixed passive/active high frequency trading strategies. We model and study both the standard price/time microstructure and the more exotic pro-rata market microstructure. We attempt to provide a complete treatment of each situation, from modelling and mathematical resolution, to calibration and *a posteriori* numerical experiments, including simulation and real-data tests when available.

From a mathematical perspective, we extend the frameworks that were recently developed in the mathematical finance litterature. We propose detailled numerical resolution of mixed regular/impulse optimal stochastic control problems. We design original numerical schemes for solving dynamic programming variationnal inequalities, that match the constraints of their high frequency purpose: we provide dimension reduction techniques, along with explicit computationnal algorithms that allows us to fasten the resolution of such problems, and therefore allow us to use the algorithms in real time. We prove the convergence of each specific scheme and provide numerical examples and illustrations.

In chapter 4, we consider the situation of an investor willing to sell a large portfolio at best price. We provide numerical solutions to an impulse control problem arising from the modelling of this problem, and we are able to take into account the effects of the bid-ask spread and market price impact penalizing speedy or large trades. The corresponding dynamic programming (DP) equation is a quasi-variational inequality (QVI) with solvency constraint satisfied by the value function in the sense of constrained viscosity solutions. We propose a tractable numerical solution based on an explicit backward numerical scheme for

the DPQVI. The convergence of this discrete-time scheme is shown by viscosity solutions arguments. We also provide numerical results both that show the behavior of the numerical scheme, the typical shape of the optimal strategy, and comparative performance analysis with respect to some benchmark execution strategies.

In chapter 5, we propose a framework for studying optimal high frequency trading strategies in the standard price/time microstructure, and we propose an application to european cash equities. We consider an (high-frequency) investor whose objective is to maximize her expected utility from revenue over a short term horizon, given that they are able to trade both with limit and market orders, while controlling their inventory position. This is formulated as a mixed regime switching regular/impulse control problem that we characterize in terms of quasi-variational system by dynamic programming methods. Calibration procedures are derived for the model, along with practical example for real-data fitting. We provide an explicit backward splitting scheme for solving the problem, and show how dimension can be reduced to a system of simple equations involving only the inventory and spread variables. Several computational tests are performed both on simulated and real data, and we perform detailled benchmarked performance analysis for the resulting algorithm.

In chapter 6, we propose a framework to study optimal high frequency trading (HFT) strategy in an exotic market microstructure, the so-called pro-rata microstructure, and propose an application to HFT on short-term interest rate futures contracts. Here again, the high-frequency trader has the choice to trade via market orders or limit orders, which are represented respectively by impulse controls and regular controls. We discuss and model the main risks specific to this microstructure, which are linked to the fact the size of the HF trades is not controlled. We assess the consequences of this specific fact in the context of optimal liquidation. The optimal trading problem is studied by stochastic control and dynamic programming methods, which lead to a characterization of the value function in terms of an integro quasi-variational inequality. We then provide the associated numerical resolution procedure, and convergence of this computational scheme is proved. We also propose algorithm simplifications for specific cases of practical interest: in particular we demonstrate a high frequency trading strategy in the case where a (predictive) directional information on the price is available. Each of the resulting strategies are illustrated by performance tests.

Let us now introduce the financial context that led to such models, along with a qualitative presentation of the financial objects and mechanisms that are at the core of high frequency trading. In 2.2.1 we propose an overview of the context of high frequency trading. In 2.2.2, we present the general vocabulary that we will use throughout this thesis, and we recall general results and qualitatives observations about market microstructure. In 2.2.3 we propose to sum up the different issues faced in the financial industry, where a solution involving high-frequency trading is available. Finally, we provide an outline of this thesis,

along with the summary of our main results in section .

2.2 Qualitative observations and context

2.2.1 General presentation

High-frequency trading is the use of automated strategies to trade securities such as cash equities, currencies or derivatives, with the distinguishing feature that positions are held for a very short period of time, ranging from a few seconds to a few hours. The term encompass several distinct trading techniques, that are often associated with the use of highly quantitative or data-intensive decision methods, heavy technology infrastructure, and no overnight position.

However, due to the recent increased availability of electronic trading technologies, as well as regulatory changes, a large range of investors are now able to implement high frequency trading strategies. The main regulatory frameworks that recently impacted high-frequency trading are MiFID in Europe (Market in Financials Instruments Directive, implemented 1 November 2007) and RegNMS in the United States (Regulation National Market System, 2007). They both aim at fostering competition between marketplaces, and promoting fair price formation processes. The practical results of these framework is the development of alternatives marketplaces (such as BATS or Chi-X, for example), and coincidently new needs in liquidity provision, orders routing and arbitrage.

In most of modern public security markets, the price formation process, or price discovery, results from competition between several market agents that take part in a public auction. In particular, day trading sessions, which are also called continuous trading phases, consist of continuous double auctions. High-frequency trading takes place in the continuous trading sessions, and therefore the precise study and modelling of actual mechanisms implementing this continuous double auction is of central importance when designing a high frequency trading strategy. This is precisely the subject of market microstructure: from [56], market microstructure theory is "the study of the process and outcomes of exchanging assets under a specific set of rules. While much of economics abstracts from the mechanics of trading, microstructure theory focuses on how specific trading mechanisms affect the price formation process." In the next subsection, we present the main mechanisms involved in price formation process.

2.2.2 The different types of limit order books

In this subsection, we present the mechanisms for order peering in the continuous trading phase, along with the general vocabulary that we will be using throughout this thesis.

The continuous trading phase is implemented in the general setup of continuous dou-

ble auctions. This means that the marketplace (for example the London Stock Exchange) displays publicly at least partial information about offered selling and buying prices. The liquidity providers are market agents that offers those prices, waiting for a counterpart market agent to take their offer, therefore leading to a trade. Liquidity providers compete in an auction on both buy side (called the bid side) and sell side (called the ask side). Practionners often distinguish between price-driven markets and order-driven markets. Although the definition of those notions may vary depending on the author, the general distinction consist in the following:

- Price-driven markets are markets where liquidity providers offer a price for any transaction volume. Generally speaking, there is a small number of dedicated market agents that act as liquidity providers. In actual markets however, e.g. FX markets, the price offered by the liquidity providers often depends on the volume wanted by their counterparts. This microstructure can also be encountered on more rudimentary markets, as for example real-time online betting markets, where a monopolistic market-maker set prices for a bet game so that the number of bets is balanced on both side of the game.

- Order-driven markets are markets where liquidity providers offer a given quantity at a given price, either to buy or to sell. Contrary to the previous organization, any market participant is able to act as a liquidity provider, thanks to the use of limit order trading (see below). This mechanism is the most common microstructure on electronic financial markets, for example it can be found on European cash equities, commodities or interest rates derivatives. This is implemented by the use of a limit order book (LOB), an object that we will describe in the following paragraphs.

In this thesis, we will focus on order-driven markets, since this is the mainstream market organization. Let us now define what is a limit order book, and examine two different orders peering rules.

We mention the complete survey article [33] about the limit order book, from where we adapted the following definitions. The role of a marketplace is to gather and to match the order to trade, originated from market participants, that can be submitted at any time during the continuous trading phase. They are of two types:

Definition. A market order of size m is an order to buy (sell) m units of the asset being traded at the lowest (highest) available price in the market.

Definition. A limit order of size ℓ at price p is an order to buy (sell) ℓ units of the asset being traded at the specified price p.

A limit order can be submitted to the market, updated in price or quantity or cancelled at any time, and therefore we call:

Definition. An active limit order at time t is a limit order that has been submitted at

some time $t_0 < t$, but has not been fully filled or cancelled by time t.

It is precisely the active limit orders in a market that make up the limit order book:

Definition. The limit order book (LOB), for a given asset, is the set of all active limit orders in the market at time t for this asset.

In addition to that, each marketplace follow its own policy about what information is publicly displayed. For example, the most transparent marketplaces (*lit microstructure*) display the aggregated volumes offered at each prices, and the usual data presentation looks as follows 2.1:

	Ask		Bid	
	Price	Quantity	Price	Quantity
Level 1	50.01	80	49.98	120
Level 2	50.02	53	49.97	89
Level 3	50.03	81	49.96	64
Level 4	50.04	112	49.95	163
Level 5	50.05	44	49.94	101

Table 2.1: Schematic representation of a snapshot at time t of a *lit* limit order book. In this example, at time t, the tick size is 0.01 the ask price is 50.01, the bid price is 49.98, and the spread is 0.03

Every electronic market allows only a discrete set of possible limit prices, where the *tick size* is the minimum increment between two possible prices. Therefore, we are able to use the definition of best prices, and the bid/ask spread:

Definition. The bid price at time t is equal to the highest stated price among buy limit orders in the limit order book.

Definition. The ask price at time t is equal to the lowest stated price among sell limit orders in the limit order book.

Definition. The bid-ask spread at time t is the difference between the ask price at time t and the bid price at time t.

Moreover, some more complex notions are often associated with the LOB, in particular some concepts coming from dynamic models of the LOB. Among them, the concept of *market impact* refers to the phenomenon of the ask (bid) price being deteriorated, i.e. being higher (lower), after a market order to buy (sell) has consumed several levels of the LOB at once. The opposite concept is the *order book resilience*, which stands for the fact that after a market order causing market impact, the emptied levels tends to re-populate with new limit orders.

Finally, let us give two practical examples of such limit order books. First, the most

common implementation of the LOB is the price/time microstructure. This microstructure is the most common microstructure in modern exchanges. It can be found e.g. on all European cash equities, with varying level of data confidentiality depending on the exchange. Its principle is very straightforward: an incoming market order is matched with the oldest active limit order among the best priced limit orders in the LOB. A detailled description of this microstructure can be found e.g. in [21] and [37].

Another important microstructure, however more exotic, is the pro-rata microstructure. The pro-rata microstructure (see [43] for extensive presentation and discussion) can be schematically described as follows: when a market order comes in the pro-rata limit order book, its volume is dispatched among all active limit orders at best prices, proportionnally to each limit orders volumes, and therefore create several transactions (see Figure 2.1).

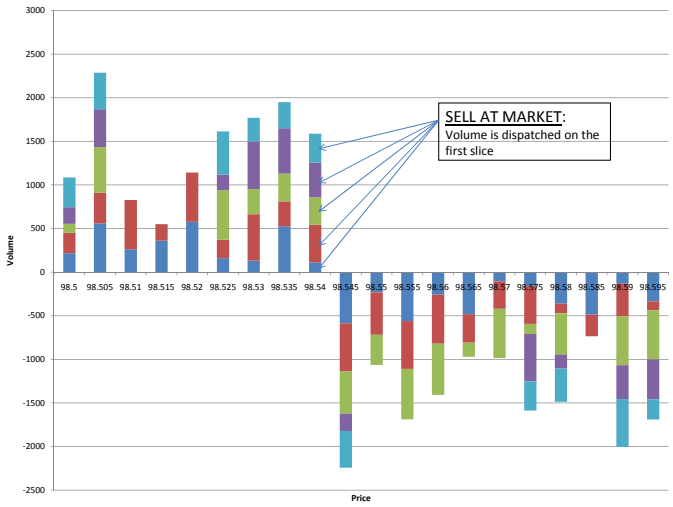

Figure 2.1: Schematic view of the pro-rata Limit Order Book.

This pro-rata microstructure is in use in some derivatives markets (e.g. London International Financial Futures and options Exchange, or Chicago Mercantile Exchange), and will be the subject of a whole chapter of this thesis.

2.2.3 Issues faced in high-frequency trading industry

In this subsection, we sum up the main industrial issues where high-frequency trading applies. We focus on the strategic stakes of high-frequency trading, and we put aside the technology issues such as latency minimization, direct market access or hardware speed improvement, which are however crucial aspects of the high frequency trading practice.

Indeed, our aim in to provide coverage for several distinct use of high frequency trading strategies, which are listed and summarized below.

Indirect trading costs minimization

Indirect trading costs minimization consists in obtaining the highest possible price from a sell trade, or obtaining the lowest possible price for a buy trade.

This problem naturally arises when the traded volume is large, due to finite liquidity offering in the LOB (see the above section) : indeed, a large single transaction at market price can desequilibrate the LOB by consuming several levels at once. For example, if an investor sends a market order to buy e.g. 200 shares in the book represented in table 2.1, the result of that transaction is:

- 80 shares at 50.01
- 53 shares at 50.02
- 67 shares at 50.03

therefore, the ask price at the end of this transaction is 50.03 with a volume offered of 14. Then, the Volume Weighted Average Price of this single transaction is $(80 \times 50.01 + 53 \times 50.02 + 67 \times 50.03)/200 = 50.0193$ which is about one tick greater than the ask price before the transaction, which leads to a loss of 2 bp. This effect is known as market impact. To give a comparison point, a strategy that trades on a daily basis, and that is expected to make a 5% return a year, have a daily expected return of 2 bp, and this is wiped out by the market impact. Moreover, several other costs, as the cost of crossing the spread, the brokers' fee or latency-related issues can penalize a single trade. Therefore we see that it is of crucial importance for portfolios managers to ensure the best possible execution of their trades.

Actors involved in the indirect trading costs optimization are both investors such as large hedge funds or investment banks, that develops their proprietary solution to this problem, and brokers, that typically have a large daily volume to trade on behalf of their clients. The brokers are moreover bound by the MIFiD regulations in Europe, and RegNMS act in the US, that force them to operate best execution algorithms. Some estimates that about $70\% - 80\%$ of the european equities [34] traded volume is done by execution algorithms, and other algorithmic trading.

Classical solutions to this problem can be classified around two central ideas: the space-optimization methods, and the time-optimization methods.

The space optimization procedure has received little focus from from academic literature, but some works are available, e.g. [48]. The idea underlying this method is to profit from the fact that an asset can often be traded on several distincts marketplaces. Therefore, by splitting a large parent order into smaller children orders, and dispatching them on several marketplaces, the investor is able to take more liquidity at the same time, hence to

be less exposed market impact. This technique is known as smart order routing (SOR), and is extensively implemented by numerous brokers in the industry. The optimization procedure in such tools typically involve latency considerations [49], along with high-frequency trading tools to be able to update quickly the trade schedule.

On the contrary, the time optimization procedure received extensive academic coverage, for example [3], [31] or [35]. The idea underlying this method is to split a large parent order into smaller children orders, and to pass the children orders on a extended time period. One can see the optimization procedure here as finding a balance in the following trade-off: if the investor trades quickly, they will face no market risk, but will have a large market impact ; on the contrary, if they trade slowly, they will face a large market risk, due to price movements, but will have reduced market impact. Several solutions to this problem have been proposed, with different assumptions, and the general technique is to trade according to a predefined schedule (optimal trading pattern) that arises when balancing the above mentionned trade-off under simplifying assumptions. We will give a lot more precisions on this topic in the following sections.

Finally, from an industrial perspective, some issues remains in that topic. Firstly, the detectability of trade optimization techniques is central to brokers and portfolio managers. Indeed, the massive use of execution algorithms is know to be at the source of autocorrelation in trade signs (see [18]) or lagged correlation in the trade data of the same asset on two distinct marketplaces. Therefore, such algorithms are very sensitive to the response of the LOB they trade onto, and therefore are less efficient when easily detected by competitors. Secondly, mixed market/limit orders execution strategies have so far received less focus from academical literature (see [67] or [37]), although the use of limit order trading is much cheaper than market order trading, and therefore extensively used in the industry in optimal execution strategies.

Pure alpha strategies

Now, let us focus on *pure alpha* strategies, which is a jargon term that refers to profit maximisation strategies that are largely irrespective of market conditions. This category includes the following strategies:

- *Market-making* strategies. This class of strategies are based on the idea that using limit orders trading, one can buy at the bid price, and sell at the ask price, and therefore gain the bid/ask spread. Such a strategy typically involve continuously providing bid and ask quotes, along with optimally chosing the prices and quantities of these quotes. The market maker will aim at balancing their inventory, i.e. keeping their position on the risky asset close to zero at all times, and therefore reducing their market risk.

- *Statistical arbitrage* strategies. This class of strategies are based on the idea that one can exploit the statistical relationship between asset prices (e.g. the cointegration

structure of a market sector, or the relationship between an index and its components) to profit from transient inefficiencies. Such strategies are typically data-intensive, they are directionnal over a short-term horizon and repeat a large number of times the same bet in order to reduce the variance of the outcome. Very often, such strategies are aggressive strategies, meaning that they take liquidity in the LOB (*hit orders*). They are also critically dependent on the latency of the trading infrastructure, due to competition between actors running the same strategy.

- *Mixed strategies*, that are the combination of the two above strategies classes.

Actors involved in such strategies include investment banks, hedge funds, proprietary trading firms and dedicated market-makers. The advantages of running these types of strategies is that their performance is very stable accross market conditions, and therefore the investor is not exposed to market risk. On the contrary, shortcomings of running pure alpha strategies is of two kinds: first, the absolute performance of the strategy is bounded most of the time, due to the fact that arbitrage opportunities are rare, and second, the operational risk is high, since technological performance is of crucial importance in this activity.

This class of strategies was studied in academic litterature, with an emphasis on market-making strategies.

Firstly, the market-making strategies have been succesfully presented as an inventory management problem since the pionner works of Amihud and Mendelsohn in 1980 [5] and Ho and Stoll in 1981 [42], and this approach was modernised in the work of Avellaneda and Stoikov in 2008 [7]. The underlying idea in this approach is take a risk/reward approach: the market-maker objective is to *make the spread*, i.e. to buy an asset at the bid, and sell it at the ask price, and therefore gain the bid/ask spread as a revenue. When doing this, the market-maker is subject to the market risk, i.e. the risk of holding a non-zero position in the risky asset, subject to price change. Therefore, the limit orders trading operated by the market-maker has two opposite goals: on one hand, they seek at maximizing the number of trades in which they participate, in order to maximize revenue from making the spread, and on the other hand, they need to keep their position on the risky asset close to zero at all time, in order to keep the market risk low, and this constraint leads to offering a more aggressive price at ask when they hold a long inventory, and conversely. This subject recently received sustained interest in academic works, with for example the works [16], [35] and [37].

Secondly, statistical arbitrage strategies have received less academic interest despite of their wide popularity among high frequency traders. The general idea of such strategies is to build a predictive price indicator based on market phenomena observation, and then trade accordingly. Let us illustrate this principle with two examples. In the work [6], the authors developed a generalized pairs trading approach: they perform a principal compo-

nent analysis on stocks returns, and then obtain a *market portfolio* that explains the stocks returns. Then, the main idea is to assume that the residual between one single stock and the market portfolio should revert to its mean, and trade accordingly. Another example is in the work [21], where the authors propose a simple statistical arbitrage strategy to illustrate the relevance of a predictive price indicator based on a poissonian model for a LOB. Based on the current state of the LOB, they are able to compute the probability of price going up or down in the next milliseconds, and they propose a HF strategy to exploit this information. Finally, in chapter 6, we propose a way to include such predictive price indicator to a mixed limit/market orders strategy.

The next section is devoted to outlining the main results of this thesis.

2.3 Thesis outline and main results

2.3.1 Optimal execution problem

In chapter 4, we consider the problem of an investor willing to unwind a large position on a risky asset. This situation is presented as a trade-off between market risk and market impact. Indeed, trading slowly has a small impact on the market price, but the investor keeps a non-zero position for a longer time, therefore bears more market risk. On the contrary, trading quickly has a large impact on the market price, but reduces market risk.

More precisely, we aim at controlling the difference between the marked to market value (or book value) of a portfolio, and the realized revenue when actually selling this portfolio. This shortfall is due to illiquidity effects including the bid/ask spread, the broker's fees and the market impact. We discuss the notion of market impact, presented as an adverse market reaction, actually resulting from finite liquidity offering in the market. This modelling was suggested by the seminal papers [10] and [3] that first introduced the concept of market impact in a discrete-time model. Applying an optimal control approach to the order execution problem was already documented in [63] and [28] with continuous controls (approximation of continuous trading), and in [44] with an impulse control approach. We use this last approach since it provides a more realistic modelling and still leads to tractable solutions. Our goal is to find optimal trading schedule and associated quantities.

Let us provide a brief overview of the model and our contributions.

Market model and trading strategies

We consider a financial market where an investor has to liquidate an initial position of $y > 0$ shares of risky asset by time T. We consider the following processes:

- $(P_t)_{t\in[0,T]}$ the market price of the risky asset
- $(X_t)_{t\in[0,T]}$ the cash holdings

- $(Y_t)_{t\in[0,T]}$ the number of stock shares held by the investor
- $(\Theta_t)_{t\in[0,T]}$ the time interval between t and the last trade before t

Trading strategies are considered to be made of impulse controls, in the form:

$$\alpha = (\tau_n, \xi_n)_{n\in\mathbb{N}}$$

where (τ_n), representing the trading dates, are $\mathbb{F}$-stopping times and (ξ_n), representing the traded quantities, are $\mathcal{F}_{\tau_n}$-measurable $\mathbb{R}$-valued variables. Dynamics for the shares and lag processes are under α:

$$\begin{aligned}
\Theta_t &= t-\tau_n,\ \tau_n \le t < \tau_{n+1} \\
\Theta_{\tau_{n+1}} &= 0,\ n \ge 0. \\
Y_s &= Y_{\tau_n},\quad \tau_n \le s < \tau_{n+1} \\
Y_{\tau_{n+1}} &= Y_{\tau_n} + \xi_{n+1} \quad n \ge 0.
\end{aligned}$$

We assume that market price of risky asset process follows a geometric Brownian motion:

$$dP_t = P_t(bdt + \sigma dW_t)$$

Suppose now that the investor decides to trade the quantity e. If the current market price is p, and the time lag from the last order is θ, then the price they actually get for the order e is:

$$Q(e,p,\theta) = pf(e,\theta)$$

we do allow a large set of admissible functions f, but we take the following example for our impact function:

$$f(e,\theta) = \exp\left(\lambda|\frac{e}{\theta}|^{\beta}\mathrm{sgn}(e)\right).\left(\kappa_a \mathbf{1}_{e>0} + \mathbf{1}_{e=0} + \kappa_b \mathbf{1}_{e<0}\right),$$

In this expression, $\kappa_a > 1$ and $\kappa_b < 1$ so that $\left(\kappa_a \mathbf{1}_{e>0} + \mathbf{1}_{e=0} + \kappa_b \mathbf{1}_{e<0}\right)$ represents the effect of crossing the bid/ask spread. The exponential part $\exp\left(\lambda|\frac{e}{\theta}|^{\beta}\mathrm{sgn}(e)\right)$ represents the non-linear effect of finite liquidity offering, i.e. the fact that a large market order will consume several slices of the order book at the same time. Reflexions about the shape of such function can be found for example in [50].

Then cash holdings have the following dynamics:

$$\begin{aligned}
X_t &= X_{\tau_n},\quad \tau_n \le t < \tau_{n+1},\ n \ge 0. \\
X_{\tau_{n+1}} &= X_{\tau_{n+1}^-} - \xi_{n+1} P_{\tau_{n+1}} f(\xi_{n+1}, \Theta_{\tau_{n+1}^-}) - \epsilon,\quad n \ge 0.
\end{aligned}$$

PDE characterization

We choose a constant relative risk aversion utility function $U(x) = x^\gamma$ with $\gamma \in (0,1)$ and denote $U_L(.) = U(L(.))$, where $L(.)$ is the liquidation function, which is the revenue obtained for selling the portfolio. The value function is defined by (we denoted $z = (x, y, p)$):

$$v(t,z,\theta) = \sup_{\alpha \in \mathcal{A}(t,z,\theta)} \mathbb{E}\big[U_L(Z_T)\big], \quad (t,z,\theta) \in [0,T] \times \mathcal{S}$$

where $\mathcal{A}(t,z,\theta)$ is a suitable set of admissible controls and $\mathcal{S} \subset \mathbb{R}^3$ the solvency region where the state variables lives.

From [44] v is a unique viscosity solution to a quasi-variational inequality (QVI) written as:

$$\begin{aligned} \min\left[-\frac{\partial}{\partial t}v - \mathcal{L}v \ , \ v - \mathcal{H}v\right] &= 0, \quad \text{on} \quad [0,T) \times \mathcal{S}, \\ \min\left[v - U_L, v - \mathcal{H}v\right] &= 0, \quad \text{on} \quad \{T\} \times \mathcal{S}. \end{aligned}$$

where $\mathcal{L}$ is the infinitesimal generator associated to the process (X, Y, P, Θ) in a no trading period:

$$\mathcal{L}\varphi = \frac{\partial}{\partial \theta}\varphi + bp\frac{\partial}{\partial p}\varphi + \frac{1}{2}\sigma^2 p^2 \frac{\partial^2}{\partial p^2}\varphi$$

and $\mathcal{H}$ is the impulse operator:

$$\mathcal{H}\varphi(t,z,\theta) = \sup_{e \in \mathcal{C}(t,z,\theta)} \varphi(t, \Gamma(z,\theta,e), 0)$$

with

$$\Gamma(z,\theta,e) = (x - epf(e,\theta) - \epsilon, y + e, p), \ \ z = (x,y,p) \in \mathcal{S}, \ e \in \mathbb{R}$$

indeed, during a no-trading period, the state process evolve only with the price P and lag Θ variables, in a diffusive fashion. When an impulse control occurs, the state variables jumps under the effect of a transaction, with a net loss of marked-to-market value (or book value), due to the presence of indirect trading costs.

From now, our goal is to solve numerically this HJBQVI.

Explicit numerical scheme

The choice of the numerical scheme is of crucial importance since it will impact the computing time. We choosed to use an explicit backward scheme by using a specific property of our problem. We start by considering the standard time discretization scheme:

$$S^h(t,z,\theta,v^h(t,z,\theta),v^h) = 0, \quad (t,z,\theta) \in [0,T] \times \bar{\mathcal{S}},$$

with

$$\begin{aligned} & S^h(t,z,\theta,r,\varphi) \\ := & \begin{cases} \min\left[r - \mathbb{E}\big[\varphi(t+h, Z_{t+h}^{0,t,z}, \Theta_{t+h}^{0,t,\theta})\big] \ , \ r - \mathcal{H}\varphi(t,z,\theta)\right] & \text{if } t \in [0, T-h] \\ \min\left[r - \mathbb{E}\big[\varphi(T, Z_T^{0,t,z}, \Theta_T^{0,t,\theta})\big] \ , \ r - \mathcal{H}\varphi(t,z,\theta)\right] & \text{if } t \in (T-h, T) \\ \min\left[r - U_L(z,\theta) \ , \ r - \mathcal{H}\varphi(t,z,\theta)\right] & \text{if } t = T. \end{cases} \end{aligned}$$

which can be formulated equivalently as an implicit backward scheme:

$$\begin{aligned} v^h(T,z,\theta) &= \max\big[U_L(z,\theta)\ ,\ \mathcal{H}v^h(T,z,\theta)\big], \\ v^h(t,z,\theta) &= \max\Big[\mathbb{E}\big[v^h(t+h,Z_{t+h}^{0,t,z},\theta+h)\big]\ ,\mathcal{H}v^h(t,z,\theta)\Big], \quad 0\le t\le T-h, \end{aligned}$$

and $v^h(t,z,\theta)=v^h(T-h,z,\theta)$ for $T-h<t<T$.

The usual way to treat implicit backward scheme is to solve by iterations a sequence of optimal stopping problems:

$$\begin{aligned} v^{h,n+1}(T,z,\theta) &= \max\big[U_L(z,\theta)\ ,\ \mathcal{H}v^{h,n}(T,z,\theta)\big], \\ v^{h,n+1}(t,z,\theta) &= \max\Big[\mathbb{E}\big[v^{h,n+1}(t+h,Z_{t+h}^{0,t,z},\theta+h)\big]\ ,\mathcal{H}v^{h,n}(t,z,\theta)\Big], \end{aligned}$$

starting from $v^{h,0}=\mathbb{E}[U_L(Z_T^{0,t,z},\Theta_T^{0,t,\theta})]$. Due to the effect of the lag variable Θ_t in the market impact function, it is not optimal to trade immediately after a trade. Therefore we are able to write equivalently this scheme as an explicit backward scheme:

$$\begin{aligned} v^h(T,z,\theta) &= \max\big[U_L(z,\theta)\ ,\ \mathcal{H}U_L(z,\theta)\big], \\ v^h(t,z,\theta) &= \max\Big[\mathbb{E}\big[v^h(t+h,Z_{t+h}^{0,t,z},\theta+h)\big]\ ,\ \sup_{e\in\mathcal{C}_\varepsilon(z,\theta)}\mathbb{E}\big[v^h(t+h,Z_{t+h}^{0,t,z_\theta^e},h)\big]\Big], \end{aligned}$$

where $z_\theta^e=\Gamma(z,\theta,e)$

Convergence of the numerical scheme

We prove the stability, monotonicity and consistency properties for the numerical scheme, and therefore it is convergent, thanks to an argument adapted from Barles-Souganidis [8].

Performance analysis

We provide some numerical results that we obtain from our implementation. We tested the optimal strategy against a benchmark of two other strategies. We test several aspect of the optimal strategy, as for example the terminal performance, as shown in figure 2.2, with detailled methodology and comments.

2.3.2 Optimal high-frequency trading with limit and market orders

In chapter 5, we move to another important aspect of high-frequency trading, the market-making strategies. Market-making is the action of continuously providing liquidity to the market by trading with limit orders. In this work, we consider an investor who is able to trade with limit orders, but also with market orders, and therefore we consider a slightly larger class of strategies than strict market-making. The investor's objective is to maximize the utility of their profit over a finite time horizon. Our goal is to obtain a simple and tractable market model, with a precise modelling of the underlying microstructure. We chose the context of the price/time microstructure, which is the most standard market

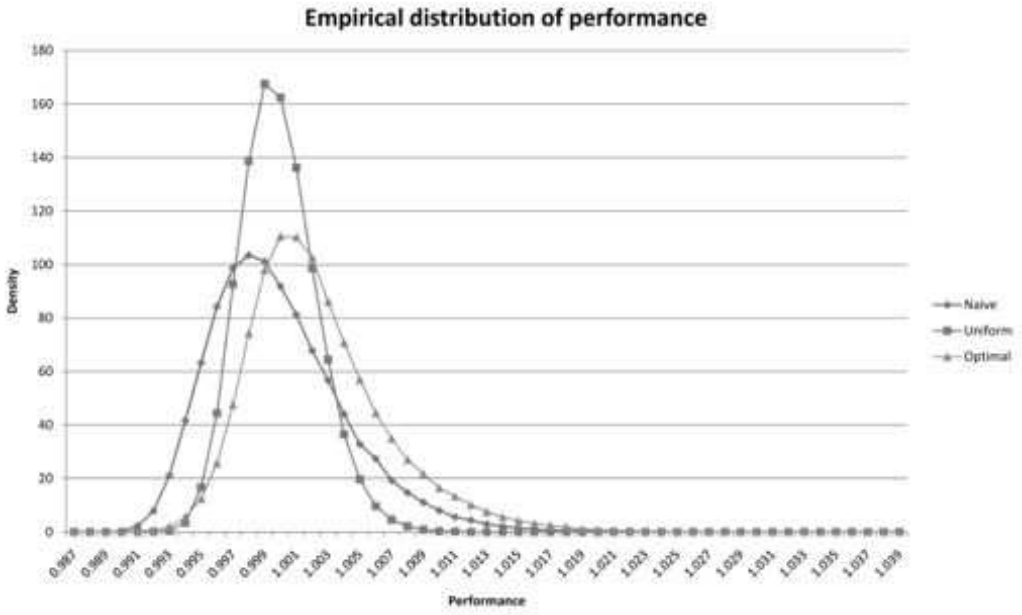

Figure 2.2: Optimal strategy perfrmance empirical distribution.

microstructure, and can be encountered on most cash equities, for example. We propose an easy to calibrate model that reflects some crucial elements of the price/time microstructure: in particular, we are able to fit very general behaviour for the bid/ask spread, and we also take into account the fact that the market can react the investor's actions, thanks to a control-dependent modelisation of the trades intensities. We represent this situation as a mixed stochastic control problem, that we study by dynamic programming means, and we provide a fast numerical scheme to solve it, thanks to a dimension reduction technique. We prove that this scheme is convergent, and we provide detailled numerical results along with precise performance analysis.

Most of modern equities exchanges are organized through a mechanism of Limit Order Book (LOB) which is the central element in market microstructure. In such mechanism, quoted prices are discrete, separated by the tick size which is typically of order 0.01 EUR per share. Market makers are liquidity providers in the LOB in the sense that they trade with limit orders sending buying orders at the bid, selling orders at the ask. Limit orders strategies are usually referred to as passive trading, since they are executed only when they meet incoming counterpart market orders. This uncertainty in execution is compensated

by the profit one can do by making the bid-ask spread (i.e. selling at the ask price, and buying at the bid price). Yet, passive trading is subject to a series of strategic risks:

- Inventory risk: risk exposure for holding a position in the stock due to price fluctuations also called market risk or volatility risk
- Execution risk: uncertainty in limit orders execution. For example in the price/time microstructure (FIFO rule):
  - an incoming market order is executed against the best priced, first arrived limit order (queuing system)
  - a market maker must be fast enough to profit from this priority rule for catching the market order flow.
- Adverse selection risk: market reacts adversely to the investor's quotes

Our objective is to address these three strategic risks in our market-making strategy. We adopt the inventory management perspective that have been successfully developped by the seminal paper by Avellaneda and Stoikov [7]: the market maker can submit bid and ask quotes with unit orders anywhere around a mid price, and the arrival of incoming counterpart market orders is modelled by a Poisson process with intensity depending on the distance of the quote to the mid price. This model leads to keeping the position in the risky asset close to zero at all times. Other recent litterature in line with this approach includes e.g. [35] and [16] .

Market model and trading strategies

We assume that the stock (mid)-price is a Markov process P with generator $\mathcal{P}$ and state space $\mathbb{P}$. The number of price updates, the so-called tick time clock is assumed to be a Poisson process $(N_t)_t$ with deterministic intensity $\lambda(t)$. Now under this tick time clock, the spread is assumed to be a stationary Markov chain $(\hat{S}_n)_{n\in\mathbb{N}}$ valued in $\mathbb{S} = \delta\mathbb{I}_m$, $\mathbb{I}_m = \{1, \ldots, m\}$, where δ is the tick size. We also define its transition matrix $(\rho_{ij})_{ij}$: $\rho_{ij} = \mathbf{P}[\hat{S}_{n+1} = j\delta | \hat{S}_n = i\delta]$, $i, j \in \mathbb{I}_m$, $\rho_{ii} = 0$. In regular time, i.e. calendar time, the spread is therefore: $S_t = \hat{S}_{N_t}$ and assumed to be independent of P. Then the best bid and best ask prices are simply defined by:

$$P_t^b \;=\; P_t - \frac{S_t}{2}, \qquad P_t^a \;=\; P_t + \frac{S_t}{2}.$$

Let us turn now to the trading strategies. First, limit orders (make strategy) are modelled as continuous-time predictable control process:

$$\alpha_t^{make} \;\;=\;\; \{(Q_t^b, L_t^b), (Q_t^a, L_t^a)\}$$

where Q_t^b represents the bid quote valued in $\mathcal{Q}^b = \{Bb, Bb_+\}$, which means:

- Bb: Best bid price, and Bb_+: Best bid price + one tick (to get priority in order execution)
- L^b: size of the limit buy order valued in $[0, \bar{\ell}]$

and Q^a_t represents the ask quote valued in $\mathcal{Q}^a = \{Ba, Ba_-\}$, which means:

- Ba: Best ask price, and Ba_-: Best ask price − one tick (to get priority in order execution)
- L^a: size of the limit sell order valued in $[0, \bar{\ell}]$

In this context, we can write how variables describing the investor's portfolio evolve. By applying a limit order strategy $\alpha^{make}_t = \{(Q^b_t, L^b_t), (Q^a_t, L^a_t)\}$, inventory Y and cash X evolve as:

$$\begin{aligned} dY_t &= L^b_t dN^b_t - L^a_t dN^a_t, \\ dX_t &= \pi^a(Q^a_t, P_{t^-}, S_{t^-}) L^a_t dN^a_t - \pi^b(Q^b_t, P_{t^-}, S_{t^-}) L^b_t dN^b_t. \end{aligned}$$

where

$$\begin{aligned} \pi^a(q^a, p, s) &= p + \frac{s}{2} - \delta 1_{q^b = Ba_-} \\ \pi^b(q^b, p, s) &= p - \frac{s}{2} + \delta 1_{q^b = Bb_+}, \end{aligned}$$

and where we introduced the trade processes N^a and N^b, counting the trades occurring at ask and bid sides respectively, which are, more precisely:

- N^a_t: arrival of market buy orders matching the limit sell orders $\sim \text{Cox}(\lambda^a(Q^a_t, S_t))$: $\lambda^a(Ba, s) < \lambda^a(Ba_-, s)$
- N^b_t: arrival of market sell orders matching the limit buy orders $\sim \text{Cox}(\lambda^b(Q^b_t, S_t))$: $\lambda^b(Bb, s) < \lambda^b(Bb_+, s)$

Note that the intensity of these trade processes depends on the investor's limit orders controls (Q^a_t, Q^b_t), which is relevant to model a market reaction to the investor actions, but also on other market variable, in our case the bid/ask spread.

Now, market orders strategy is modelled as impulse controls $\alpha^{take} = (\tau_n, \zeta_n)_{n \geq 0}$ where $(\tau_n)_n$ is an increasing sequence of stopping times representing market order decision times and ζ_n are $\mathcal{F}_{\tau_n}$-measurable, representing the number of stocks bought at best ask (if $\zeta_n \geq 0$), and sold at best bid (if $\zeta_n < 0$). Thos market orders are immediately executed, are therefore their effect on portfolio variables is:

$$\begin{aligned} Y_{\tau_n} &= Y_{\tau_n^-} + \zeta_n \\ X_{\tau_n} &= X_{\tau_n^-} - c(\zeta_n, P_{\tau_n}, S_{\tau_n}), \end{aligned}$$

where

$$c(e,p,s) \;=\; ep + |e|\frac{s}{2} + \varepsilon,$$

with $\varepsilon > 0$ denotes a fixed fee.

Estimation

The next section is devoted to model calibration. First, we show how to estimate the parameters involved in spread dynamics. We assume that the continuous-time Markov chain spread (S_t) is observable. We observe the following quantities:

- The tick times $(\theta_n)_n$ defined by:

$$\theta_{n+1} \;=\; \inf\left\{t > \theta_n : S_t \neq S_{t^-}\right\}, \qquad \theta_0 \;=\; 0.$$

- The associated Point process:

$$N_t \;=\; \#\left\{\theta_j > 0 \,:\, \theta_j \leq t\right\} \;, t \geq 0,$$

- The spread in tick time:

$$\hat{S}_n \;=\; S_{\theta_n} \;, n \geq 0.$$

Then, the transition probability $\rho_{ij} = \mathbf{P}[\hat{S}_{n+1} = j\delta | \hat{S}_n = i\delta]$ of the stationary Markov chain $(\hat{S}_n)$ is estimated from K samples of $\hat{S}_n$, $n = 1, \ldots, K$ with the standard estimator. For estimating the intensity of the tick time clock (which is a proxy for market activity) we propose a straightforward method, based on simplifying assumptions, valid for high-frequency data.

We go on presenting a method to fit the Cox processes N^a and N^b intensities. If we focus on N^b for example, this process represent arrivals of markets orders matching bid quote. Assuming that we can observe the following triplet: (Q_t^b, N_t^b, S_t), $t \geq 0$, we aim at estimating the intensity function of the Cox process N^b:

$$\lambda_i^b(q^b) \;:=\; \lambda^b(q^b, s), \;\; q^b \in \{Bb, Bb_+\}, s = i\delta, i = 1, \ldots, m.$$

Estimating this execution intensity is equivalent to estimating $2m$ scalars, which provides flexibility for model fitting, but requires a specific method. Let us define:

$$\begin{aligned} N_t^{b,q^b,i} &= \int_0^t 1_{\{Q_u^b = q^b, S_{u-} = i\delta\}} dN_u^b, \\ \mathcal{T}_t^{b,q^b,i} &= \int_0^t 1_{\{Q_u^b = q^b, S_{u-} = i\delta\}} du. \end{aligned}$$

here, $N_t^{b,q^b,i}$ counts the number of bid market orders that arrives when the spread is $i\delta$, and $\mathcal{T}^{b,q^b,i}$ is the time spent in the state $i\delta$ and then we propose the following estimator of $\lambda_i^b(q^b)$:

$$\hat{\lambda}_i^b(q^b) \quad = \quad \frac{N_T^{b,q^b,i}}{\mathcal{T}_T^{b,q^b,i}}$$

which is a consistent estimator once $\mathcal{T}_T^{b,q^b,i} >> 1/\lambda_i^b(q^b)$. Indeed $N_t^{b,q^b,i}$ has intensity $\lambda_i^b(q^b)1_{\{Q_t^b=q^b,S_{t^-}=i\delta\}}$ and we apply law of large numbers for the compensated martingale. Figure 2.3 illustrate this estimation procedure on real data.

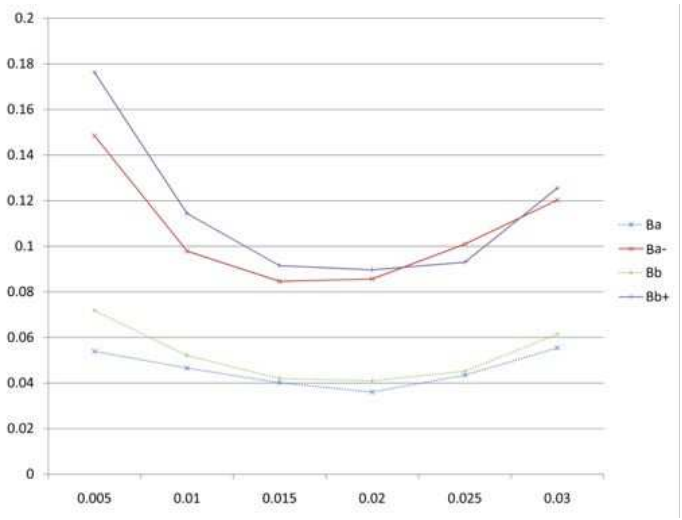

Figure 2.3: Plot of execution intensities for the stock SOGN.PA on April 18, 2011, expressed in s^{-1} (affine interpolation) as a function of the spread.

Optimization

We propose to optimize the terminal utility of profit of the market-maker, over a finite time horizon, with two example of utility function: the exponential utility and mean-variance utility. In this outline, for conciseness, we focus on the mean-variance criterion:

$$\text{maximize} \quad \mathbf{E}\big[X_T - \gamma \int_0^T Y_t^2 d < P >_t \big]$$

over all limit/market order strategies $\alpha = (\alpha^{make}, \alpha^{take}) \in \mathcal{A}$ such that $Y_T = 0$. Therefore, our objective is to maximize the terminal cash, given that we hold no risky position by time T, and we penalize holding a large inventory during $[0;T]$ by penalizing the integrated variance of the investor's portfolio. $\gamma > 0$ penalizes the quadratic risk of holding an inventory of Y shares in stock P. We can easily get rid of the terminal constraint $Y_T = 0$ by

introducing the liquidation function:

$$L(x,y,p,s) \;=\; x - c(-y,p,s) \;=\; x + yp - |y|\frac{s}{2} - \varepsilon.$$

and now we define the value function:

$$v(t,x,y,p,s) \;=\; \sup_{\alpha\in\mathcal{A}} \mathbf{E}_{t,x,y,p,s}\Big[L(X_T,Y_T,P_T,S_T) - \gamma\int_t^T Y_u^2\varrho(Y_u)du\Big],$$

where we assumed $d < P >_t = \varrho(P_t)dt$.

Since the spread takes discrete values, $s = i\delta$, $i \in \mathbb{I}_m$, we denote

$$v_i(t,x,y,p) \;=\; v(t,x,y,p,i\delta)$$

and we identify v with $(v_i)_{i=1,\dots,m}$: $\mathbb{R}^m$-vector valued function on $[0,T]\times\mathbb{R}\times\mathbb{R}\times\mathbb{P}$. We use similar notations L_i, c_i, π_i^a, π_i^b, λ_i^a, λ_i^b.

And we characterize v_i as the unique viscosity solution of a 3-dimensional QVI, that we will simplify.

Dimension reduction

In order to fasten numerical resolution of the HJB-QVI, we are now interested in reducing the dimensions of the state space. If we assume that P is a Lévy process, we have:

$$\mathcal{P}I_{\mathbb{P}} \;=\; c_P, \qquad d < P >_t = \; \varrho dt,$$

where $I_{\mathbb{P}}$ is the identity, for some constants c_P, ϱ. In this case, we obtain the following reduction. The value function $v = (v_i)_{i=1,\dots,m}$ is in the form:

$$v_i(t,x,y,p) \;=\; L_i(x,y,p) + \phi_i(t,y).$$

Moreover, there exists some constant κ s.t.

$$0 \;\le\; \phi_i(t,y) \;\le\; (T-t)\kappa,$$

for all $(t,y,i) \in [0,T]\times\mathbb{R}\times\mathbb{I}_m$.

Finally, the simplified problem reads as a system of unidimensionnal QVI:

$$\begin{aligned}
&\min\Big[-\frac{\partial\phi_i}{\partial t} - yc_P + \gamma\varrho y^2 - \lambda(t)\sum_{j=1}^m \rho_{ij}\big[\phi_j - \phi_i + |y|(j-i)\frac{\delta}{2}\big] \\
&\quad - \sup_{(q^b,\ell^b)\in\mathcal{Q}_i^b\times[0,\bar\ell]} \lambda_i^b(q^b)\big[\phi_i(t,y+\ell^b) - \phi_i(t,y) + \frac{i\delta}{2}(|y| + \ell^b - |y+\ell^b|) - \delta\ell^b 1_{q^b = Bb_+}\big] \\
&\quad - \sup_{(q^a,\ell^a)\in\mathcal{Q}_i^a\times[0,\bar\ell]} \lambda_i^a(q^a)\big[\phi_i(t,y-\ell^a) - \phi_i(t,y) + \frac{i\delta}{2}(|y| + \ell^a - |y-\ell^a|) - \delta\ell^a 1_{q^a = Ba_-}\big]\;; \\
&\qquad \phi_i(t,y) - \sup_{e\in\mathbb{R}}\big[\phi_i(t,y+e) - \frac{i\delta}{2}(|y+e| + |e| - |y|) - \varepsilon\big] \;=\; 0,
\end{aligned}$$

for $(t, y, i) \in [0, T) \times \mathbb{R} \times \{1, \dots, m\}$, together with the terminal condition:

$$\phi_i(T, y) \quad = \quad 0, \qquad \forall y \in \mathbb{R},\ i = 1, \dots, m.$$

Numerical scheme and results

We solve the QVI numerically, by providing an explicit backward numerical scheme. We first discretize the time line, by introducing a simple time grid on $[0, T]$: $\mathbb{T}_n = \{t_k = kh, k = 0, \dots, n\}$, $h = T/n$. Then, we discretize and localize the inventory domain: $\mathbb{Y}_{R,M} = \{\ell \frac{R}{M}, \ell = -M, \dots, M\}$. On the boundaries, $\ell = \pm M$, orders are place on only one side of the book.

$(\phi_i)_{i=1,\dots,m}$ approximated by $(\phi_i^{h,R,M})_{i=1,\dots,m}$, starting from the terminal condition: $\phi_i^{h,R,M}(t_n, y) = 0$, and we obtain the numerical scheme $\mathcal{S}^{h,R,M}$ by replacing the following quantities in the system of non local differential equations:

$$\frac{\partial \phi_i}{\partial t}(t_k, y) \quad \sim \quad \frac{\phi_i^{h,R,M}(t_k + h, y) - \phi_i^{h,R,M}(t_k, y)}{h}$$

the non local terms at (t_k, z, i) computed at time $t_k + h$ with:

$$\phi_i(t_k, z) \quad \sim \quad \phi_i^{h,R,M}(t_k + h, \mathrm{Proj}_{[-R,R]}(z))$$

So that we can write the explicit backward scheme:

$$\begin{array}{rl} & \phi_i^{h,R,M}(t_k, y) \\ = & \mathcal{S}^{h,R,M}\Big(t_k, y, \phi_i^{h,R,M}(t_k + h, .), \big(\phi_j^{h,R,M}(t_k + h, y)\big)_{j=1,\dots,m}\Big), \end{array}$$

and we prove that $\mathcal{S}^{h,R,M}$ is stable, and monotone provided that:

$$\Big[\max_{i \in \mathbb{I}_m, q^b \in \mathcal{Q}_i^b} \lambda_i^b(q^b) + \max_{i \in \mathbb{I}_m, q^a \in \mathcal{Q}_i^a} \lambda_i^a(q^a) + \sup_{t \in [0,T]} \lambda(t)\Big] h < 1,$$

Moreover $\mathcal{S}^{h,R,M}$ is consistent (when $h \to 0,\ M, N \to \infty$), hence convergent by using Barles-Souganidis [8] arguments.

Finally, we provide detailled numerical tests, along with a backtest and performance analysis on simulated data, and we produce here the main figures: figure 2.4 represents two views of the optimal policy, at two different dates, and table 2.2 is a synthesis our benchmarked performance analysis. We also plotted here the empirical distribution of the performance in figure 2.3.2 and the efficient frontier, obtained by varying the arbitrary parameter γ, in figure 2.6

(a) near date 0 (b) near date T

Figure 2.4: Stylized shape of the optimal policy sliced in YS.

		optimal $\alpha^\star$	WoMO α^w	constant α^c	random α^r
Terminal wealth	$m(X_T)/\sigma(X_T)$	2.117	1.999	0.472	0.376
	$m(X_T)$	26.759	25.19	24.314	24.022
	$\sigma(X_T)$	12.634	12.599	51.482	63.849
Num. of exec. at bid	$m(N_T^b)$	18.770	18.766	13.758	21.545
	$\sigma(N_T^b)$	3.660	3.581	3.682	4.591
Num. of exec. at ask	$m(N_T^a)$	18.770	18.769	13.76	21.543
	$\sigma(N_T^a)$	3.666	3.573	3.692	4.602
Num. of exec. at market	$m(N_T^{market})$	6.336	0	0	0
	$\sigma(N_T^{market})$	2.457	0	0	0
Maximum Inventory	$m(\sup_{s\in[0;T]} \lvert Y_s\rvert)$	241.019	176.204	607.913	772.361
	$\sigma(\sup_{s\in[0;T]} \lvert Y_s\rvert)$	53.452	23.675	272.631	337.403

Table 2.2: Performance analysis: synthesis of benchmarked backtest (10^5 simulations).

2.3.3 Optimal high-frequency trading in a pro-rata microstructure with predictive information

In chapter 6, we investigate a mixed market-making strategy in a exotic microstructure, called the pro-rata microstructure. This microstructure can be encountered for example on short-term interest rates futures. Here again, we consider the situation of an investor willing to maximize their terminal profit over a finite time horizon, who is able to trade with limit and market orders. We adopt the perspective of inventory management, which

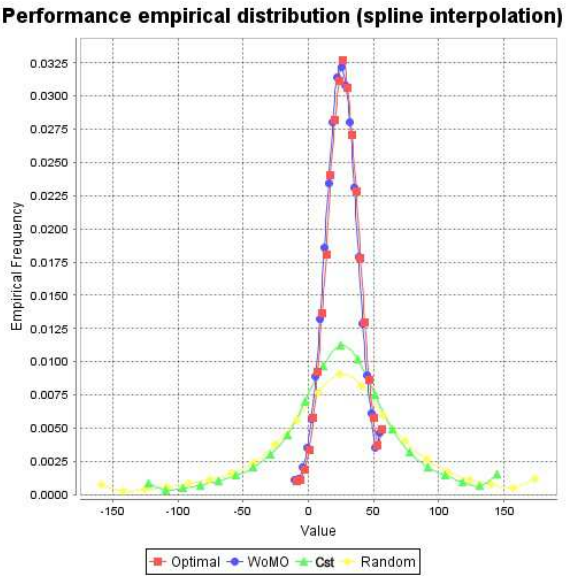

Figure 2.5: Empirical distribution of terminal wealth X_T (spline interpolation).

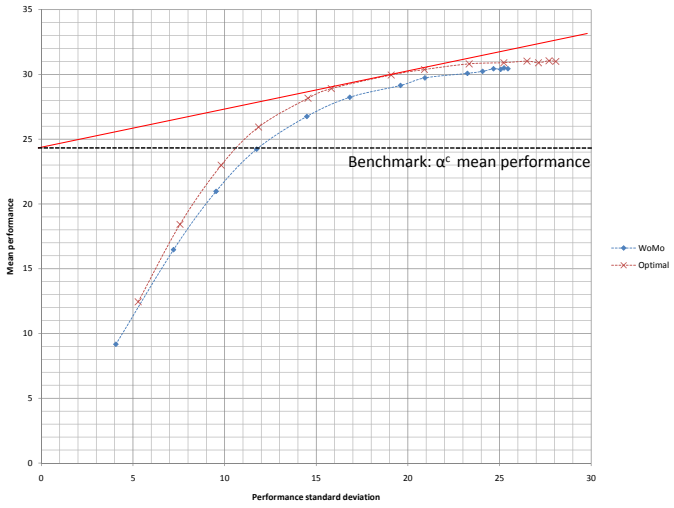

Figure 2.6: Efficient frontier plot

means that the investor primary objective is to keep their position on the risky asset close to zero at all times, in order to avoid being exposed to market risk. In this particular microstructure, we are able to define and address two other types of risk: the overtrading risk, which is the risk of large variations in the investor inventory, due to the fact that they

do not control the quantity they trade at limit ; and the adverse selection risk, which is the risk of market reacting unfavorably to the investor quotes. For this last purpose, we introduce a new state variable, that we interpret as a predictive price indicator, that allows us to balance our position before the price changes. This last feature also provides an extra performance on our empirical tests.

We are interested in the so-called "vanilla Pro-Rata microstructure", which can be described succintly the following way: each incoming market order is dispatched on all active limit orders on the best priced slice of the LOB, proportionnally to each limit order's volume. Figures 2.7 describes the peering of a market order with pre-existing limit orders in the LOB.

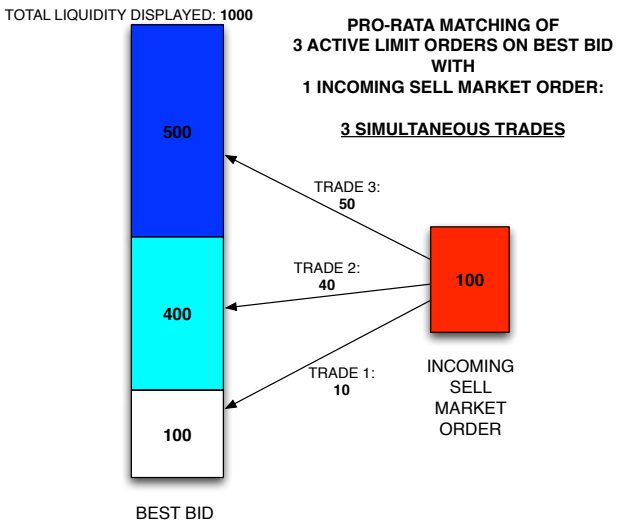

Figure 2.7: Simultaneous trades triggered by a market order.

This type of microstructure, along with characteristic tick size leads to 2 particularities, as shown on figure 2.8 (top left instruments), reproduced from [25]:

- First, the bid/ask spread is most of the time equal to 1 tick
- Second, the liquidity offered on the best priced slices of the LOB is largely oversized compared to the average transaction size.

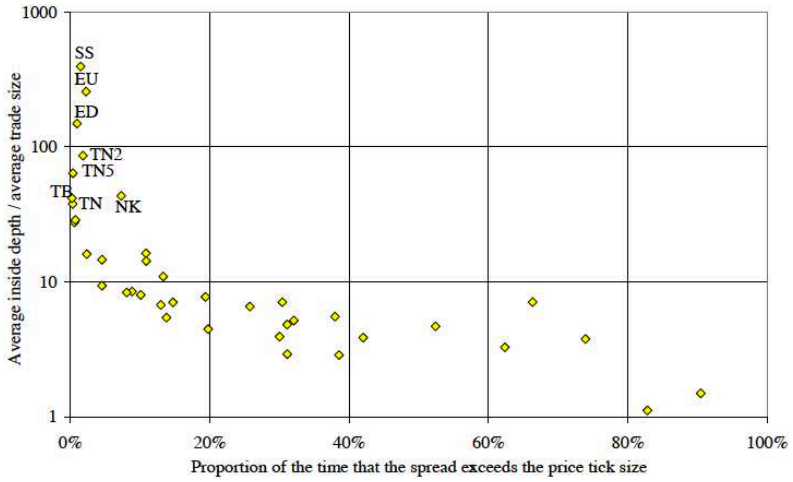

Figure 2.8: Market exploration

Our work is based on the inventory management approach as developped by Avelaneda and Stoikov (2008) [7]. We also used some methods from [35], [37] or [68]. Moreover, the idea of using a predictive price indicator comes from [21]. Finally, we matched our empirical results with the work [25] that is among the few that are dedicated to such market.

Market model

We use the following simple price model:

- P the mid-price (observable: *lit* microstructure): a Markov process of generator $\mathcal{P}$ valued in $\mathbb{P}$. It is assumed to be a special semimartingale.
- δ the tick size, generally on STIR: 12.5 EUR per contract
- P^a (resp. P^b) the ask (resp. bid) price (*one-tick* microstructure):

$$P^a := P + \delta/2 \ , \ P^b := P - \delta/2$$

Now, we consider mixed trading strategies, i.e. made of limit orders and market orders, that are modelled respectively as continuous controls and impulse controls. Indeed, limit order submission, update or cancel is free of charge, therefore limit orders are modelled as time-continuous controls. On the contrary, execution is costly, therefore market orders,

leading to immediate execution, are modelled as impulse controls. More precisely, a trading strategy is a pair $\alpha := (\alpha^{make}, \alpha^{take})$ of regular/impulse controls:

$$\alpha^{make} := (L^a_t, L^b_t)_{t\geq 0} \ , \ \alpha^{take} := (\tau_n, \xi_n)_{n\in\mathbb{N}}$$

where L^a and L^b are predictable process, valued in $\{0,1\}$, representing the *make* regimes. $L^a = 1$ (resp. $L^b = 1$) means that the high frequency trader has active limit order on ask (resp. bid) side. Also, (τ_n) is sequence of non-decreasing stopping times, and ξ_n is a $\mathcal{F}_{\tau_n}$-measurable random variable, valued in $[-\bar{e}, \bar{e}]$, representing the quantity purchased (if $\xi_n > 0$) or sold (if $\xi_n < 0$) by the high frequency trader. The set of such strategies is denoted by $\mathcal{A}$.

Let us now describe our model for trade processes. Due to pro-rata rule, liquidity providers must oversize their limit orders: they post orders with much higher volume than they really intend to trade. This is a way to catch a larger incoming market order volume. Therefore they do not control the size of trades in which they participate: this is the overtrading risk. This differs from the price-time microstructure. The incoming market volume at ask (resp. bid), in which the high frequency trader participate, will be modelled by a random Poisson measure ν^a (resp. ν^b) of intensity $\lambda dt \times \mu(dz)$ on $\mathbb{R}^+ \times \mathbb{R}^+$. $\lambda > 0$ represent the trade clock intensity and μ is a finite measure representing the distribution of a single trade's volume. We also define:

- The number of trades in which the HFT participated:

$$N^a_t := \int_0^t \int_{z\geq 0} \nu^a(dt, dz) \ , \ N^b_t := \int_0^t \int_{z\geq 0} \nu^b(dt, dz)$$

- The cumulative volume executed by the HFT:

$$\vartheta^a_t := \int_0^t \int_{z\geq 0} z\nu^a(dt, dz) \ , \ \vartheta^b_t := \int_0^t \int_{z\geq 0} z\nu^b(dt, dz)$$

In this situation, we are able to describe the evolution of variables describing our portfolio. Inventory Y and cash X evolve under the following dynamics under control α:

$$\begin{aligned}
dY_t &= L^b_t d\vartheta^b_t - L^a_t d\vartheta^a_t \ , \ \tau_n \leq t < \tau_{n+1} \\
dX_t &= L^a_t (P_t + \frac{\delta}{2}) d\vartheta^a_t - L^b_t (P_t - \frac{\delta}{2}) d\vartheta^b_t \ , \ \tau_n \leq t < \tau_{n+1} \\
Y_{\tau_n} - Y_{\tau_n -} &= \xi_n \\
X_{\tau_n} - X_{\tau_n -} &= -\xi_n P_{\tau_n} - |\xi_n| (\frac{\delta}{2} + \epsilon) - \varepsilon_0
\end{aligned}$$

where $\epsilon > 0$ is a per share trading cost and $\varepsilon_0 > 0$ is a fixed trading cost. Remark that the marked to market value of portfolio (or book value, liquidative value) evaluated at

mid-price, $V := X + YP$ evolve under the following dynamics:

$$\begin{aligned} dV_t &= \frac{\delta}{2}(L_t^a d\vartheta_t^a + L_t^b d\vartheta_t^b) + Y_{t-} dP_t \\ V_{\tau_n} - V_{\tau_n -} &= -(\frac{\delta}{2} + \epsilon)|\xi_n| - \varepsilon_0 \end{aligned}$$

Optimization

The system is completely determined by the state variables (X, Y, P) controlled by the limit/market orders strategy $\alpha \in \mathcal{A}$. Let $T > 0$ be a finite time horizon. We choose to :

$$\text{maximize } \mathbf{E}\left[X_T - \gamma \int_0^T Y_t^2 d < P >_t\right] \text{ over all } \alpha \in \mathcal{A} \text{ s.t. } Y_T = 0$$

where $\gamma > 0$ is a penalization parameter. This is equivalent to:

$$\text{maximize } \mathbf{E}\left[L(X_T, Y_T, P_T) - \gamma \int_0^T Y_t^2 \varrho(P_t) dt\right] \text{ over all } \alpha \in \mathcal{A}$$

where it is assumed that $d < P >_t = \varrho(P_t)dt$, with ϱ positive, continuous on $\mathbb{R}$. The liquidation function L is equal to:

$$L(x, y, p) = x + yp - |y|(\frac{\delta}{2} + \epsilon) - \varepsilon_0$$

We now define the value function:

$$v(t, x, y, p) := \sup_{\alpha \in \mathcal{A}} \mathbf{E}_{t,x,y,p}\left[L(X_T, Y_T, P_T) - \gamma \int_t^T Y_s^2 \varrho(P_s) ds\right]$$

and we have some bounds on the value function (Proposition 6.3.1): there exist a constant $K_P \in \mathbb{R}$ s.t.: $L(x, y, p) \leq v(t, x, y, p) \leq x + yp + \delta\lambda\bar{\mu}(T - t) + K_P$ where $\bar{\mu}$ is the mean of μ.

Now we introduce the operators involved in the DPP. For any $(\ell^a, \ell^b) \in \{0, 1\}^2$ we define the non-local operator associated with the limit order control:

$$\mathcal{L}^{\ell^a, \ell^b} := \mathcal{P} + \ell^a \Gamma^a + \ell^b \Gamma^b$$

where

$$\begin{aligned} \Gamma^a \phi(t, x, y, p) &:= \lambda \int_0^\infty \left[\phi(t, x + z(p + \delta/2), y - z, p) - \phi(t, x, y, p)\right] \mu(dz) \\ \Gamma^b \phi(t, x, y, p) &:= \lambda \int_0^\infty \left[\phi(t, x - z(p - \delta/2), y + z, p) - \phi(t, x, y, p)\right] \mu(dz) \end{aligned}$$

We also define the impulse operator (obstacle) associated with the market order control:

$$\mathcal{M}\phi(t, x, y, p) := \sup_{e \in [-\bar{e}; \bar{e}]} \phi(t, x - ep - |e|(\delta/2 + \epsilon) - \varepsilon_0, y + e, p)$$

The dynamic programming equation associated to this problem is a QVI:

$$\min\left\{-\frac{\partial v}{\partial t}-\sup_{(\ell^a,\ell^b)\in\{0,1\}^2}\mathcal{L}^{\ell^a,\ell^b}v+\gamma g\ ;\ v-\mathcal{M}v\right\}=0\ ,\quad \text{on } [0,T)\times\mathbb{R}^2\times\mathbb{P}$$

together with terminal condition:

$$v(T,.)=L\ ,\quad \text{on } \mathbb{R}^2\times\mathbb{P}$$

where we denoted $g(y,p)=y^2\varrho(p)$. This last equation can be expressed explicitly (see chapter 5).

Dimension reduction

We are able to simplify this last QVI in the case where the mid-price is a Lévy process so that:

$$\mathcal{P}\mathbb{I}_{\mathbb{P}}=c_P \text{ and } \varrho \text{ is constant.}$$

where $\mathbb{I}_{\mathbb{P}}$ is the identity function on $\mathbb{P}$ i.e. $\mathbb{I}_{\mathbb{P}}(p)=p$ and c_P is a constant depending only on the characteristic triplet of P.
In this Lévy context, the value function v is decomposed into the form:

$$v(t,x,y,p)=L(x,y,p)+w(t,y)$$

so we see this decomposition makes the liquidative value of the portfolio apparent. With this simplification, we have w is solution to the integral variational inequality:

$$\min\Big[-\frac{\partial w}{\partial t}-yc_{_P}+\gamma\varrho y^2-\mathcal{I}^a w-\mathcal{I}^b w\ ,\ w-\tilde{\mathcal{M}}w\Big]\ =\ 0,\ \text{ on }\ [0,T)\times\mathbb{R},$$

together with the terminal condition:

$$w(T,y)\ =\ 0,\quad \forall y\in\mathbb{R},$$

where $\mathcal{I}^a$ and $\mathcal{I}^b$ are the nonlocal integral operators:

$$\begin{aligned}\mathcal{I}^a w(t,y)\ &=\ \lambda^a\Big(\int_0^\infty\Big[w(t,y-z)-w(t,y)+z\frac{\delta}{2}+(\frac{\delta}{2}+\varepsilon)(|y|-|y-z|)\Big]\mu^a(dz)\Big)_+\\ \mathcal{I}^b w(t,y)\ &=\ \lambda^b\Big(\int_0^\infty\Big[w(t,y+z)-w(t,y)+z\frac{\delta}{2}+(\frac{\delta}{2}+\varepsilon)(|y|-|y+z|)\Big]\mu^b(dz)\Big)_+,\end{aligned}$$

and $\tilde{\mathcal{M}}$ is the nonlocal operator:

$$\tilde{\mathcal{M}}w(t,y)\ =\ \sup_{e\in[-|y|,|y|]}\Big[w(t,y+e)-(\frac{\delta}{2}+\varepsilon)(|y+e|+|e|-|y|)-\varepsilon_0\Big].$$

Finally, we have bounds and symmetry properties for w.

- We have the bounds (comparison principle):

$$0 \;\leq\; w(t,y) \;\;\leq\;\; (T-t)\Big[\frac{c_P^2}{4\gamma\rho} + \lambda^a(\delta+\epsilon)\bar{\mu}^a + \lambda^b(\delta+\epsilon)\bar{\mu}^b\Big],$$

- Stressing the dependence in c_P, we have that

$$w(t,y,c_P) = w(t,-y,-c_P)$$

Numerical scheme

We provide an explicit backward computational scheme for the integral variational inequality. Let us define a regular time grid:

$$\mathbb{T}_N := \{t_k = kh\,,\; k = 0,\ldots,N\}$$

and a regular discretization/truncation of the state space:

$$\mathbb{Y}_M \;\;=\;\; \big\{y_i = i\Delta_Y,\; i = -N_Y,\ldots,N_Y\big\}.$$

Finally, we denote by $\mathrm{Proj}_M(y) := -M \vee (y \wedge M)$, and consider the discrete approximating distribution of μ^a and μ^b, defined by:

$$\hat{\mu}^a \;=\; \sum_{i\in\mathbb{Z}^+} \mu^a([i\Delta_Y;(i+1)\Delta_Y))\delta_{i\Delta_Y} \;\;, \qquad \hat{\mu}^b \;=\; \sum_{i\in\mathbb{Z}^+} \mu^b([i\Delta_Y;(i+1)\Delta_Y))\delta_{i\Delta_Y},$$

with δ_x the Dirac measure at x. For any real-valued function φ on $[0,T]\times\mathbb{R}$, $t\in[0,T]$, and $y\in\mathbb{R}$, we define:

$$\mathcal{S}^{h,\Delta_Y,M}(t,y,\varphi) \;\;=\;\; \max\Big[\mathcal{T}^{h,\Delta_Y,M}(t,y,\varphi)\;;\;\tilde{\mathcal{M}}^{h,\Delta_Y,M}(t,y,\varphi)\Big],$$

where

$$\begin{aligned}
&\mathcal{T}^{h,\Delta_Y,M}(t,y,\varphi) = \\
&\varphi(t,y) - h\gamma\varrho y^2 + hyc_P \\
&+\lambda^a h\Big(\int_0^\infty \big[\varphi(t,\mathrm{Proj}_M(y-z)) - \varphi(t,y)\big]\hat{\mu}^a(dz) + \big[\frac{\delta}{2}z + (\frac{\delta}{2}+\varepsilon)(|y|-|y-z|)\big]\mu^a(dz)\Big)_+ \\
&+\lambda^b h\Big(\int_0^\infty \big[\varphi(t,\mathrm{Proj}_M(y+z)) - \varphi(t,y)\big]\hat{\mu}^b(dz) + \big[\frac{\delta}{2}z + (\frac{\delta}{2}+\varepsilon)(|y|-|y+z|)\big]\mu^b(dz)\Big)_+,
\end{aligned}$$

and

$$\begin{aligned}
&\tilde{\mathcal{M}}^{h,\Delta_Y,M}(t,y,\varphi) \\
&= \sup_{e\in\mathbb{Y}_M\cap[-|y|,|y|]} \big[\varphi(t,\mathrm{Proj}_M(y+e)) - (\frac{\delta}{2}+\varepsilon)(|y+e|+|e|-|y|) - \varepsilon_0\big].
\end{aligned}$$

Finally, we prove that this numerical scheme is monotonous, stable and consistent (Proposition 6.4.1-6.4.2-6.4.3) and therefore the solution $w^{h,\Delta_Y,M}$ to the numerical scheme converges locally uniformly to w on $[0,T)\times\mathbb{R}$, as (h,Δ_Y,M) goes to $(0,0,\infty)$ (Theorem 6.4.1).

Application: HFT with information on price trend

Finally, we made numerical tests with the mid price assumed to be a Lévy process, on which we have directionnal information. More precisely, we assume that:

- The mid-price P is a pure jump process valued in the discrete grid $\delta\mathbb{Z}$.
- We have:

$$\begin{aligned}\mathbf{P}\left(P_{t+h}-P_t=\delta\,|\mathcal{F}_t\right) &= \pi^+h+o(h)\\ \mathbf{P}\left(P_{t+h}-P_t=-\delta\,|\mathcal{F}_t\right) &= \pi^-h+o(h)\\ \mathbf{P}\left(|P_{t+h}-P_t|>\delta\,|\mathcal{F}_t\right) &= o(h)\end{aligned}$$

with fixed $\pi^+,\pi^->0$ and we denote $\varpi:=\pi^+-\pi^-$

Therefore $\mathcal{P}\mathbb{I}_{\mathbb{P}}=c_P=\varpi\delta$ and $\varrho(.)\equiv(\pi^++\pi^-)\delta^2$.

In this context we are able to compute the value function and the optimal policy (figure 2.9).

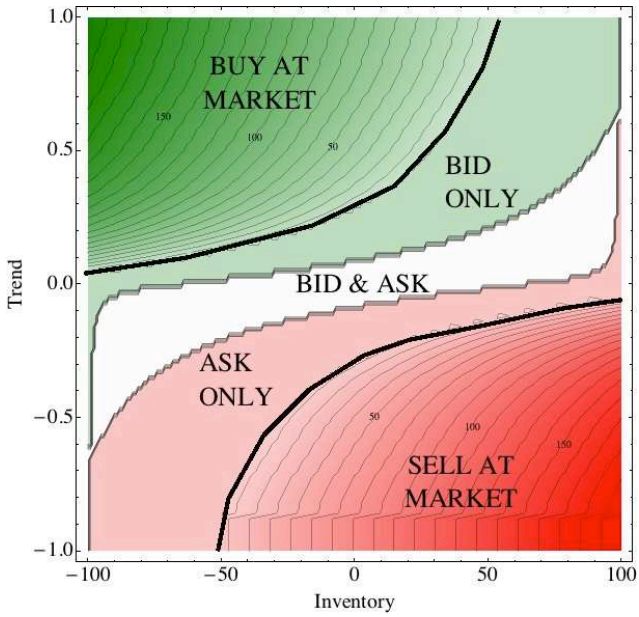

Figure 2.9: Policy $\alpha^\star$ at date $t=0$ for the trend information case.

As a numerical illustration, we performed a detailled performance analysis on simulated data, against a benchmark strategy made of constant controls. The main results are gathered in table 2.10.

Quantity	Definition	$\alpha^\star$	α^{WoMO}	α^{cst}
Info ratio over T	$m(\hat{V}_T^{\cdot})/\sigma(\hat{V}_T^{\cdot})$	3.67	0.89	0.18
Profit per trade	$m(\hat{V}_T^{\cdot})/m(\hat{Q}^{total,\cdot})$	8.06	16.31	5.57
Risk per trade	$\sigma(\hat{V}_T^{\cdot})/m(\hat{Q}^{total,\cdot})$	2.19	18.31	29.56
Mean performance	$m(\hat{V}_T^{\cdot})$	31446.4	28246.3	21737.2
Standard deviation of perf	$\sigma(\hat{V}_T^{\cdot})$	8555.46	31701.2	115312
Skew of perf	skew$(\hat{V}_T^{\cdot})$	0.64	0.16	-0.007
Kurtosis of perf	kurt$(\hat{V}_T^{\cdot})$	3.82	3.31	7.02
Mean total executed volume	$m(\hat{Q}^{total,\cdot})$	3900.68	1730.82	3900.61
Mean at market volume	$m(\hat{Q}^{market,\cdot})$	1932.29	0	0
Ratio market over total exec	$m(\hat{Q}^{market,\cdot})/m(\hat{Q}^{total,\cdot})$	0.495	0	0

Figure 2.10: Synthetis table for backtest. Categories are, from top to bottom: relative performance metrics, period-adjusted performance metrics, absolute performance metrics and absolute activity metrics.

Chapter 3

Literature survey: quantitative methods in high-frequency trading

This chapter surveys the current literature on quantitative methods in high-frequency trading. In a first part, we review some recent developments in market impact and execution costs modelling. We sum up the original linear impact framework, then we describe the limit order book resilience-based optimization frameworks. We also review continuous-time trading models and impulse control models, and finally we expose a recent work on Smart Order Routing.

Then, in a second part, we survey the literature on market-making and mixed trading strategies. We expose the classical inventory management framework with the associated linear market-making strategies. Then, we review recent market model enhancements, including mixed strategies, and high-dimensional modelling of the trades processes. We conclude by presenting a recent work on a Poissonian model for the limit order book dynamics.

3.1 Introduction

In this chapter, we propose an overview of quantitative literature on high frequency trading. From the modelling perspective, our objective is to compare and to expose the contributions of several models for each topic. From the financial application perspective, we are interested in the two main application of high-frequency trading: trading costs optimization, and market-making. Many of the topics that we present in that section will be further detailled in the rest of that thesis, and we present the context and motivations for some of our modelling choices.

On one hand, an extensive literature is devoted to optimizing the trading costs at the scale of a single transaction, since the seminal work of Almgren and Chriss [3]. Indeed, a genuine application of high-frequency trading is to systematically improve the outcome of an order sent by a human investor, by minimizing direct and indirect trading costs. The general mechanism of such costs minimization strategy is to split the original transaction, ordered by the investor, in multiple children transactions and dispatch them on an extended time period and on several trading venues.

We present several distinct approaches that solves a time optimization problem, where the goal is to reduce market impact while dynamically managing market risk. Market impact is a key factor when executing large orders. It is the difference between the reference price and the realized transaction price: due to finite liquidity available in the limit order book, large transactions induce a change in the quoted price that is unfavourable to the investor. Therefore, the investor faces a trade-off between trading quickly and being exposed to market impact, and trading gradually during a certain time period and being exposed to price fluctuation, i.e. market risk. We also mention approaches that solves a space optimization problem, where the goal is to find the best available price among several trading venues, usually including displayed and hidden liquidity, the so-called Smart Order Routing (SOR) strategies.

On the other hand, market-making strategies, and mixed passive/active high-frequency strategies received large academic interest, since the reference work of [7], that modernized the framework of [5]. Such papers examine a direct application of high-frequency trading: an investor who continuously submit bid and ask quotes in a limit order book wants to control its exposure to market risk, by keeping its position on the risky asset close to zero at all time. The standard approach is to present this problem as an inventory management problem, where the investor seeks at maximizing their terminal profit over a short-term horizon, while balancing their inventory by choosing their position in the limit order book.

We first detail the inventory management framework, along with the main concepts associated to market-making strategies, in the context of pure passive strategies, and with a very simple market model. We describe the concepts of uncertain execution, and of inventory risk, and we recall the main steps in the Avellaneda and Stoikov solution. We

also present the main results of the recent extension of this framework in [35].

Then, we present recent developments in market-making strategies and richer market models, and especially two specific tools: the use of both market and limit orders in such strategies, and the use of predictive information as an input of the market-making strategy. We detail the mixed stochastic control approach to the market/limit order strategies with recent example of application. We describe the main evolutions in market models underlying such strategies, with a specific focus on trades flows modelling with point processes. We also present recent models for the limit order book dynamics.

This chapter is organized as follows. In the first section 3.2, we propose an overview of indirect trading costs minimizations strategies. We sum up the original Almgren and Chriss framework, then go on explaining optimization framework that are based on the resilience of the limit order book. We also review continuous-time trading models and impulse control models, and finally we expose a recent work on Smart Order Routing. Then, in 3.3, we survey the literature on market-making and mixed trading strategies. We expose the classical Avellaneda and Stoikov framework with the associated linear market-making strategies. Then, we review recent market model enhancements, including mixed strategies, and high-dimensional modelling of the trades processes. We conclude by presenting a recent work on a Poissonian model for the limit order book dynamics.

3.2 Costs optimization strategies

In this first section, we propose an overview of model-based costs minimization strategies.

Almgren and Chriss framework

The two seminal papers [3] and [10] first provided a framework to manage market impact. In a discrete-time setup, the goal of their work is to minimize the expected costs of trading for an execution strategy: in [10], sole the expectation of the costs are subject to minimization, whereas in [3], the authors take into account a variance criterion in addition to costs minimization. More precisely, [3] considers the following scheme: the investor has X units of a risky asset at time 0 and wants to liquidate the whole portfolio before time T by trading at regularly spaced discrete dates $t_k = k\tau = k\frac{T}{N}$ the quantities n_k with the conditions:

$$x_k = X - \sum_{i=0}^{k} n_i \ , \ x_0 = X \ , \ x_N = 0,$$

and the asset price is a random walk with volatility σ:

$$S_k = S_{k-1} + \sigma\sqrt{\tau}\xi_k - \tau g(\frac{n_k}{\tau}).$$

In this last expression, (ξ_k) are i.i.d standard normal random variables, and the last term

$\tau g(\frac{n_k}{\tau})$ represents the *permanent impact* from the investor's trading on market price. Here, $\frac{n_k}{\tau}$ is the rate of trading. Moreover, the investor is also penalized by a *temporary impact*, that other agents in the market do not face, modelled by :

$$\bar{S}_k = S_k - h(\frac{n_k}{\tau}).$$

In this setup, the cash obtained for the selling operation is:

$$\sum_{k=1}^{N} n_k \bar{S}_k = XS_0 + \sum_{k=1}^{N} \left(\sigma \tau^{\frac{1}{2}} \xi_k - \tau g(\frac{n_k}{\tau}) \right) x_k - \sum_{k=1}^{N} n_k h(\frac{n_k}{\tau}).$$

The authors define thus the following quantities:

$$\text{costs due to market impact} \quad E(x) = \sum_{k=1}^{N} \tau g(\frac{n_k}{\tau}) x_k + \sum_{k=1}^{N} n_k h(\frac{n_k}{\tau}),$$

$$\text{risk due to price variations} \quad V(x) = \sigma^2 \sum_{k=1}^{N} \tau \xi_k^2,$$

and propose the following mean-variance criterion (using the Lagrange multipliers technique, where λ represents the risk aversion parameter):

$$C_\lambda = \min_x E(x) + \lambda V(x).$$

By solving the first order conditions, the authors are able to provide a closed form for the optimal strategy, which is static in the sense that it does not depends on the price path:

$$x_j = \frac{\sinh(\kappa(T - t_j))}{\sinh(\kappa T)} X \ , \ j = 0...N,$$

where κ depends on λ , η and σ. By varying the λ, they are able to provide the efficient frontier.

This method is quite simple to implement and can be easily extended to the case of multiple correlated assets. Moreover, the tractability of computations allows to have a complete numerical description of the strategy behaviour in terms of return/risk ratio, which is useful for the financial intuition. Yet, the optimal solution is static and does not take as an input the price of a risky asset during the liquidation period. This last point is counter-intuitive since one would expect the optimal strategy to sell quickly when the price is high, and more slowly when the asset price is low. The optimal strategy, due to the choice of price dynamics and risk measure, is here a fixed pattern.

Limit Order Book dynamics based framework

Another type of modelling appears in the papers [2],[55], [64], [30], [61] and [31] where the form of the impact is directly inferred from a stylized dynamics of the limit order book (LOB). We will present here the general mechanics of this type of models, irrespective to the particular features of each work. For this purpose and for sake of simplicity, we follow the presentation of [2], that is set up in discrete time framework. A continuous time version of this model is studied in [30].

In these models, the asset's market price is assumed to be a martingale process (or risk-neutral process) (S_t^0) in absence of any trading. When the investor trades, the market price is a perturbation of this martingale price due to the resilience of the order book. Indeed, when the investor decides to buy (resp. sell) a quantity ξ of the risky asset at time t, he will consume the liquidity offered at the bid (resp. ask) side of the LOB, and therefore this will shift down (resp. shift up) the best bid (resp. best ask) quotes, according to the volume traded and the shape of the order book at time t. After time t, the LOB will regenerate by the effect of incoming limit orders that provide liquidity. This regeneration is viewed as a resilience of the LOB around the price (S_t^0). In this type of models, the resulting strategy typically consists in a large trade at first date, a large trade at terminal date, and constant trading in-between.

More precisely, trading strategies are represented on a discrete (regular) time grid $t_k = k\frac{T}{N}$, $k = 0...N$ and, as is Almgren and Chriss framework we have:

$$\sum_{n=0}^{N} \xi_{t_n} = X_0.$$

Here the trade is a purchase of X_0 shares, and the quantity traded at date t_n is ξ_{t_n}. The market price of asset evolves under the following dynamics:

$$S_t = S_t^0 + \sum_{t_n<t} \xi_{t_n} G(t - t_n),$$

where S^0 is the *reference price*, assumed to be a continuous martingale, and the function G is the *resilience function* or *resilience kernel*, and is non-increasing on $[0, \infty)$. This kernel describes the time structure of the impact, indeed:

- The *instantaneous impact* is $\xi_{t_n}(G(0) - G(0+))$, where $G(0+)$ denotes the righthand limit of G at $t = 0$. It only affects the execution cost of ξ_{t_n} and not any subsequent orders.
- The *permanent impact* is $\xi_{t_n} G(\infty)$. It affects all current and future trades equally and vanishes eventually.
- The remaining part, $\xi_{t_n}(G(0+) - G(\infty))$, is called the *transient impact*. Its effect on future trades decays over time.

After setting this model, the authors prove that the expected execution costs for the strategy ξ can be expressed as follows, in the case where the liquidity supply has a constant shape:

$$\mathcal{C}(\xi) := X_0 S_0 + \mathbb{E}(C(\xi)),$$

where C is the quadratic form:

$$C(x) = \sum_{i,j} x_i x_j G(|t_i - t_j|).$$

Notice that this form comes from the fact that the reference price S^0 is a martingale. The optimality criterion is defined by these expected execution costs, and an optimal solution exists and is unique if C is strictly definite positive. In this case, the authors provide a closed form for the optimal solution (which is static):

$$\xi* = \frac{X_0}{1^* M^{-1} 1} M^{-1} 1,$$

where $M := (G(|t_i - t_j|))_{ij}$. For the case when C is not strictly definite positive, the authors introduce and discuss the concept of price manipulation strategies, with several numerical examples.

The advantages of this approach is to start from a natural modelling of the order book, and to derive a closed-form optimal strategy. Moreover, it gives an intuitive interpretation and insights about dynamic arbitrage of the LOB. The optimal strategy has several typical patterns that vary with the shape of the resilience kernel. Yet, in this case again, due to the choice of the dynamics for the reference price, the optimal strategy does not depend on the price path, i.e. it is static and so induces a fixed pattern for all price realizations.

Dynamic strategies

In a recent paper [61], the authors propose an extension of the LOB-based model where the order book can have a general shape, and they analyze the optimal strategy both in discrete time and continuous time. In this paper, the optimality criterion is again the total expected costs of trading, expressed as:

$$\mathcal{C}(X) = \mathbb{E}(C(X)),$$

where X denotes the (continuous-time) strategy and

$$C(X) := \tilde{C}(X) + \int_0^T S_t^0 dX_t,$$

with S^0 the reference price as defined in the above section, and $\tilde{C}(X)$ is a function representing costs of trading due to market impact, and independent on S^0. The integral term

represents the costs of trading due to price variations. Using the integration by parts formula on this integral term, the authors show, under the assumption that martingale price process S^0, that $\mathcal{C}(X)$ does not anymore depend on S^0, so that there is no more source of randomness in the minimization problem of the total expected trading costs. Therefore, the authors deduce that they can restrict their search for optimal strategy to non-random functions of time. This argument can also be formulated in the following way: in this setup, the *statically optimal* strategy is also *dynamically optimal*.

We then see that the both the shape (pattern of execution) and the nature (static or dynamic) of the optimal strategy is fundamentally related to the choice of the setup, and in particular is determined by those two elements:

- Dynamics of the reference price (i.e. price without intervention of the big investor)
- Choice of the optimality criterion.

In the following section, we will present frameworks where the choice of this two elements makes dynamic the optimal strategies.

Continuous-time models

A recent paper by Forsyth [28] proposes to solve the optimal execution problem in a continuous-time framework, formulating the problem as a mean-variance problem leading to an optimal stochastic control problem. The market impact is linear in the trading rate. The resulting optimal strategy is *dynamic* in the sense that it depends both on the time, the price of the risky asset, the cash amount and the quantity of shares in the portfolio.

The computation of the optimal strategy in terms of trading rate involves the numerical resolution of an Hamilton Jacobi Bellman PDE. More precisely, the market model is set up as follows:

Market price of risky asset :	$dS = (\eta + g(v))\, Sdt + \sigma S dW$
Number of shares of underlying asset :	α
Rate of trading :	$v = \dfrac{d\alpha}{dt}$
Cash amount :	$dB = (rB - vSf(v))\, dt,$

where W is a standard Brownian motion. We see that the market price of risky asset is assumed to be a geometrical Brownian motion with drift η and volatility σ, but penalized by a *permanent impact* function g that is linear with slope κ_p. The cash amount of the investor evolves under a risk free return r and the cash obtained from trading is penalized by a *temporary impact* function f which is non-linear. In the paper [28], they assume an exponential form for f but the general methodology of this article still work under more general assumptions on f.

Informally, the execution problem is to reach a fixed number of shares α_T at time T, starting from α_0 at time 0. Therefore, a trading strategy is viewed as a function of state

variables S, B, α and $t \in [0,T]$. It is denoted $v(S,B,\alpha,t)$ and the expected gain from the strategy $v(.)$ is denoted $\mathbb{E}^{t=0}_{v(.)}(B_T)$, with associated variance $\mathbb{V}^{t=0}_{v(.)}(B_T)$. The optimality criterion is chosen to be, for a given expectation d:

$$\begin{cases} & \min \mathbb{V}^{t=0}_{v(.)}(B_T) \\ \text{subject to} & \mathbb{E}^{t=0}_{v(.)}(B_T) = d, \end{cases}$$

together with additionnal admissibility constraints. By using Lagrange multipliers technique, and interpreting γ as the risk aversion of the investor, the problem is reduced to:

$$\min_{v(.)} \mathbb{E}^{t=0}_{v(.)}((B_T - \frac{\gamma}{2})^2).$$

Therefore, by solving this last problem for all γ the author is able to provide the *efficient frontier* for this setup. At this point, the rest of the paper consists in reducing the dimension of the problem, and describing a numerical procedure to solve it. The minimization problem is expressed by means of the dynamic programming principle under the form of an HJB equation: denote by $\mathcal{B}_T = B_T - \frac{\gamma}{2}$, $\tau = T - t$ and $\mathcal{L}V = \dfrac{\sigma^2 S^2}{2} V_{SS} + \eta S V_S$, so that:

$$\begin{cases} V_\tau = \mathcal{L}V + r\mathcal{B}V_{\mathcal{B}} + \min_{v(.)} \left[-VSf(v)V_{\mathcal{B}} + vV_\alpha + g(v)SV_S \right] \\ V(S,\mathcal{B},\alpha,\tau=0) = \mathcal{B}_T^2. \end{cases}$$

There is no closed-form solution to the HJB, but it is possible to map the space of the state variables (S,B,α,t) to the optimal control in terms of rate of trading by solving numerically a PDE associated to the HJB on a discretized grid. The method presented in this paper is the finite difference method, with improvements on the differentiation approximations and on complexity of the computation. This finite difference method is well-suited for solving PDE on domain that have a simple shape, but it is not suitable for complex-shaped domains.

The advantages of this method is that the optimal solution is dynamic and takes into account both market price of the risky asset and cash amount in the portfolio. Moreover, the problem can be reduced to a two-dimensionnal problem, which is quite useful for computing the optimal strategy. This numerical tractability allows the author to obtain a whole risk-return characterization of the optimal strategy by computing the efficient frontier. Yet, one can consider that is not realistic to assume that the investor can trade continuously, in particular if the overall problem is to schedule trading decisions. To address this last scheduling issue, a suitable formulation to the execution problem is provided by impulse control approach as described in the next section.

Impulse control formulation

As seen in previous sections, there exists both continuous-time models and discrete time model to solve an optimal liquidation problem. The principal advantage of a continuous-time model is the use of the powerful stochastic calculus theory, which provides tractable

computations. Yet, it may not be realistic to assume continuous-time trading, especially in presence of transaction costs and illiquidity effects. On the other hand, discrete time models are more readily implementable, but suffer from two shortfalls: first, it may be less easy to have a complete computational treatment of the problem because of the need of ad-hoc resolution method for complex discrete systems; second, the time discretization structure if often chosen *exogenously* or even, in many cases, arbitrarily. Therefore, a discrete model may not be suitable for building a *trading agenda* since in this case, the goal of the investor is to *endogenously* determine the optimal trading times.

The approach of the best execution problem by means of *impulse control* combines the advantages of both continuous-time and discrete-time framework. In this setup, the investor is able to choose discrete-time controls in a continuous-time system: typically, a trading strategy will be the choice of a discrete number of dates τ_n associated with trade quantities ξ_n, which control a state process Z evolving in some diffusive regime. This approach has the advantage of the tractability of stochastic calculus, together with the possibility of a direct implementation. Moreover, the computation of the optimal strategy provides endogenously both the dates and the quantities to trade. This formulation can be seen as a sequence of optimal stopping problems. Therefore, it is possible to use classical optimal stopping theory as the main ingredient for the resolution method. Finally, we will show in later sections that the optimal strategy is *dynamic*, in the sense that it depends both on the market price but also on a set of variables describing entirely the investor portfolio. We mention the papers [39], [51], [12], which address the optimal liquidation problem in terms of impulse control formulation. In chapter 4 we use this last approach, and we detail its resolution.

Smart order routing techniques

Finally, let us conclude this section by mentionning the work [48], that are concerned with the situation of an investor (or actually a broker) that wants to trade an asset on several distincts venues. These works are original in the sense that most of existing solutions to this routing problem (known as Smart Order Routing) are technology based, and does not rely on precise mathematical modelling.

The difficulty of such situation is that the liquidity offered on each marketplace is not publicly displayed, indeed, those works tackle the problem of trading simultaneously on multiple *dark* trading venues, and illustrate the results on *dark pools*. Therefore, the investor does not know if an order sent to the venue i will be executed or not. Moreover, the fees structure (i.e. the amount of money paid by the investor to the marketplace to place an order) differs from one venue to another, and must be taken into account in the optimization. This is the general setup of *dark pools*, where the trade price is stuck to the market mid-price, but liquidity available is not known pre-trade.

In [48], the goal of the investor is to dispatch a large order to several of such marketplaces, with the objective of minimizing the trading costs. The proposed approach involves

the use of recursive stochastic algorithms, and the authors proves the optimality of the resulting strategy.

3.3 Market-making and mixed strategies

In this section, we are interested in describing both standard approach to the market-maker problem, and recent extensions of this framework. In a first part, we put aside market modelling issues, and focus and the optimization framework developped in [5] and [7], and further extensions and observations in recent studies. In a second part, we propose an overview of some rich features recently developped in order to make this optimization framework more suitable to industrial needs, along with popular limit order books models.

The standard inventory management approach and the linear market-making strategy

Pricing strategies of market-makers have received extensive coverage in the microstructure litterature, while quantitative approaches were taken more recently. Survey of such results in microstructure theory can be found in [11] or [56]. Historically, quantitative approaches to market-making policies aimed to address the inventory risk, which is the market risk associated with holding a non-zero portfolio.

The pionnering work in developping "automated" market-making strategies was made in 1980 by Amihud and Mendelsohn [5] and another related work is [42]. They propose to examine a monopolistic market-maker that sets bid and ask prices for some asset. The incoming market order flow (i.e. counterparts of the monopolistic market-maker) is modelled as a price-dependent Poisson process, so that the aggregated buying flow is greater when price in low, and conversely, aggregated selling flow is greater when price is high. In this setup, they show that the bid and ask prices provided by the monopolistic market-maker depends on their inventory, i.e. their position on the risky asset.

More precisely, they study the optimal market-making policy in this context, where the objective of the market-maker is to maximize their average profit per unit-time. They prove (Theorem 3.2 of [5]) that the optimal bid and ask quotes resulting from this optimization are monotone decreasing functions of the inventory held by the market-maker. They establish (Theorem 3.7 of [5]) that "the market-maker adopts a pricing policy which produces a preffered inventory position", in the sense that the optimal strategy consists in choosing bid and ask price in order to favor sell trades when the inventory is positive, and conversely to favor buy trades when the inventory is negative. The rest of the paper is concerned with finer results about this market dynamics.

The idea of presenting market-making as an inventory management problem have been made successful in the more recent work of Avellaneda and Stoikov [7]. In this work, the market-maker pricing is influenced not only by the price-dependent nature of counterpart

order flows (although presented in a slightly different way), but also by market risk. Indeed, in this work, the market-maker is no more monopolistic, and therefore cannot choose the price of the risky asset based on their own objective. Another risk factor is added to the market model that drives the price.

Let us have a brief explanation of this model. They consider the situation of an investor trading with limit orders only, on an asset whose mid-price S is a Brownian motion (Bachelier model) with volatility $\sigma > 0$:

$$dS_u = \sigma dW_u$$

Then, the agent continuously quotes the bid price p^b and the ask price p^a (continuous controls), which means that they are committed to respectively buy and sell one share of stock at these prices when a market order comes in. Then the cash X and the inventory q of the market-maker evolve according to the following dynamics:

$$\begin{aligned} dX_t &= p_t^a dN_t^a - p_t^b dN_t^b \\ q_t &= N_t^b - N_t^a \end{aligned}$$

where N^a and N^b are Cox processes, whose jumps represent respectively trades at ask and bid, and whose intensity depends respectively on (decrease with) $\delta_t^a := p_t^a - S_t$ and $\delta_t^b := S_t - p_t^b$. This decreasing dependence on the distance to mid-price is the modern equivalent of the price-dependent Poisson process appearing in Amihud and Mendelsohn and is chosen to exponential:

$$\begin{aligned} &N^a \sim \mathrm{Cox}\left(\lambda(\delta_t^a)\right) \quad , \quad N^b \sim \mathrm{Cox}\left(\lambda(\delta_t^b)\right) \\ &\lambda(d) := A\exp(-kd) \end{aligned}$$

Where A and k are constants to be fitted, representing characteristics of execution probability. Based on that simple dynamics, the objective of the market-maker is to maximize their profit utility over a finite time-horizon. More precisely, the value function is defined by:

$$u(s,x,q,t) = \sup_{p^a,p^b} \mathbf{E}_t\left[-\exp\left(-\gamma\left(X_T + q_T S_T\right)\right)\right]$$

here, $T > 0$ is a finite time horizon, γ is an arbitrary risk aversion parameter, the utility is chosen to be exponential (for tractability) and $X_T + q_T S_T$ is the terminal marked-to-market value (or book value) of the investor's portfolio.

Applying the dynamic programming principle, the authors obtain the Hamilton-Jacobi-

Bellman equation:

$$\begin{aligned} 0 &= u_t + \frac{1}{2}\sigma^2 u_{ss} \\ &\quad + \max_{\delta^b} \lambda(\delta^b) \left[u(s, x - s + \delta^b, q + 1, t) - u(s, x, q, t) \right] \\ &\quad + \max_{\delta^a} \lambda(\delta^a) \left[u(s, x + s + \delta^a, q - 1, t) - u(s, x, q, t) \right] \\ u(s, x, q, T) &= -\exp(-\gamma(x + qs)) \end{aligned}$$

In the second and third lines, one can identify the infinitesimal generators of N^b and N^a respectively, applied to u, which make such equation non-local. Thanks to a variable change, the authors are able to obtain explicit approximating formulas for the optimal quotes, and they perform numerical tests. The paper [35] provides detailled analytical resolution and experiments, along with several observations about that model that we present in what follows. Indeed, using the same model, authors of [35] show with a variable change that the HJB equation of [7] can be reduced to a system of linear ordinary differential equations. In these conditions, they are able to provide a close-form approximation formula to the optimal quotes.

Indeed, they observe numerically that the behavior of the optimal quotes is almost time-independent when far from the terminal date T, and they argue that this steady-state market-making policy is more relevant than the time-varying one, because of the arbitrary nature of T. In figure 3.1, that we reproduced from [35], representing optimal bid quote as a function of inventory and time, one can see that this optimal quote is mainly time-invariant and linear when far from T.

The authors go on proposing a linear approximation (actually an asymptotics as $T \to \infty$) for the optimal bid and ask quotes in the Avellaneda and Stoikov model. They read (Theorem 2 and proposition 3 of [35]):

$$\begin{aligned} \delta^{b\star}(q) &= Cq + d \\ \delta^{a\star}(q) &= -Cq - d \end{aligned}$$

where C and d are constants which are explicitly given in [35], exhibiting dependence on market volatility, execution probability and risk aversion.

To sum up, this type of model is easily tractable and allows us to obtain closed form linear solutions for optimal quotes, and is parsimonous, since very few market effects are specifically taken into account, and therefore can be fitted to a large class of real-world data. Indeed, a direct data-oriented approach can use such a results to look for the best-performing linear market-making strategy based on backtests results.

Mixed market/limit orders strategies

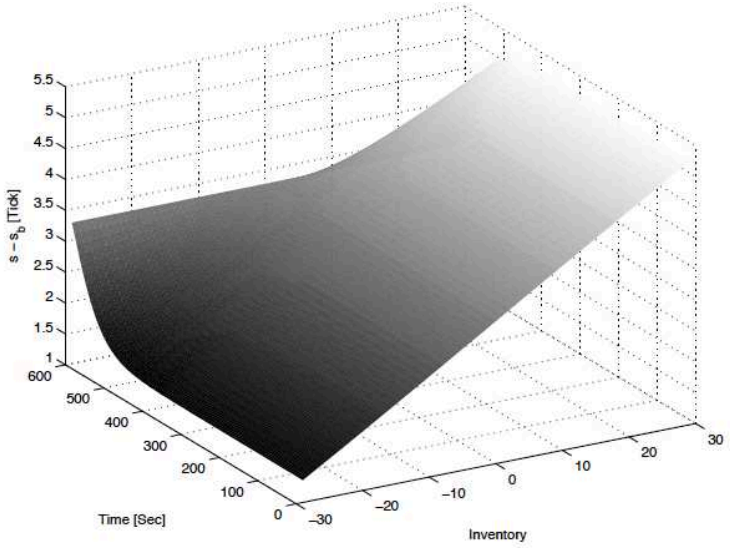

Figure 3.1: Optimal bid quote as a function of inventory and time in the Avellaneda and Stoikov model.

A natural extension of this framework is presented in [67], [68], [37] and [38], and the two last references are adapted in chapter 5 and 6 of this thesis. The idea in these paper is to consider a market orders strategy that will superpose to the limit order placement strategy explained in Avellaneda and Stoikov. The market orders used in these works are "hit orders", which means that they are actually marketable limit orders, i.e. limit orders that hits the opposite side of the LOB, therefore leading to an immediate execution.

Let us propose a brief recall of the approach that we use in this thesis, along with impact of this new perspective on the HJB equation. Mixed strategies are represented as a pair:

$$\alpha := (\alpha^{make}, \alpha^{take})$$

where, on one hand, α^{make}, represents the limit order strategy, directly corresponding to the Avellaneda and Stoikov model. For simplicity sake, let us assume here that it is the pair (p^a, p^b) of limit orders prices as explained in the previous paragraph, represented as predictible continuous-time process. On the other hand, α^{take} has the following structure:

$$\alpha^{take} := (\tau_n, \xi_n)_{n\in\mathbb{N}}$$

where τ_n is a stopping time in the underlying stochastic basis, representing the date where the investor decides to send a market order of size ξ_n (which a mesurable variable at date τ_n). $\xi_n > 0$ represents a buy market order, and $\xi_n < 0$ represents a sell market order.

Now let us show the effect of such extension on the corresponding HJB equation. If we re-write the Avellaneda and Stoikov equation in a less explicit form, in order to abstract from the specific features of the model, in can be written in the following way:

$$0 \quad = \quad u_t + \mathcal{P}u + \mathcal{N}^a u + \mathcal{N}^b u$$

along with some terminal condition at date T. Here $\mathcal{P}$ is the infinitesimal generator of the price process, $\mathcal{N}^a$ is the infinitesimal generator of the trade process at ask (here chosen to be a Cox process), and $\mathcal{N}^b$ the infinitesimal generator of the trade process at bid. Now, if we had the possibility to trade with market orders in addition to the limit orders strategy, this "diffusive part" is embedded in a quasi-variationnal inequality (QVI), where the obstacle part correspond to the market order optimization:

$$0 \quad = \quad \min\left\{-u_t - \mathcal{P}u - \mathcal{N}^a u - \mathcal{N}^b u \; ; \; u - \mathcal{M}u\right\}$$

Here the the operator $\mathcal{M}$ represents the variation of the state variables when trading via market orders. Typically, it will include the costs of crossing the spread, and a proportionnal, per share or fixed fee. For practical example of such operators, we refer to [37] for example.

To sum up, adding a market orders strategy in addition to the limit orders strategy leads to adding an obstacle part the resulting HJB equation. The resulting optimal strategy will be represented as a mapping that associate the optimal control to the current state variable process, including an obstacle region, where it is optimal to trade via a market order.

Enriching the market model

Recent developments ([38], [16], [27]) or in high-frequency trading strategies included building up richer market models. The objective of such work is to take into account in the HF trading strategy such features as: partial execution of the investor's limit orders, more precise dynamics for the trade process, or predictive information on the price trend. Indeed, in practice, the performance of a high-frequency trading strategy depends on the accuracy of the investor's views on short-term evolution of the market, which in turn depends on the accuracy of their market model.

These short-term predictions on the market evolution usually come from three distinct types of arguments. The first and most commonly used type of argument are the so-called statistical arbitrage arguments, that are typically cross-assets. For example, in [6], the authors propose an extension of the pairs trading technique, which means that they exploit the covariance structure of a market sector to trade one stock against the sector. Other famous techniques includes trading one index against sectors ETF (Exchange Traded Funds). The second type of argument comes from limit order books models, as the one

presented in [21]. In such works, the objective is to infer the evolution of price at a very short timescale, typically a few ticks, from the current state of the limit order book. Indeed, by analyzing limit orders data, one is able to compute such quantity as the probability that the price will go up or down at the next tick. Finally, the third type of argument comes from trade processes models, as for example presented in [13], [41], [54] and [16]. These works typically use superior information coming from the detection of autocorrelations, or cross-correlation in trades occuring on a given market, and they use spot estimation of time-varying buying and selling intensities for a given stock. In such models, the market-maker will adapt their quotes not only to control their inventory, but also in function of a dynamic supply/demand process estimated dynamically on the market. In this part, we describe the general framework of such strategy, based on the presentation of [16].

The framework presented by Cartea, Jaimungal and Ricci [16] is very similar to the one of Avellaneda and Stoikov regarding the optimization procedure, but it differs largely when it comes to the market model. Ou goal is to present the modelling ingredients, and how they impact the resulting strategy. First of all, the authors assume that the mid-price of the risky asset is an arithmetic Brownian motion with an adverse selection term:

$$dS_t = \alpha_t dt + \sigma dW_t$$

Where α_t, representing a predictive information on short-term reward, has known dynamics derived from market variables. From this point, they observe that market activity, i.e. the number of trades per second, exhibit burst periods (also called the trade clustering effect). This means that there are short time period where market activity is intense, and this effects quickly reverts to a normal behavior (in a few seconds timescale). We reproduce in figure 3.2 their observation on the stock IBM for a time period of 3 minutes.

Their goal is therefore to provide a point process model, whose jumps will represent trades, and whose intensity will fit such historical process. To do so, they introduce a qualitative distinction between market orders. The first kind are influential orders which excite the state of the market and induce other traders to increase their trading activity. The second type of orders are non-influential orders which are viewed as arising from players who do not excite the state of the market. The proportion of influential market orders is $\rho \in [0; 1]$. Such a distinction is to compare with the asymmetric information model as developped in [56]. In this last model, a certain proportion of the trades come from *informed traders*, who have more information than the market maker, and therefore induces an adverse selection risk from the market maker point of view.

They propose the following "piece-wise exponential" dynamics for the intensities of trades process (i.e. overall market orders counting process) respectively for sell and buy

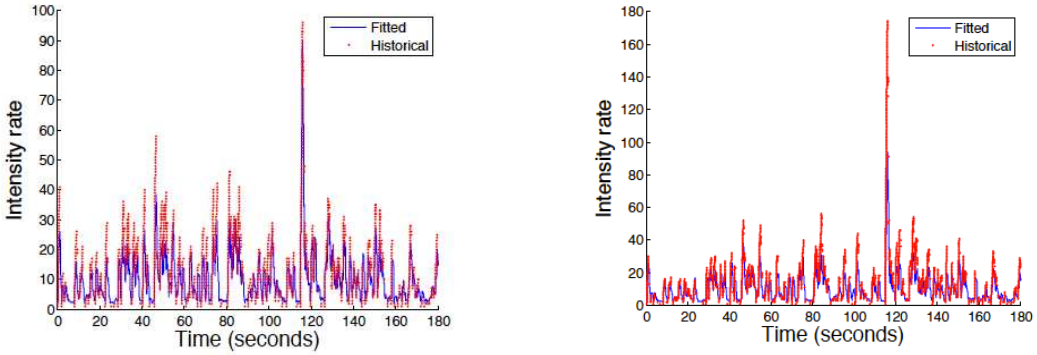

(a) Number of BUY orders per second

(b) Number of SELL orders per second

Figure 3.2: IBM market orders. Historical running intensity using a 1 second sliding window for IBM for a 3 minute period, between 3.30 and 3.33 pm, February 1 2008. Reproduced from [16].

market orders:

$$d\lambda_t^- = \beta(\theta - \lambda_t^-)dt + \eta d\bar{M}_t^- + \nu d\bar{M}_t^+ + \tilde{\eta} dZ_t^- + \tilde{\nu} dZ_t^+$$
$$d\lambda_t^+ = \beta(\theta - \lambda_t^+)dt + \eta d\bar{M}_t^+ + \nu d\bar{M}_t^- + \tilde{\eta} dZ_t^+ + \tilde{\nu} dZ_t^-$$

where Z^+ and Z^- are independent Poisson processes with constant intensity, which represent news events, and M_t^+ and M_t^- are the total number of influential buy and sell orders up until time t. Moreover, η, ν, $\tilde{\eta}$, $\tilde{\nu}$, β, θ are non-negative constants satisfying some constraint.

This choice is a simple version of the (symmetric) Hawkes process model as presented in [41] or [54]. It has the advantadge of providing a tractable SDE while still exhibiting auto- and cross-excitation effects of the trades. We also mention the recent work [24] for useful insights on modelling with self-exciting point processes. We reproduce in figure 3.3 their simulation of the resulting (λ^+, λ^-).

Now, the high-frequency trader only participates in a fraction of trades occuring in the market. Indeed, processes counting the number of trades in which the high-frequency trader participated are denoted N^+ and N^- and their intensities are expressed as functions of λ^+ and λ^-. In [16], they allow a several form for these function, but let us focus on the exponential form, which is closest to the Avellaneda and Stoikov model:

$$N^+ \sim \text{Cox}(\Lambda^+) \quad , \quad N^- \sim \text{Cox}(\Lambda^-)$$
$$\Lambda^+ := \lambda^+ \exp(-\kappa_t^+ \delta^+) \quad , \quad \Lambda^- := \lambda^- \exp(-\kappa_t^- \delta^-)$$

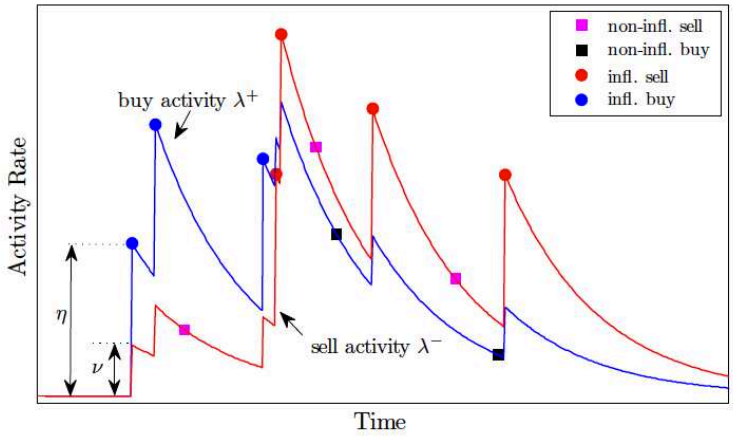

Figure 3.3: Sample path for (λ^+, λ^-). Reproduced from [16].

where δ^+ and δ^- are the distance of the market maker (ask and bid respectively) quotes to the mid-price (continuous controls), and κ^+ and κ^- are the so-called execution intensities at ask and bid respectively. Note that the processes κ^+ and κ^- are the time-varying equivalent of the parameter k in the Avellaneda and Stoikov model. They parametrize the probability that the market maker receive an execution on their bid or ask limit order. They have their own (piecewise exponential) dynamics that reads as follows:

$$\begin{aligned} d\kappa_t^- &= \beta_\kappa(\theta_\kappa - \kappa_t^-)dt + \eta_\kappa d\bar{M}_t^- + \nu_\kappa d\bar{M}_t^+ \\ d\kappa_t^- &= \beta_\kappa(\theta_\kappa - \kappa_t^+)dt + \eta_\kappa d\bar{M}_t^+ + \nu_\kappa d\bar{M}_t^- \end{aligned}$$

The final elements of the model are the dynamics of the portfolios variables, that are exactly similar those of Avellaneda and Stoikov:

$$\begin{aligned} q_t &= N_t^- - N_t^+ \\ dX_t &= (S_t + \delta_t^+)dN_t^+ - (S_t - \delta_t^-)dN_t^- \end{aligned}$$

where q is the inventory process, and X is the cash process. Now the market-maker faces a 6-dimensional optimization problem. Indeed, the value function associated to the market-maker problem reads:

$$\Phi(t, x, S, \lambda^+, \lambda^-, \kappa^+, \kappa^-) = \sup_{\delta^+, \delta^-} \mathbf{E}_t \left(X_T + q_T(S_T + \alpha q_T) - \phi\sigma^2 \int_t^T q_s^2 ds \right)$$

The rest of the paper is devoted to the resolution of this control problem. The associated HJB equation is a non-local variationnal equality, somewhat similar as the one observed in Avellaneda and Stoikov, however more sophisticated due to the presence of varying market orders intensities and execution probabilities. The authors are able to provide an explicit form for the optimal controls as function of the state variables. They also provide a brief procedure for model calibration.

To sum up, in this setup, the author performs a 6-dimensional optimization procedure, in which they input rich information about trade process and execution probabilities. This model is strongly related to self-exciting point process models of trades, similar to those that appears in [54] and [41] for example. This type of model have been proved to perform better than the Poisson model in empirical studies.

Limit order book models

Finally, let us mention the papers [21]. In this work, the authors build up a simple stochastic model for the dynamics of a limit order book, in which arrivals of market order, limit orders and order cancellations occurs at jump times of a Poisson process. Although it has been shown (e.g.[54]) that this Poissonian framework performs poorly when it comes to fitting real-world trades processes, the tractability of this model allows the authors to compute analytically various quantities related to the LOB such as the distribution of the duration between price changes, the distribution and autocorrelation of price changes, and the probability of an upward move in the price, conditional on the state of the order book.

Another objective of this work is to study the relationship between the price volatility, as defined on a macro timescale, and micro characteristics (arrival intensities) of the order flow in this model by studying the diffusion limit of the resulting price process. For example, the authors show that the volatility of the macro price process can be expressed:

$$\sigma^2 = \frac{\delta^2 \lambda}{D(f)}$$

where δ is the tick size, λ is the intensity of orders arrival and $D(f)$ is some measure of market depth.

This stylized model is an example of what can be done to assess future prices movements at the high-frequency timescale, based on the current state of the order book. It is further developed in the work [22], where the authors can apply their results to wide class of stochastic models proposed for order book dynamics, including models based on Poisson point processes, self-exciting point processes.

Chapter 4

Numerical methods for an optimal order execution problem

This chapter deals with numerical solutions to an impulse control problem arising from optimal portfolio liquidation with bid-ask spread and market price impact penalizing speedy execution trades. The corresponding dynamic programming (DP) equation is a quasi-variational inequality (QVI) with solvency constraint satisfied by the value function in the sense of constrained viscosity solutions. By taking advantage of the lag variable tracking the time interval between trades, we can provide an explicit backward numerical scheme for the time discretization of the DPQVI. The convergence of this discrete-time scheme is shown by viscosity solutions arguments. An optimal quantization method is used for computing the (conditional) expectations arising in this scheme. Numerical results are presented by examining the behaviour of optimal liquidation strategies, and comparative performance analysis with respect to some benchmark execution strategies. We also illustrate our optimal liquidation algorithm on real data, and observe various interesting patterns of order execution strategies. Finally, we provide some numerical tests of sensitivity with respect to the bid/ask spread and market impact parameters.

Note: this chapter is adapted from the article: [36] Guilbaud F., Mnif M. and H. Pham (2010): "Numerical methods for an optimal order execution problem", to appear in *Journal of Computational Finance*.

4.1 Introduction

Portfolios managers define "implementation shortfall" as the difference in performance between a theoretical trading strategy and the implemented portfolio. In a theoretical strategy, the investor observes price displayed by the market and assumes that trades will actually be executed at this price. Implementation shortfall measures the distance between the realized transaction price and the pre-trade decision price. Indeed, the investor has to face several adverse effects when executing a trading strategy, usually referred to as trading costs. Let us describe the three main components of these illiquidity effects: the bid/ask spread, the broker's fees and the market impact. The best bid (resp. best ask) price is the best offer to buy (resp. to sell) the asset, and the bid/ask spread is the difference (always positive in the continuous trading session) between the best ask price and best bid price. The broker's fees are the amount paid to the broker for executing the order. The market impact refers to the following phenomenon: any buy or sell market order passed by an investor induces an adverse market reaction that will penalize quoted price from the investor point of view.

Market impact is a key factor when executing large orders since price impact may noticeably affect a trading strategy. On April 29, 2010, Reuters agency reports that Citadel Investment Group sold 170M shares of the E*Trade stock, and raised about 301M$: this operation led to a price fall of 7.1%. These example explain why measurement and efficient management of market impact is a key issue for financial institutions, and the research of low-touch trading strategies has found a great interest among academics.

Most of market places and brokers offer several common tools to reduce market impact. We can cite as an example the simple time slicing (we will refer to this example later as the *uniform strategy*): a large order is split up in multiple *children* orders of the same size, and these children orders are sent to the market at regular time intervals. Brokers also propose more sophisticated tools as *smart order routing* (SOR) or volume weighted average price (VWAP) based algorithmic strategies. Indeed, one basic observation is that market impact can be reduced by splitting up a large order into several children orders. Then the investor has to face the following trade-off: if he chooses to trade immediately, he will penalize his performance due to market impact; if he trades gradually, he is exposed to price variation on the period of the operation. Our goal in this article is to provide a numerical method to find optimal schedule and associated quantities for the children orders.

Recently, there has been considerable interest for this problem in the academic literature. The seminal papers [10] and [3] first provided a framework for managing market impact in a discrete-time model. The optimality is determined according to a mean-variance criterion, and this leads to a static strategy, in the sense that it is independent of the stock price. Models of market impact based on stylized order book dynamics were proposed in [55], [64] and [31]. There also has been several optimal control approaches to the order execution problem, using a penalizing function to model price impact: the papers [63] and

[28] assume continuous-time trading, and use an Hamilton-Jacobi-Bellman approach for the mean-variance criterion, while [39], [51], and [44] consider real trading taking place in discrete-time by using an impulse control approach. This last approach combines the advantages of realistic modelling of portfolio liquidation and the tractability of continuous-time stochastic calculus. In these papers, the optimal liquidation strategies are price-dependent in contrast with static strategies.

In this article, we adopt the model investigated in [44]. Let us describe the main features of this model. The stock price process is assumed to follow a geometrical Brownian motion. The price impact is modelled via a nonlinear transaction costs function, that depends both on the quantity traded, and on a lag variable θ tracking the time spent since the investor's last trade. This lag variable will penalize rapid execution trades, and ensures in particular that trading times are strictly increasing, which is consistent with market practice in limit order books. In this context, we consider the problem of an investor seeking to unwind an initial position in stock shares over a finite horizon. Risk aversion of the investor is modelled through a utility function, and we use an impulse control approach for the optimal order execution problem, which consists in maximizing the expected utility from terminal liquidation wealth, under a natural economic solvency constraint involving the liquidation value of portfolio. The theoretical part of this impulse control problem is studied in [44], and the solution is characterized through dynamic programming by means of a quasi-variational inequality (QVI) satisfied by the value function in the (constrained) viscosity sense. The aim of this paper is to solve numerically this optimal order execution problem. There are actually few papers dealing with a complete numerical treatment of impulse control problems, see [19], [52], or [20]. In these papers, the domain has a simple shape, typically rectangular, and a finite-difference method is used. In contrast, our domain is rather complex due to the solvency constraint naturally imposed by the liquidation value under market impact, and we propose a suitable probabilistic numerical method for solving the associated impulse control problem. Our main contributions are the following:

- We provide a numerical scheme for the QVI associated to the impulse control problem and prove that this method is monotone, consistent and stable, hence converges to the viscosity solution of the QVI. For this purpose, we adapt a proof from [8].

- We take advantage of the lag variable θ to provide an explicit backward scheme and then simplify the computation of the solution. This contrasts with the classical approach by iterative sequence of optimal stopping problems, see e.g. [19].

- We provide the detailed computational probabilistic algorithm with an optimal quantization method for the approximation of conditional expectations arising in the backward scheme.

- We provide several numerical tests and statistics, both on simulated and real data,

and compare the optimal strategy to a benchmark of two other strategies: the uniform strategy and the naive one consisting in the liquidation of all shares in one block at the terminal date. We also provide some sensitivity numerical analysis with respect to the bid/ask spread and market impact parameters.

This paper is organized as follows: Section 2 recalls the problem formulation and main properties of the model, in particular the PDE characterization of the impulse control problem by means of constrained viscosity solutions to the QVI, as stated in [44]. Section 3 is devoted to the time discretization and the proof of convergence of the numerical scheme. Section 4 provides the numerical algorithm and numerical methods to solve the DPQVI. We also address the convergence of the numerical scheme when approximating the exact expectation by the quantized expectation, discuss the complexity of the algorithm, and compare with the finite-difference scheme methods. Section 5 presents the results obtained with our implementation, both on simulated and historical data.

4.2 Problem formulation

4.2.1 The model of portfolio liquidation

We consider a financial market where an investor has to liquidate an initial position of $y > 0$ shares of risky asset by time T. He faces the following risk/cost tradeoff: if he trades rapidly, this results in higher costs due to market impact; if he splits the order into several smaller blocks, he is exposed to the risk of price depreciation during the trading horizon. We adopt the recent continuous-time framework of [44], who proposed a modeling where trading takes place at discrete random times through an impulse control formulation, and with a temporary price impact depending on the time interval between trades, and including a bid-ask spread.

Let us recall the details of the model. We set a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with a filtration $\mathbb{F} = (\mathcal{F}_t)_{0 \leq t \leq T}$ supporting a one-dimensional Brownian motion W on a finite horizon $[0, T]$, $T < \infty$. We denote by P_t the market price of the risky asset, by X_t the cash holdings, by Y_t the number of stock shares held by the investor at time t and by Θ_t the time interval between t and the last trade before t.

Trading strategies. We assume that the investor can only trade at discrete time on $[0, T]$. This is modelled through an impulse control strategy $\alpha = (\tau_n, \zeta_n)_{n \geq 1}$ where $\tau_1 \leq \ldots \tau_n \leq \ldots \leq T$ are stopping times representing the trading times and ζ_n, $n \geq 1$, are $\mathcal{F}_{\tau_n}$-measurable random variables valued in $\mathbb{R}$ and giving the quantity of stocks purchased if $\zeta_n \geq 0$ or selled if $\zeta_n < 0$ at these times. A priori, the sequence (τ_n, ζ_n) may be finite or infinite. We introduce the lag variable tracking the time interval between trades, which

evolves according to

$$\Theta_t = t - \tau_n,\ \tau_n \leq t < \tau_{n+1}, \quad \Theta_{\tau_{n+1}} = 0,\ n \geq 0. \tag{4.2.1}$$

The dynamics of the number of stock shares Y is then given by :

$$Y_t \ = \ Y_{\tau_n}, \quad \tau_n \leq t < \tau_{n+1}, \qquad Y_{\tau_{n+1}} \ = \ Y_{\tau_n} + \zeta_{n+1}, \quad n \geq 0. \tag{4.2.2}$$

Cost of illiquidity. The market price of the risky asset process follows a geometric Brownian motion:

$$dP_t \ = \ P_t(bdt + \sigma dW_t), \tag{4.2.3}$$

with constant b and $\sigma > 0$. We focus here on the temporary price impact that penalizes the price at which an investor will trade the asset. Suppose now that the investor decides at time t to trade the quantity e. If the current market price is p, and the time lag from the last order is θ, then the price he actually gets for the order e is:

$$Q(e, p, \theta) \ = \ pf(e, \theta), \tag{4.2.4}$$

where f is a temporary price impact function from $\mathbb{R} \times [0, T]$ into $\mathbb{R}_+ \cup \{\infty\}$. Actually, in the rest of the paper, we consider a function f in the form

$$f(e, \theta) \ = \ \exp\big(\lambda|\frac{e}{\theta}|^{\beta}\mathrm{sgn}(e)\big).\big(\kappa_a \mathbf{1}_{e>0} + \mathbf{1}_{e=0} + \kappa_b \mathbf{1}_{e<0}\big), \tag{4.2.5}$$

where $\beta > 0$ is the price impact exponent, $\lambda > 0$ is the temporary price impact factor, $\kappa_b < 1$, and $\kappa_a > 1$ are the bid and ask spread parameters. The impact of liquidity modelled in (4.2.4) is like a transaction cost combining nonlinearity and proportionality effects. The nonlinear costs come from the dependence of the function f on e, but also on θ. On the other hand, this transaction cost function f can be determined implicitly from the impact of a market order placed by a large trader in a limit order book, as explained in [55], [64] or [63]. Moreover, the dependence of f in θ in (4.2.5) means that rapid trading has a larger temporary price impact than slower trading. Such kind of assumption is also made in the seminal paper [3], and reflects stylized facts on limit order books. The form (4.2.5) was suggested in several empirical studies, see [50], [60], [4], and used also in [28], [44].

Remark 4.2.1 We could consider a permanent price impact, i.e. the lasting effect of large trade, in our modelling by introducing a jump in the market price P at a trading time (as in [39] or [51]), which depends on the order size and time lag from the last order size. Alternatively, one can introduce a permanent price impact in the spirit of [3], [28] or [2] by modelling the rate of return $b = (b_t)$ of the market price as a state variable process following the dynamics:

$$db_t \ = \ \rho(\theta_t)(\bar{b} - b_t)dt, \quad \tau_n \leq t < \tau_{n+1}, \qquad b_{\tau_{n+1}} \ = \ b_{\tau_{n+1}^-} + g\big(\frac{\zeta_{n+1}}{\tau_{n+1} - \tau_n}\big),$$

where g is the permanent price impact function, e.g. in the linear form $g(\eta) = \kappa_p \eta$, with a factor $\kappa_p > 0$, and ρ is an increasing positive resilience function, e.g. in the linear form $\rho(\theta) = \kappa_r \theta$, $\kappa_r > 0$, measuring the reversion rate of the return process to a reference constant value $\bar{b}$.

Cash holdings. We assume a zero risk-free return, so that the cash holdings are constant between two trading times:

$$X_t \;=\; X_{\tau_n}, \quad \tau_n \le t < \tau_{n+1}, \;\; n \ge 0. \tag{4.2.6}$$

When a discrete trading $\Delta Y_t = \zeta_{n+1}$ occurs at time $t = \tau_{n+1}$, this results in a variation of the cash amount given by $\Delta X_t := X_t - X_{t^-} = -\Delta Y_t . Q(\Delta Y_t, P_t, \Theta_{t^-})$ due to the illiquidity effects. Moreover, there is a fixed cost $\varepsilon \ge 0$ to be paid at each transaction. In other words, we have

$$X_{\tau_{n+1}} \;=\; X_{\tau_{n+1}^-} - \zeta_{n+1} P_{\tau_{n+1}} f(\zeta_{n+1}, \tau_{n+1} - \tau_n) - \varepsilon, \quad n \ge 0. \tag{4.2.7}$$

Remark 4.2.2 Notice that since $f(e, 0) = 0$ if $e < 0$ and $f(e, 0) = \infty$ if $e > 0$, an immediate sale does not increase the cash holdings, i.e. $X_{\tau_{n+1}} = X_{\tau_{n+1}^-} = X_{\tau_n}$, while an immediate purchase leads to a bankruptcy i.e. $X_{\tau_{n+1}} = -\infty$.

Liquidation value and solvency constraint. The solvency constraint is a key issue in portfolio choice problem. The point is to define in an economically meaningful way what is the portfolio value of a position in cash and stocks. In our context, we first impose a no-short selling constraint on the trading strategies, i.e.

$$Y_t \;\ge\; 0, \quad 0 \le t \le T.$$

Next, we introduce the liquidation function $L_\varepsilon(x, y, p, \theta)$ representing the value that an investor would obtain by liquidating immediately his stock position y by a single block trade, when the pre-trade price is p and the time lag from the last order is θ. It is defined on $\mathbb{R} \times \mathbb{R}_+ \times (0, \infty) \times [0, T]$ by

$$L_\varepsilon(x, y, p, \theta) \;=\; \max[x, x + ypf(-y, \theta) - \varepsilon].$$

The interpretation of this liquidation function is the following. Due to the presence of the transaction fee at each trading, it may be advantageous for the investor not to liquidate his position in stock shares (which would give him $x + ypf(-y, \theta) - \varepsilon$), and rather bin his stock shares, by keeping only his cash amount (which would give him x). Hence, the investor chooses the best of these two possibilities, which induces a liquidation value $L_\varepsilon(z, \theta)$.

We thus constrain the portfolio's liquidative value to satisfy the solvency criterion:

$$L_\varepsilon(X_t, Y_t, P_t, \Theta_t) \;\ge\; 0, \quad 0 \le t \le T.$$

We then naturally introduce the solvency region:

$$\mathcal{S}_\varepsilon \;=\; \{(z,\theta)=(x,y,p,\theta)\in\mathbb{R}\times\mathbb{R}_+\times(0,\infty)\times[0,T]\;:L_\varepsilon(z,\theta)>0\}\,.$$

and we denote its boundary and its closure by

$$\partial\mathcal{S}_\varepsilon \;=\partial_y\mathcal{S}_\varepsilon\cup\partial_L\mathcal{S}_\varepsilon \quad \text{and} \quad \bar{\mathcal{S}}_\varepsilon \;=\; \mathcal{S}_\varepsilon\cup\partial\mathcal{S}_\varepsilon.$$

where

$$\begin{aligned}\partial_y\mathcal{S}_\varepsilon &= \{(z,\theta)=(x,y,p,\theta)\in\mathbb{R}\times\mathbb{R}_+\times(0,\infty)\times[0,T]:\; y=0 \text{ and } x=L_\varepsilon(z,\theta)\geq 0\}\,,\\ \partial_L\mathcal{S}_\varepsilon &= \{(z,\theta)=(x,y,p,\theta)\in\mathbb{R}\times\mathbb{R}_+\times(0,\infty)\times[0,T]:\; L_\varepsilon(z,\theta)=0\}\,.\end{aligned}$$

In the sequel, we also introduce the corner lines in $\partial\mathcal{S}_\varepsilon$:

$$D_0 \;=\; \{(0,0)\}\times(0,\infty)\times[0,T] \;=\; \partial_y\mathcal{S}_\varepsilon\cap\partial_L\mathcal{S}_\varepsilon.$$

Admissible trading strategies. Given $(t,z,\theta)\in[0,T]\times\bar{\mathcal{S}}_\varepsilon$, we say that the impulse control strategy $\alpha=(\tau_n,\zeta_n)_{n\geq 0}$ is admissible, denoted by $\alpha\in\mathcal{A}_\varepsilon(t,z,\theta)$, if $\tau_0=t-\theta$, $\tau_n\geq t$, $n\geq 1$, and the process $\{(Z_s,\Theta_s)=(X_s,Y_s,P_s,\Theta_s),t\leq s\leq T\}$ solution to (4.2.1)-(6.2.5)-(4.2.3)-(4.2.6)-(4.2.7), with an initial state $(Z_{t^-},\Theta_{t^-})=(z,\theta)$ (and the convention that $(Z_t,\Theta_t)=(z,\theta)$ if $\tau_1>t$), satisfies $(Z_s,\Theta_s)\in[0,T]\times\bar{\mathcal{S}}$ for all $s\in[t,T]$. As usual, to alleviate notations, we omit the dependence of (Z,Θ) in (t,z,θ,α), when there is no ambiguity.

Portfolio liquidation problem. We consider a utility function U from $\mathbb{R}_+$ into $\mathbb{R}$, strictly increasing, concave and w.l.o.g. $U(0)=0$, and s.t. there exists $K\geq 0$, $\gamma\in[0,1)$:

$$U(w) \;\leq\; Kw^\gamma, \quad \forall w\geq 0.$$

The problem of optimal portfolio liquidation is formulated as

$$v_\varepsilon(t,z,\theta) \;=\; \sup_{\alpha\in\mathcal{A}_\varepsilon(t,z,\theta)}\mathbb{E}\big[U_{L_\varepsilon}(Z_T,\Theta_T)\big], \quad (t,z,\theta)\in[0,T]\times\bar{\mathcal{S}}_\varepsilon, \tag{4.2.8}$$

where $U_{L_\varepsilon}(z,\theta)=U(L_\varepsilon(z,\theta))$ is the terminal liquidation utility function.

Remark 4.2.3 The function $z\to v_\varepsilon(t,z,0)$ is strictly increasing in the argument of cash holdings x, for $(z=(x,y,p),0)\in\bar{\mathcal{S}}_\varepsilon$, and fixed $t\in[0,T]$. Indeed, for $x<x'$, and $z=(x,y,p)$, $z'=(x',y,p)$, any strategy $\alpha\in\mathcal{A}_\varepsilon(t,z,\theta)$ with corresponding state process $(Z_s=(X_s,Y_s,P_s),\Theta_s)_{s\geq t}$, is also in $\mathcal{A}_\varepsilon(t,z',\theta)$, and leads to an associated state process $(Z'_s=(X_s+x'-x,Y_s,P_s),\Theta_s)_{s\geq t}$. Using the fact that the utility function is strictly increasing, we deduce that $v_\varepsilon(t,x,y,p,0)<v_\varepsilon(t,x',y,p,0)$. Moreover, the function $z\to v_\varepsilon(t,z,0)$ is nondecreasing in the argument of number of shares y. Indeed, fix $z=(x,y,p)$, and $z'=(x,y',p)$ with $y\leq y'$. Given any arbitrary $\alpha=(\tau_n,\zeta_n)_n\in\mathcal{A}_\varepsilon(t,z,0)$, consider

the strategy $\alpha' = (\tau'_n, \zeta'_n)$, starting from (x, y', p) at time t, which consists in trading again immediately at time t by selling $y' - y$ shares (which does not change the cash holdings, see Remark 4.2.2), and then follow the same strategy than α. The corresponding state process satisfies $(Z'_s, \Theta'_s) = (Z_s, \Theta_s)$ a.s. for $s \geq t$, and in particular $\alpha' \in \mathcal{A}_\varepsilon(t, z', 0)$, together with $\mathbb{E}[U_{L_\varepsilon}(Z'_T, \Theta'_T)] = \mathbb{E}[U_{L_\varepsilon}(Z_T, \Theta_T)] \leq v(t, z', \theta)$. Since α is arbitrary in $\mathcal{A}_\varepsilon(t, z, 0)$, this shows that $v(t, x, y, p, 0) \leq v(t, x, y', p, 0)$.

We recall from [44] that v_ε is in the set $\mathcal{G}([0,T] \times \bar{\mathcal{S}}_\epsilon)$ of functions satisfying the growth condition:

$$\mathcal{G}([0,T] \times \bar{\mathcal{S}}_\epsilon) \;=\; \Big\{ \varphi : [0,T] \times \bar{\mathcal{S}}_\epsilon \longrightarrow \mathbb{R} \text{ s.t. } \sup_{[0,T] \times \bar{\mathcal{S}}_\epsilon} \frac{|\varphi(t,z,\theta)|}{\big(1 + (x + yp)^\gamma\big)} < \infty \Big\}.$$

In the sequel, we shall denote by $\mathcal{G}_+([0,T] \times \bar{\mathcal{S}}_\varepsilon)$ the set of functions φ in $\mathcal{G}([0,T] \times \bar{\mathcal{S}}_\varepsilon)$ such that $\varphi(t, x, y, p, 0)$ is strictly increasing in x and nondecreasing in y.

4.2.2 PDE characterization

The dynamic programming Hamilton-Jacobi-Bellman (HJB) equation corresponding to the stochastic control problem (4.2.8) is a quasi-variational inequality written as

$$\min\big[-\frac{\partial v}{\partial t} - \mathcal{L}v \;,\; v - \mathcal{H}_\varepsilon v \big] \;=\; 0, \quad \text{on} \quad [0,T) \times \bar{\mathcal{S}}_\varepsilon, \tag{4.2.9}$$

together with the relaxed terminal condition

$$\min\big[v - U_{L_\varepsilon}, v - \mathcal{H}_\varepsilon v\big] \;=\; 0, \quad \text{on} \quad \{T\} \times \bar{\mathcal{S}}_\varepsilon. \tag{4.2.10}$$

Here, $\mathcal{L}$ is the infinitesimal generator associated to the process $(Z = (X, Y, P), \Theta)$ in a no-trading period:

$$\mathcal{L}\varphi \;=\; \frac{\partial \varphi}{\partial \theta} + bp\frac{\partial \varphi}{\partial p} + \frac{1}{2}\sigma^2 p^2 \frac{\partial^2 \varphi}{\partial p^2},$$

$\mathcal{H}_\varepsilon$ is the impulse operator defined by

$$\mathcal{H}_\varepsilon \varphi(t,z,\theta) \;=\; \sup_{e \in \mathcal{C}_\varepsilon(z,\theta)} \varphi(t, \Gamma_\varepsilon(z,\theta,e), 0), \quad (t,z,\theta) \in [0,T] \times \bar{\mathcal{S}}_\varepsilon,$$

Γ_ε is the impulse transaction function defined from $\bar{\mathcal{S}}_\varepsilon \times \mathbb{R}$ into $\mathbb{R} \times \mathbb{R} \times (0, \infty)$:

$$\Gamma_\varepsilon(z,\theta,e) \;=\; (x - epf(e,\theta) - \varepsilon, y + e, p), \;\; z = (x,y,p) \in \bar{\mathcal{S}}_\varepsilon, \; e \in \mathbb{R},$$

and $\mathcal{C}_\varepsilon(z,\theta)$ the set of admissible transactions :

$$\mathcal{C}_\varepsilon(z,\theta) \;=\; \Big\{ e \in \mathbb{R} \;:\; \Big(\Gamma_\varepsilon(z,\theta,e), 0\Big) \;\in\; \bar{\mathcal{S}}_\varepsilon \Big\}.$$

Remark 4.2.4 Fix $t \in [0,T]$. For $\theta = 0$, and $z = (x,y,p)$ s.t. $(z,0) \in \bar{\mathcal{S}}_\varepsilon$, the set of admissible transactions $\mathcal{C}_\varepsilon(z,0) = [-y,0]$ (and $\Gamma_\varepsilon(z,0,e) = (x-\varepsilon, y+e, p)$ for $e \in \mathcal{C}_\varepsilon(z,0)$) if $x \geq \varepsilon$, and is empty otherwise. Thus, $\mathcal{H}_\varepsilon w(t,z,0) = \sup_{e\in[-y,0]} w(t,x-\varepsilon,y+e,p,0)$ if $x \geq \varepsilon$, and is equal to $-\infty$ otherwise. This implies in particular that

$$\mathcal{H}_\varepsilon w(t,z,0) \quad < \quad w(t,z,0), \tag{4.2.11}$$

for any $w \in \mathcal{G}_+([0,T] \times \bar{\mathcal{S}}_\varepsilon)$, which is the case of v_ε (see Remark 4.2.3). Therefore, due to the market impact function f in (4.2.5) penalizing rapid trades, it is not optimal to trade again immediately right after some trade, i.e. the optimal trading times are strictly increasing.

A main result in [44] is to provide a unique PDE characterization of the value functions v_ε, $\varepsilon > 0$, and to prove that the sequence $(v_\varepsilon)_\varepsilon$ converges, as ε goes to zero, to the value function v_0 in the model without transaction fee, i.e. when $\varepsilon = 0$.

Theorem 4.2.1 *(1) The sequence $(v_\varepsilon)_\varepsilon$ is nonincreasing, and converges pointwise on $[0,T]\times(\bar{\mathcal{S}}_0 \setminus \partial_{L_0}\mathcal{S}_0)$ towards v_0 as ε goes to zero, with $v_\varepsilon \leq v_0$.*
(2) For any $\varepsilon > 0$, the value function v_ε is continuous on $[0,T)\times\mathcal{S}_\varepsilon$, and is the unique (in $[0,T)\times\mathcal{S}_\varepsilon$) constrained viscosity solution to (4.2.9)-(4.2.10)*, satisfying the growth condition in $\mathcal{G}([0,T]\times\bar{\mathcal{S}}_\varepsilon)$, and the boundary condition:*

$$\begin{aligned} \lim_{(t',z',\theta')\to(t,z,\theta)} v_\varepsilon(t',z',\theta') &= v_\epsilon(t,z,\theta) \\ &= U(0), \quad \forall (t,z=(0,0,p),\theta) \in [0,T]\times D_0. \end{aligned} \tag{4.2.12}$$

The rest of this paper is devoted to the numerical analysis and resolution of the QVI (4.2.9)-(4.2.10) characterizing the optimal portfolio liquidation problem with fixed transaction fee. On the other hand, this also provide an ε-approximation of the optimal portfolio liquidation problem without fixed transaction fee.

4.3 Time discretization and convergence analysis

In this section, we fix $\varepsilon > 0$, and we study time discretization of the QVI (4.2.9)-(4.2.10) characterizing the value function v_ε. For a time discretization step $h > 0$ on the interval $[0,T]$, let us consider the following approximation scheme:

$$S^h(t,z,\theta,v^h(t,z,\theta),v^h) \quad = \quad 0, \quad (t,z,\theta) \in [0,T]\times\bar{\mathcal{S}}_\varepsilon, \tag{4.3.1}$$

where $S^h : [0,T] \times \bar{\mathcal{S}}_\varepsilon \times \mathbb{R} \times \mathcal{G}_+([0,T] \times \bar{\mathcal{S}}_\varepsilon) \to \mathbb{R}$ is defined by

$$
\begin{aligned}
& S^h(t,z,\theta,r,\varphi) \qquad\qquad (4.3.2)\\
:= & \begin{cases}
\min\Big[r - \mathbb{E}\big[\varphi(t+h, Z^{0,t,z}_{t+h}, \Theta^{0,t,\theta}_{t+h})\big]\ ,\ r - \mathcal{H}_\epsilon\varphi(t,z,\theta)\Big] & \text{if } t \in [0, T-h]\\
\min\Big[r - \mathbb{E}\big[\varphi(T, Z^{0,t,z}_{T}, \Theta^{0,t,\theta}_{T})\big]\ ,\ r - \mathcal{H}_\epsilon\varphi(t,z,\theta)\Big] & \text{if } t \in (T-h, T)\\
\min\Big[r - U_{L_\epsilon}(z,\theta)\ ,\ r - \mathcal{H}_\epsilon\varphi(t,z,\theta)\Big] & \text{if } t = T.
\end{cases}
\end{aligned}
$$

Here, $(Z^{0,t,z}, \Theta^{0,t,\theta})$ denotes the state process starting from (z,θ) at time t, and without any impulse control strategy: it is given by

$$
\Big(Z^{0,t,z}_s, \Theta^{0,t,\theta}_s\Big) \ = \ (x, y, P^{t,p}_s, \theta + s - t), \quad s \geq t,
$$

with $P^{t,p}$ the solution to (4.2.3) starting from p at time t. Notice that (4.3.1) is formulated as a backward scheme for the solution v^h through:

$$
\begin{aligned}
v^h(T,z,\theta) &= \max\big[U_{L_\epsilon}(z,\theta)\ ,\ \mathcal{H}_\epsilon v^h(T,z,\theta)\big], \qquad (4.3.3)\\
v^h(t,z,\theta) &= \max\Big[\mathbb{E}\big[v^h(t+h, Z^{0,t,z}_{t+h}, \theta+h)\big]\ , \mathcal{H}_\epsilon v^h(t,z,\theta)\Big], \quad 0 \leq t \leq T-h, \qquad (4.3.4)
\end{aligned}
$$

and $v^h(t,z,\theta) = v^h(T-h,z,\theta)$ for $T-h < t < T$. This approximation scheme seems a priori implicit due to the nonlocal obstacle term $\mathcal{H}_\varepsilon$. This is typically the case in impulse control problems, and the usual way (see e.g. [19], [52]) to circumvent this problem is to iterate the scheme by considering a sequence of optimal stopping problems:

$$
\begin{aligned}
v^{h,n+1}(T,z,\theta) &= \max\big[U_{L_\epsilon}(z,\theta)\ ,\ \mathcal{H}_\epsilon v^{h,n}(T,z,\theta)\big],\\
v^{h,n+1}(t,z,\theta) &= \max\Big[\mathbb{E}\big[v^{h,n+1}(t+h, Z^{0,t,z}_{t+h}, \theta+h)\big]\ , \mathcal{H}_\epsilon v^{h,n}(t,z,\theta)\Big],
\end{aligned}
$$

starting from $v^{h,0} = \mathbb{E}[U_{L_\varepsilon}(Z^{0,t,z}_T, \Theta^{0,t,\theta}_T)]$. Here, we shall make the numerical scheme (4.3.1) explicit, i.e. without iteration, by taking effect of the state variable θ in our model. Recall indeed from Remark 4.2.4 that it is not optimal to trade again immediately right after some trade. Thus, for $v^h \in \mathcal{G}_+([0,T] \times \bar{\mathcal{S}}_\varepsilon)$, and any $(z',0) \in \bar{\mathcal{S}}_\varepsilon$, we have from (4.2.11) and (4.3.3)-(4.3.4):

$$
\begin{aligned}
v^h(T,z',0) &= U_{L_\epsilon}(z',0)\\
v^h(t,z',0) &= \mathbb{E}\big[v^h(t+h, Z^{0,t,z'}_{t+h}, h)\big].
\end{aligned}
$$

Therefore, by using again the definition of $\mathcal{H}_\varepsilon$ in the relations (4.3.3)-(4.3.4), we see that the scheme (4.3.1) is written equivalently as an explicit backward scheme:

$$
\begin{aligned}
v^h(T,z,\theta) &= \max\big[U_{L_\epsilon}(z,\theta)\ ,\ \mathcal{H}_\epsilon U_{L_\epsilon}(z,\theta)\big], \qquad (4.3.5)\\
v^h(t,z,\theta) &= \max\Big[\mathbb{E}\big[v^h(t+h, Z^{0,t,z}_{t+h}, \theta+h)\big]\ ,\ \sup_{e \in \mathcal{C}_\varepsilon(z,\theta)} \mathbb{E}\big[v^h(t+h, Z^{0,t,z^e_\theta}_{t+h}, h)\big]\Big], \qquad (4.3.6)
\end{aligned}
$$

for $0 \leq t \leq T-h$, and $v^h(t,z,\theta) = v^h(T-h,z,\theta)$ for $T-h < t < T$, where we denote $z^e_\theta = \Gamma_\varepsilon(z,\theta,e)$ in (4.3.6) to alleviate notations. Notice that at this stage, this approximation scheme is not yet fully implementable since it requires an approximation method for the expectations arising in (4.3.6). This is the concern of the next section.

We focus now on the convergence (when h goes to zero) of the solution v^h to (4.3.1) towards the value function v_ε solution to (4.2.9)-(4.2.10). Following [8], we have to show that the scheme S^h in (4.3.2) satisfies monotonicity, stability and consistency properties. As usual, the monotonicity property follows directly from the definition (4.3.2) of the scheme.

Proposition 4.3.1 *(Monotonicity)*
For all $h > 0$, $(t,z,\theta) \in [0,T] \times \bar{\mathcal{S}}_\epsilon$, $r \in \mathbb{R}$, and φ, $\psi \in \mathcal{G}_+([0,T] \times \bar{\mathcal{S}}_\epsilon)$ s.t. $\varphi \leq \psi$, we have

$$S^h(t,z,\theta,r,\varphi) \quad \geq \quad S^h(t,z,\theta,r,\psi).$$

We next prove the stability property.

Proposition 4.3.2 *(Stability)*
For all $h > 0$, there exists a unique solution $v^h \in \mathcal{G}_+([0,T] \times \bar{\mathcal{S}}_\varepsilon)$ to (4.3.1), and the sequence $(v^h)_h$ is uniformly bounded in $\mathcal{G}([0,T] \times \bar{\mathcal{S}}_\varepsilon)$: there exists $w \in \mathcal{G}([0,T] \times \bar{\mathcal{S}}_\varepsilon)$ s.t. $|v^h| \leq |w|$ for all $h > 0$.

Proof. The uniqueness of a solution $\in \mathcal{G}_+([0,T] \times \bar{\mathcal{S}}_\varepsilon)$ to (4.3.1) follows from the explicit backward scheme (4.3.5)-(4.3.6). For $t \in [0,T]$, denote by $N_{t,h}$ the integer part of $(T-t)/h$, and $\mathbf{T}_{t,h} = \{t_k = t + kh, k = 0, \ldots, N_{t,h}\}$ the partition of the interval $[t,T]$ with time step h. For $(t,z,\theta) \in [0,T] \times \bar{\mathcal{S}}_\varepsilon$, we denote by $\mathcal{A}^h_\varepsilon(t,z,\theta)$ the subset of elements $\alpha = (\tau_n, \zeta_n)_n$ in $\mathcal{A}_\varepsilon(t,z,\theta)$ such that the trading times τ_n are valued in $\mathbf{T}_{t,h}$. Let us then consider the impulse control problem

$$v^h(t,z,\theta) \quad = \quad \sup_{\alpha \in \mathcal{A}^h_\varepsilon(t,z,\theta)} \mathbb{E}\big[U_{L_\varepsilon}(Z^\varepsilon_T, \Theta_T)\big], \quad (t,z,\theta) \in [0,T] \times \bar{\mathcal{S}}_\varepsilon. \tag{4.3.7}$$

It is clear from the representation (4.3.7) that for all $h > 0$, $0 \leq v^h \leq v_\varepsilon$, which shows that the sequence $(v^h)_h$ is uniformly bounded in $\mathcal{G}([0,T] \times \bar{\mathcal{S}}_\varepsilon)$. Moreover, similarly as for v_ε, and by the same arguments as in Remark 4.2.3, we see that $v^h(t,z,0)$ is strictly increasing in x and nondecreasing in y for $(z,0) = (x,y,p,0) \in \bar{\mathcal{S}}_\varepsilon$. Finally, we observe that the numerical scheme (4.3.1) is the dynamic programming equation satisfied by the value function v^h. This proves the required stability result. □

We now move on the consistency property.

Proposition 4.3.3 *(Consistency)*

(i) For all $(t,z,\theta) \in [0,T) \times \bar{\mathcal{S}}_\epsilon$ *and* $\phi \in C^{1,2}([0,T) \times \bar{\mathcal{S}}_\epsilon)$*, we have*

$$\begin{aligned} \limsup_{\substack{(h,t',z',\theta') \to (0,t,z,\theta) \\ (t',z',\theta') \in [0,T) \times \mathcal{S}_\epsilon}} & \min\left\{ \frac{\phi(t',z',\theta') - \mathbb{E}\Big[\phi(t'+h, Z^{0,t',z'}_{t'+h}, \Theta^{0,t',\theta'}_{t'+h})\Big]}{h}, \Big(\phi - \mathcal{H}_\epsilon \phi\Big)(t',z',\theta') \right\} \\ \leq & \min\Big\{ \Big(-\frac{\partial \phi}{\partial t} - \mathcal{L}\phi\Big)(t,z,\theta), \Big(\phi - \mathcal{H}_\epsilon \phi\Big)(t,z,\theta) \Big\} \end{aligned} \tag{4.3.8}$$

and

$$\begin{aligned} \liminf_{\substack{(h,t',z',\theta') \to (0,t,z,\theta) \\ (t',z',\theta') \in [0,T) \times \mathcal{S}_\epsilon}} & \min\left\{ \frac{\phi(t',z',\theta') - \mathbb{E}\Big[\phi(t'+h, Z^{0,t',z'}_{t'+h}, \Theta^{0,t',z'}_{t'+h})\Big]}{h}, \Big(\phi - \mathcal{H}_\epsilon \phi\Big)(t',z',\theta') \right\} \\ \geq & \min\Big\{ \Big(-\frac{\partial \phi}{\partial t} - \mathcal{L}\phi\Big)(t,z,\theta), \Big(\phi - \mathcal{H}_\epsilon \phi\Big)(t,z,\theta) \Big\} \end{aligned} \tag{4.3.9}$$

(ii) For all $(z,\theta) \in \bar{\mathcal{S}}_\epsilon$ *and* $\phi \in C^{1,2}([0,T] \times \bar{\mathcal{S}}_\epsilon)$*, we have*

$$\begin{aligned} \limsup_{\substack{(t',z',\theta') \to (T,z,\theta) \\ (t',z',\theta') \in [0,T) \times \mathcal{S}_\epsilon}} & \min\Big\{ \phi(t',z',\theta') - U_{L_\epsilon}(z',\theta'), \Big(\phi - \mathcal{H}_\epsilon \phi\Big)(t',z',\theta') \Big\} \\ \leq & \min\Big\{ \phi(T,z,\theta) - U_{L_\epsilon}(z,\theta), \Big(\phi - \mathcal{H}_\epsilon \phi\Big)(T,z,\theta) \Big\} \end{aligned} \tag{4.3.10}$$

and

$$\begin{aligned} \liminf_{\substack{(t',z',\theta') \to (T,z,\theta) \\ (t',z',\theta') \in [0,T) \times \mathcal{S}_\epsilon}} & \min\Big\{ \phi(t',z',\theta') - U_{L_\epsilon}(z',\theta'), \Big(\phi - \mathcal{H}_\epsilon \phi\Big)(t',z',\theta') \Big\} \\ \geq & \min\Big\{ \Big(\phi(T,z,\theta) - U_{L_\epsilon}(z,\theta)\Big), \Big(\phi - \mathcal{H}_\epsilon \phi\Big)(T,z,\theta) \Big\} \end{aligned} \tag{4.3.11}$$

Proof. The arguments are standard, and can be adapted e.g. from [19] or [20]. We sketch the proof, and only show the inequality (4.3.8) since the other ones are derived similarly. Fix $t \in [0,T)$. Since the minimum of two upper-semicontinous (usc) functions is also usc

and using the caracterization of usc functions, we have

$$
\begin{aligned}
&\limsup_{\substack{(h,t',z',\theta')\to(0,t,z,\theta)\\(t',z',\theta')\in[0,T)\times\mathcal{S}_\epsilon}} \min\Big\{\Big(\phi-\mathcal{H}_\epsilon\phi\Big)(t',z',\theta'),\frac{\phi(t',z',\theta')-\mathbb{E}\Big[\phi(t'+h,Z^{0,t',z'}_{t'+h},\Theta^{0,t',\theta'}_{t'+h})\Big]}{h}\Big\}\\
\leq\ &\limsup_{\substack{(h,t',z',\theta')\to(0,t,z,\theta)\\(t',z',\theta')\in[0,T)\times\mathcal{S}_\epsilon}} \min\Big\{\limsup_{\substack{(h,t'',z'',\theta'')\to(0,t',z',\theta')\\(t'',z'',\theta'')\in[0,T)\times\mathcal{S}_\epsilon}}\Big(\phi-\mathcal{H}_\epsilon\phi\Big)(t'',z'',\theta''),\\
&\qquad\limsup_{\substack{(h,t'',z'',\theta'')\to(0,t',z',\theta')\\(t'',z'',\theta'')\in[0,T)\times\mathcal{S}_\epsilon}}\frac{\phi(t'',z'',\theta'')-\mathbb{E}\Big[\phi(t''+h,Z^{0,t'',z''}_{t''+h},\Theta^{0,t'',\theta''}_{t''+h})\Big]}{h}\Big\}\\
\leq\ &\min\Big\{\limsup_{\substack{(h,t',z',\theta')\to(0,t,z,\theta)\\(t',z',\theta')\in[0,T)\times\mathcal{S}_\epsilon}}\Big(\phi-\mathcal{H}_\epsilon\phi\Big)(t',z',\theta'),\\
&\qquad\limsup_{\substack{(h,t',z',\theta')\to(0,t,z,\theta)\\(t',z',\theta')\in[0,T)\times\mathcal{S}_\epsilon}}\frac{\phi(t',z',\theta')-\mathbb{E}\Big[\phi(t'+h,Z^{0,t',z'}_{t'+h},\Theta^{0,t',\theta'}_{t'+h})\Big]}{h}\Big\}\\
\leq\ &\min\Big\{\phi(t,z,\theta)-\mathcal{H}_\epsilon\phi(t,z,\theta)\\
&\qquad\limsup_{\substack{(h,t',z',\theta')\to(0,t,z,\theta)\\(t',z',\theta')\in[0,T)\times\mathcal{S}_\epsilon}}\frac{\phi(t',z',\theta')-\mathbb{E}\Big[\phi(t'+h,Z^{0,t',z'}_{t'+h},\Theta^{0,t',\theta'}_{t'+h})\Big]}{h}\Big\},
\end{aligned}
\tag{4.3.12}
$$

where the last inequality follows from the continuity of ϕ and the lower semicontinuity of $\mathcal{H}_\varepsilon$. Moreover, by Itô's formula applied to $\phi(s,Z^{0,t',z'}_s,\Theta^{0,t',\theta'}_s)$, and standard arguments of localization to remove in expectation the stochastic integral, we get

$$
\limsup_{\substack{(h,t',z',\theta')\to(0,t,z,\theta)\\(t',z',\theta')\in[0,T)\times\mathcal{S}_\epsilon}}\frac{\phi(t',z',\theta')-\mathbb{E}\Big[\phi(t'+h,Z^{0,t',z'}_{t'+h},\Theta^{0,t',\theta'}_{t'+h})\Big]}{h} \;=\; -\Big(\frac{\partial\phi}{\partial t}+\mathcal{L}\phi\Big)(t,z,\theta)
$$

Substituting into (4.3.12), we obtain the desired inequality (4.3.8). □

Since the numerical scheme (4.3.1) is monotone, stable and consistent, we can follow the viscosity solutions arguments as in [8] to prove the convergence of v^h to v_ε, by relying on the PDE characterization of v_ε in Theorem 4.2.1 (2), and the strong comparison principle for (4.2.9)-(4.2.10) proven in [44].

Theorem 4.3.1 *(Convergence) The solution v^h of the numerical scheme* (4.3.1) *converges locally uniformly to v_ε on $[0,T)\times\mathcal{S}_\epsilon$.*

Proof. Let $\overline{v_\epsilon}$ and $\underline{v_\epsilon}$ be defined on $[0,T]\times\bar{\mathcal{S}}_\varepsilon$ by

$$\begin{aligned}\overline{v_\epsilon}(t,z,\theta) &= \limsup_{\substack{(h,t',z',\theta')\to(0,t,z,\theta)\\(t',z',\theta')\in[0,T)\times\mathcal{S}_\epsilon}} v^h(t',z',\theta')\\ \underline{v_\epsilon}(t,z,\theta) &= \liminf_{\substack{(h,t',z',\theta')\to(0,t,z,\theta)\\(t',z',\theta')\in[0,T)\times\mathcal{S}_\epsilon}} v^h(t',z',\theta')\end{aligned}$$

We first see that $\overline{v_\epsilon}$ and $\underline{v_\epsilon}$ are respectively viscosity subsolution and supersolution of (4.2.9)-(4.2.10). These viscosity properties follow indeed, by standard arguments as in [8] (see also [19] or [20] for impulse control problems), from the monotonicity, stability and consistency properties. Details can be obtained upon request to the authors. Moreover, from (4.3.7), we have the inequality: $U(0)\le v^h\le v_\varepsilon$, which implies by (4.2.12):

$$\liminf_{\substack{(t',z',\theta')\to(t,z,\theta)\\(t',z',\theta')\in[0,T)\times\mathcal{S}_\epsilon}} \underline{v_\epsilon}(t',z',\theta') \;=\; U(0) \;=\; \overline{v_\epsilon}(t,z,\theta),\;\; \forall\,(t,z,\theta)\in[0,T]\times D_0 \quad (4.3.13)$$

Thus, by using the strong comparison principle for (4.2.9)-(4.2.10) stated in Theorem 5.2 [44], we deduce that $\overline{v_\epsilon}\le\underline{v_\epsilon}$ on $[0,T]\times\mathcal{S}_\epsilon$ and so $\overline{v_\epsilon}=\underline{v_\epsilon}=v_\epsilon$ on $[0,T]\times\mathcal{S}_\epsilon$. This proves the required convergence result. □

4.4 Numerical Algorithm

Let us consider a time step $h=T/m$, $m\in\mathbb{N}\setminus\{0\}$, and denote by $\mathbb{T}_m=\{t_i=ih, i=0,\dots,m\}$ the regular grid over the interval $[0,T]$. We recall from the previous section that the time discretization of step h for the QVI (4.2.9)-(4.2.10) leads to the convergent explicit backward scheme:

$$v^h(t_m,z,\theta) = \begin{cases} U_{L_\varepsilon}(z,\theta) & \text{if } \theta=0\\ \max\Big[U_{L_\epsilon}(z,\theta)\,, \sup_{e\in\mathcal{C}_\varepsilon(z,\theta)} v^h(t_m,\Gamma_\varepsilon(z,\theta,e),0)\Big], & \text{if } \theta>0, \end{cases} \quad (4.4.1)$$

$$v^h(t_i,z,\theta) = \begin{cases} \mathbb{E}\big[v^h(t_{i+1},Z^{0,t_i,z}_{t_{i+1}},\theta+h)\big] & \text{if } \theta=0\\ \max\Big[\mathbb{E}\big[v^h(t_{i+1},Z^{0,t_i,z}_{t_{i+1}},\theta+h)\big]\,, \sup_{e\in\mathcal{C}_\varepsilon(z,\theta)} v^h(t_i,\Gamma_\varepsilon(z,\theta,e),0)\big]\Big], & \text{if } \theta>0 \end{cases} \quad (4.4.2)$$

for $i=0,\dots,m-1$, $(z=(x,y,p),\theta)\in\bar{\mathcal{S}}_\varepsilon$. Recall that the variable θ represents the time lag between the current time t and the last trade. Thus, it suffices to consider at each time step t_i of $\mathbb{T}_m$, a discretization for θ valued in the time grid

$$\mathbb{T}_i \;=\; \big\{\theta_j=jh,\;\; j=0,\dots,i\big\},\quad i=0,\dots,m.$$

On the other hand, the above scheme involves nonlocal terms in the variable z for the solution v^h in relation with the supremum over $e \in \mathcal{C}_\varepsilon(z,\theta)$ and the expectations in (4.4.1)-(4.4.2), and thus the practical implementation requires a discretization of the set of admissible transactions $\mathcal{C}_\varepsilon(z,\theta)$ and a computational approximation for the above expectations. Moreover, since the state space $\bar{\mathcal{S}}_\varepsilon$ is unbounded, we also need to localize the domain on which computations are done. For any $\theta_j \in \mathbb{T}_i$, let us denote by

$$\begin{aligned} \mathcal{Z}^j &= \big\{z=(x,y,p)\in\mathbb{R}\times\mathbb{R}_+\times\mathbb{R}_+ : (z,\theta_j)\in\bar{\mathcal{S}}_\varepsilon\big\}, \\ \mathcal{Z}^j_{loc} &= \mathcal{Z}^j \cap \big([x_{min},x_{max}]\times[0,y_{max}]\times[0,p_{max}]\big), \end{aligned}$$

where $x_{min} < x_{max}$ in $\mathbb{R}$, $0 < y_{max}$, $0 < p_{max}$ are fixed constants.

Let us first discretize the set of admissible transactions $\mathcal{C}_\varepsilon(z,\theta_j)$ over which the supremum in (4.4.2) is taken, for any $\theta_j \in \mathbb{T}_i$, $z \in \mathcal{Z}^j_{loc}$. Recall from [44] that $\mathcal{C}_\varepsilon(z,\theta_j)$ is compact in the form $[\underline{e}(z,\theta_j), \bar{e}(z,\theta_j)]$. We then consider the discrete set of admissible transactions of size M:

$$\mathcal{C}_\varepsilon^{M,loc}(z,\theta_j) = \Big\{e = \underline{e}(z,\theta_j) + \frac{i}{M}(\bar{e}(z,\theta_j) - \underline{e}(z,\theta_j,e)), i=0,\dots,M : \Gamma_\varepsilon(z,\theta_j)\in\mathcal{Z}^0_{loc}\Big\},$$

and define the associated discrete impulse operator:

$$\mathcal{H}_\varepsilon^{M,loc}v^h(t_i,z,\theta_j) = \sup_{e\in\mathcal{C}_\varepsilon^{M,loc}(z,\theta_j)} v^h(t_i,\Gamma_\varepsilon(z,\theta_j,e),0).$$

Optimal quantization method and truncation. Let us now describe the numerical procedure for computing the expectations arising in (4.4.2). Recalling that $Z^{0,t,z} = (x,y,P^{t,p})$, this involves only the expectation with respect to the price process, assumed here to follow a Black-Scholes model (4.2.3). We shall then use an optimal quantization for the standard normal random variable $\mathcal{U}$, which consists in approximating the distribution of $\mathcal{U}$ by the discrete law of a random variable $\hat{\mathcal{U}}$ of support $(u_k)_{1\le k\le N} \in \mathbb{R}^N$, and defined as the projection of $\mathcal{U}$ on the grid $(u_k)_{1\le k\le N}$ according to the closest neighbour. The grid $(u_k)_{1\le k\le N}$ is optimized in order to minimize the distorsion error, i.e. the quadratic norm between $\mathcal{U}$ and $\hat{\mathcal{U}}$. This optimal grid and the associated weights $(\pi_k)_{1\le k\le N}$ are downloaded from the website: "http://www.quantize.maths-fi.com/downloads". We refer to the survey article [58] for more details on the theoretical and computational aspects of optimal quantization methods. From (4.4.2), we have to compute at any time step $t_i \in \mathbb{T}_m$, and for any $\theta_j \in \mathbb{T}_i$, $z=(x,y,p)\in\mathcal{Z}^j_{loc}$, expectations in the form:

$$\mathbb{E}\big[v^h(t_i+h, Z^{0,t_i,z}_{t_i+h}, \theta_j+h)\big] = \mathbb{E}\big[v^h(t_i+h, x, y, p\exp\big(\bar{b}h+\sigma\sqrt{h}\,\mathcal{U}\big), \theta_j+h)\big].$$

where we set $\bar{b} = (b - \frac{\sigma^2}{2})$. The optimal quantization method consists then in approximating

the above exact expectation by the discrete expectation operator:

$$\begin{aligned}\mathcal{E}^N\big[v^h(t_i+h, Z^{0,t_i,z}_{t_i+h}, \theta_j+h)\big] &:= \sum_{k=1}^N \pi_k\, v^h(t_i+h, x, y, p\exp\big(\bar{b}h+\sigma\sqrt{h}u_k\big), \theta_j+h)\\ &= \mathbb{E}\big[v^h(t_i+h, \hat{Z}^{0,t_i,z}_{t_i+h}, \theta_j+h)\big],\end{aligned}$$

where $\hat{Z}^{0,t_i,z}_{t_i+h} = (x, y, \hat{P}^{t_i,p}_{t_i+h})$, and $\hat{P}^{t_i,p}_{t_i+h} = p\exp\big(\bar{b}h+\sigma\sqrt{h}\,\hat{\mathcal{U}}\big)$ is the discrete random variable valued in $\hat{P}^{t_i,p,u_k}_{t_i+h} = p\exp\big(\bar{b}h+\sigma\sqrt{h}\;u_k\big)$, with weights π_k, $k=1,\dots,N$. Actually, since for $0\le p\le p_{max}$, the discrete positive random variable $\hat{P}^{t_i,p}_{t_i+h}$ can take values above p_{max}, we truncate to the nearest neighbour of p_{max}, and consider the approximate expectation operator:

$$\begin{aligned}\mathcal{E}^N_{loc}\big[v^h(t_i+h, Z^{0,t_i,z}_{t_i+h}, \theta_j+h)\big] &:= \sum_{k=1}^N \pi_k\, v^h(t_i+h, x, y, \mathrm{Proj}_{[0,p_{max}]}(P^{t_i,p,u_k}_{t_i+h}), \theta_j+h)\\ &= \mathbb{E}\big[v^h(t_i+h, x, y, \hat{P}^{loc,t_i,p}_{t_i+h}, \theta_j+h)\big], \qquad (4.4.3)\end{aligned}$$

where $\mathrm{Proj}_{[0,p_{max}]}(p) = p1_{p\le p_{max}} + p_{max}1_{p>p_{max}}$ for $p\ge 0$, and $\hat{P}^{loc,t_i,p}_{t_i+h} = \mathrm{Proj}_{[0,p_{max}]}(\hat{P}^{t_i,p}_{t_i+h})$.

We may then rewrite the actual numerical scheme used as:

$$S^{h,M,N}_{loc}(t_i, z, \theta_j, v^h(t,z,\theta_j), v^h) = 0, \quad t_i\in\mathbb{T}_m, \theta_j\in\mathbb{T}_i, z\in\mathcal{Z}^j_{loc}, \qquad (4.4.4)$$

for $i=0,\dots,m$, $j=0,\dots,i$, where $S^{h,M,N}_{loc}$ is defined by

$$\begin{aligned}&S^{h,M,N}_{loc}(t_i, z, \theta_j, r, \varphi) \qquad (4.4.5)\\ :=\ &\begin{cases}\min\Big[r-\mathcal{E}^N_{loc}\big[\varphi(t_i+h, Z^{0,t_i,z}_{t_i+h}, \theta_j+h)\big]\ ,\ r-\mathcal{H}^{M,loc}_\epsilon\varphi(t_i,z,\theta)\Big] & \text{for } i=0,\dots,m-1\\ \min\Big[r-U_{L_\epsilon}(z,\theta_j)\ ,\ r-\mathcal{H}^{M,loc}_\epsilon\varphi(T,z,\theta_j)\Big] & \text{for } i=m.\end{cases}\end{aligned}$$

Let us now address the convergence proofs of this computational scheme by adapting the arguments in Section 4.3. The monotonicity in the sense of Proposition 4.3.1 easily follows since the weights $(\pi_k)_{k=1,\dots,N}$ appearing in the definition of $\mathcal{E}^N_{loc}$ are nonnegative. In order to get the stability, we notice that the numerical scheme (4.4.4) is actually the dynamic programming equation for the following discrete impulse control problem:

$$v^{h,N,M}_{loc}(t_i,z,\theta_j) = \sup_{\alpha\in\mathcal{A}^{h,M,loc}_\varepsilon(t_i,z,\theta_j)} \mathbb{E}\big[U_{L_\varepsilon}(\hat{Z}^{loc}_{t_m}, \hat{\Theta}_{t_m})\big],$$

where $\mathcal{A}^{h,M,loc}_\varepsilon(t_i,z,\theta_j)$ is the set of elements $\alpha=(\tau_n,\zeta_n)_n$ s.t. the trading times τ_n are $\hat{\mathbb{F}} = (\hat{\mathcal{F}}_{t_\ell})$-stopping times, valued in $\mathbb{T}_{i,m} = \{t_\ell = \ell h, \ell = i,\dots,m\}$, and ζ_n is $\hat{\mathcal{F}}_{\tau_n}$-measurable,

valued in the discrete set of admissible transactions $\mathcal{C}_\varepsilon^{M,loc}(\hat{Z}^{loc}_{\tau_n}, \tau_n - \tau_{n-1})$, where the discrete time controlled process $\{\hat{Z}^{loc}_{t_\ell} = (\hat{X}_{t_\ell}, \hat{Y}_{t_k}, \hat{P}^{loc}_{t_\ell}), \hat{\Theta}_{t_\ell}, \ell = i, \ldots, m\}$ is governed by $\hat{Z}^{loc}_{t_i} = z$, and

$$\begin{aligned}
\hat{X}_{t_\ell} &= \hat{X}_{\tau_n}, \qquad \hat{Y}_{t_\ell} = \hat{Y}_{\tau_n}, \quad \hat{\Theta}_{t_\ell} = t_\ell - \tau_n, \quad \tau_n \le t_\ell < \tau_{n+1}, \\
\hat{P}^{loc}_{t_\ell} &= \mathrm{Proj}_{[p_{min}, p_{max}]}\Big(\hat{P}^{loc}_{t_{\ell-1}} \exp\big((b - \frac{\sigma^2}{2})h + \sigma\sqrt{h}U_\ell\big)\Big), \\
(\hat{Z}^{loc}_{\tau_{n+1}}, \hat{\Theta}_{\tau_{n+1}}) &= \Big(\Gamma_\varepsilon\big(\hat{Z}^{loc}_{\tau_n}, \tau_{n+1} - \tau_n, \zeta_{n+1}\big), 0\Big),
\end{aligned}$$

where U_ℓ, $\ell = i+1, \ldots, m$ are i.i.d. discrete random variables with support $(u_k)_{k=1,\ldots,N}$ and weights $(\pi_k)_{k=1,\ldots,N}$, and $\hat{\mathcal{F}}_{t_\ell}$ is the σ-algebra generated by U_j, $j \le \ell$. Assuming for simplicity that the utility function U is bounded, we then see that the solution $v^{h,N,M}_{loc}$ to the numerical scheme is pointwise bounded uniformly in (h, N, M) and the localization parameters $(x_{min}, x_{max}, y_{max}, p_{max})$. For proving the (pointwise) consistency in the line of Proposition 4.3.3, we have to estimate, for any fixed $t_i \in \mathbb{T}_m$, $\theta_j \in \mathbb{T}_i$, $z \in \mathcal{Z}^j$, any smooth test function ϕ, the accuracy of the approximate expectation $\mathcal{E}^N_{loc}\big[\phi(t_i + h, Z^{0,t_i,z}_{t_i+h}, \theta_j + h)\big]$ with respect to the exact expectation $\mathbb{E}\big[\phi(t_i + h, Z^{0,t_i,z}_{t_i+h}, \theta_j + h)\big]$, when h goes to zero, N goes to infinity, and $R := \min[|x_{min}|, |x_{max}|, y_{max}, p_{max}]$ goes to infinity. Assuming that the smooth test function is uniformly Lipschitz in p, we have:

$$\begin{aligned}
&\Big|\mathbb{E}\big[\phi(t_i + h, Z^{0,t_i,z}_{t_i+h}, \theta_j + h)\big] - \mathcal{E}^N_{loc}\big[\phi(t_i + h, Z^{0,t_i,z}_{t_i+h}, \theta_j + h)\big]\Big| \\
= \;& \Big|\mathbb{E}\big[\phi(t_i + h, x, y, P^{t_i,p}_{t_i+h}, \theta_j + h)\big] - \mathbb{E}\big[\phi(t_i + h, x, y, \hat{P}^{loc,t_i,p}_{t_i+h}, \theta_j + h)\big]\Big| \\
= \;& \Big|\mathbb{E}\big[\phi(t_i + h, x, y, pe^{\bar{b}h + \sigma\sqrt{h}\,\mathcal{U}}, \theta_j + h)\big] \\
&\qquad - \mathbb{E}\big[\phi(t_i + h, x, y, \mathrm{Proj}_{[0,p_{max}]}(pe^{\bar{b}h + \sigma\sqrt{h}\,\hat{\mathcal{U}}})), \theta_j + h)\big]\Big| \\
\le \;& C\mathbb{E}\big|pe^{\bar{b}h + \sigma\sqrt{h}\,\mathcal{U}} - \mathrm{Proj}_{[0,p_{max}]}\big(pe^{\bar{b}h + \sigma\sqrt{h}\,\hat{\mathcal{U}}}\big)\big| \\
\le \;& C\mathbb{E}\big|pe^{\bar{b}h + \sigma\sqrt{h}\,\mathcal{U}} - \mathrm{Proj}_{[0,p_{max}]}\big(pe^{\bar{b}h + \sigma\sqrt{h}\mathcal{U}}\big)\big| + Cp\,\mathbb{E}\big|e^{\bar{b}h + \sigma\sqrt{h}\,\mathcal{U}} - e^{\bar{b}h + \sigma\sqrt{h}\,\hat{\mathcal{U}}}\big| \\
\le \;& Cp\,\mathbb{E}\Big[e^{\bar{b}h + \sigma\sqrt{h}\,\mathcal{U}} 1_{pe^{\bar{b}h + \sigma\sqrt{h}\,\mathcal{U}} > p_{max}}\Big] + Cp\sqrt{h}\mathbb{E}\Big[\Big(e^{\bar{b}h + \sigma\sqrt{h}\mathcal{U}} + e^{\bar{b}h + \sigma\sqrt{h}\hat{\mathcal{U}}}\Big)|\mathcal{U} - \hat{\mathcal{U}}|\Big],
\end{aligned}$$

where C denotes a generic constant independent of h, N, R, and we used in the second inequality the fact that the projection on $[0, p_{max}]$ is a Lipschitz function, and in the third inequality the relation: $|e^x - e^y| \le \frac{e^x + e^y}{2}|x - y|$. Then, by Cauchy-Schwarz inequality, we obtain:

$$\begin{aligned}
&\Big|\mathbb{E}\big[\phi(t_i + h, Z^{0,t_i,z}_{t_i+h}, \theta_j + h)\big] - \mathcal{E}^N_{loc}\big[\phi(t_i + h, Z^{0,t_i,z}_{t_i+h}, \theta_j + h)\big]\Big| \\
\le \;& Cp\sqrt{E\Big[e^{2\bar{b}h + 2\sigma\sqrt{h}\,\mathcal{U}}\Big]}\sqrt{\mathbb{P}\big[pe^{\bar{b}h + \sigma\sqrt{h}\,\mathcal{U}} > p_{max}\big]} \\
&\quad + Cp\sqrt{h}\sqrt{\mathbb{E}\Big[e^{2\bar{b}h + 2\sigma\sqrt{h}\,\mathcal{U}}\Big] + \mathbb{E}\Big[e^{2\bar{b}h + 2\sigma\sqrt{h}\,\hat{\mathcal{U}}}\Big]}\sqrt{\mathbb{E}|\mathcal{U} - \hat{\mathcal{U}}|^2}.
\end{aligned}$$

Now, since $\hat{\mathcal{U}}$ is an optimal quantization of $\mathcal{U}$, we have the stationary property, meaning that $\mathbb{E}[\mathcal{U}|\hat{\mathcal{U}}] = \hat{\mathcal{U}}$ (see [58]), which implies from Jensen's inequality applied to the convex function $u \to e^{2\bar{b}h+2\sigma\sqrt{h}u}$, and the law of iterated conditional expectations:

$$\mathbb{E}\Big[e^{2\bar{b}h+2\sigma\sqrt{h}\,\hat{\mathcal{U}}}\Big] \;\leq\; \mathbb{E}\Big[e^{2\bar{b}h+2\sigma\sqrt{h}\,\mathcal{U}}\Big] \;=\; e^{(2b+\sigma^2)h}.$$

Denoting by Φ the distribution function of $\mathcal{U}$, we then have:

$$\begin{aligned}&\Big|\mathbb{E}\big[\phi(t_i+h, Z^{0,t_i,z}_{t_i+h}, \theta_j+h)\big] - \mathcal{E}^N_{loc}\big[\phi(t_i+h, Z^{0,t_i,z}_{t_i+h}, \theta_j+h)\big]\Big| \\ \leq\;& Cpe^{(b+\frac{\sigma^2}{2})h}\Big\{\sqrt{1-\Phi\Big(\frac{\ln\big(\frac{p_{max}}{pe^{\bar{b}h}}\big)}{\sigma\sqrt{h}}\Big)} + \sqrt{h}\sqrt{\mathbb{E}|\mathcal{U}-\hat{\mathcal{U}}|^2}\Big\}.\end{aligned}$$

From Zador's theorem (see [58]), the asymptotic distorsion error for the optimal quantization satisfies: $\lim_{N\to\infty} N\sqrt{\mathbb{E}|\mathcal{U}-\hat{\mathcal{U}}|^2} \in (0,\infty)$, and so $\sqrt{\mathbb{E}|\mathcal{U}-\hat{\mathcal{U}}|^2} = O(1/N)$. Recalling the well known estimate: $1-\Phi(d) \sim \varphi(d)/d$, as d goes to infinity, where $\varphi = \Phi'$ is the density of $\mathcal{U}$, we obtain by taking N s.t. $N\sqrt{h} \to \infty$, e.g. $N = O(1/h^{\frac{1}{2}+\varepsilon})$, with $\varepsilon > 0$, and $p_{max} > p$, the pointwise estimation:

$$\mathcal{E}^N_{loc}\big[\phi(t_i+h, Z^{0,t_i,z}_{t_i+h}, \theta_j+h) \;=\; \mathbb{E}\big[\phi(t_i+h, Z^{0,t_i,z}_{t_i+h}, \theta_j+h)\big] + o(h),$$

where the notation $o(h)$ means that $o(h)/h$ goes to zero as h goes to zero. This yields

$$\begin{aligned}&\lim_{\substack{h\,\to\,0\\ N\sqrt{h},\,R\,\to\,\infty}} \frac{\phi(t_i,z,\theta_j) - \mathcal{E}^N_{loc}\big[\phi(t_i+h, Z^{0,t_i,z}_{t_i+h}, \theta_j+h)\big]}{h} \\ =\;& \lim_{h\to 0} \frac{\phi(t_i,z,\theta_j) - \mathbb{E}\big[\phi(t_i+h, Z^{0,t_i,z}_{t_i+h}, \theta_j+h)\big]}{h}. \qquad (4.4.6)\end{aligned}$$

On the other hand, for fixed $t_i \in \mathbb{T}_m$, $\theta_j \in \mathbb{T}_i$, $z \in \mathcal{Z}^j$, we notice that $\cup_{M,R=1}^{\infty} \mathcal{C}^{M,loc}_\varepsilon(z,\theta_j)$ is dense in $\mathcal{C}_\varepsilon(z,\theta_j)$. Hence, by continuity of ϕ, Γ_ε, and compacity of $\mathcal{C}_\varepsilon(z,\theta_j)$, we deduce that

$$\lim_{M,R\to\infty} \mathcal{H}^{M,loc}_\varepsilon \phi(t_i,z,\theta_j) \;=\; \mathcal{H}_\varepsilon\phi(t_i,z,\theta_j).$$

Together with (4.4.6), we then obtain similarly as in Proposition 4.3.3:

$$\begin{aligned}&\lim_{\substack{h\,\to\,0\\ N\sqrt{h},\,M,\,R\,\to\,\infty}} \min\Big\{\frac{\phi(t_i,z,\theta_j) - \mathcal{E}^N_{loc}\big[\phi(t_i+h, Z^{0,t_i,z}_{t_i+h}, \theta_j+h)\big]}{h}\;,\; (\phi - \mathcal{H}^{M,loc}_\varepsilon\phi)(t_i,z,\theta_j)\Big\} \\ =\;& \min\Big\{-\Big(\frac{\partial\phi}{\partial t} + \mathcal{L}\phi\Big)(t_i,z,\theta_j)\;,\; (\phi - \mathcal{H}_\varepsilon\phi)(t_i,z,\theta_j)\Big\},\end{aligned}$$

which then proves the convergence of the numerical scheme $S^{h,M,N}_{loc}$.

Algorithm description. In summary, our numerical scheme provides an algorithm for computing approximations v^h of the value function, and ζ^h of the optimal trading strategy at each time step $t_i \in \mathbb{T}_m$, and each point (z, θ) of the grid $(\mathbb{X}_n \times \mathbb{Y}_n \times \mathbb{P}_n \times \mathbb{T}_i) \cap \bar{\mathcal{S}}_\varepsilon$, where $\mathbb{X}_n$ is the uniform grid with n nodes on $[x_{min}, x_{max}]$, i.e. of step $(x_{max} - x_{min})/n$, and similarly for $\mathbb{Y}_n$, $\mathbb{P}_n$. Let us also denote by $\mathbb{Z}_n^j = \{z \in \mathbb{Z}_n : (z, \theta_j) \in \bar{\mathcal{S}}_\varepsilon\}$. The parameters in the algorithm are:

- T the maturity
- b and σ the Black and Scholes parameters of the stock price
- λ the impact parameter, β the impact exponent in the market impact function (4.2.5)
- κ_a, κ_b the spread parameters in percent, ε the transactions costs fee
- We take by default a CRRA utility function: $U(x) = x^\gamma$
- $x_{min}, x_{max} \in \mathbb{R}$, $0 \leq y_{min} < y_{max}$, $0 \leq p_{min} < p_{max}$, the boundaries of the localized domain
- m number of steps in time discretization, n the number of steps in space discretization
- N number of points for optimal quantization of the normal law, M number of points used in the static supremum in e

The algorithm is described explicitly in backward induction as follows:

▶ *Initialization step at time $t_m = T$:*

- *(s:0)* For $j = 0$, set $v^h(t_m, z, 0) = U_{L_\varepsilon}(z, 0)$, $\zeta^h(t_m, z, 0) = 0$ on $\mathbb{Z}_n^0$, and interpolate $v^h(t_m, z, 0)$ on $\mathcal{Z}_{loc}^0$.
- *(s:j)* For $j = 1, \dots, m$,
  - for $z \in \mathbb{Z}_n^j$, compute $v := \sup_{e \in \mathcal{C}_\varepsilon^{M,loc}(z,\theta_j)} U_{L_\varepsilon}(\Gamma_\varepsilon(z, \theta_j, e), 0)$ and denote by $\hat{e}$ the argument maximum:
  - if $v > U_{L_\epsilon}(z, \theta_j)$, then set $v^h(t_m, z, \theta_j) = v$ and $\zeta^h(t_m, z, \theta_j) = \hat{e}$,
  - else set $v^h(t_m, z, \theta_j) = U_{L_\varepsilon}(z, \theta_j)$, and $\zeta^h(t_m, z, \theta_j) = 0$.
  - Interpolate $z \to v^h(t_m, z, \theta_j)$ on $\mathcal{Z}_{loc}^j$.

▶ *From time step t_{i+1} to t_i, $i = m-1, \dots, 0$:*

- *(s:0)* For $j = 0$, compute $\mathcal{E}_{loc}^N\big[v^h(t_i + h, Z_{t_i+h}^{0,t_i,z}, \theta_j + h)\big]$ from (4.4.3) and *(s:1)* of time step t_{i+1}, and set $v^h(t_i, z, 0) = \mathcal{E}_{loc}^N\big[v^h(t_i + h, Z_{t_i+h}^{0,t_i,z}, \theta_j + h)\big]$, $\zeta^h(t_i, z, 0) = 0$ on $\mathbb{Z}_n^0$; interpolate $v^h(t_i, z, 0)$ on $\mathcal{Z}_{loc}^0$.
- *(s:j)* For $j = 1, \dots, i$,

- for $z \in \mathbb{Z}_n^j$, compute $\mathcal{E}_{loc}^N\big[v^h(t_i+h, Z_{t_i+h}^{0,t_i,z}, \theta_j+h)\big]$ from (4.4.3) and *(s:j+1)* of time step t_{i+1}, $v := \sup_{e \in \mathcal{C}_\varepsilon^{M,loc}(z,\theta_j)} v^h(t_i, \Gamma_\varepsilon(z,\theta_j,e),0)$ from *(s:0)*, and denote by $\hat{e}$ the argument maximum:
- if $v > \mathcal{E}_{loc}^N\big[v^h(t_i+h, Z_{t_i+h}^{0,t_i,z}, \theta_j+h)\big]$, then set $v^h(t_i,z,\theta_j) = v$, $\zeta^h(t_i,z,\theta_j) = \hat{e}$,
- else set $v^h(t_i,z,\theta_j) = \mathcal{E}_{loc}^N\big[v^h(t_i+h, Z_{t_i+h}^{0,t_i,z}, \theta_j+h)\big]$, and $\zeta^h(t_i,z,\theta_j) = 0$.
- Interpolate $z \to v^h(t_i,z,\theta_j)$ on $\mathcal{Z}_{loc}^j$.

Complexity of the algorithm. Due to the high dimension of the grid

$$\mathbb{S} \;=\; \mathbb{T}_m \times \bigcup_{i=1\dots m} \big((\mathbb{X}_n \times \mathbb{Y}_n \times \mathbb{P}_n \times \mathbb{T}_i) \cap \bar{\mathcal{S}}_\varepsilon\big)\,,$$

the computation of the optimal policy on the entire grid has an expensive computational cost. Indeed, this grid contains $O(m^2n^3)$ points, and at each point $(t_i,z,\theta_j) \in \mathbb{S}$, one has to compute:

- The approximation of conditional expectation $\mathcal{E}_{loc}^N\big[v^h(t_i+h, Z_{t_i+h}^{0,t_i,z}, \theta_j+h)\big]$ that costs $O(N)$ unitary operations.
- The approximation of the static supremum $\sup_{e \in \mathcal{C}_\varepsilon^{M,loc}(z,\theta_j)} v^h(t_i, \Gamma_\varepsilon(z,\theta_j,e),0)$, together with its argument maximum, that costs $O(M)$ unitary operations when using linear search[1].
- The localization procedure and the interpolation procedure has constant computational cost $O(1)$.

Therefore, we obtain a complexity of:

$$\text{Complexity} \;=\; O(m^2n^3 \max(N,M)).$$

Actually, denoting by $K = \max(n,m,N,M)$, the complexity of the algorithm is $O(K^6)$. Yet, practical implementation of the algorithm can achieve quite better performance. First, in the optimal quantization for the computation of the expectations in the numerical algorithm, we can choose $N = O(m^{1/2+\varepsilon})$ for all $\epsilon > 0$. Assuming that we are able to use a dichotomy-based method for computing the static supremum, which has logarithmic complexity, the main computational costs are due to the computation of the approximate

[1]Note that the supremum computation can be improved by the use of dichotomy-based search instead of linear search if we are able to use a concavity argument on $e \mapsto v(t,\Gamma(x,y,p,\theta,e),0)$ which would lead to a complexity of $O(\ln(M))$. From numerical experiments, this dichotomy search method leads to acceptable results.

conditional expectation, and we can neglect the cost of computing the static supremum. In this case, the complexity is reduced to:

$$\text{Complexity} \quad = \quad O(m^{5/2+\varepsilon}n^3) \ , \forall \varepsilon > 0,$$

which is satisfactory when considering that there is $O(m^2n^3)$ points to compute in the grid. Second, the grid computation algorithm can be parallelized easily, which is a very desirable property when targeting an industrial application. Indeed, at each date t_i the computation of $\mathcal{E}^N_{loc}\big[v^h(t_i+h,x,y,P^{0,t_i,p}_{t_i+h},\theta_j+h)\big]$ and $\sup_{e\in\mathcal{C}^{M,loc}_\varepsilon(z,\theta_j)} v^h(t_i,\Gamma_\varepsilon(x,y,p,\theta_j,e),0)$ can be done independently for each quadruplet (x,y,p,θ_j) provided that $\theta_j > 0$.

Finally, the complexity displayed above represents the amount of computations needed to build up the optimal policy. When targeting a live trading application, one can compute off-line and store optimal policies for a given set of market parameters, and when actually trading, one does only need to read (with constant cost) the optimal policy corresponding to current market state.

Comparison with finite difference scheme. In order to motivate our numerical scheme proposal, let us compare it with usual finite difference scheme. Let us briefly introduce the class of theta-schemes. We refer to [47] for complete discussion about this class of schemes. We will assume that the value function is sufficiently smooth, and we focus in this paragraph on the diffusive part of the QVI, so that our target equation to solve is:

$$\frac{\partial}{\partial t} + \mathcal{L}v \quad = \quad 0 \text{ on } \bar{\mathcal{S}}_\varepsilon \times [0,T),$$

together with a terminal condition on $\bar{\mathcal{S}}_\varepsilon \times \{T\}$. To solve numerically this Kolmogorov parabolic equation with finite time horizon, we can discretize it using a theta-scheme of parameter a according to [47]. This approximation consists in the following:

$$\left(\frac{\partial}{\partial t}v + \mathcal{L}v\right)(t,z,\theta) \simeq \mathcal{P}^a_{h,\delta}v(t,z,\theta)$$

where

$$\mathcal{P}^a_{h,\delta}v(t,z,\theta) = \frac{v(t+h,z,\theta+h)-v(t,z,\theta)}{h} + aL_\delta v(t,z,\theta) + (1-a)L_\delta v(t+h,z,\theta+h)$$

and L_δ is the finite difference approximation of $\tilde{\mathcal{L}} := bp\dfrac{\partial}{\partial p} + \frac{1}{2}\sigma^2p^2\dfrac{\partial^2}{\partial p^2}$ of (space) step δ and $a \in [0,1]$. The discretized equation is:

$$\mathcal{P}^a_{h,\delta}v(t,z,\theta) \quad = \quad 0 \text{ on } \mathcal{O}_\delta \cap \bar{\mathcal{S}} \times [0,T),$$

where $\mathcal{O}_\delta$ is a suitable regular grid of (space) step δ. From the finite differences approximation, we have the following precision:

$$\left(\frac{\partial}{\partial t}v + \mathcal{L}v\right)(t,z,\theta) \quad = \quad \mathcal{P}^a_{h,\delta}v(t,z,\theta) + o(h^p+\delta^q),$$

where p and q depends on the choice of a: if $a \neq 1/2$ we obtain that $p = 1$, and if $a = 1/2$ we obtain that $p = 2$, which corresponds to the Crank-Nicholson scheme. Due to the second order derivative in $\mathcal{L}$, and by using standard finite difference approximation, the rate of convergence for the spatial approximation is $q = 1$, $\forall a \in [0, 1]$. Therefore, in our case, we see that theta-schemes have order 1 in time and order 1 in space, except for the Crank-Nicholson scheme, which gives an order 2 in time and order 1 in space. For comparison purpose, the optimally quantized scheme that we use has order 1 in time provided that $N = O(h^{-(1/2+\varepsilon)})$ where N is the number of points in the optimal quantization grid:

$$\left(\frac{\partial}{\partial t}v + \mathcal{L}v\right)(t, z, \theta) = 0 \ , \ N = O(h^{-(1/2+\varepsilon)}) \Longrightarrow v^h(t, z, \theta) = \mathcal{E}^N_{loc}\big[v^h(t+h, Z^{0,t,z}_{t+h}, \theta+h)\big] + o(h).$$

This raises two comments. First, we see that in contrast with finite difference scheme, the precision of the optimally quantized scheme is controlled by the number of points N of the optimal quantization grid, and not by the space step δ, provided that interpolation procedure is sufficiently efficient. Therefore, one can improve the precision by increasing N and without increasing the size of the grid, which is very interesting when dealing with high-dimension state space. Second, the above result allows us to choose $n = O(m^{1/2})$, while keeping a precision of $o(1/m)$, whereas if using a finite-difference scheme, the precision would be $o(1/m^{1/2})$ due to spatial approximation. Therefore, by using an optimally quantized scheme, we can obtain a satisfactory precision, while managing efficiently the size of the grid, and subsequently the memory needed to achieve computation, which is quite relevant when dealing with high-dimensional state space.

Yet, two other theta-schemes may be good candidates for solving numerically our QVI, the Crank-Nicholson scheme due to its higher order in time, and the fully-implicit scheme, corresponding to $a = 1$ because it has the property of being stable without restriction on the choice of time step versus space step.

4.5 Numerical Results

4.5.1 Procedure

For each of the numerical tests, we used the same procedure consisting in the following steps:

(1) Set the parameters according to the parameter table described in the first subsection of each test

(2) Compute and save the grids representing value function and optimal policy according to the optimal liquidation algorithm

(3) Generate Q paths for the stock price process following a geometrical Brownian motion: we choose parameters b and σ that allows us to observe several empirical facts on the

performance and the behavior of optimal liquidation strategy. These parameters can also be estimated from historical observations on real data by standard statistical methods.

(4) Consider the portfolio made of X_0 dollars and Y_0 shares of risky asset

(5) For each price path realization, update the portfolio along time and price path accordingly to the policy computed in the second step

(6) Save each optimal liquidation realization

(7) Compute statistics

In the sequel, we shall use the following quantities as descriptive statistics:

- The performance of the i-th realization of the optimal strategy is defined by

$$L_{opt}^{(i)} = \frac{L_\epsilon(Z_T^{(i),\alpha^{opt}}, \Theta_T^{(i),\alpha^{opt}})}{X_0 + Y_0 P_0}$$

where $(Z_T^{(i),\alpha^{opt}}, \Theta_T^{(i),\alpha^{opt}})$ is the state process, starting at date 0 at $(X_0, Y_0, P_0, 0)$, evolving under the i-th price realization and the optimal control α^{opt}. This quantity can be interpreted as the ratio between the cash obtained from the optimal liquidation strategy and the ideal Merton liquidation. We define in the same way the quantities $L_{naive}^{(i)}$ and $L_{uniform}^{(i)}$ respectively associated with the controls α^{naive} and $\alpha^{uniform}$ of the naive and uniform strategy, referred to as benchmark strategies. Recall that the naive strategy consists in liquidating the whole portfolio in one block at the last date, and the uniform strategy consists in liquidating the same quantity of asset at each predefined date until the last date. Notice that the score 1 corresponds to the strategy, which consists in liquidating the whole portfolio immediately in an ideal Merton market.

When denoting by Q the number of paths of our simulation, we define:

- The mean utility $\hat{V}_\cdot = \frac{1}{Q}\sum_{i=1}^{Q} U(L_\cdot^{(i)})$
- The mean performance $\hat{L}_\cdot = \frac{1}{Q}\sum_{i=1}^{Q} L_\cdot^{(i)}$
- The standard deviation of the strategy $\hat{\sigma}_\cdot = \sqrt{\frac{1}{Q}\sum_{i=1}^{Q}(L_\cdot^{(i)})^2 - \hat{L}_\cdot^{\,2}}$

Here the dot . stands for *opt*, *naive* or $uniform$. We will also compute the third and fourth standardized moments for the series $(L_\cdot^{(i)})_i$.

4.5.2 Test 0: Convergence of the numerical scheme

In order to experiment numerically the convergence of the scheme, we performed two series of convergence tests. First, we computed a reference value function with a fine discretization grid, and computed for various sizes of grids the difference to this reference result. Second, we backtested the optimal policy obtained with various discretization grid sizes, using the the procedure described in 5.1, and compared the results.

Due to the high dimension of the problem, we restricted our convergence analysis to reasonably sized discretization grids, except for the reference computation, and therefore missing values in tables 4.3 and 4.4 corresponds either to grids that required too much memory space or too much time to compute. When targeting industrial applications, one can avoid these restrictions by using a suitable parallel algorithm, as we did for computing the reference value function. Yet, with a reasonable size of grid, for example ($m = 64, n = 32$) one can achieve satisfactory results (see table 4.3).

Convergence of the value function First, we computed a reference value function that we will denote v^∞ using a parallelized version of our algorithm with parameters shown in table 4.1. We ran the computations on two SGI Altix ICE 8200EX supercomputers made of 256 computing cores 64-bit at 2.83 GHz with 512 GB of distributed RAM. Computations took 11 hours and 36 minutes to complete and size of computer representation of v^∞ was 0.991 TB.

Parameter	Value	Parameter	Value
Maturity	1 day	m	256
λ	0.02	n	128
β	0.2	N	30
γ	0.5		
κ_A	1.001		
κ_B	0.999		
ϵ	0.001		

Table 4.1: Test 0: parameters for the reference computation v^∞

Second, we computed value functions for different values of n and m, (see parameters in table 4.2) that we will denote $v^{n,m}$, and we computed the relative error we made compared to v^∞, i.e. $\dfrac{\|v^\infty - v^{n,m}\|_2}{\|v^\infty\|_2}$. The results are reported on table 4.3.

As a consequence of this convergence test, we will use in the following tests the following values: $m \in [|30...60|]$ and $n \in [|20...60|]$. Indeed these sizes of grid are a good compromise between computational complexity and precision.

Parameter	Value	Parameter	Value
Maturity	1 day	X_0	20
λ	0.02	Y_0	250
β	0.2	P_0	1.0
γ	0.5	N	$\sqrt{m}$
κ_A	1.001	Q	10^5
κ_B	0.999		
ϵ	0.001		

Table 4.2: Test 0: parameters

n / **m**	4	8	16	32	64
16	0.2251	0.1043	0.0668	0.0582	0.0563
32	0.2231	0.0997	0.0567	0.0445	0.0416
64	0.2210	0.0970	0.0501	0.0343	
128	0.2207	0.0968	0.0498		
256	0.2204	0.0967	0.0498		

Table 4.3: Test 0: convergence of the value function. Quantity displayed is $\frac{\|v^\infty - v^{n,m}\|_2}{\|v^\infty\|_2}$.

Backtesting the optimal strategy We compared the fully implicit scheme to our optimally quantized scheme, following the procedure described in section 5.1. The fully implicit scheme corresponds to a theta-scheme with parameter $a = 1$, and has the property of inducing no restriction on the choice of timestep. Therefore we use it as a benchmark for our optimally quantized scheme. Parameters are reported in table 4.2 and results in table 4.4.

In table 4.5 we display the same convergence test measured in terms of the statistics $\frac{\hat{L}_{Quantized} - \hat{L}_{Implicit}}{\hat{\sigma}_{Quantized}}$ where $\hat{L}_{Quantized}$ (resp. $\hat{L}_{Implicit}$) is the estimate of performance for the initial portfolio (X_0, Y_0, P_0) using the optimally quantized scheme (resp. the fully implicit scheme) for computing the optimal policy and $\hat{\sigma}_{Quantized}$ its standard deviation. This quantity is more intuitive from the financial point of view, and can be interpreted as the gain in mean performance when using the optimally quantized scheme compared to using the fully implicit scheme, measured with the standard deviation as unit. We remark that the optimally quantized scheme performs better for most values of (m, n), especially for small-sized time grids. When increasing the size of the time grid, the difference of performance for these two scheme seems to vanish, in terms of the above statistics, but we

scheme	**n** / **m**	4	8	16	32	64
Quantized	16	0.5238	0.8603	0.8667	0.8752	0.8749
Implicit		0.5410	0.8478	0.8574	0.8574	0.8593
Quantized	32	0.5420	0.8486	0.8676	0.8747	0.8743
Implicit		0.5411	0.8458	0.8589	0.8607	0.8619
Quantized	64	0.5411	0.8465	0.8578	0.8601	
Implicit		0.5410	0.8465	0.8609	0.8603	
Quantized	128	0.5405	0.8417			
Implicit		0.5411	0.8456			
Quantized	256	0.5193				
Implicit		0.5278				

Table 4.4: Test 0: Convergence of the numerical algorithm: table of value function estimated by Monte-Carlo simulation $\hat{V}$ with initial portfolio (X_0, Y_0, P_0) when varying grid size (m is number of time steps, n the number of space steps, with boundaries fixed). We display results for the optimally quantized scheme (referred to as "Quantized" scheme in the table) against the benchmark made of the theta-scheme of parameter $a = 1$ and usual finite difference approximation (referred to as "Implicit" scheme in the table).

need more precise tests to conclude.

n / **m**	4	8	16	32	64
16	0.1191	0.1179	0.1085	0.1662	0.1482
32	0.0890	0.0612	0.0902	0.1175	0.1065
64	0.0109	0.0367	0.0328	0.0521	
128	-0.0439	0.0127			
256	-0.3437				

Table 4.5: Test 0: Convergence of the numerical algorithm: table for the statistics $(\hat{L}_{Quantized} - \hat{L}_{Implicit})/\hat{\sigma}_{Quantized}$ when varying grid size.

4.5.3 Test 1: A toy example

The goal of this test is to show the main characteristics of our results. We choose a set of parameters that is unrealistic but that has the advantage of emphasizing the typical

behavior of the optimal liquidation strategy.

Parameters We choose the set of parameters shown in table 4.6.

Parameter	Value	Parameter	Value
Maturity	1 year	X_0	2000
λ	5.00E-07	Y_0	2500
β	0.5	P_0	5.0
γ	0.5	x_{min}	-30000
κ_A	1.01	x_{max}	80000
κ_B	0.99	y_{min}	0
ϵ	0.001	y_{max}	5000
b	0.1	p_{min}	0
σ	0.5	p_{max}	20
		m	40
		n	20
		N	100
		Q	10^5

Table 4.6: Test 1: parameters

Execution statistics The results were computed using Intel® Core 2 Duo at 2.93Ghz CPU with 2.98 Go of RAM. Statistics are shown in table 4.7.

Quantity	Evaluation
Time Elapsed for grid computation in seconds	7520
Number Of Available Processors	2
Estimated Memory Used (Upper bound)	953MB
Time Elapsed for statistics Computation in seconds	21

Table 4.7: Test 1: Execution statistics

Shape of policy In this paragraph we plotted the shape of the policy sliced in the plane (x, y), i.e. the (cash, shares) plane, for a fixed (t, θ, p) (figure 4.1). The color of the map at (x_0, y_0) on the graph represents the action one has to take when reaching the state (t, θ, x_0, y_0, p). We can see three zones: a buy zone (denoted BUY on the graph), a sell

zone (denoted SELL on the graph) and a no trade zone (denoted NT on the graph). Note that the bottom left zone on the graph is outside the domain $\bar{\mathcal{S}}$. These results have the intuitive financial interpretation: when x is big and y is small, the investor has enough cash to buy shares of the risky asset and tries to profit from an increased exposure. When y is large and x is small, the investor has to reduce exposure to match the terminal liquidation constraint.

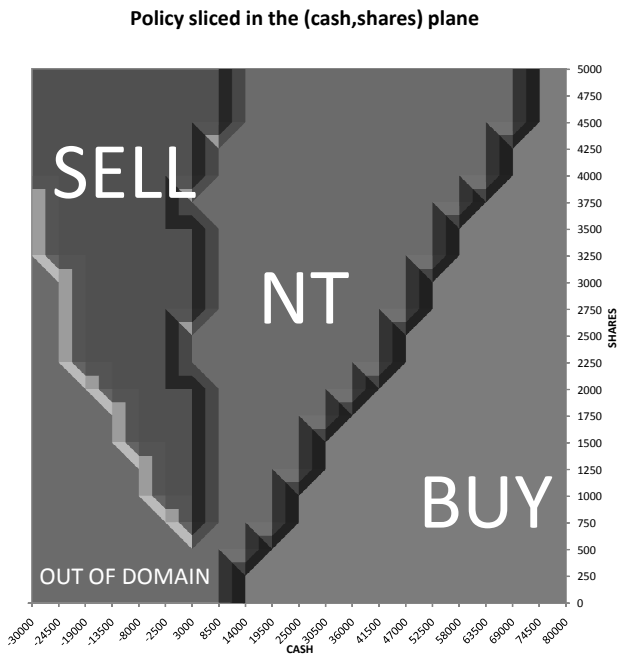

Figure 4.1: Test 1: Typical shape of the policy sliced in XY near date 0

We also plotted the shape of the policy sliced in the plane (y, p), i.e. the (shares,price) plane, for a fixed (t, θ, x) (figure 4.2). As before, the color of the map at (y_0, p_0) on the graph represents the action one has to take when reaching the state (t, θ, x, y_0, p_0). Again, we can distinguish the three zones: buy, sell and no trade.

Remark 4.5.1 In our modelling, we allow buying to occur during liquidation. This may be a priori undesirable in practice, and one could easily enforce a no-buying constraint in our model by requiring that the strategies (ζ_n) should be nonpositive, so that the shape policy is reduced to two zones instead of three zones as above: a no-trade and a sell zone. However, by giving more flexibility to the investor, we allow him to take advantage of a drop of the asset price, as illustrated in Figure 4.7, and so to realize a better performance.

Figure 4.2: Test 1: Typical shape of the policy sliced in YP

Shape of value function Figure 4.3 shows the value function sliced in the (x, y) plane. This figure is a typical pattern of the value function. Recall from Proposition 3.1 in [44] the following Merton theoretical bound for the value function:

$$v(t,z,\theta) \;\leq\; v_M(t,x,y,p) \;=\; e^{\rho(T-t)}(x+yp)^\gamma, \quad \text{with } \rho \;=\; \frac{\gamma}{1-\gamma}\frac{b^2}{2\sigma^2}.$$

In the figure 4.4 we plotted the difference between the value function and this theoretical bound. We observe that this difference is increasing with the number of shares, and decreasing with the cash. This result is interpreted as follows: the price impact increases with the number of shares, but this can be reduced by the liquidation strategy whose efficiency is greater if the investor can sustain bigger cash variations.

4.5.4 Test 2: Short term liquidation

The goal of this test is to show the behavior of the algorithm on a realistic set of parameters and real data. We used Reuters data fed by OneTick TimeSeries Database. We used the spot prices (Best Bid and Best Ask) for the week starting 04/19/2010 on BNP.PA. We

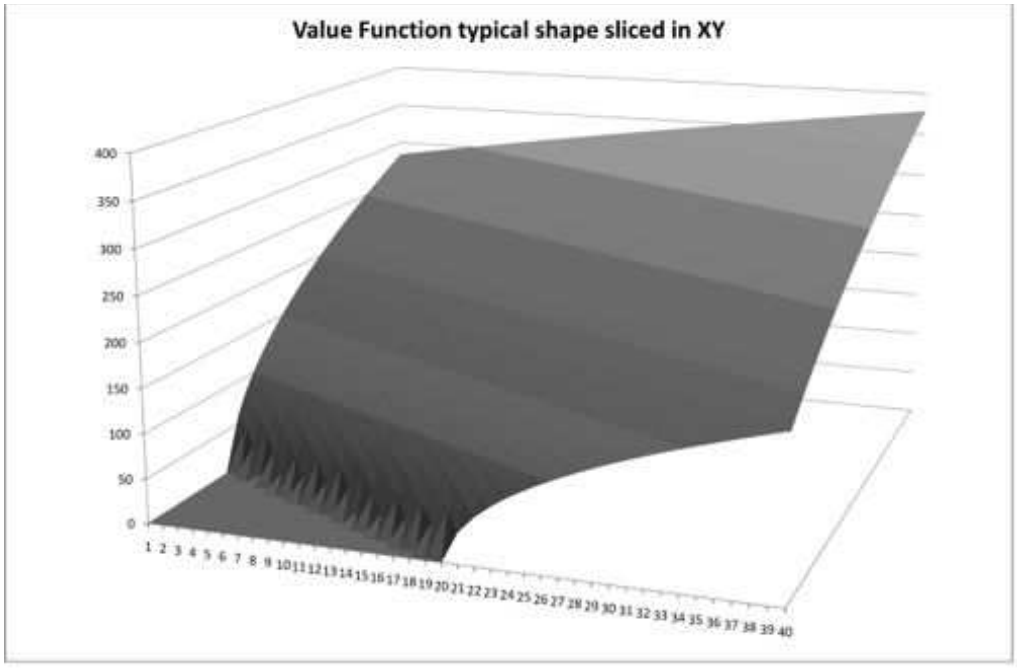

Figure 4.3: Test 1: Typical shape of the value function sliced in XY

computed mid-price that is the middle between best bid and best ask price. We choose the impact parameter λ in order to penalize by approximately 1% the immediate liquidation of the whole portfolio compared to Merton liquidation. In other words, we take λ so that: $\lambda|\frac{Y_0}{T}|^\beta \simeq 0.01$.

Parameters We computed the strategy with parameters shown in table 4.8.

Execution statistics We obtained the results using Intel® Core 2 Duo at 2.93Ghz CPU with 2.98 Go of RAM, the computations statistics are gathered in table 4.9.

Performance Analysis We computed the mean utility and the first four moments of the optimal strategy and the two benchmark strategies in table 4.10 and plotted the empirical distribution of performance in figure 4.5. It is remarkable that the optimal strategy gives an empirical performance that is above the immediate liquidation at date 0 in the Merton ideal market (represented by performance $\hat{L} = 1$, and usually refered to as *reference price* benchmark). This is due to the fact that the optimal strategy has an "opportunistic" behavior: indeed, an optimal trading strategy is embedded with the liquidation constraint: in this example, this feature not only compensates the trading costs, but also provides an

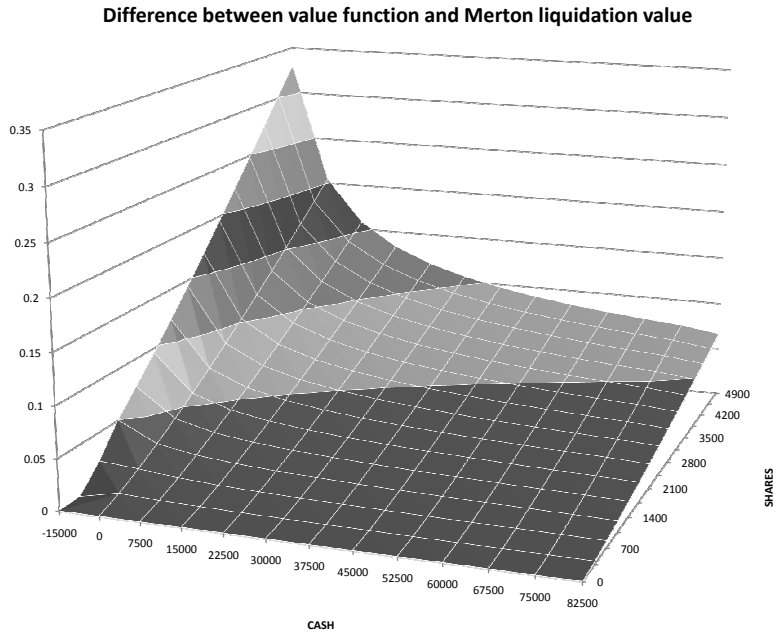

Figure 4.4: Test 1: Difference between value function and Merton theoretical bound

extra performance compared to an ideal immediate liquidation at date 0. Still, the Merton case is a theoretical upper bound in the following sense: the optimal value function with trading costs is below the optimal value function without trading costs, recall the figure 4.4. As expected, the empirical distribution is between the distributions of the two other benchmark strategies. We also notice that the optimal strategy outperforms the two others by approximatively 0.25% in utility and in performance. We also computed other statistics in table 4.11.

Behavior Analysis In this paragraph, we analyze the behaviour of the strategy as follows: first, we plotted in figure 4.6 the empirical distribution of the number of trades for one trading session. Secondly, we plotted trades realizations for three days of the BNPP.PA stock for the week starting on 04/19/2010.

The three following graphs represent three days of market data for which we computed the mid-price (lines) with associated trades realizations for the optimal strategy (vertical bars). A positive quantity for the vertical bar means a buying operation, while a negative quantity means a selling operation.

Figure 4.7 shows the trade realizations of the optimal strategy for the day 04/19/2010

Parameter	Value	Parameter	Value
Maturity	1 Day	X_0	20000
λ	5.00E-04	Y_0	2500
β	0.2	P_0	52.0
γ	0.5	x_{min}	-30000
κ_A	1.0001	x_{max}	200000
κ_B	0.9999	y_{min}	0
ϵ	0.001	y_{max}	5000
b	0.005	p_{min}	50.0
σ	0.25	p_{max}	54.0
		m	30
		n	40
		N	100
		Q	10^5

Table 4.8: Test 2: Parameters

Quantity	Evaluation
Time Elapsed for grid computation in seconds	8123
Number Of Available Processors	2
Estimated Memory Used (Upper bound)	573MB

Table 4.9: Test 2: Execution statistics

Strategy	Utility $\hat{V}$	Mean $\hat{L}$	Standard Dev.	Skewness	Kurtosis
Naive	0.99993	0.99986	0.00429	0.94584	4.68592
Uniform	0.99994	0.99988	0.00240	0.42788	3.34397
Optimal	1.00116	1.00233	0.00436	1.03892	4.89161

Table 4.10: Test 2: Utility and first four moments for the optimal strategy and the two benchmark strategies

on the BNPP.PA stock. The interesting feature in this first graph is that we see two buying decisions when the price goes down through the 54.5 EUR barrier, and which corresponds roughly to a daily minimum. The following selling decision can be viewed as a failure. On the contrary, the two last selling decisions correspond quite precisely to local maxima.

Figure 4.8 (resp.4.9) shows the trade realizations of the optimal strategy for the day

Quantity	Formula	Value
Winning percentage	$\frac{1}{Q}\sum_{i=1}^{Q} 1_{\{L_{opt}^{(i)} > \max(L_{naive}^{(i)}, L_{uniform}^{(i)})\}}$	58.8%
Relative Optimal Utility	$\frac{\hat{V}_{opt} - \max(\hat{V}_{naive}, \hat{V}_{uniform})}{\hat{V}_{opt}}$	0.00238
Relative Optimal Performance	$\frac{\hat{L}_{opt} - \max(\hat{L}_{naive}, \hat{L}_{uniform})}{\hat{L}_{opt}}$	0.00244
Utility Sharpe Ratio	$\frac{\hat{V}_{opt} - \max(\hat{V}_{naive}, \hat{V}_{uniform})}{\hat{\sigma}_{opt}}$	0.28017
Performance Sharpe Ratio	$\frac{\hat{L}_{opt} - \max(\hat{L}_{naive}, \hat{L}_{uniform})}{\hat{\sigma}_{opt}}$	0.56140
VaR 95% Naive Strategy	$\sup\left\{x \mid \frac{1}{Q}\sum_{i=1}^{Q} 1_{\{L_{naive}^{(i)} > x\}} \geq 0.95\right\}$	0.994
VaR 95% Uniform Strategy	$\sup\left\{x \mid \frac{1}{Q}\sum_{i=1}^{Q} 1_{\{L_{uniform}^{(i)} > x\}} \geq 0.95\right\}$	0.996
VaR 95% Optimal Strategy	$\sup\left\{x \mid \frac{1}{Q}\sum_{i=1}^{Q} 1_{\{L_{opt}^{(i)} > x\}} \geq 0.95\right\}$	0.997
VaR 90% Naive Strategy	$\sup\left\{x \mid \frac{1}{Q}\sum_{i=1}^{Q} 1_{\{L_{naive}^{(i)} > x\}} \geq 0.90\right\}$	0.995
VaR 90% Uniform Strategy	$\sup\left\{x \mid \frac{1}{Q}\sum_{i=1}^{Q} 1_{\{L_{uniform}^{(i)} > x\}} \geq 0.90\right\}$	0.997
VaR 90% Optimal Strategy	$\sup\left\{x \mid \frac{1}{Q}\sum_{i=1}^{Q} 1_{\{L_{opt}^{(i)} > x\}} \geq 0.90\right\}$	0.998

Table 4.11: Test 2: Other statistics on performance of optimal strategy

04/22/2010 (resp. 04/23/2010) on the BNPP.PA stock.on. Note that in figure 4.9, the naive strategy was overperforming the optimal strategy, due to an unexpected price increase. Despite this, it is satisfactory to see that there are only three trades, which is less than on April 19 and 22, 2010, and that trading occurs when price conditions are favourable.

4.5.5 Test 3: Sensitivity to Bid/Ask spread

In this last section, we are interested in the sensitivity of the results to the bid/ask spread, determined here by the two parameters κ_a and κ_b. More precisely, we look at the dominant effect between the spread and the multiplicative price impact through the parameter λ.

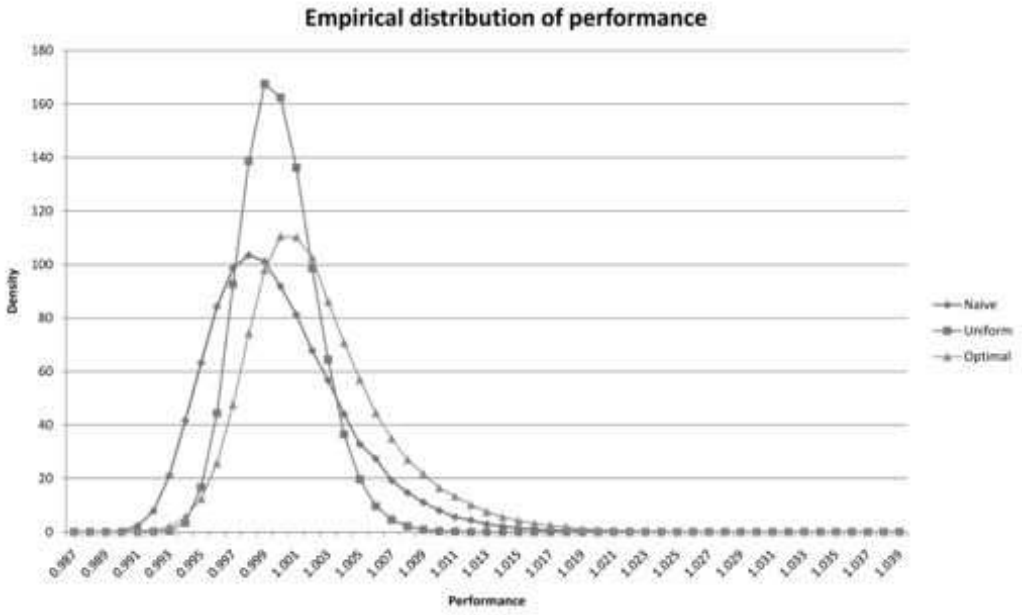

Figure 4.5: Test 2: Strategy empirical distribution

We proceeded to two tests here: one without bid/ask spread, i.e. $\kappa_a = \kappa_b = 1$ and with $\lambda = 5.10^{-4}$ as before, and one with a spread of 0.2% and a price impact parameter $\lambda = 0$.

Parameters The table 4.12 shows the parameters of the two tests. We only changed the impact and spread parameters and let the others be identical.

Performance Analysis In table 4.13 we computed several statistics on the results. In figure 4.10 we plotted the empirical distribution of performance in the two tests, with the test 2 distribution (Cf. figure 4.5) serving as a reference. In figure 4.11 we plotted the empirical distribution of the number of trades in the two tests, which is helpful for interpreting the results. Indeed, we observe from figure 4.11 that increasing the spread reduces the number of trades of the optimal strategy. Intuitively, the more frequently a strategy trades, the smaller its standard deviation: for example, the limiting case of the uniform strategy achieve the smallest standard deviation in our benchmark, and the naive strategy, that trades only once, the biggest. Qualitatively speaking, the standard deviation increases when the number of trades decreases: this help us explain qualitatively why the

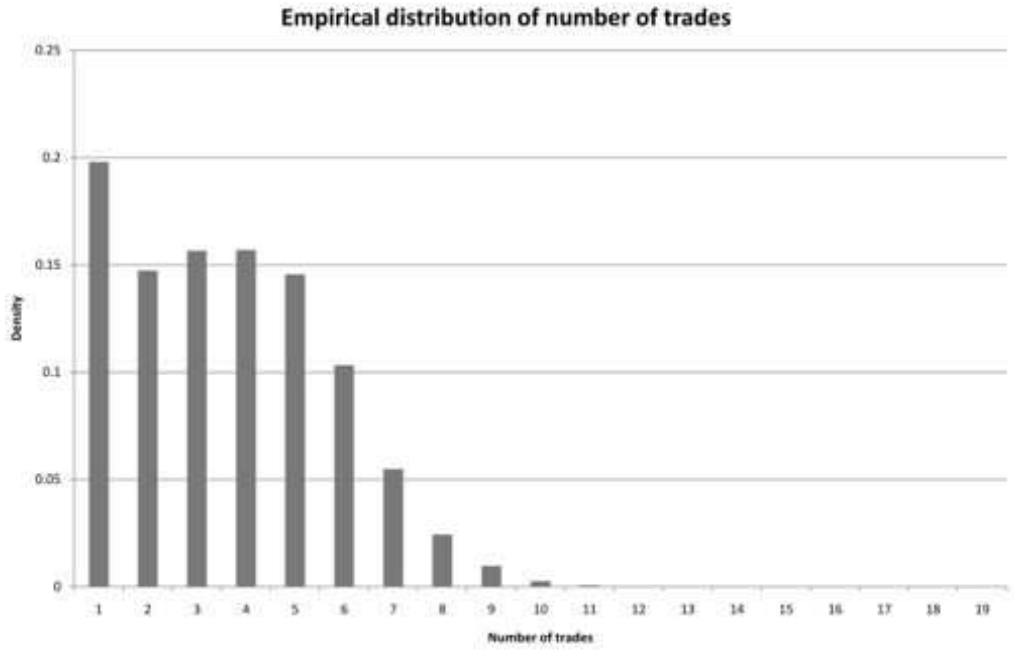

Figure 4.6: Test 2: Empirical distribution of the number of trades

standard deviation is higher in the case of a large spread (we used $\kappa_a - \kappa_b = 20\ bps$, which is much larger than usually observed in equity markets). Now, to provide an interpretation of why the optimal strategy trades less frequently when the spread is large, we can note two facts. First, in the large spread test, we considered that $\lambda = 0$, in other words that there is no market impact. Therefore, any trading rate ξ/θ will lead to same transaction price: this explain the clustering effect: the optimal strategy tends to trade a bigger quantity of assets at the same time to match terminal liquidation constraint. Second, a large spread will penalize strategies that can both buy and sell, and in particular the optimal strategy. Indeed, let us consider the typical scale of quantities involved in our optimization: we expect the optimal strategy to profit from price variation at the scale of 1 EUR in our example; if the spread is about 0.1 EUR, like in our last example, and if we usually do about 10 trades on the liquidation period, the effect of the spread (10×0.1 EUR$= 1$ EUR) is at the same scale as the price fluctuation. Therefore, the larger the spread, the more the optimal strategy tends to be one-sided, i.e. trading quantities (ξ_n) tends to be negative. Due to this phenomenon, the profit from optimal trading reduces with the spread, and the optimization becomes less efficient in this one-sided setup. This is consistent with the

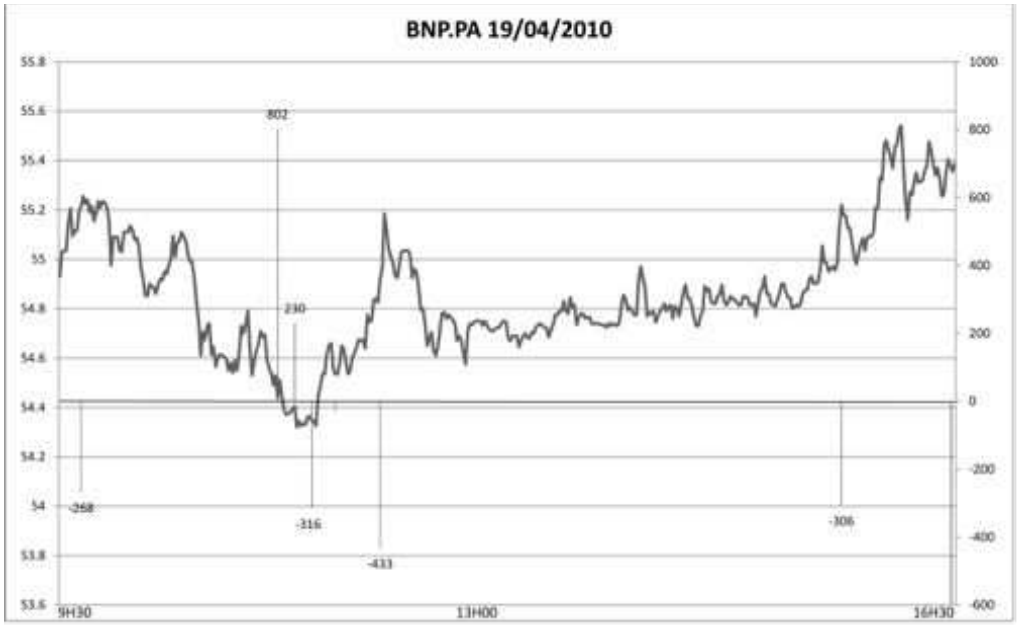

Figure 4.7: Test 2: Strategy realization on the BNP.PA stock the 04/19/2010.

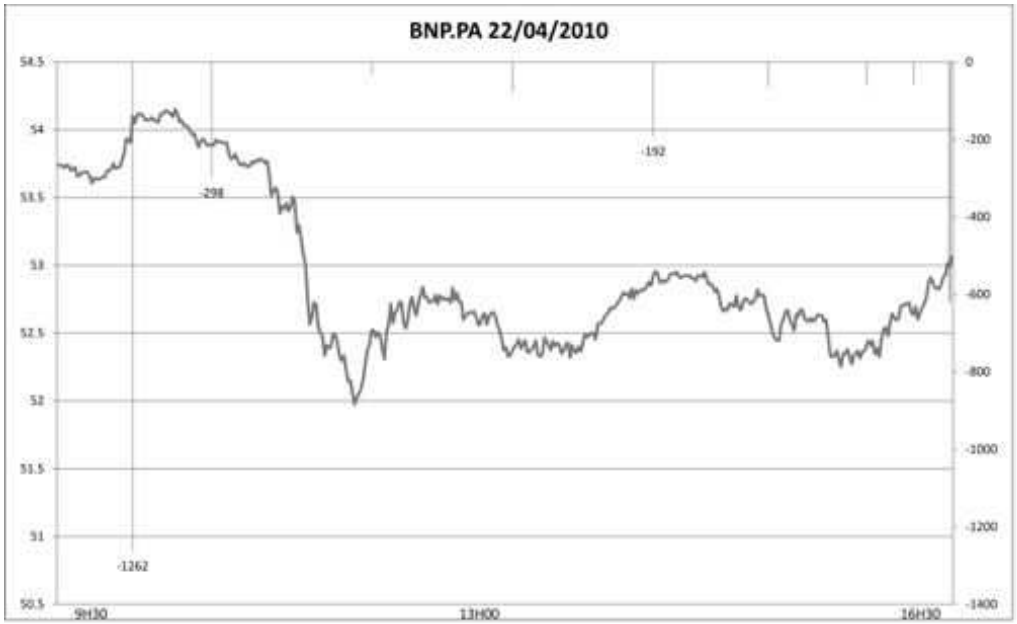

Figure 4.8: Test 2: Strategy realization on the BNP.PA stock the 04/22/2010.

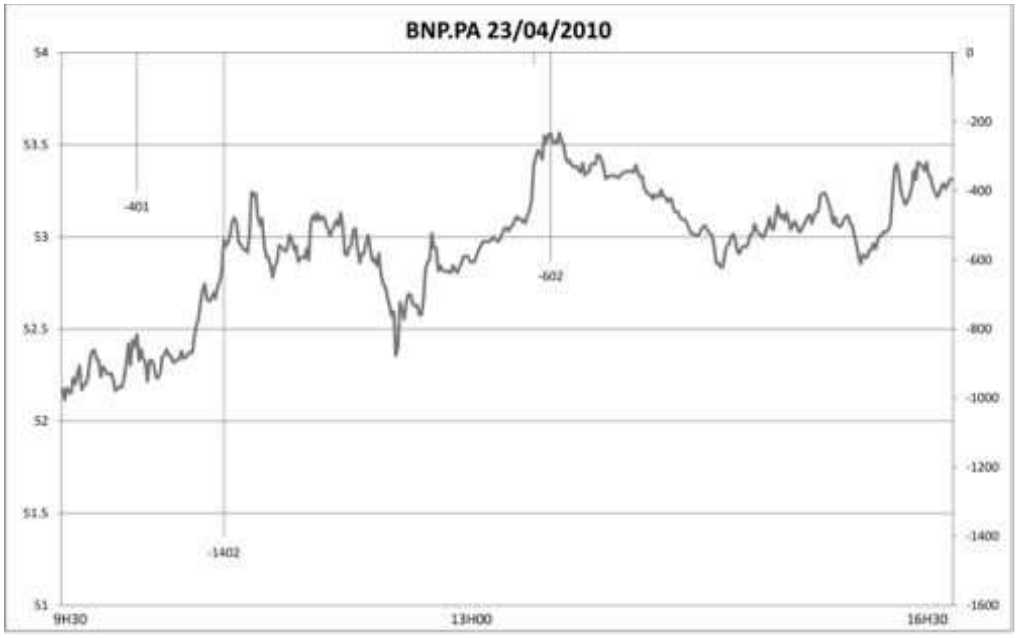

Figure 4.9: Test 2: Strategy realization on the BNP.PA stock the 04/23/2010.

Parameter	No spread test	No impact test	Parameter	No spread test	No impact test
Maturity	1 Day	1 Day	X_0	20000	20000
λ	5.00E-04	0	Y_0	2500	2500
β	0.2	0	P_0	51	51
γ	0.5	0.5	x_{min}	-20000	-20000
κ_a	1	1.001	x_{max}	200000	200000
κ_b	1	0.999	y_{min}	0	0
ϵ	0.001	0.001	y_{max}	5000	5000
b	0.01	0.01	p_{min}	49	49
σ	0.25	0.25	p_{max}	53	53
			n	30	30
			m	40	40
			N	100	100
			Q	10^5	10^5

Table 4.12: Test 3: Parameters

financial viewpoint: an investor that can both buy and sell have opportunities to profit from price fluctuations, whereas an investor that can only sell may only have opportunities

to sell at high price; therefore the number of trades decreases as the optimal strategy tends to be one-sided. Finally, we observe that both spread and non-linear impact influence the trading schedule. We also expect that the optimal quantity ξ_n to trade at date τ_n is influenced directly by the non-linear impact parameter λ.

Quantity	No spread test	No impact test	No spread vs. T2	No impact vs. T2
Mean Utility	1.00113	1.00025	$-3.00.10^{-5}$	$-9.08.10^{-4}$
Mean Performance	1.00227	1.00053	$-5.98.10^{-5}$	$-1.80.10^{-3}$
Standard Deviation	0.00432	0.00906	$-9.17.10^{-3}$	1.078

Table 4.13: Test 3: Statistics. In the two last columns "No spread vs. T2" (resp."No impact vs. T2") are shown the relative values of "No spread" test (resp. "No impact" test) against the values of test 2 of the preceding section.

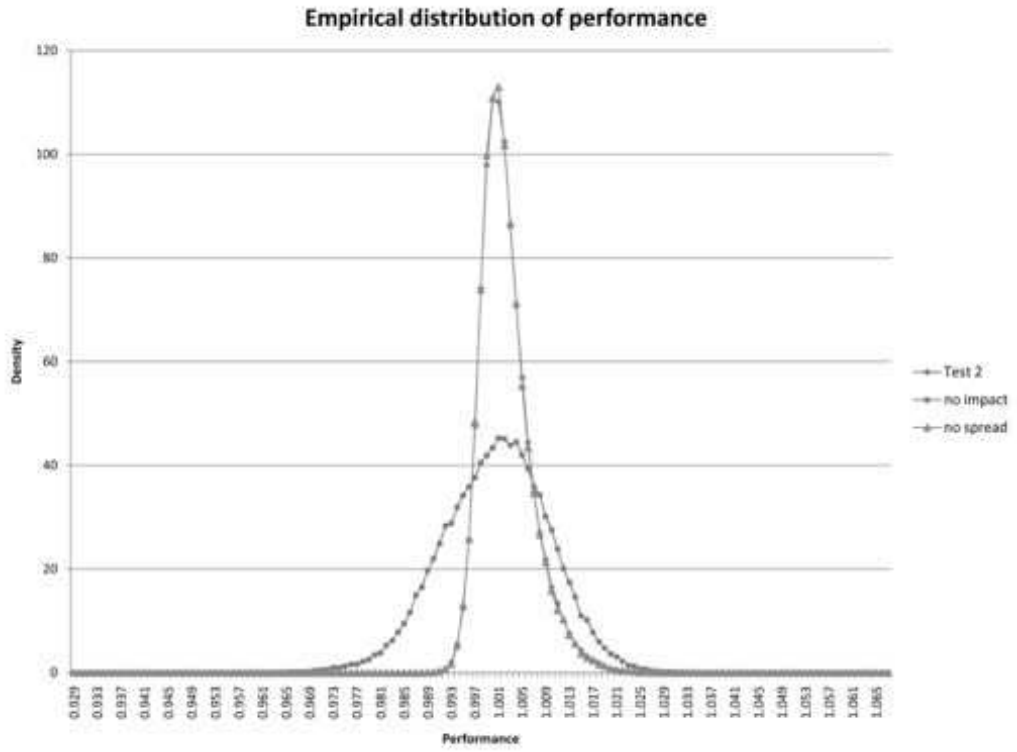

Figure 4.10: Test 3: Empirical distributions of performance

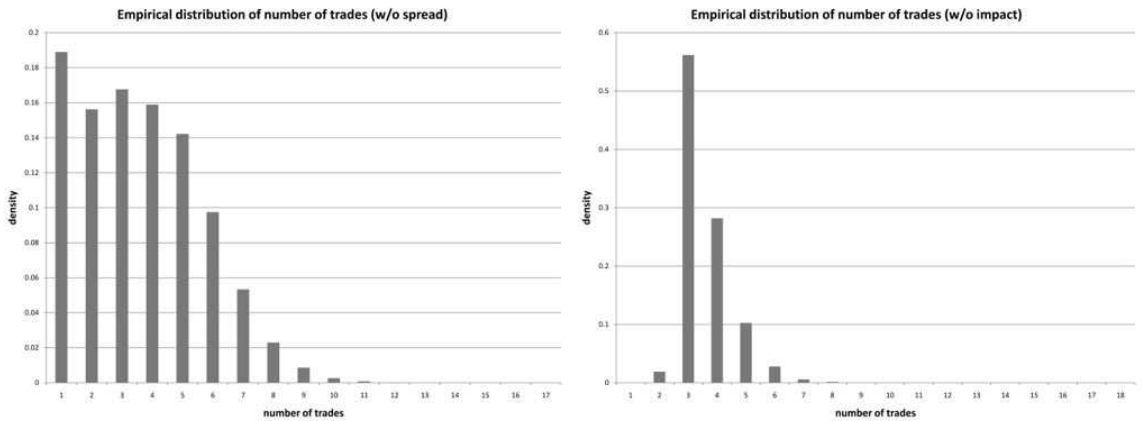

(a) No spread (b) No impact

Figure 4.11: Test 3: Empirical distributions of number of trades

Chapter 5

Optimal high frequency trading with limit and market orders

We propose a framework for studying optimal market making policies in a limit order book (LOB). The bid-ask spread of the LOB is modelled by a tick-valued continuous time Markov chain. We consider an agent who continuously submits limit buy/sell orders at best bid/ask quotes, and may also set limit orders at best bid (resp. ask) plus (resp. minus) a tick for getting the execution priority. The agent faces an execution risk since her limit orders are executed only when they meet counterpart market orders. She is also subject to the inventory risk due to price volatility when holding the risky asset. Then the agent can also choose to trade with market orders, and therefore get immediate execution, but at a less favorable price.

The objective of the market maker is to maximize her expected utility from revenue over a finite horizon, while controlling her inventory position. This is formulated as a mixed regime switching regular/impulse control problem that we characterize in terms of quasi-variational system by dynamic programming methods.

Calibration procedures are derived for fitting the market model. We provide an explicit backward splitting scheme for solving the problem, and show how it can be reduced to a system of simple equations involving only the inventory and spread variables. Several computational tests are performed both on simulated and real data.

Note: This chapter is adapted from the article [37] Guilbaud F. and H. Pham (2011): "Optimal high frequency trading with limit and market orders", to appear in *Quantitative Finance*.

5.1 Introduction

Most of modern equity exchanges are organized as *order driven* markets. In such type of markets, the price formation exclusively results from operating a *limit order book* (LOB), an order crossing mechanism where *limit orders* are accumulated while waiting to be matched with incoming *market orders*. Any market participant is able to interact with the LOB by posting either market orders or limit orders[1].

In this context, *market making* is a class of strategies that consists in simultaneously posting limit orders to buy and sell during the continuous trading session. By doing so, market makers provide counterpart to any incoming market orders: suppose that an investor A wants to sell one share of a given security at time t and that an investor B wants to buy one share of this security at time $t' > t$; if both use market orders, the economic role of the market maker C is to buy the stock as the counterpart of A at time t, and carry until date t' when she will sell the stock as a counterpart of B. The revenue that C obtains for providing this service to final investors is the difference between the two quoted prices at ask (limit order to sell) and bid (limit order to buy), also called the market maker's spread. This role was traditionally fulfilled by specialist firms, but, due to widespread adoption of electronic trading systems, any market participant is now able to compete for providing liquidity. Moreover, as pointed out by empirical studies (e.g. [53],[40]) and in a recent review [34] from AMF, the French regulator, this renewed competition among liquidity providers causes reduced effective market spreads, and therefore reduced indirect costs for final investors.

Empirical studies (e.g. [53]) also described stylized features of market making strategies. First, market making is typically not directional, in the sense that it does not profit from security price going up or down. Second, market makers keep almost no overnight position, and are unwilling to hold any risky asset at the end of the trading day. Finally, they manage to maintain their *inventory*, i.e. their position on the risky asset close to zero during the trading day, and often equilibrate their position on several distinct marketplaces, thanks to the use of high-frequency order sending algorithms. Estimations of total annual profit for this class of strategy over all U.S. equity market were around 10 G$ in 2009 [34]. Another important aspect of empirical litterature is high-frequency data modelling and estimation, a field surveyed in the forthcoming volume [26]. Typically, this literature investigates such topics as designing methodologies to discover elasticity and plasticity of price evolution [14],

[1] A market order of size m is an order to buy (sell) m units of the asset being traded at the lowest (highest) available price in the market, its execution is immediate; a limit order of size ℓ at price q is an order to buy (sell) ℓ units of the asset being traded at the specified price q, its execution is uncertain and achieved only when it meets a counterpart market order. Given a security, the *best bid* (resp. *ask*) price is the highest (resp. lowest) price among limit orders to buy (resp. to sell) that are active in the LOB. The *spread* is the difference, expressed in numéraire per share, of the best ask price and the best bid price, positive during the continuous trading session (see [33]).

and therefore allowing HFT to use persistence of some price properties; risk management in a high frequency setup; but also constructing microstructure simulation models [66], which are relevant for HFT strategies design or backtesting.

Popular models of market making strategies were set up using a risk-reward approach. Two distinct sources of risk are usually identified: the inventory risk, and the execution risk. In the early 1980's, the paper [5] contributes to electronic market design, attempting to allow the marketplace to provide liquidity automatically, and suggests that market-making can be seen as an inventory management problem. The *inventory risk* [7] is comparable to the market risk, i.e. the risk of holding a long or short position on a risky asset. Moreover, due to the uncertain execution of limit orders, market makers only have partial control on their inventory, and therefore the inventory has a stochastic behavior. The *execution risk* is the risk that limit orders may not be executed, or be partially executed [45]. Indeed, given an incoming market order, the matching algorithm of LOB determines which limit orders are to be executed according to a price/time priority[2], and this structure fundamentally impacts the dynamics of executions. We also mention a third type of risk, the so-called *adverse selection risk*, popular in economic and econometric litterature. This is the risk that market price unfavourably deviates for a short time period, from the market maker point of view, after their quote was taken. This type of risk appears naturally in models where the market orders flow contains information about the fundamental asset value (e.g. [29]).

Some of these risks were studied in previous works. The seminal work [7] provided a framework to manage inventory risk in a stylized LOB. The market maker objective is to maximize the expected utility of her terminal profit, in the context of limit orders executions occurring at jump times of Poisson processes. This model shows its efficiency to reduce inventory risk, measured via the variance of terminal wealth, against the symmetric strategy. Several extensions and refinement of this setup can be found in recent litterature: [35] provides simplified solution to the backward optimization, an in-depth discussion of its characteristics and an application to the liquidation problem. In [9], the authors develop a closely related model to solve a liquidation problem, and study continuous limit case. The paper [16] provides a way to include more precise empirical features to this framework by embedding a hidden Markov model for high frequency dynamics of LOB. Some aspects of the execution risk were also studied previously, mainly by considering the trade-off between passive and aggressive execution strategies. In [45], the authors solve the Merton's portfolio optimization problem in the case where the investor can choose between market orders or limit orders; in [67], [68], the possibility to use market orders in addition to limit orders is also taken into account, in the context of market making in the foreign exchange market.

[2]A different type of LOB operates under *pro-rata* priority, e.g. for some futures on interest rates. In this paper, we do not consider this case and focus on the main mechanism used in equity market.

Yet the relation between execution risk and the microstructure of the LOB, and especially the price/time priority is, so far, poorly investigated.

In this paper we develop a new model to address the execution and inventory risks in market making. The stock mid-price is driven by a general Markov process, and we model the market spread as a discrete Markov chain that jumps according to a stochastic clock. Therefore, the spread takes discrete values in the price grid, multiple of the tick size. We allow the market maker to trade both via limit orders, whose execution is uncertain, and via market orders, whose execution is immediate but costly. The market maker can post limit orders at best quote or improve this quote by one tick. In this last case, she hopes to capture market order flow of agents who are not yet ready to trade at the best bid/ask quote. Therefore, she faces a trade off between waiting to be executed at the current best price, or improve this best price, and then be more rapidly executed but at a less favorable price. We model the limit orders strategy as continuous controls, due to the fact that these orders can be updated at high frequency at no cost. On the contrary, we model the market orders strategy as impulse controls that can only occur at discrete dates. We also include fixed, per share or proportional fees or rebates coming with each execution. Execution processes, counting the number of executed limit orders, are modelled as Cox processes with intensity depending both on the market maker's controls and on the bid/ask spread: therefore, we consider that execution intensities are conditional to the state of the LOB, in the same vein as in [21], and we assume that the main variable of interest is the spread. In this context, we optimize the expected utility from profit over a finite time horizon, by choosing optimally between limit and market orders, while controlling the inventory position. We study in detail classical frameworks including mean-variance criterion and exponential utility criterion.

The outline of this paper is as follows. In section 2, we detail the model, and comment its features. We also provide direct calibration methods for all quantities involved in our model. We formulate in Section 3 the optimal market making control problem and derive the associated Hamilton-Jacobi-Bellman quasi variational inequality (HJBQVI) from dynamic programming principle. Section 4 is devoted to the numerical scheme for solving the HJBQVI and computing the optimal policy. We also examine several situations, where we are able simplify this algorithm by reducing the number of state variables to the inventory and spread. In section 5, we provide some numerical results and an empirical performance analysis for our computational scheme.

5.2 A market-making model

5.2.1 Mid price and spread process

Let us fix a probability space $(\Omega, \mathcal{F}, \mathbf{P})$ equipped with a filtration $\mathbb{F} = (\mathcal{F}_t)_{t\geq 0}$ satisfying the usual conditions. It is assumed that all random variables and stochastic processes are defined on the stochastic basis $(\Omega, \mathcal{F}, \mathbb{F}, \mathbf{P})$.

The mid-price of the stock is an exogenous Markov process P, with infinitesimal generator $\mathcal{L}_P$ and state space $\mathbb{P}$. For example, P is a Lévy process or an exponential of Lévy process (including Black-Scholes-Merton model with jumps). In the limit order book (LOB) for this stock, we consider a stochastic bid-ask spread resulting from the behaviour of market participants, taking discrete values, which are finite multiple of the tick size $\delta > 0$, and jumping at random times. This is modelled as follows: we first consider the tick time clock associated to a Poisson process $(N_t)_t$ with deterministic intensity $\lambda(t)$, for taking into account intra-day seasonnality, and representing the random times where the buy and sell orders of participants in the market affect the bid-ask spread. We next define a discrete-time stationary Markov chain $(\hat{S}_n)_{n\in\mathbb{N}}$, valued in the finite state space $\mathbb{S} = \delta\mathbb{I}_m$, $\mathbb{I}_m := \{1, \ldots, m\}$, $m \in \mathbb{N}\setminus\{0\}$, with probability transition matrix $(\rho_{ij})_{1\leq i,j\leq M}$, i.e. $\mathbf{P}[\hat{S}_{n+1} = j\delta | \hat{S}_n = i\delta] = \rho_{ij}$, s.t. $\rho_{ii} = 0$, independent of N, and representing the random spread in tick time. The spread process $(S_t)_t$ in calendar time is then defined as the time-change of $\hat{S}$ by N, i.e.

$$S_t \;=\; \hat{S}_{N_t}, \quad t \geq 0. \tag{5.2.1}$$

Hence, $(S_t)_t$ is a continuous time (inhomogeneous) Markov chain with intensity matrix $R(t) = (r_{ij}(t))_{1\leq i,j\leq m}$, where $r_{ij}(t) = \lambda(t)\rho_{ij}$ for $i \neq j$, and $r_{ii}(t) = -\sum_{j\neq i} r_{ij}(t)$. We assume that S and P are independent. The best-bid and best-ask prices are defined by: $P^b_t = P_t - \frac{S_t}{2}$, $P^a_t = P_t + \frac{S_t}{2}$.

5.2.2 Trading strategies in the limit order book

We consider an agent (market maker), who trades the stock using either limit orders or market orders. She may submit limit buy (resp. sell) orders specifying the quantity and the price she is willing to pay (resp. receive) per share, but will be executed only when an incoming sell (resp. buy) market order is matching her limit order. Otherwise, she can post market buy (resp. sell) orders for an immediate execution, but, in this case obtain the opposite best quote, i.e. trades at the best-ask (resp. best bid) price, which is less favorable.

Limit orders strategies. The agent may submit at any time limit buy/sell orders at the current best bid/ask prices (and then has to wait an incoming counterpart market order matching her limit), but also control her own bid and ask price quotes by placing buy (resp. sell) orders at a marginal higher (resp. lower) price than the current best bid (resp. ask),

i.e. at $P_t^{b+} := P_t^b + \delta$ (resp. $P_t^{a-} := P_t^a - \delta$). Such an alternative choice is used in practice by a market maker to capture market orders flow of undecided traders at the best quotes, hence to get priority in the order execution w.r.t. limit order at current best/ask quotes, and can be taken into account in our modelling with discrete spread of tick size δ.

There is then a tradeoff between a larger performance for a quote at the current best bid (resp. ask) price, and a smaller performance for a quote at a higher bid price, but with faster execution. The submission and cancellation of limit orders are for free, as they provide liquidity to the market, and are thus stimulated. Actually, market makers receive some fixed rebate once their limit orders are executed. The agent is assumed to be small in the sense that she does not influence the bid-ask spread. The limit order strategies are then modelled by a continuous time predictable control process:

$$\alpha_t^{make} \quad = \quad (Q_t^b, Q_t^a, L_t^b, L_t^a), \quad t \geq 0,$$

where $L = (L^b, L^a)$ valued in $[0, \bar{\ell}]^2$, $\bar{\ell} > 0$, represents the size of the limit buy/sell order, and $Q = (Q^b, Q^a)$ represent the possible choices of the bid/ask quotes either at best or at marginally improved prices, and valued in $\mathcal{Q} = \mathcal{Q}^b \times \mathcal{Q}^a$, with $\mathcal{Q}^b = \{Bb, Bb_+\}$, $\mathcal{Q}^a = \{Ba, Ba_-\}$:

- Bb: best-bid quote, and Bb_+: bid quote at best price plus the tick
- Ba: best-ask quote, and Ba_-: ask quote at best price minus the tick

Notice that when the spread is equal to one tick δ, a bid quote at best price plus the tick is actually equal to the best ask, and will then be considered as a buy market order. Similarly, an ask quote at best price minus the tick becomes a best bid, and is then viewed as a sell market order. In other words, the limit orders $Q_t = (Q_t^b, Q_t^a)$ take values in $\mathcal{Q}(S_{t^-})$, where $\mathcal{Q}(s) = \mathcal{Q}^b \times \mathcal{Q}^a$ when $s > \delta$, $\mathcal{Q}(s) = \{Bb\} \times \{Ba\}$ when $s = \delta$. We shall denote by $\mathcal{Q}_i^b = \mathcal{Q}^b$ for $i > 1$, and $\mathcal{Q}_i^b = \{Bb\}$ for $i = 1$, and similarly for $\mathcal{Q}_i^a$ for $i \in \mathbb{I}_m$.

We denote at any time t by $\pi^b(Q_t^b, P_t, S_t)$ and $\pi^a(Q_t^a, P_t, S_t)$ the bid and ask prices of the market maker, where the functions π^b (resp. π^a) are defined on $\mathcal{Q}^b \times \mathbb{P} \times \mathbb{S}$ (resp. $\mathcal{Q}^a \times \mathbb{P} \times \mathbb{S}$) by:

$$\pi^b(q^b, p, s) \quad = \quad \begin{cases} p - \frac{s}{2}, & \text{for } q^b = Bb \\ p - \frac{s}{2} + \delta & \text{for } q^b = Bb_+. \end{cases}$$

$$\pi^a(q^a, p, s) \quad = \quad \begin{cases} p + \frac{s}{2}, & \text{for } q^a = Ba \\ p + \frac{s}{2} - \delta & \text{for } q^a = Ba_-. \end{cases}$$

We shall denote by $\pi_i^b(q^b, p) = \pi^b(q^b, p, s)$, $\pi_i^a(q^a, p) = \pi^a(q^a, p, s)$ for $s = i\delta$, $i \in \mathbb{I}_m$.

Remark 5.2.1 One can take into account proportional rebates received by the market makers, by considering; $\pi^b(q^b, p, s) = (p - \frac{s}{2} + \delta 1_{q^b = Bb_+})(1 - \rho)$, $\pi^a(q^a, p, s) = (p + \frac{s}{2} -$

$\delta 1_{q^a=Ba_-})(1+\rho)$, for some $\rho \in [0,1)$, or per share rebates with: $\pi^b(q^b,p,s) = p - \frac{s}{2} + \delta 1_{q^b=Bb_+} - \rho$, $\pi^a(q^a,p,s) = p + \frac{s}{2} - \delta 1_{q^a=Ba_-} + \rho$, for some $\rho > 0$.

The limit orders of the agent are executed when they meet incoming counterpart market orders. Let us then consider the arrivals of market buy and market sell orders, which completely match the limit sell and limit buy orders of the small agent, and modelled by independent Cox processes N^a and N^b. The intensity rate of N^a_t is given by $\lambda^a(Q^a_t, S_t)$ where λ^a is a deterministic function of the limit quote sell order, and of the spread, satisfying $\lambda^a(Ba,s) < \lambda^a(Ba_-,s)$. This natural condition conveys the price/priority in the order execution in the sense that an agent quoting a limit sell order at ask price P^{a-} will be executed before traders at the higher ask price P^a, and hence receive more often market buy orders. Typically, one would also expect that λ^a is nonincreasing w.r.t. the spread, which means that the larger is the spread, the less often the market buy orders arrive. Likewise, we assume that the intensity rate of N^b_t is given by $\lambda^b(Q^b_t, S_t)$ where λ^b is a deterministic function of the spread, and $\lambda^b(Bb,s) < \lambda^b(Bb_+,s)$. We shall denote by $\lambda^a_i(q^a) = \lambda^a(q^a,s)$, $\lambda^b_i(q^b) = \lambda^b(q^b,s)$ for $s = i\delta$, $i \in \mathbb{I}_m$.

For a limit order strategy $\alpha^{make} = (Q^b, Q^a, L^b, L^a)$, the cash holdings X and the number of shares Y hold by the agent (also called inventory) follow the dynamics

$$
\begin{aligned}
dY_t &= L^b_t dN^b_t - L^a_t dN^a_t, && (5.2.2)\\
dX_t &= -\pi^b(Q^b_t, P_{t^-}, S_{t^-}) L^b_t dN^b_t + \pi^a(Q^a_t, P_{t^-}, S_{t^-}) L^a_t dN^a_t. && (5.2.3)
\end{aligned}
$$

Market order strategies. In addition to market making strategies, the investor may place market orders for an immediate execution reducing her inventory. The submissions of market orders, in contrast to limit orders, take liquidity in the market, and are usually subject to fees. We model market order strategies by an impulse control:

$$
\alpha^{take} = (\tau_n, \zeta_n)_{n\geq 0},
$$

where (τ_n) is an increasing sequence of stopping times representing the market order decision times of the investor, and ζ_n, $n \geq 1$, are $\mathcal{F}_{\tau_n}$-measurable random variables valued in $[-\bar{e}, \bar{e}]$, $\bar{e} > 0$, and giving the number of stocks purchased at the best-ask price if $\zeta_n \geq 0$, or selled at the best-bid price if $\zeta_n < 0$ at these times. Again, we assumed that the agent is small so that her total market order will be executed immediately at the best bid or best ask price. In other words, we only consider a linear market impact, which does not depend on the order size. When posting a market order strategy, the cash holdings and the inventory jump at times τ_n by:

$$
\begin{aligned}
Y_{\tau_n} &= Y_{\tau_n^-} + \zeta_n, && (5.2.4)\\
X_{\tau_n} &= X_{\tau_n^-} - c(\zeta_n, P_{\tau_n}, S_{\tau_n}) && (5.2.5)
\end{aligned}
$$

where

$$c(e,p,s) \quad = \quad ep + |e|\frac{s}{2} + \varepsilon$$

represents the (algebraic) cost function indicating the amount to be paid immediately when passing a market order of size e, given the mid price p, a spread s, and a fixed fee $\varepsilon > 0$. We shall denote by $c_i(e,p) = c(e,p,s)$ for $s = i\delta$, $i \in \mathbb{I}_m$.

Remark 5.2.2 One can also include proportional fees $\rho \in [0,1)$ paid at each market order trading by considering a cost function in the form: $c(e,p,s) = (e+\varepsilon|e|)p + (|e|+\rho e)\frac{s}{2} + \varepsilon$, or fixed fees per share with $c(e,p,s) = ep + |e|(\frac{s}{2}+\rho) + \varepsilon$.

In the sequel, we shall denote by $\mathcal{A}$ the set of all limit/market order trading strategies $\alpha = (\alpha^{make}, \alpha^{take})$.

5.2.3 Market making problem

The objective of the market maker is the following. She wants to maximize over a finite horizon T the profit from her transactions in the LOB, while keeping under control her inventory (usually starting from zero), and getting rid of her inventory at the terminal date:

$$\text{maximize} \quad \mathbf{E}\big[U(X_T) - \gamma \int_0^T g(Y_t)dt\big] \tag{5.2.6}$$

over all limit/market order trading strategies $\alpha = (\alpha^{make}, \alpha^{take})$ in $\mathcal{A}$ such that $Y_T = 0$. Here U is an increasing reward function, γ is a nonnegative constant, and g is a nonnegative convex function, so that the last integral term $\int_0^T g(Y_t)dt$ penalizes the variations of the inventory. Typical frameworks include the two following cases:

Mean-quadratic criterion: $U(x) = x$, $\gamma > 0$, $g(y) = y^2$.

Exponential utility maximization: $U(x) = -\exp(-\eta x)$, $\gamma = 0$.

We shall investigate in more detail these two important cases, which lead to nice simplifications for the numerical resolution.

5.2.4 Parameters estimation

In most order-driven markets, available data are made up of *Level 1 data* that contain transaction prices and quantities at best quotes, and of *Level 2 data* containing the volume updates for the liquidity offered at the L first order book slices (L usually ranges from 5 to 10). In this section, we propose some direct methods for estimating the intensity of the spread Markov chain, and of the execution point processes, based only on the observation

of *Level 1 data*. This has the advantage of low computational cost, since we do not have to deal with the whole volume of *Level 2 data*.

Estimation of spread parameters. Assuming that the spread S is observable, let us define the jump times of the spread process:

$$\theta_0 \;=\; 0, \qquad \theta_{n+1} \;=\; \inf\left\{t > \theta_n \,:\, S_t \neq S_{t-}\right\} \;, \forall n \geq 1.$$

From these observable quantities, one can reconstruct the processes:

$$\begin{aligned} N_t &= \#\left\{\theta_j > 0 \,:\, \theta_j \leq t\right\} \;, t \geq 0, \\ \hat{S}_n &= S_{\theta_n} \;,\; n \geq 0. \end{aligned}$$

Then, our goal is to estimate the deterministic intensity of the Poisson process $(N_t)_t$, and the transition matrix of the Markov chain $(\hat{S}_n)_n$ from a path realization with high frequency data of the tick-time clock and spread in tick time over a finite trading time horizon T, typically of one day. From the observations of K samples of $\hat{S}_n$, $n = 1, \ldots, K$, and since the Markov chain $(\hat{S}_n)$ is stationary, we have a consistent estimator (when K goes to infinity) for the transition probability $\rho_{ij} := \mathbf{P}[\hat{S}_{n+1} = j\delta | \hat{S}_n = i\delta] = \mathbf{P}[(\hat{S}_{n+1}, \hat{S}_n) = (j\delta, i\delta)]/\mathbf{P}[\hat{S}_n = i\delta]$ given by:

$$\hat{\rho}_{ij} \;=\; \frac{\displaystyle\sum_{n=1}^{K} 1_{\{(\hat{S}_{n+1}, \hat{S}_n) = (j\delta, i\delta)\}}}{\displaystyle\sum_{n=1}^{K} 1_{\{\hat{S}_n = i\delta\}}} \tag{5.2.7}$$

For the estimation of the deterministic intensity function $\lambda(t)$ of the (non)homogeneous Poisson process (N_t), we shall assume in a first approximation a simple natural parametric form. For example, we may assume that λ is constant over a trading day, and more realistically for taking into account intra-day seasonality effects, we consider that the tick time clock intensity jumps e.g. every hour of a trading day. We then assume that λ is in the form:

$$\lambda(t) \;=\; \sum \lambda_k 1_{\{t_k \leq t < t_{k+1}\}}$$

where $(t_k)_k$ is a fixed and known increasing finite sequence of $\mathbb{R}_+$ with $t_0 = 0$, and $(\lambda_k)_k$ is an unknown finite sequence of $(0, \infty)$. In other words, the intensity is constant equal to λ_k over each period $[t_k, t_{k+1}]$, and by assuming that the interval length $t_{k+1} - t_k$ is large w.r.t. the intensity λ_k (which is the case for high frequency data), we have a consistent estimator of λ_k, for all k, and then of $\lambda(t)$ given by:

$$\hat{\lambda}_k \;=\; \frac{N_{t_{k+1}} - N_{t_k}}{t_{k+1} - t_k}. \tag{5.2.8}$$

We performed these two estimation procedures (5.2.7) and (5.2.8) on Société Générale (SOGN.PA) stock on April 18, 2011 between 9:30 and 16:30 in Paris local time. We used tick-by-tick level 1 data provided by Quanthouse, and fed via a OneTick Timeseries database. Number of data point were roughly 10^5.

In table 5.1, we display the estimated transition matrix: first row and column indicate the spread value $s = i\delta$ and the cell ij shows $\hat{\rho}_{ij}$. For this stock and at this date, the tick size was $\delta = 0.005$ euros, and we restricted our analysis to the first 6 values of the spread ($m = 6$) due to the small number of data outside this range: indeed, in our set, less than 1% of datapoints corresponded to a spread above 0.03. Note that this observation is valid, on Euronext Paris, only for stocks priced less than 50 EUR, since the tick size doubles (to 0.01 EUR) for stocks priced higher than 50 EUR. After truncating to $m \leq 6$ we performed a re-normalization in order to obtain a transition matrix.

spread	0.005	0.01	0.015	0.02	0.025	0.03
0.005	0	0.410	0.220	0.160	0.142	0.065
0.01	0.201	0	0.435	0.192	0.103	0.067
0.015	0.113	0.221	0	0.4582	0.147	0.059
0.02	0.070	0.085	0.275	0	0.465	0.102
0.025	0.068	0.049	0.073	0.363	0	0.446
0.03	0.077	0.057	0.059	0.112	0.692	0

Table 5.1: Estimation of the transition matrix (ρ_{ij}) for the underlying spread of the stock SOGN.PA on April 18, 2011.

In figure 5.1, we plot the tick time clock intensity by using an affine interpolation, and observed a typical U-pattern. This is consistent with the empirical observation that trading activity is more important in the beginning and at the end of the day trading session, and less active around noon, see [18]. A further step for the estimation of the intensity could be to specify a parametric form for the intensity function fitting U pattern, e.g. parabolic functions in time, and then use a maximum likelihood method for estimating the parameters.

Estimation of execution parameters. When performing a limit order strategy α^{make}, we suppose that the market maker permanently monitors her execution point processes N^a and N^b, representing respectively the number of arrivals of market buy and sell orders matching the limit orders for quote ask Q^a and quote bid Q^b. We also assume that there is no latency so that the observation of the execution processes is not noisy. Therefore, observable variables include the quintuplet:

$$(N_t^a, N_t^b, Q_t^a, Q_t^b, S_t) \in \mathbb{R}^+ \times \mathbb{R}^+ \times \mathcal{Q}^a \times \mathcal{Q}^b \times \mathbb{S}\,,\ t \in [0, T]$$

Moreover, since N^a and N^b are assumed to be independent, and both sides of the order

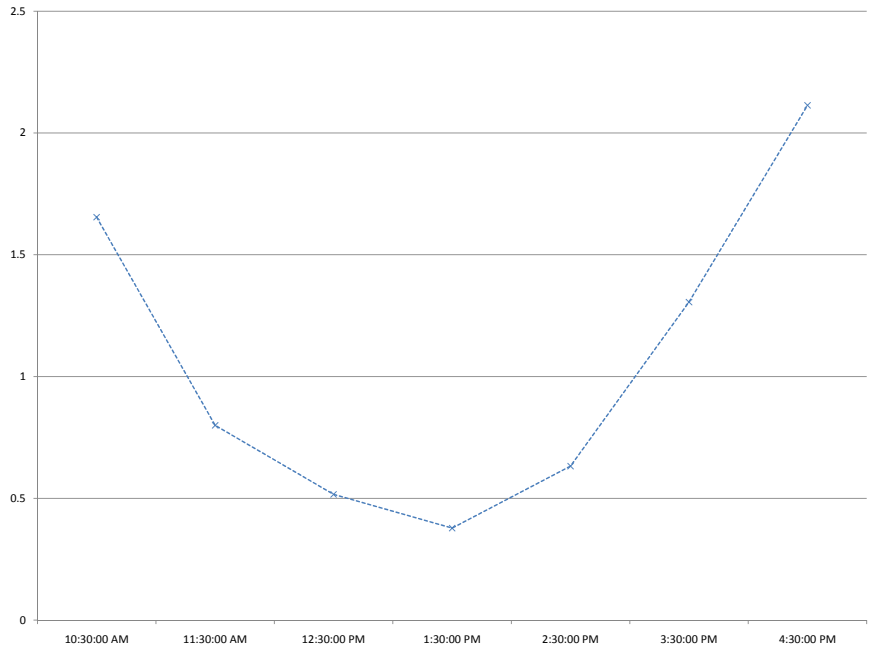

Figure 5.1: Plot of tick time clock intensity estimate for the stock SOGN.PA on April 18, 2011 expressed in second^{-1} (affine interpolation).

book can be estimated using the same procedure, we shall focus on the estimation for the intensity function $\lambda^b(q^b, s)$, $q^b \in \mathcal{Q}^b = \{Bb, Bb_+\}$, $s \in \mathbb{S} = \delta\mathbb{I}_m$, of the Cox process N^b.

The estimation procedure for $\lambda^b(q^b, s)$ basically matchs the intuition that one must count the number of executions at bid when the system was in the state (q^b, s) and normalize this quantity by the time spent in the state (q^b, s). This is justified mathematically as follows. For any $(q^b, s = i\delta) \in \mathcal{Q}^b \times \mathbb{S}$, let us define the point process

$$N_t^{b,q^b,i} \;=\; \int_0^t 1_{\{Q_u^b = q, S_{u-} = i\delta\}} dN_u^b, \quad t \geq 0,$$

which counts the number of jumps of N^b when (Q^b, S) was in state $(q^b, s = i\delta)$. Then, for any nonnegative predictable process (H_t), we have

$$\begin{aligned}
\mathbf{E}\Big[\int_0^\infty H_t dN_t^{b,q^b,i}\Big] &= \mathbf{E}\Big[\int_0^\infty H_t 1_{\{Q_t^b = q^b, S_{t-} = i\delta\}} dN_t^b\Big] \\
&= \mathbf{E}\Big[\int_0^\infty H_t 1_{\{Q_t^b = q^b, S_{t-} = i\delta\}} \lambda^b(Q_t^b, S_t) dt\Big] \\
&= \mathbf{E}\Big[\int_0^\infty H_t 1_{\{Q_t^b = q^b, S_{t-} = i\delta\}} \lambda_i^b(q^b) dt\Big], \qquad (5.2.9)
\end{aligned}$$

where we used in the second equality the fact that $\lambda^b(Q_t^b, S_t)$ is the intensity of N^b. The relation (5.2.9) means that the point process $N^{b,q^b,i}$ admits for intensity $\lambda_i^b(q^b)1_{\{Q_t^b=q^b,S_{t-}=i\delta\}}$. By defining

$$\mathcal{T}_t^{b,q^b,i} \;=\; \int_0^t 1_{\{Q_u^b=q,S_{u-}=i\delta\}}du$$

as the time that (Q^b, S) spent in the state $(q^b, s = i\delta)$, this means equivalently that the process $M_t^{b,q^b,i} = N^{b,q^b,i}_{A_t^{b,q^b,i}}$, where $A_t^{b,q^b,i} = \inf\{u \geq 0 : \mathcal{T}_u^{b,q^b,i} \geq t\}$ is the càd-làg inverse of $\mathcal{T}^{b,q^b,i}$, is a Poisson process with intensity $\lambda_i(q^b)$. By assuming that $\mathcal{T}_T^{b,q^b,i}$ is large w.r.t. $\lambda_i(q^b)$, which is the case when $(\hat{S}_n)$ is irreducible (hence recurrent), and for high-frequency data over $[0, T]$, we have a consistent estimator of $\lambda_i^b(q^b)$ given by:

$$\hat{\lambda}_i^b(q^b) \;=\; \frac{N_T^{b,q^b,i}}{\mathcal{T}_T^{b,q^b,i}}. \tag{5.2.10}$$

Similarly, we have a consistent estimator for $\lambda_i^a(q^a)$ given by:

$$\hat{\lambda}_i^a(q^a) \;=\; \frac{N_T^{a,q^a,i}}{\mathcal{T}_T^{a,q^a,i}}, \tag{5.2.11}$$

where $N_T^{a,q^a,i}$ counts the number of executions at ask quote q^a and for a spread $i\delta$, and $\mathcal{T}_T^{a,q^a,i}$ is the time that (Q^a, S) spent in the state $(q^a, s = i\delta)$ over $[0, T]$.

Let us now illustrate this estimation procedure on real data, with the same market data as above, i.e. tick-by-tick level 1 for SOGN.PA on April 18, 2011, provided by Quanthouse via OneTick timeseries database. Actually, since we did not perform the strategy on this real-world order book, we could not observe the real execution processes N^b and N^a. We built thus simple proxies $\tilde{N}^{b,q^b,i}$ and $\tilde{N}^{a,q^a,i}$, for $q^b = Bb, Bb_+$, $q^a = Ba, Ba_-$, $i = 1, \ldots, m$, based on the following rules. Let us also assume that in addition to $(S_{\theta_n})_n$, we observe at jump times θ_n of the spread, the volumes $(V_{\theta_n}^a, V_{\theta_n}^b)$ offered at the best ask and best bid price in the LOB together with the cumulated market order quantities $\vartheta_{\theta_{n+1}}^{BUY}$ and $\vartheta_{\theta_{n+1}}^{SELL}$ arriving between two consecutive jump times θ_n and θ_{n+1} of the spread, respectively at best ask price and best bid price. We finally fix an arbitrarily typical volume V_0, e.g. $V_0 = 100$ of our limit orders, and define the proxys $\tilde{N}^{b,q^b,i}$ and $\tilde{N}^{a,q^a,i}$ at times θ_n by:

$$\begin{aligned}
\tilde{N}_{\theta_{n+1}}^{b,Bb_+,i} &= \tilde{N}_{\theta_n}^{b,Bb_+,i} + 1_{\left\{V_0<\vartheta_{\theta_{n+1}}^{SELL},S_{\theta_n}=i\delta\right\}},\ \tilde{N}_0^{b,Bb_+,i} = 0\\
\tilde{N}_{\theta_{n+1}}^{b,Bb,i} &= \tilde{N}_{\theta_n}^{b,Bb,i} + 1_{\left\{V_0+V_{\theta_n}^b<\vartheta_{\theta_{n+1}}^{SELL},S_{\theta_n}=i\delta\right\}},\ \tilde{N}_0^{b,Bb,i} = 0\\
\tilde{N}_{\theta_{n+1}}^{a,Aa_-,i} &= \tilde{N}_{\theta_n}^{a,Aa_-,i} + 1_{\left\{V_0<\vartheta_{\theta_{n+1}}^{BUY},S_{\theta_n}=i\delta\right\}},\ \tilde{N}_0^{a,Aa_-,i} = 0\\
\tilde{N}_{\theta_{n+1}}^{a,Aa,i} &= \tilde{N}_{\theta_n}^{b,Aa,i} + 1_{\left\{V_0+V_{\theta_n}^a<\vartheta_{\theta_{n+1}}^{BUY},S_{\theta_n}=i\delta\right\}},\ \tilde{N}_0^{a,Aa,i} = 0,
\end{aligned}$$

together with a proxy for the time spent in spread $i\delta$:

$$\tilde{\mathcal{T}}^i_{\theta_{n+1}} \;=\; \tilde{\mathcal{T}}^i_{\theta_n} + (\theta_{n+1}-\theta_n)1_{\{S_{\theta_n}=i\delta\}}, \quad \tilde{\mathcal{T}}^i_0 = 0.$$

The interpretation of these proxies is the following: we consider the case where the (small) market maker instantaneously updates her quote Q^b (resp. Q^a) and volume $L^b \leq V_0$ (resp. $L^a \leq V_0$) only when the spread changes exogenously, i.e. at dates (θ_n), so that the spread remains constant between her updates, not considering her own quotes. If she chooses to improve best price i.e $Q^b_{\theta_n} = Bb_+$ (resp. $Q^a_{\theta_n} = Ba_-$) she will be in top priority in the LOB and therefore captures all incoming market order flow to sell (resp. buy). Therefore, an unfavourable way for (under)-estimating her number of executions is to increment $\tilde{N}^b$ (resp. $\tilde{N}^a$) only when total traded volume at bid $\xi^{SELL}_{\theta_{n+1}}$ (resp. total volume traded at ask $\xi^{BUY}_{\theta_{n+1}}$) was greater than V_0. If the market maker chooses to add liquidity to the best prices i.e. $Q^b_{\theta_n} = Bb$ (resp. $Q^a_{\theta_n} = Ba$), she will be ranked behind $V^b_{\theta_n}$ (resp. $V^a_{\theta_n}$) in LOB priority queue. Therefore, we increment $\tilde{N}^b$ (resp. $\tilde{N}^a$) only when the total traded volume at bid $\vartheta^{SELL}_{\theta_{n+1}}$ (resp. total volume traded at ask $\vartheta^{BUY}_{\theta_{n+1}}$) was greater than $V_0 + V^b_{\theta_n}$ (resp. $V_0 + V^a_{\theta_n}$). We then provide a proxy estimate for $\lambda^b_i(q^b)$, $\lambda^a_i(q^a)$ by:

$$\tilde{\lambda}^b_i(q^b) \;=\; \frac{\tilde{N}^{b,q^b,i}_{\theta_n}}{\tilde{\mathcal{T}}^i_{\theta_n}}, \qquad \tilde{\lambda}^a_i(q^a) \;=\; \frac{\tilde{N}^{a,q^a,i}_{\theta_n}}{\tilde{\mathcal{T}}^i_{\theta_n}}. \tag{5.2.12}$$

We performed the estimation procedure (5.2.12), by computing $\tilde{\lambda}^a_i(q^a)$ and $\tilde{\lambda}^b_i(q^b)$, for $i = 1, \ldots, 6$, and limit order quotes $q^b = Bb_+, Bb$, $q^a = Ba, Ba_-$. Due to the lack of data, estimate for large values of the spread are less robust. In figure 5.2, we plotted this estimated intensity as a function of the spread, i.e. $s = i\delta \to \tilde{\lambda}^b_i(q^b)$, $\tilde{\lambda}^a_i(q^a)$ for $q^b \in \mathcal{Q}^b$, and $q^a \in \mathcal{Q}^a$. As one would expect, $(\tilde{\lambda}^a_i(.), \tilde{\lambda}^b_i(.))$ are decreasing functions of i for the small values of i which matches the intuition that the higher are the (indirect) costs, the smaller is market order flow. Surprisingly, for large values of i this function becomes increasing, which can be due either to an estimation error, caused by the lack of data for this spread range, or a "gaming" effect, in other word liquidity providers increasing their spread when large or autocorrelated market orders come in.

5.3 Optimal limit/market order strategies

5.3.1 Value function

We shall study the market making problem (5.2.6) by stochastic control methods. This problem is determined by the state variables (X, Y, P, S) controlled by the limit/marker order trading strategies $\alpha \in \mathcal{A}$. Let us first remove mathematically the terminal constraint

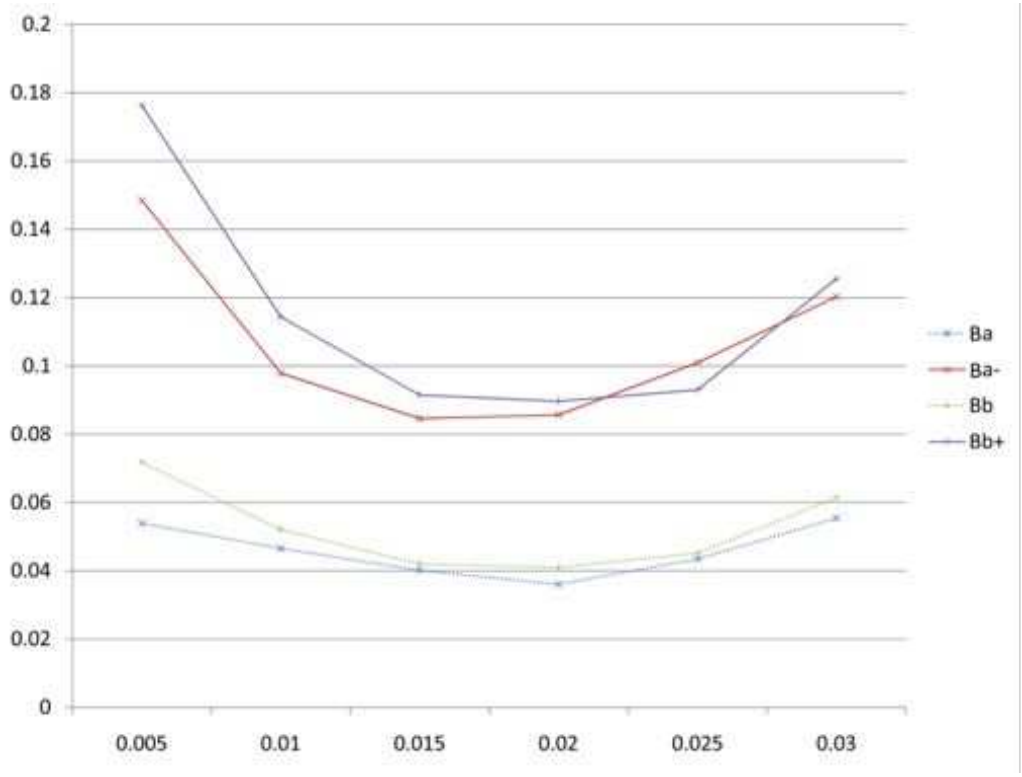

Figure 5.2: Plot of execution intensities estimate as a function of the spread for the stock SOGN.PA on the 18/04/2011, expressed in s^{-1} (affine interpolation).

on the inventory: $Y_T = 0$, by introducing the liquidation function $L(x, y, p, s)$ defined on $\mathbb{R}^2 \times \mathbb{P} \times \mathbb{S}$ by:

$$L(x, y, p, s) \;\; = \;\; x - c(-y, p, s) \;=\; x + yp - |y|\frac{s}{2} - \varepsilon.$$

This represents the value that an investor would obtained by liquidating immediately by a market order her inventory position y in stock, given a cash holdings x, a mid-price p and a spread s. Then, problem (5.2.6) is formulated equivalently as

$$\text{maximize} \quad \mathbf{E}\big[U(L(X_T, Y_T, P_T, S_T)) - \gamma \int_0^T g(Y_t)dt\big] \tag{5.3.1}$$

over all limit/market order trading strategies $\alpha = (\alpha^{make}, \alpha^{take})$ in $\mathcal{A}$. Indeed, the maximal value of problem (5.2.6) is clearly smaller than the one of problem (5.3.1) since for any $\alpha \in \mathcal{A}$ s.t. $Y_T = 0$, we have $L(X_T, Y_T, P_T, S_T) = X_T$. Conversely, given an arbitrary $\alpha \in \mathcal{A}$, let us consider the control $\tilde{\alpha} \in \mathcal{A}$, coinciding with α up to time T, and to which one add at the terminal date T the market order consisting in liquidating all the inventory Y_T.

The associated state process $(\tilde{X}, \tilde{Y}, P, S)$ satisfies: $\tilde{X}_t = X_t$, $\tilde{Y}_t = Y_t$ for $t < T$, and $\tilde{X}_T = L(X_T, Y_T, P_T, S_T)$, $\tilde{Y}_T = 0$. This shows that the maximal value of problem (5.3.1) is smaller and then equal to the maximal value of problem (5.2.6).

We then define the value function for problem (5.3.1) (or (5.2.6)):

$$v(t,z,s) = \sup_{\alpha\in\mathcal{A}} \mathbf{E}_{t,z,s}\Big[U(L(Z_T,S_T)) - \gamma\int_t^T g(Y_u)du\Big] \tag{5.3.2}$$

for $t \in [0,T]$, $z = (x,y,p) \in \mathbb{R}^2\times\mathbb{P}$, $s \in \mathbb{S}$. Here, given $\alpha \in \mathcal{A}$, $\mathbf{E}_{t,z,s}$ denotes the expectation operator under which the process $(Z,S) = (X,Y,P,S)$ solution to (5.2.1)-(6.2.5)-(6.2.4)-(6.2.7)-(6.2.6), with initial state $(Z_{t^-}, S_{t^-}) = (z,s)$, is taken. Problem (5.3.2) is a mixed regular/impulse control problem in a regime switching jump-diffusion model, that we shall study by dynamic programming methods. Since the spread takes finite values in $\mathbb{S} = \delta\mathbb{I}_m$, it will be convenient to denote for $i \in \mathbb{I}_m$, by $v_i(t,z) = v(t,z,i\delta)$. By misuse of notation, we shall often identify the value function with the $\mathbb{R}^m$-valued function $v = (v_i)_{i\in\mathbb{I}_m}$ defined on $[0,T]\times\mathbb{R}^2\times\mathbb{P}$.

5.3.2 Dynamic programming equation

For any $q = (q^b, q^a) \in \mathcal{Q}$, $\ell = (\ell^b, \ell^a) \in [0,\bar{\ell}]^2$, we consider the second-order nonlocal operator:

$$\begin{aligned}\mathcal{L}^{q,\ell}\varphi(t,x,y,p,s) &= \mathcal{L}_P\varphi(t,x,y,p,s) + R(t)\varphi(t,x,y,p,s)\\ &\quad + \lambda^b(q^b,s)\big[\varphi(t,\Gamma^b(x,y,p,s,q^b,\ell^b),p,s) - \varphi(t,x,y,p,s)\big]\\ &\quad + \lambda^a(q^a,s)\big[\varphi(t,\Gamma^a(x,y,p,s,q^a,\ell^a),p,s) - \varphi(t,x,y,p,s)\big]\end{aligned} \tag{5.3.3}$$

for $(t,x,y,p,s) \in [0,T]\times\mathbb{R}^2\times\mathbb{P}\times\mathbb{S}$, where

$$R(t)\varphi(t,x,y,p,s) = \sum_{j=1}^m r_{ij}(t)\big[\varphi(t,x,y,p,j\delta) - \varphi(t,x,y,p,i\delta)\big], \quad \text{for } s = i\delta,\ i\in\mathbb{I}_m,$$

and Γ^b (resp. Γ^a) is defined from $\mathbb{R}^2\times\mathbb{P}\times\mathbb{S}\times\mathcal{Q}^b\times\mathbb{R}_+$ (resp. $\mathbb{R}^2\times\mathbb{P}\times\mathbb{S}\times\mathcal{Q}^a\times\mathbb{R}_+$ into $\mathbb{R}^2$) by

$$\begin{aligned}\Gamma^b(x,y,p,s,q^b,\ell^b) &= (x - \pi^b(q^b,p,s)\ell^b, y+\ell^b)\\ \Gamma^a(x,y,p,s,q^a,\ell^a) &= (x + \pi^a(q^a,p,s)\ell^a, y-\ell^a).\end{aligned}$$

The first term of $\mathcal{L}^{q,\ell}$ in (5.3.3) corresponds to the infinitesimal generator of the diffusion mid-price process P, the second one is the generator of the continuous-time spread Markov chain S, and the two last terms correspond to the nonlocal operator induced by the jumps of the cash process X and inventory process Y when applying an instantaneous limit order control $(Q_t, L_t) = (q,\ell)$.

Let us also consider the impulse operator associated to market order control, and defined by

$$\mathcal{M}\varphi(t,x,y,p,s) \;=\; \sup_{e\in[-\bar{e},\bar{e}]} \varphi(t,\Gamma^{take}(x,y,p,s,e),p,s),$$

where Γ^{take} is the impulse transaction function defined from $\mathbb{R}^2\times\mathbb{P}\times\mathbb{S}\times\mathbb{R}$ into $\mathbb{R}^2$ by:

$$\Gamma^{take}(x,y,p,s,e) \;=\; \Big(x-c(e,p,s),y+e\Big),$$

The dynamic programming equation (DPE) associated to the control problem (5.3.2) is the quasi-variational inequality (QVI):

$$\min\big[-\frac{\partial v}{\partial t} - \sup_{(q,\ell)\in\mathcal{Q}(s)\times[0,\bar{\ell}]^2} \mathcal{L}^{q,\ell}v \;+\; \gamma g\;,\; v-\mathcal{M}v\big] \;=\; 0, \tag{5.3.4}$$

on $[0,T)\times\mathbb{R}^2\times\mathbb{P}\times\mathbb{S}$, together with the terminal condition:

$$v(T,x,y,p,s) \;=\; U(L(x,y,p,s)),\ \forall (x,y,p)\in\mathbb{R}^2\times\mathbb{P}\times\mathbb{S}. \tag{5.3.5}$$

This is also written explicitly in terms of system of QVIs for the functions v_i, $i\in\mathbb{I}_m$:

$$\begin{aligned}
\min\Big[-\frac{\partial v_i}{\partial t} - \mathcal{L}_{_P}v_i - \sum_{j=1}^m r_{ij}(t)[v_j(t,x,y,p)-v_i(t,x,y,p)] \\
- \sup_{(q^b,\ell^b)\in\mathcal{Q}_i^b\times[0,\bar{\ell}]} \lambda_i^b(q^b)[v_i(t,x-\pi_i^b(q^b,p)\ell^b,y+\ell^b,p)-v_i(t,x,y,p)] \\
- \sup_{(q^a,\ell^a)\in\mathcal{Q}_i^a\times[0,\bar{\ell}]} \lambda_i^a(q^a)[v_i(t,x+\pi_i^a(q^a,p)\ell^a,y-\ell^a,p)-v_i(t,x,y,p)] \;+\gamma g(y)\;; \\
v_i(t,x,y,p)-\sup_{e\in[-\bar{e},\bar{e}]} v_i(t,x-c_i(e,p),y+e,p)\Big] \;=\; 0,
\end{aligned}$$

for $(t,x,y,p)\in[0,T)\times\mathbb{R}^2\times\mathbb{P}$, together with the terminal condition:

$$v_i(T,x,y,p) \;=\; U(L_i(x,y,p)), \qquad \forall (x,y,p)\in\mathbb{R}^2\times\mathbb{P},$$

where we set $L_i(x,y,p)=L(x,y,p,i\delta)$.

By the dynamic programming principle, one can show by standard arguments that the value function v is a viscosity solution to the QVI (6.3.7)-(6.3.8), see e.g. Chapter 4, sec. 3 in [59] or Chap. 9, sec. 3 in [57]. Uniqueness of viscosity solution to (6.3.7)-(6.3.8) can also be proved by standard arguments as presented in the seminal reference [23] (see also [44] for an impulse control problem arising in optimal liquidation), and are stated within a class of functions depending on the growth conditions on the utility function U and penalty function g. The next section is devoted to numerical schemes for the resolution of the dynamic programming equation DPE (6.3.7)-(6.3.8), and to some particular cases of interest for reducing remarkably the number of states variables in the DPE.

5.4 Numerical scheme

We study a time discretization of the QVI (6.3.7)-(6.3.8). For a time step $h = T/n$, and a regular time grid $\mathbb{T}_n = \{t_k = kh, k = 0, \ldots, n\}$ over the interval $[0, T]$, we consider the following operators: for any real-valued function φ on $[0, T] \times \mathbb{R}^2 \times \mathbb{P} \times \mathbb{S}$, identified with the $\mathbb{R}^m$-valued function $(\varphi_i)_{i=1,\ldots,m}$ on $[0, T] \times \mathbb{R}^2 \times \mathbb{P}$ through $\varphi_i(t, x, y, p) = \varphi(t, x, y, p, i\delta)$, we define

$$\mathcal{D}_i^h(t, x, y, p, \varphi) \;=\; \max\big[\mathcal{T}_i^h(t, x, y, p, \varphi), \mathcal{M}_i^h(t, x, y, p, \varphi)\big],$$

where

$$\begin{aligned}\mathcal{T}_i^h(t, x, y, p, \varphi) \;=\; & -h\gamma g(y) + \frac{1}{4}\Big\{\mathbf{E}\big[\varphi_i(t + h, x, y, P_{t+4h}^{t,p})\big] + \mathbf{E}[\varphi(t + h, x, y, p, S_{t+4h}^{t,i\delta})\big] \\ & + \sup_{(q^b,\ell^b)\in\mathcal{Q}_i^b\times[0,\bar{\ell}]} \mathbf{E}\big[\varphi_i(t + h, x - \pi_i^b(q^b, p)\ell^b \Delta N_{4h}^{i,q^b}, y + \ell^b \Delta N_{4h}^{i,q^b}, p)\big] \\ & + \sup_{(q^a,\ell^a)\in\mathcal{Q}_i^a\times[0,\bar{\ell}]} \mathbf{E}\big[\varphi_i(t + h, x + \pi_i^a(q^a, p)\ell^a \Delta N_{4h}^{i,q^a}, y - \ell^a \Delta N_{4h}^{i,q^a}, p)\big]\Big\},\end{aligned}$$

and

$$\mathcal{M}_i^h(t, x, y, p, \varphi) \;=\; \sup_{e\in[-\bar{e},\bar{e}]} \varphi_i(t_{k+1}, x - c_i(e, p), y + e, p),$$

for $t \in [0, T]$, $(x, y, p) \in \mathbb{R}^2 \times \mathbb{P}$, $i \in \mathbb{I}_m$. Here, $P^{t,p}$ denotes the Markov price process of generator $\mathcal{L}_P$ starting from p at time t, $S^{t,i\delta}$ is the Markov chain of generator R starting from $i\delta$ at time t, $\Delta N_h^{i,q^b}$ is the increment over a period h of a Poisson process with intensity $\lambda_i(q^b)$, and similarly for $\Delta N_h^{i,q^a}$.

We then consider an approximation of the value function $v = (v_i)_{i\in\mathbb{I}_m}$ by $v^h = (v_i^h)_{i\in\mathbb{I}_m}$ through the explicit backward scheme:

$$\begin{aligned} v_i^h(t_n, x, y, p) \;&=\; U(L_i(x, y, p)), \quad i \in \mathbb{I}_m, (x, y, p) \in \mathbb{R}^2 \times \mathbb{P}, \\ v_i^h(t_k, x, y, p) \;&=\; \mathcal{D}_i^h(t_k, x, y, p, v^h), \quad k = 0, \ldots, n - 1,\ i \in \mathbb{I}_m, (x, y, p) \in \mathbb{R}^2 \times \mathbb{P}. \end{aligned} \tag{5.4.1}$$

Here, we identified again the real-valued function v^h on $\mathbb{T}_n \times \mathbb{R}^2 \times \mathbb{P} \times \mathbb{S}$ with the $\mathbb{R}^m$-valued function $(v_i^h)_{i\in\mathbb{I}_m}$ on $\mathbb{T}_n \times \mathbb{R}^2 \times \mathbb{P}$ via $v_i^h(t, x, y, p) = v^h(t, x, y, p, i\delta)$.

Remark 5.4.1 The convergence of the above numerical scheme can be shown formally as follows. First, it is monotone in the sense that the operator $\mathcal{D}_i^h$ is nondecreasing in φ, i.e. for any $t \in [0, T]$, $(x, y, p) \in \mathbb{R}^2 \times \mathbb{P}$, $i \in \mathbb{I}_m$, and real-valued functions φ, ψ on $[0, T] \times \mathbb{R}^2 \times \mathbb{P} \times \mathbb{S}$ s.t. $\varphi \le \psi$:

$$\mathcal{D}_i^h(t, x, y, p, \varphi) \;\le\; \mathcal{D}_i^h(t, x, y, p, \psi).$$

Secondly, by observing that the scheme (5.4.1) can be written as:

$$\min\Big[\frac{v_i^h(t,x,y,p)-\mathcal{T}_i^h(t,x,y,p,v^h)}{h}\ ,\ v_i^h(t,x,y,p)-\mathcal{M}_i^h(t,x,y,p,v^h)\Big] \ = \ 0,$$

it is consistent in the sense that

$$\begin{aligned}&\lim_{h\to 0}\min\Big[\frac{\varphi_i(t,x,y,p)-\mathcal{T}_i^h(t,x,y,p,\varphi)}{h}\ ,\ \varphi_i(t,x,y,p)-\mathcal{M}_i^h(t,x,y,p,\varphi)\Big]\\ =\ &\min\big[-\frac{\partial\varphi_i}{\partial t}-\sup_{(q,\ell)\in\mathcal{Q}(s)\times[0,\bar\ell]^2}\mathcal{L}^{q,\ell}\varphi_i\ +\ \gamma g\ ,\ \varphi_i-\mathcal{M}\varphi_i\big],\end{aligned}$$

which is the DPE (6.3.7) satisfied by the value function v. Thus, by the viscosity solutions arguments of [8], we obtain the convergence of v^h to v.

Remark 5.4.2 The approximation scheme (5.4.1) can be compared with another approximation of the value function $v=(v_i)_{i\in\mathbb{I}_m}$ by $\tilde v^h=(\tilde v_i^h)_{i\in\mathbb{I}_m}$ given by the standard explicit backward scheme:

$$\begin{aligned}\tilde v_i^h(t_n,x,y,p) &= U(L_i(x,y,p)),\quad i\in\mathbb{I}_m,(x,y,p)\in\mathbb{R}^2\times\mathbb{P},\\ \tilde v_i^h(t_k,x,y,p) &= \tilde{\mathcal{D}}_i^h(t_k,x,y,p,v^h),\quad k=0,\dots,n-1,\ i\in\mathbb{I}_m,(x,y,p)\in\mathbb{R}^2\times\mathbb{P},\end{aligned}$$

where $\tilde{\mathcal{D}}_i^h(t,x,y,p,\varphi)=\max\big[\tilde{\mathcal{T}}_i^h(t,x,y,p,\varphi),\mathcal{M}_i^h(t,x,y,p,\varphi)\big]$ with

$$\begin{aligned}&\tilde{\mathcal{T}}_i^h(t,x,y,p,\varphi)\\ =\ &\sup_{(q^b,q^a,\ell^b,\ell^a)\in\mathcal{Q}_i^b\times\mathcal{Q}_i^a\times[0,\bar\ell]^2}\mathbf{E}\Big[v^h(t+h,x-\pi_i^b(q^b,p)\ell^b\Delta N_h^{i,q^b}+\pi_i^a(q^a,p)\ell^a\Delta N_h^{i,q^a},\\ &\qquad y+\ell^b\Delta N_h^{i,q^b}-\ell^a\Delta N_h^{i,q^a},P_{t+h}^{t_k,p},S_{t+h}^{t_k,i\delta})\Big]\ -h\gamma g(y).\end{aligned}$$

The practical computation of the expectations in $\tilde{\mathcal{T}}_i^h(t_k,x,y,p,\varphi)$ would involve approximations of $P_{t_{k+1}}^{t_k,p}$ by a discrete random variable taking, say M values, approximations of $S_{t_{k+1}}^{t_k,i\delta}$ by a discrete random variable taking value $j\delta$, $j=1,\dots,m$, with probability $r_{ij}(t_k)h$ for $j\neq i$, and $1-\sum_{j\neq i}r_{ij}(t_k)h$ for $j=i$, and approximations of $\Delta N_h^{i,q^b}$ (resp. $\Delta N_h^{i,q^b}$) by the discrete variable taking value 1 with probability $\lambda_i(q^b)h$ (resp. $\lambda_i(q^a)h$) and 0 with probability $1-\lambda_i(q^b)h$ (resp. $1-\lambda_i(q^a)h$). Therefore, the global computation in $\tilde{\mathcal{T}}_i^h(t_k,x,y,p,\varphi)$, for each (t_k,x,y,p,i), would require a complexity of order $4\times\bar\ell^2\times M\times m$. Instead, we use in (5.4.1) a splitting scheme for computing separately the expectations in $\mathcal{T}_i^h(t_k,x,y,p,\varphi)$ w.r.t. the independent random variables $P_{t_{k+1}}^{t_k,p}$, $S_{t_{k+1}}^{t_k,i\delta}$, and $\Delta N_h^{i,q^b}$, $\Delta N_h^{i,q^a}$. This allows us to reduce the complexity to an order $M+m+2\bar\ell$.

In the two next paragraphs, we present two important cases leading to simplifications in the above explicit backward splitting scheme, actually by removing the cash and stock price variables.

5.4.1 Mean criterion with penalty on inventory

In this paragraph, we consider the case as in [65] where:

$$U(x) \;=\; x, \;\; x \in \mathbb{R}, \quad \text{and} \quad (P_t)_t \;\text{ is a martingale.} \tag{5.4.2}$$

The martingale assumption of the stock price under the historical measure under which the market maker performs her criterion, reflects the idea that she has no information on the future direction of the stock price. Moreover, by starting typically from zero endowment in stock, and by introducing a penalty function on inventory, the market maker wants to keep an inventory that fluctuates around zero.

In this case, similarly as in [9], the solution v^h to the above approximation scheme is reduced into the form:

$$v_i^h(t,x,y,p) \;\;=\;\; x + yp + \phi_i^h(t,y) \tag{5.4.3}$$

where $(\phi_i^h)_{i\in\mathbb{I}_m}$ is solution to the backward scheme:

$$\phi_i^h(t_n,y) \;\;=\;\; -|y|\frac{i\delta}{2} - \varepsilon \tag{5.4.4}$$

$$\begin{aligned}\tilde{\phi}_i(t_k,y) \;\;=\;\; & \frac{1}{4}\Big\{\phi_i^h(t_{k+1},y) + \mathbf{E}\big[\phi^h\big(t_{k+1},y,S_{t_{k+4}}^{t_k,i\delta}\big)\big] \\ & + \sup_{(q^b,\ell^b)\in\mathcal{Q}_i^b\times[0,\bar{\ell}]} \mathbf{E}\Big[\big(\frac{i\delta}{2} - \delta 1_{q^b=Bb_+}\big)\ell^b\Delta N_{4h}^{i,q^b} + \phi_i(t_{k+1},y+\ell^b\Delta N_{4h}^{i,q^b})\Big] \\ & + \sup_{(q^a,\ell^a)\in\mathcal{Q}_i^a\times[0,\bar{\ell}]} \mathbf{E}\Big[\big(\frac{i\delta}{2} - \delta 1_{q^a=Ba_-}\big)\ell^a\Delta N_{4h}^{i,q^a} + \phi_i(t_{k+1},y-\ell^a\Delta N_{4h}^{i,q^a})\Big] \\ & - h\gamma g(y)\Big\}\end{aligned} \tag{5.4.5}$$

$$\phi_i(t_k,y) \;\;=\;\; \max\Big[\tilde{\phi}_i^h(t_k,y)\;,\; \sup_{e\in[-\bar{e},\bar{e}]}\big[-\frac{i\delta}{2}|e| - \varepsilon + \phi_i^h(t_{k+1},y+e)\big]\Big] \tag{5.4.6}$$

for $k = 0,\ldots,n-1$, $i \in \mathbb{I}_m$, $y \in \mathbb{R}$. By misuse of notation, we have set $\phi^h(t,y,s) = \phi_i^h(t,y)$ for $s = i\delta$.

The reduced form (5.4.3) shows that the optimal market making strategies are price independent, and depend only on the level of inventory and of the spread, which is consistent with stylized features in the market.

Remark 5.4.3 The scheme for (ϕ_i^h) is the time discretization of the system of one-dimensional

integro-differential equations (IDEs):

$$\begin{aligned}
\min\Big[-\frac{\partial \phi_i}{\partial t} - \sum_{j=1}^m r_{ij}(t)[\phi_j(t,y)-\phi_i(t,y)] \\
- \sup_{(q^b,\ell^b)\in\mathcal{Q}_i^b\times[0,\bar\ell]} \lambda_i^b(q^b)[\phi_i(t,y+\ell^b)-\phi_i(t,y)+\big(\frac{i\delta}{2}-\delta 1_{q^b=Bb_+}\big)\ell^b] \\
- \sup_{(q^a,\ell^a)\in\mathcal{Q}_i^a\times[0,\bar\ell]} \lambda_i^a(q^a)[\phi_i(t,y-\ell^a)-\phi_i(t,y)+\big(\frac{i\delta}{2}-\delta 1_{q^a=Ba_-}\big)\ell^a] + \gamma g(y) ; \\
\phi_i(t,y) - \sup_{e\in[-\bar e,\bar e]}[\phi_i(t,y+e)-\frac{i\delta}{2}|e|-\varepsilon]\Big] &= 0,
\end{aligned}$$

together with the terminal condition:

$$\phi_i(T,y) = -|y|\frac{i\delta}{2}-\varepsilon,$$

which can be also derived from the dynamic programming system (6.3.7)-(6.3.8) for $v=(v_i)_{i\in\mathbb{I}_m}$ reduced into the form: $v_i(t,x,y,p)=x+yp+\phi_i(t,y)$. This system of IDEs also show that optimal policies do not depend on the martingale modeling of the stock price.

5.4.2 Exponential utility criterion

In this paragraph, we consider as in [7] a risk averse market marker:

$$U(x) = -\exp(-\eta x),\ x\in\mathbb{R},\ \eta>0,\quad \gamma = 0, \tag{5.4.7}$$

and assume that P is a Lévy process so that

$$P_{t+h}^{t,p} = p+\mathcal{E}_h$$

where $\mathcal{E}_h$ is a random variable, which does not depend on p. In this case, similarly as in [35], the solution v^h to the above approximation scheme is reduced into the form

$$v_i^h(t,x,y,p) = U(x+yp)\varphi_i^h(t,y), \tag{5.4.8}$$

where $(\varphi_i^h)_{i\in\mathbb{I}_m}$ is solution to the backward scheme:

$$\varphi_i^h(t_n,y) = \exp(\eta|y|\frac{i\delta}{2}) \tag{5.4.9}$$

$$\begin{aligned}
\tilde\varphi_i^h(t_k,y) &= \frac14\Big\{\mathbf{E}\big[\exp\big(-\eta y\mathcal{E}_h\big)\big]\varphi_i^h(t_{k+1},y)+\mathbf{E}\big[\varphi^h(t_{k+1},y,S_{t_{k+1}}^{t_k,i\delta}\big)\big] \\
&+ \inf_{(q^b,\ell^b)\in\mathcal{Q}_i^b\times[0,\bar\ell]}\mathbf{E}\big[\exp\big(-\eta\big(\frac{i\delta}{2}-\delta 1_{q^b=Bb_+}\big)\ell^b\Delta N_h^{i,q^b}\big)\varphi_i^h(t_{k+1},y+\ell^b\Delta N_h^{i,q^b})\big] \\
&+ \inf_{(q^a,\ell^a)\in\mathcal{Q}_i^a\times[0,\bar\ell]}\mathbf{E}\big[\exp\big(-\eta\big(\frac{i\delta}{2}-\delta 1_{q^a=Ba_-}\big)\ell^a\Delta N_h^{i,q^a}\big)\varphi_i^h(t_{k+1},y-\ell^a\Delta N_h^{i,q^a})\big]\Big\}
\end{aligned} \tag{5.4.10}$$

$$\varphi_i^h(t_k,y) = \min\Big[\tilde\varphi_i^h(t_k,y)\ ,\ \inf_{e\in[-\bar e,\bar e]}\big[\exp\big(\eta|e|\frac{i\delta}{2}+\eta\varepsilon\big)\varphi_i^h(t_{k+1},y+e)\big]\Big] \tag{5.4.11}$$

for $k = 0, \ldots, n-1$, $i \in \mathbb{I}_m$, $y \in \mathbb{R}$. Here, we set $\varphi^h(t,y,s) = \varphi^h_i(t,y)$ for $s = i\delta$.

As in the case (5.4.2), the reduced form (5.4.8) shows that the optimal market making strategies are price independent, and depend only on the level of inventory and of the spread. However, it depends on the model (typically the volatility) for the stock price through the term $\mathcal{E}_h$.

Remark 5.4.4 Let us consider the example of Lévy process: $dP_t = bdt + \sigma dW_t + \kappa(dM_t - \mu dt)$, where b, $\sigma > 0$, κ are real constants, W is a Brownian motion, and M is an independent Poisson process of intensity μ. Thus, $\mathcal{E}_h = bh + \sigma W_h + \kappa(M_h - \mu h)$, and the above scheme for $(\varphi^h_i)_{i\in\mathbb{I}_m}$ corresponds to the time discretization of the system of one-dimensional integro-differential equations:

$$\begin{aligned}
\max\Big[-\frac{\partial\varphi_i}{\partial t} + \big(b\eta y - \frac{1}{2}\sigma^2(\eta y)^2 + \mu(1 - \kappa\eta y - e^{-\eta\kappa y})\big)\varphi_i - \sum_{j=1}^m r_{ij}(t)[\varphi_j(t,y) - \varphi_i(t,y)] \\
- \inf_{(q^b,\ell^b)\in\mathcal{Q}^b_i\times[0,\bar\ell]} \lambda^b_i(q^b)[\exp\big(-\eta\big(\frac{i\delta}{2} - \delta 1_{q^b = Bb_+}\big)\ell^b\big)\varphi_i(t,y+\ell^b) - \varphi_i(t,y)] \\
- \inf_{(q^a,\ell^a)\in\mathcal{Q}^a_i\times[0,\bar\ell]} \lambda^a_i(q^a)[\exp\big(-\eta\big(\frac{i\delta}{2} - \delta 1_{q^a = Ba_-}\big)\ell^a\big)\varphi_i(t,y-\ell^a) - \varphi_i(t,y)] \\
\varphi_i(t,y) - \inf_{e\in[-\bar e,\bar e]}[\exp\big(\eta|e|\frac{i\delta}{2} + \eta\varepsilon\big)\varphi_i(t,y+e)]\Big] &= 0,
\end{aligned}$$

together with the terminal condition:

$$\varphi_i(T,y) = \exp(\eta|y|\frac{i\delta}{2}),$$

which can be also derived from the dynamic programming system (6.3.7)-(6.3.8) for $v = (v_i)_{i\in\mathbb{I}_m}$ reduced into the form: $v_i(t,x,y,p) = x + yp + \varphi_i(t,y)$.

Remark 5.4.5 We observe that numerical scheme simplifications are due to the specific form of the value function. In the case of a general utility function U (e.g. CRRA utility function, see [44]), such simplifications as (5.4.8) or (5.4.3) may not exist, and therefore the optimization is performed on a 4-dimensionnal model (plus time). From the computational point of view, this requires to solve a 4 dimensions numerical scheme, which may lead to much heavier computations, with sometimes untractable memory and time requirements, and less precise numerical results.

The main difference between the case of mean-variance criterion and exponential criterion is that the price model parameters, as the volatility σ, appears naturally in the case of exponential criterion. Indeed, the two main objects of interest in our model are the price model and the trade processes model. This last feature can be used to favour the dependence of the resulting strategy on price parameters, against the trades processes models.

When the high-frequency trader has no information on the price behavior, or when the volatility is not relevant for the timescale of trading, one may want to take greater care of characteristics of the trade processes than of the price.

In the case of the mean-variance criterion, choice of the risk aversion parameter γ is left to the decision of the high-frequency trader. As shown in figure 5.4, this parameter can be fitted a posteriori, upon results of the backtest/calibration procedure, in order to choose the relationship between the variance and the average profit of the optimal strategy. For example, the high frequency trader may want to choose γ in order to maximize the information ratio (or Sharpe ratio) against a benchmark, the example that we chose to illustrate graphically in figure 5.4. The equivalent parameter in the case of exponential criterion is η, that is also left to the HFT's choice. Note that in this last case, there is no explicit constraint on the inventory since we take $\gamma = 0$.

5.5 Computational results

In this section, we provide numerical results obtained with the optimal strategy computed with our implementation of the simplified scheme (5.4.4)-(5.4.6) in the case of a mean criterion with penalty on inventory, that we will denote within this section by $\alpha^\star$. We used parameters shown in table 5.2 together with transition probabilities $(\rho_{ij})_{1\leq i,j\leq M}$ calibrated in table 5.1 and execution intensities calibrated in Figure 5.2, slightly modified to make the bid and ask sides symmetric.

Parameter	Signification	Value
δ	Tick size	0.005
ρ	Per share rebate	0.0008
ϵ	Per share fee	0.0012
ϵ_0	Fixed fee	10^{-6}
$\lambda(t)$	Tick time intensity	$\equiv 1s^{-1}$

(a) Market parameters

Parameter	Signification	Value
$U(x)$	Utility function	x
$g(x)$	Penalty function	x^2
γ	Inventory penalization	5
$\bar{\ell}$	Max. volume make	100
$\bar{e}$	Max. volume take	100

(b) Optimization parameters

Parameter	Signification	Value
T	Length in seconds	300 s
y_{min}	Lower bound shares	-1000
y_{max}	Upper bound shares	1000
n	Number of time steps	100
m	Number of spreads	6

(c) Discretization/localization parameters

Parameter	Signification	Value
N^{MC}	Number of paths for MC simul.	10^5
Δt	Euler scheme time step	0.3 s
$\bar{\ell}_0$	B/A qty for bench. strat.	100
x_0	Initial cash	0
y_0	Initial shares	0
p_0	Initial price	45

(d) Backtest parameters

Table 5.2: Parameters

Shape of the optimal policy. The reduced form (5.4.3) shows that the optimal policy $\alpha^\star$ does only depend on time t, inventory y and spread level s. One can represent $\alpha^\star$ as a mapping $\alpha^\star : \mathbb{R}^+ \times \mathbb{R} \times \mathbb{S} \to \mathcal{A}$ with $\alpha^\star = (\alpha^{\star,make}, \alpha^{\star,take})$ thus it divides the space $\mathbb{R}^+ \times \mathbb{R} \times \mathbb{S}$ in two zones $\mathcal{M}$ and $\mathcal{T}$ so that $\alpha^\star_{|\mathcal{M}} = (\alpha^{\star,make}, 0)$ and $\alpha^\star_{|\mathcal{T}} = (0, \alpha^{\star,take})$. Therefore we plot the optimal policy in one plane, distinguishing the two zones by a color scale. For the zone $\mathcal{M}$, due to the complex nature of the control, which is made of four scalars, we only represent the prices regimes.

(a) near date 0

(b) near date T

Figure 5.3: Stylized shape of the optimal policy sliced in YS.

Figure 5.3 describes the optimal policy as a function of inventory and spread. Qualitatively, we can explain this strategy by thinking of a risk/reward trade-off. One can interpret the market order zones $\mathcal{M}$, located on the extreme right and left parts of the graph, as zones where the inventory becomes too large, and the inventory risk unsustainable. Therefore, the HF trader will need to unwind her portfolio at market, and therefore pay direct and indirect ("crossing the spread") costs. Otherwise, when spread becomes large, thus allowing more potential profit from the market-maker point of view, or when the inventory is low, the HF trader has a better bet trading passively with limit orders. In this last case, depending on the sign of her inventory, the market-maker may want to trade with asymetric limit orders, i.e. cancel the bid (resp. ask) side and keep an active limit order only on the ask (resp. bid) side.

Moreover, when using constant tick time intensity $\lambda(t) \equiv \lambda$ and in the case where $T \gg \frac{1}{\lambda}$ we can observe on numerical results that the optimal policy is mainly time invariant near date 0; on the contrary, close to the terminal date T the optimal policy has a transitory regime, in the sense that it critically depends on the time variable t. This matches the intuition that to ensure the terminal constraint $Y_T = 0$, the optimal policy tends to get rid of the inventory more aggressively when close to maturity. In figure 5.3, we plotted a

stylized view of the optimal policy, in the plane (y, s), to illustrate this phenomenon.

Benchmarked empirical performance analysis. We made a backtest of the optimal strategy $\alpha^\star$, on simulated data, and benchmarked the results with the three following strategies:

Optimal strategy without market orders (WoMO), that we denote by α^w: this strategy is computed using the same algorithm (5.4.4)-(5.4.6), but in the case where the investor is not allowed to use market orders, which is equivalent to setting $\bar{e} = 0$.

Constant strategy, that we denote by α^c: this strategy is the symmetric best bid, best ask strategy with constant quantity $\bar{\ell}_0$ on both sides, or more precisely $\alpha^c := (\alpha^{c,make}, 0)$ with $\alpha_t^{c,make} \equiv (Bb, Ba, \bar{\ell}_0, \bar{\ell}_0)$.

Random strategy, that we denote by α^r: this strategy consists in choosing randomly the price of the limit orders and using constant quantities on both sides, or more precisely $\alpha^r := (\alpha^{r,make}, 0)$ with $\alpha_t^{r,make} = (\varsigma_t^b, \varsigma_t^a, \bar{\ell}_0, \bar{\ell}_0)$ where $(\varsigma_\cdot^b, \varsigma_\cdot^a)$ is s.t. $\forall t \in [0;T]$, $\mathbf{P}(\varsigma_t^b = Bb) = \mathbf{P}(\varsigma_t^b = Bb_+) = \mathbf{P}(\varsigma_t^a = Ba) = \mathbf{P}(\varsigma_t^a = Ba_-) = \frac{1}{2}$.

Our backtest procedure is described as follows. For each strategy $\alpha \in \{\alpha^\star, \alpha^w, \alpha^c, \alpha^r\}$, we simulated N^{MC} paths of the tuple $(X^\alpha, Y^\alpha, P, S, N^{a,\alpha}, N^{b,\alpha})$ on $[0, T]$, according to eq. (5.2.1)-(6.2.5)-(6.2.4)-(6.2.7)-(6.2.6), using a standard Euler scheme with time-step Δt. Therefore we can compute the empirical mean (resp. empirical standard deviation), that we denote by $m(.)$ (resp. $\sigma(.)$), for several quantities shown in table 5.3.

		optimal $\alpha^\star$	WoMO α^w	constant α^c	random α^r
Terminal wealth	$m(X_T)/\sigma(X_T)$	2.117	1.999	0.472	0.376
	$m(X_T)$	26.759	25.19	24.314	24.022
	$\sigma(X_T)$	12.634	12.599	51.482	63.849
Num. of exec. at bid	$m(N_T^b)$	18.770	18.766	13.758	21.545
	$\sigma(N_T^b)$	3.660	3.581	3.682	4.591
Num. of exec. at ask	$m(N_T^a)$	18.770	18.769	13.76	21.543
	$\sigma(N_T^a)$	3.666	3.573	3.692	4.602
Num. of exec. at market	$m(N_T^{market})$	6.336	0	0	0
	$\sigma(N_T^{market})$	2.457	0	0	0
Maximum Inventory	$m(\sup_{s\in[0;T]} \lvert Y_s\rvert)$	241.019	176.204	607.913	772.361
	$\sigma(\sup_{s\in[0;T]} \lvert Y_s\rvert)$	53.452	23.675	272.631	337.403

Table 5.3: Performance analysis: synthesis of benchmarked backtest (10^5 simulations).

Optimal strategy $\alpha^\star$ demonstrates significant improvement of the information ratio $\mathrm{IR}(X_T) := m(X_T)/\sigma(X_T)$ compared to the benchmark, which is confirmed by the plot of the whole empirical distribution of X_T (see figure 5.4).

Even if absolute values of $m(X_T)$ are not representative of what would be the real-world performance of such strategies, these results prove that the different layers of optimization

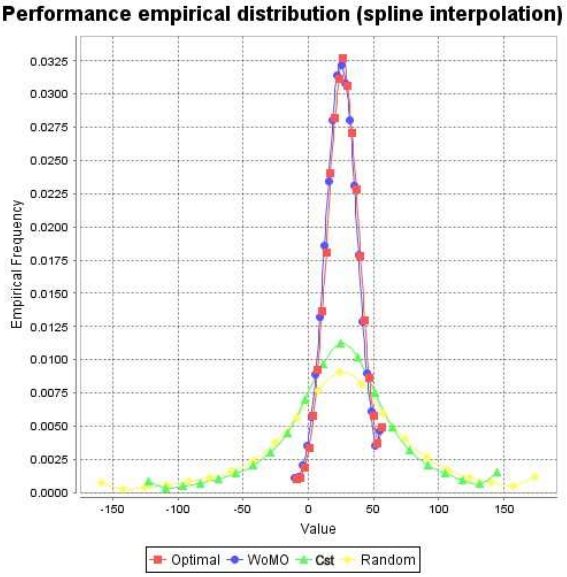

Figure 5.4: Empirical distribution of terminal wealth X_T (spline interpolation).

are relevant. Indeed, one can compute the ratios $\left[m(X_T^{\alpha^\star}) - m(X_T^{\alpha^c})\right] / \sigma(X_T^{\alpha^\star}) = 0.194$ and $\left[m(X_T^{\alpha^\star}) - m(X_T^{\alpha^w})\right] / \sigma(X_T^{\alpha^\star}) = 0.124$ that can be interpreted as the performance gain, measured in number of standard deviations, of the optimal strategy $\alpha^\star$ compared respectively to the constant strategy α^c and the WoMO strategy α^w. Another interesting statistics is the surplus profit per trade $\left[m(X_T^{\alpha^\star}) - m(X_T^{\alpha^c})\right] / \left[m(N_T^{b,\alpha^\star}) + m(N_T^{a,\alpha^\star}) + m(N_T^{market,\alpha^\star})\right] =$ 0.056 euros per trade, recalling that the maximum volume we trade is $\bar{\ell} = \bar{e} = 100$. Note that for this last statistics, the profitable effects of the per share rebates ρ are partially neutralized because the number of executions is comparable between $\alpha^\star$ and α^c; therefore the surplus profit per trade is mainly due to the revenue obtained from *making the spread*. To give a comparison point, typical clearing fee per execution is 0.03 euros on multilateral trading facilities, therefore, in this backtest, the surplus profit per trade was roughly twice the clearing fees.

We observe in the synthesis table that the number of executions at bid and ask are symmetric, which is also confirmed by the plots of their empirical distributions in figure 5.5. This is due to the symmetry in the execution intensities λ^b and λ^a, which is reflected by the symmetry around $y = 0$ in the optimal policy.

Moreover, notice that the maximum absolute inventory is efficiently kept close to zero in $\alpha^\star$ and α^w, whereas in α^c and α^r it can reach much higher values. The maximum absolute inventory is higher in the case of $\alpha^\star$ than in the case α^w due to the fact that $\alpha^\star$ can unwind any position immediately by using market orders, and therefore one may post higher volume for limit orders between two trading at market, profiting from reduced

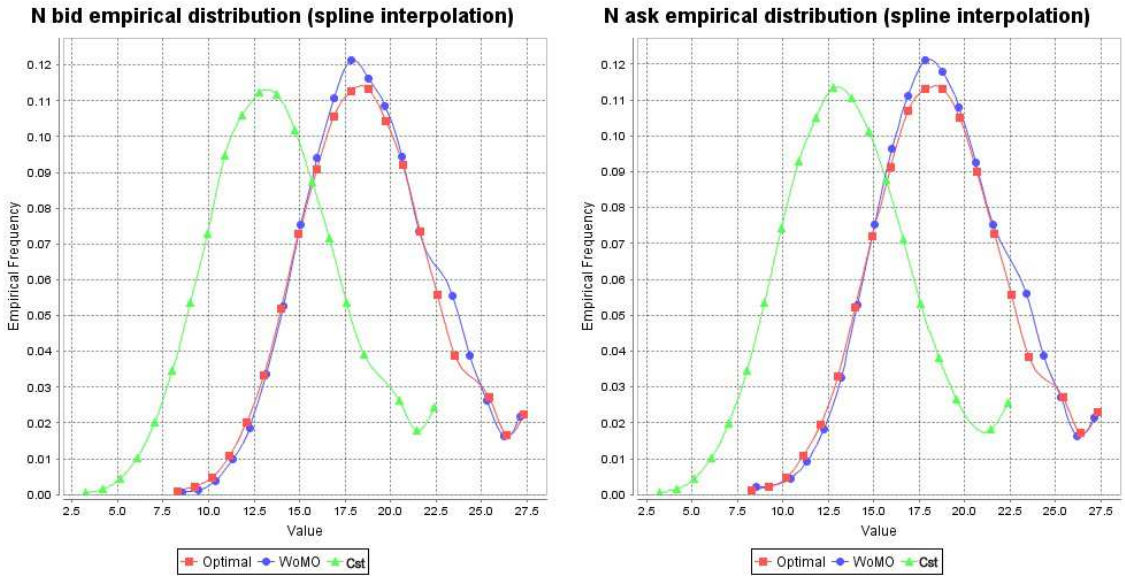

(a) N Bid empirical distribution

(b) N Ask empirical distribution

Figure 5.5: Empirical distribution of the number of executions on both sides.

execution risk.

Efficient frontier. An important feature of our algorithm is that the market maker can choose the inventory penalization parameter γ. To illustrate its influence, we varied the inventory penalization γ from 50 to 6.10^{-2}, and then build the efficient frontier for both the optimal strategy $\alpha^\star$ and for the WoMO strategy α^w. Numerical results are provided in table 5.4 and a plot of this data is in figure 5.6.

We display both the "gross" information ratio $\mathrm{IR}(X_T^{\alpha^\star}) := m(X_T^{\alpha^\star})/\sigma(X_T^{\alpha^\star})$ and the "net" information ratio $\mathrm{NIR}(X_T^{\alpha^\star}) := \left(m(X_T^{\alpha^\star}) - m(X_T^{\alpha^c})\right)/\sigma(X_T^{\alpha^\star})$ to have more precise interpretation of the results. Indeed, $m(X_T^\alpha)$ seems largely overestimated in this simulated data backtest compared to what would be real-world performance, for all $\alpha \in \{\alpha^\star, \alpha^w, \alpha^c, \alpha^r\}$. Then, to ease interpretation, we assume that α^c has zero mean performance in real-world conditions, and therefore offset the mean performance $m(X_T^{\alpha^\star})$ by the constant $-m(X_T^{\alpha^c})$ when computing the NIR. This has simple visual interpretation as shown in figure 5.6.

Observe that highest (net) information ratio is reached for $\gamma \simeq 0.8$ for this set of parameters. At this point $\gamma \simeq 0.8$, the annualized value of the NIR (obtained by simple extrapolation) is 47, but this simulated data backtest must be completed by a backtest on real data. Qualitatively speaking, the effect of increasing the inventory penalization parameter γ is to increase the zone $\mathcal{T}$ where we trade at market. This induces smaller

γ	$\sigma(X_T^{\alpha^\star})$	$m(X_T^{\alpha^\star})$	$\sigma(X_T^{w^\star})$	$\sigma(X_T^{w^\star})$	$\mathrm{IR}(X_T^{\alpha^\star})$	$\mathrm{NIR}(X_T^{\alpha^\star})$
50.000	5.283	12.448	4.064	9.165	2.356	-2.246
25.000	7.562	18.421	7.210	16.466	2.436	-0.779
12.500	9.812	22.984	9.531	20.971	2.343	-0.135
6.250	11.852	25.932	11.749	24.232	2.188	0.136
3.125	14.546	28.153	14.485	26.752	1.935	0.263
1.563	15.819	28.901	16.830	28.234	1.827	0.289
0.781	19.088	29.952	19.593	29.145	1.569	0.295
0.391	20.898	30.372	20.927	29.728	1.453	0.289
0.195	23.342	30.811	23.247	30.076	1.320	0.278
0.098	25.232	30.901	24.075	30.236	1.225	0.261
0.049	26.495	31.020	24.668	30.434	1.171	0.253
0.024	27.124	30.901	25.060	30.393	1.139	0.242
0.012	27.697	31.053	25.246	30.498	1.121	0.243
0.006	28.065	30.998	25.457	30.434	1.105	0.238

Table 5.4: Efficient frontier data

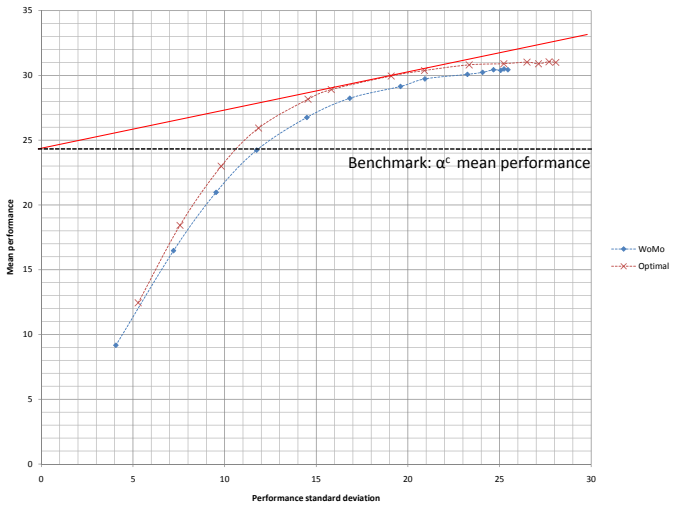

Figure 5.6: Efficient frontier plot

inventory risk, due to the fact that we unwind our position when reaching relatively small values for $|y|$. This feature can be used to enforce a soft maximum inventory constraint directly by choosing γ.

Appendix A: pseudo-code

In this appendix, we provide the pseudo-code for solving the simplified numerical schemes ((5.4.4)-5.4.6) and ((5.4.9)-(5.4.11)). In this section, given $C > 0$ and $\Delta_Y > 0$ we use the notation:

$$\mathbb{Y}_C := \{-C \vee k\Delta_Y \wedge C, k \in \mathbb{Z}\}$$

the regular grid on $\mathbb{R}$ truncated at C.

Pseudo-code for the numerical scheme in the case of mean criterion with penalty on inventory.

This algorithm is described explicitly in backward induction by the following pseudo-code:

- Timestep $t_N = T$: for each $y \in \mathbb{Y}_C$, for each $i \in \mathbb{I}_m$, set $\phi_i^h(t_n, y) = -|y|\frac{i\delta}{2} - \varepsilon$ according to eq. (5.4.4).
- For $k = N-1 \dots 0$, from timestep t_{k+1} to timestep t_k, for each $y \in \mathbb{Y}_C$, and for each $i \in \mathbb{I}_m$:
  - Compute $\tilde{\phi}_i(t_k, y)$ from 5.4.10, and store $(q^b, \ell^b)^\star$, $(q^a, \ell^a)^\star$ the argmax
  - Compute $\hat{\phi}_i(t_k, y) := \sup_{e \in [-\bar{e}, \bar{e}]} \big[-\frac{i\delta}{2}|e| - \varepsilon + \phi_i^h(t_{k+1}, y+e)\big]$, and store $e^\star$ the argmax
  - If $\tilde{\phi}_i(t_k, y) \geq \hat{\phi}_i(t_k, y)$ then set

    $$\phi_i(t_k, y) \;:=\; \tilde{\phi}_i(t_k, y)$$

    and the policy is make $(q^b, \ell^b)^\star$, $(q^a, \ell^a)^\star$. Otherwise, set

    $$\phi_i(t_k, y) \;:=\; \hat{\phi}_i(t_k, y)$$

    and the policy is take $e^\star$.

Pseudo-code for the numerical scheme in the case of exponential utility criterion.

This algorithm is described explicitly in backward induction by the following pseudo-code:

- Timestep $t_N = T$: for each $y \in \mathbb{Y}_C$, for each $i \in \mathbb{I}_m$, set $\varphi_i^h(t_n, y) = \exp(\eta|y|\frac{i\delta}{2})$ according to eq. 5.4.9.

- For $k = N-1 \dots 0$, from timestep t_{k+1} to timestep t_k, for each $y \in \mathbb{Y}_C$, and for each $i \in \mathbb{I}_m$:
  - Compute $\tilde{\varphi}_i(t_k, y)$ from 5.4.5, and store $(q^b, \ell^b)^\star$, $(q^a, \ell^a)^\star$ the argmax
  - Compute $\hat{\varphi}_i(t_k, y) := \inf_{e \in [-\bar{e}, \bar{e}]} \left[\exp\left(\eta |e| \frac{i\delta}{2} + \eta\varepsilon\right) \varphi_i^h(t_{k+1}, y+e) \right]$, and store $e^\star$ the argmax
  - If $\tilde{\varphi}_i(t_k, y) \geq \hat{\varphi}_i(t_k, y)$ then set

    $$\varphi_i(t_k, y) \quad := \quad \tilde{\varphi}_i(t_k, y)$$

    and the policy is make $(q^b, \ell^b)^\star$, $(q^a, \ell^a)^\star$. Otherwise, set

    $$\varphi_i(t_k, y) \quad := \quad \hat{\varphi}_i(t_k, y)$$

    and the policy is take $e^\star$.

Chapter 6

Optimal HF trading in a pro-rata microstructure with predictive information

We propose a framework to study optimal trading policies in a one-tick pro-rata limit order book, as typically arises in short-term interest rate futures contracts. The high-frequency trader chooses to post either market orders or limit orders, which are represented respectively by impulse controls and regular controls. We discuss the consequences of the two main features of this microstructure: first, the limit orders are only partially executed, and therefore she has no control on the executed quantity. Second, the high frequency trader faces the overtrading risk, which is the risk of brutal variations in her inventory. The consequences of this risk are investigated in the context of optimal liquidation. The optimal trading problem is studied by stochastic control and dynamic programming methods, and we provide the associated numerical resolution procedure and prove its convergence. We propose dimension reduction techniques in several cases of practical interest. We also detail a high frequency trading strategy in the case where a (predictive) directional information on the price is available. Each of the resulting strategies are illustrated by numerical tests.

Note: This chapter is adapted from the article : [38] Guilbaud F. and H. Pham (2012): "Optimal high frequency trading in a pro-rata microstructure with predictive information", available at SSRN: http://ssrn.com/abstract=2040867.

6.1 Introduction

In most of modern public security markets, the price formation process, or price discovery, results from competition between several market agents that take part in a public auction. In particular, day trading sessions, which are also called continuous trading phases, consist of continuous double auctions. In these situations, liquidity providers[1] continuously set bid and ask prices for the considered security, and the marketplace publicly displays a (possibly partial) information about these bid and ask prices, along with transactions prices. The action of continuously providing bid and ask quotes during day trading sessions is called *market making*, and this role was tradionnally performed by specialist firms. However, due to the recent increased availability of electronic trading technologies, as well as regulatory changes, a large range of investors are now able to implement such market making strategies. These strategies are part of the broader category of high frequency trading (HFT) strategies, which are characterized by the fact that they facilitate a larger number of orders being sent to the market per unit of time. HFT takes place in the continuous trading phase, and therefore in the double continuous auction context, but actual mechanisms that implement this general continuous double auction set-up directly influence the price formation process and, as a consequence, HFT strategies.

In this work, we shall focus on the case where the continuous double auction is implemented by a limit order book (LOB), operated under the pro-rata microstructure, see [43] and [1]. This microstructure can be encountered on some derivatives markets, and especially in short-term interest rate (STIR) futures markets, also known as financial futures, traded e.g. on LIFFE (London International Financial Futures and options Exchange) or on CME (Chicago Mercantile Exchange). This differs from the usual price/time microstructure found on most cash equity markets, and governed by the FIFO (first in first out) rule where limit orders are executed according to the first arrival at the best price. We will describe in detail the prorata microstructure in Section 2, but the general mechanism of this microstructure is as follows: an incoming market order is dispatched on all active limit orders at the best price, with each limit order contributing to execution in proportion to its volume. In particular, we will discuss the two main consequences of this microstructure on HFT strategies which are the oversizing of the best priced slices of the LOB and the overtrading risk.

Our main goal is to construct an HFT strategy, by means of optimal stochastic control, that targets the pro-rata microstructure. We allow both limit orders and market orders in this HFT strategy, modelled respectively as continuous and impulse controls, due to considerations about direct trading costs. From a modelling point of view, the key novelty is that we take into account partial execution for limit orders, which is crucial in the pro-rata case. For this purpose we introduce a compound Poisson model for trades pro-

[1] In this paper, we call liquidity provider any investor that currently trades with limit orders

cesses, that can be fitted to a large class of real-world execution processes, since we make few assumptions about the distributions of execution volumes. From a practical trading point of view, we allow the HFT to input predictive information about price evolution into the strategy, so that our algorithm can be seen as an information-driven HFT strategy (this situation is sometimes called HFT with superior information, see [16]). We derive the dynamic programming equation corresponding to this mixed impulse/regular control problem. Moreover, we are able to reduce the number of relevant state variables to one in two situations of practical interest: first, in the simple case where the mid-price is a martingale, and second, in the case where the mid-price is a Lévy process, in particular when the HFT has predictive information on price trend, in line with recent studies [21]. We provide a computational algorithm for the resolution of the dynamic programming equation, and prove the convergence of this scheme. We illustrate numerically the behavior of the strategy and perform a simulated data benchmarked backtest.

High-frequency trading has recently received sustained academic interest, mostly in a price/time microstructure model. The reference work for inventory-based high frequency trading is Avellaneda and Stoikov (2008) [7] following early work by Ho and Stoll [42]. The authors present the HFT problem as an inventory management problem and define inventory risk as the risk of holding a non-zero position in a risky asset. They also provide a closed-form approximate solution in a stylized market model where the controls are continuous. Several works are available that describe optimal strategies for HFT on cash equities or foreign exchange, e.g. [45], [15], [35], [37] or [68]. Guéant, Tapia and Lehalle ([35]) provide extensive analytical treatment of the Avellaneda and Stoikov model. Veraart ([68]) includes market orders (that are modelled as impulse controls) as well as limit order in the context of FX trading. Guilbaud and Pham ([37]) study market/limit orders HFT strategies on stocks with a focus on the price/time priority microstructure and the bid/ask spread modelling. Cartea, Jaimungal and Ricci ([16]) consider a HFT strategy that takes into account influence of trades on the LOB, and give the HFT superior information about the security price evolution. A growing literature is dedicated to modelling the dynamics of the limit order book itself, and its consequences for the price formation process. A popular approach is the Poisson Limit Order Book model as in Cont and de Larrard ([21]). These authors are able to retrieve a predictive information on price behavior (together with other LOB features) based on the current state of the order book. Finally, in empirical literature, much work is available for cash equities e.g. [33], but very few is dedicated to markets operating under pro-rata microstructure. We would like to mention the work by Field and Large ([25]), which provides a detailed empirical description of such microstructure.

This paper is organized as follows: in Section 2, we detail the prorata microstructure model and explain the high frequency trading strategy in this context. In Section 3, we formulate the control problem, derive the corresponding dynamic programming equation (DPE) for the value function, and state some bounds and symmetry properties. We also

simplify the DPE in two cases of practical interest, namely the case where the price is a martingale, and the case where the investor has predictive information on price trend available. In Section 4, we provide the numerical algorithm to solve the DPE, and we study the convergence of the numerical scheme, by proving the monotonicity, stability and consistency for this scheme. We also provide numerical tests including computations of the optimal policies and performance analysis on a simulated data backtest. Finally, in Section 5, we show how to extend our model in the optimal liquidation case, i.e. when the investor's objective is to minimize the trading costs for unwinding her portfolio.

6.2 Market model

Let us fix a probability space $(\Omega, \mathcal{F}, \mathbf{P})$ equipped with a filtration $\mathbb{F} = (\mathcal{F}_t)_{0 \leq t \leq T}$ satisfying the usual conditions. It is assumed that all random variables and stochastic processes are defined on the stochastic basis $(\Omega, \mathcal{F}, \mathbb{F}, \mathbf{P})$.

Prices in a one-tick microstructure. We denote by P the midprice, defined as a Markov process with generator $\mathcal{P}$ valued in $\mathbb{P}$. We shall assume that P is a special semimartingale with locally integrable quadratic variation process $[P]$, so that its dual predictable projection (also called sharp bracket) $< P >$ exists (see [62]). We assume that $< P >_T$ is integrable, and that the predictable finite variation term A of the special semimartingale P satisfies the canonical structure: $dA_t \ll d < P >_t$, with a bounded density process:

$$\theta_t \;=\; \frac{dA_t}{d < P >_t}, \tag{6.2.1}$$

and the sharp bracket process $< P >$ is absolutely continuous with respect to the Lebesgue measure:

$$d < P >_t \;=\; \varrho(P_t)dt, \tag{6.2.2}$$

for some positive continuous function ϱ on P. We denote by $\delta > 0$ the tick size, and we shall assume that the spread is constantly equal to δ, i.e. the best ask (resp. bid) price is $P^a := P + \frac{\delta}{2}$ (resp. $P^b := P - \frac{\delta}{2}$). This assumption corresponds to the case of the so-called *one-tick microstructure* [25], which can be encountered e.g. on short term interest rates futures contracts.

Trading strategies. For most of investors, the brokerage costs are paid when a transaction occurs, but new limit order submission, update or cancel are free of charge. Therefore, the investor can submit or update her quotes at any time, with no costs associated to this operation: it is then natural to model the limit order strategy (*make* strategy) as a continuous time predictable control process. On the contrary, market orders lead to immediate execution, and are costly, so that continuous submission of market orders would

lead to bankruptcy. Therefore, we choose to model the market order strategy (*take* strategy) as impulse controls. More precisely, we model trading strategies by a pair $\alpha = (\alpha^{make}, \alpha^{take})$ in the form:

$$\alpha^{make} \;=\; \big(L^a_t, L^b_t\big)_{t\geq 0}, \qquad \alpha^{take} \;=\; \big(\tau_n, \xi_n\big)_{n\in\mathbb{N}}.$$

The predictable processes L^a and L^b, valued in $\{0,1\}$ represent the possible *make regimes*: when $L^a_t = 1$ (resp. $L^b_t = 1$) this means that the investor has active limit orders at the best ask price (resp. best bid price) at time t, else, if $L^a_t = 0$ (resp. $L^b_t = 0$) this means that the investor has no active order at the best ask price (resp. best bid price) at time t. Practical implementation of such rule would be, for example, to send a limit order with a fixed quantity, when the corresponding control is 1, and cancel it when it turns to 0. Another practical implementation of the rule would be to post a constant proportion of the available volume at best prices: for example, if $V^a_M(t)$ is the current offered volume at best ask, and if $L^a_t = 1$, the practical action in this situation is to post a limit order of volume $v^a(t)$ s.t. $\dfrac{v^a(t)}{v^a(t) + V^a_M(t)} = const$ at the best ask price. Choice of practical implementation of the limit order controls will impact the outcome of the high frequency trader's strategy in term of executed volumes, and therefore we propose at the next paragraph an approach suitable in both cases. On the other hand, $(\tau_n)_{n\in\mathbb{N}}$ is an increasing sequence of stopping times, representing the times when the investor chooses to trade at market, and ξ_n, $n \geq 0$ are $\mathcal{F}_{\tau_n}$-measurable random variables valued in $\mathbb{R}$, representing the quantity purchased if $\xi_n \geq 0$ or sold if $\xi_n < 0$.

Execution processes in a pro-rata microstructure. The pro-rata microstructure (see [43] for extensive presentation and discussion) can be schematically described as follows[2]: when a market order comes in the pro-rata limit order book, its volume is dispatched among all active limit orders at best prices, proportionally to each limit orders volumes, and therefore create several transactions (see Figure 6.1).

This pro-rata microstructure fundamentally differs from price-time microstructure [37] for two reasons: first, several limit orders at the best prices receive incoming market order flow, regardless of the time priority, and second, market makers tend to oversize their liquidity offering (that is, posting limit order with much higher volume than they actually want to trade) in order to increase their transaction volume. For example, on the three-months EURIBOR futures contracts, the liquidity available at the best prices is 200 times higher than the average transaction size.

Let us examine more precisely the outcomes of the two practical implementations of limit orders posting mentionned in the last paragraph. We consider the two cases where

[2]For a detailled description of actual trading rules, and a general overview of STIR futures trading, we refer to [1] and references therein.

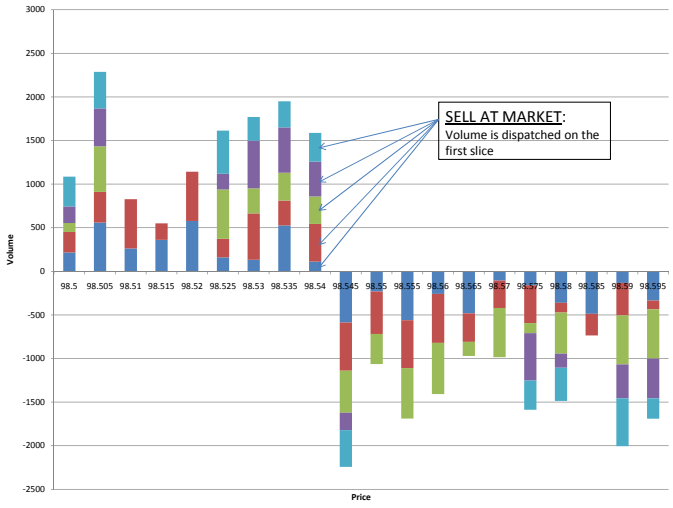

Figure 6.1: Schematic view of the pro-rata market microstructure.

the high frequency trader posts 1) limit orders with a fixed volume, say $V_0 = 100$ contracts, and 2) limits orders with volumes:

$$v^a(t) \text{ s.t. } \frac{v^a(t)}{v^a(t) + V_M^a(t)} = 10\% \;\; ; \;\; v^b(t) \text{ s.t. } \frac{v^b(t)}{v^b(t) + V_M^b(t)} = 10\%$$

where $V_M^a(t)$ (resp. $V_M^b(t)$) is the volume available at best ask (resp. bid) at time t. Considering an incoming market order of size V on the ask side, the high frequency trader receives:

- in case 1) $\min\left(V, V_0 + V_M^a(t)\right) \dfrac{V_0}{V_0 + V_M^a(t)} \leq V_0$
- and in case 2) $10\% \;\min\left(V, v^a(t) + V_M^a(t)\right) \leq v^a(t)$.

Note that in these two cases, the volume offered by the market maker is fully executed if and only if the market order's volume V is greater or equal to the total volume offered at ask $V_0 + V_M^a(t)$, resp. $v^a(t) + V_M^a(t)$. Therefore, the probability that the high frequency trader volume is fully executed is equal to the probability that the market order consume the first slice of the LOB in integrality. In other words, the volume $\dfrac{V_0}{V_0 + V_M^a(t)}$, resp. $\dfrac{v^a(t)}{v^a(t) + V_M^a(t)}$, that the HFT receives, never reaches the bound V_0, resp. $v^a(t)$, unless the market order consume the first slice of the LOB in integrality.

For illustration purposes, and in this discussion only, we assume that the volume of incoming market orders has a gamma distribution with shape 4 and scale 7.5 (which makes an average market order volume of 30 contracts, consistent with observations on the front 3-M EURIBOR contract, see [25]). In figure 6.2 we plot the probability of the HFT's limit order to be fully executed as a function of $V_0 + V_M^a(t)$, resp. $v^a(t) + V_M^a(t)$.

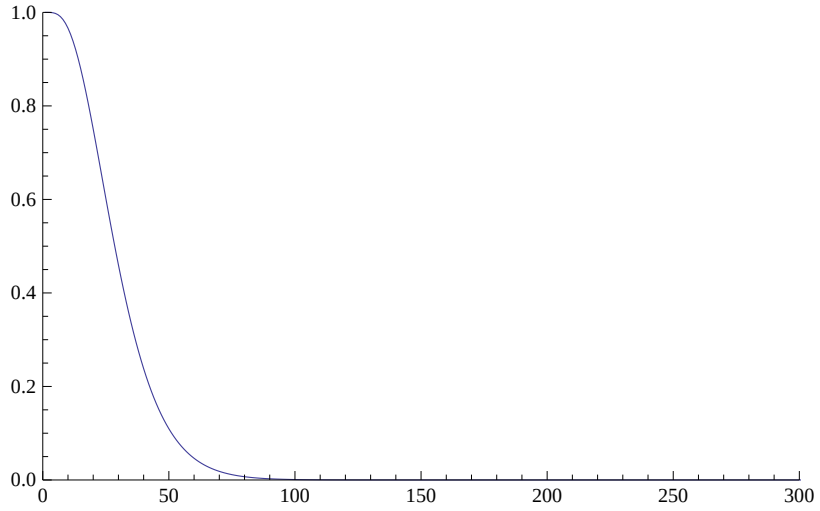

Figure 6.2: Probability of the HFT ask limit order V_0, resp. $v^a(t)$, to be fully executed as a function of total offered volume $V_0 + V_M^a(t)$, resp. $v^a(t) + V_M^a(t)$, when a market order of size $V \sim \text{Gamma}(4, 7.5)$ comes in the LOB at time t.

In this example, we see that the probability of the HFT limit order to be fully executed drops to negligible values once the total offered volume is greater than 100, which is about 3 times the average transaction size. Yet, in actual market, the average offered volume at the best priced slice is about 200 times larger than the average transaction size [25], and therefore, the probability that the HFT limit orders are fully executed is negligible. For example, if we use the average volume offered on best prices on the front EURIBOR future, 6000 contracts, the probability of such a market order consuming the first slice is 3×10^{-340}.

Therefore, our approach is to assume that the HFT's limit orders are never fully executed, and instead we model the executed volume as a random variable on which the market maker has no control. Indeed, the distribution of the volume of a single trade can be fitted directly on market data resulting from running our strategy. This approach combines the advantages of abstracting from practical details of the strategy implementation while keeping precise information on executed volumes. In other words, we assume that the outcome of the practical implementation of the strategy, in terms of executed volume distribution, is known and can be measured in market data, and especially in post-production data, i.e.

by examinating the real outcomes of trading with a given implementation of the strategy.

More precisely, let N^a (resp. N^b) be a Poisson process of intensity $\lambda^a > 0$ (resp. λ^b), whose jump times represent the times when execution by a market order flow occurs at best ask (resp. best bid), and we assume that N^a and N^b are independent. Let $(\zeta_n^a)_{n\in\mathbb{N}^*}$ and $(\zeta_n^b)_{n\in\mathbb{N}^*}$ be two independent sequences of i.i.d integrable random variables valued in $(0,\infty)$, of distribution laws μ^a and μ^b, which represent the transacted volume of the n^{th} execution at best ask and best bid. We denote by $\nu^a(dt,dz)$ (resp. $\nu^b(dt,dz)$) the Poisson random measure associated to the marked point process $(N^a,(\zeta_n^a)_{n\in\mathbb{N}^*})$ (resp. $(N^b,(\zeta_n^b)_{n\in\mathbb{N}^*})$) of intensity measure $\lambda^a\mu^a(dz)dt$ (resp. $\lambda^b\mu^b(dz)dt$), which is often identified with the compound Poisson processes

$$\vartheta_t^a \;=\; \sum_{n=1}^{N_t^a}\zeta_n^a \;=\; \int_0^t\int_0^\infty z\;\nu^a(dt,dz), \qquad \vartheta_t^b \;=\; \sum_{n=1}^{N_t^b}\zeta_n^b \;=\; \int_0^t\int_0^\infty z\;\nu^b(dt,dz). \tag{6.2.3}$$

representing the cumulative volume of transaction at ask, and bid, assumed to be independent of the mid-price process P. Notice that these processes model only the trades in which the investor has participated.

Cash holdings and inventory. The cash holdings process X and the cumulated number of stocks Y (also called inventory) hold by the investor evolve according to the following dynamics:

$$dX_t \;=\; L_t^a\big(P_{t^-}+\frac{\delta}{2}\big)d\vartheta_t^a - L_t^b\big(P_{t^-}-\frac{\delta}{2}\big)d\vartheta_t^b, \quad \tau_n\le t<\tau_{n+1} \tag{6.2.4}$$

$$dY_t \;=\; L_t^b d\vartheta_t^b - L_t^a d\vartheta_t^a, \qquad \tau_n\le t<\tau_{n+1} \tag{6.2.5}$$

$$X_{\tau_n}-X_{\tau_n-} \;=\; -\xi_n P_{\tau_n} - |\xi_n|\big(\frac{\delta}{2}+\varepsilon\big) - \varepsilon_0 1_{\xi_n\neq 0}, \tag{6.2.6}$$

$$Y_{\tau_n}-Y_{\tau_n-} \;=\; \xi_n. \tag{6.2.7}$$

The equations (6.2.4)-(6.2.5) model the evolution of the cash holdings and inventory under a limit order (make) strategy, while equations (6.2.6)-(6.2.7) describe the jump on the cash holdings and inventory when posting a market order (take) strategy, subject to a per share fee $\varepsilon > 0$ and a fixed fee $\varepsilon_0 > 0$. In the sequel, we impose the admissibility condition that the inventory should remain within a bounded interval $[-M_{_Y}, M_{_Y}]$, $M_{_Y} > 0$, after the trade at market, i.e. $\xi_n \in [-M_{_Y} - Y_{\tau_n-}, M_{_Y} - Y_{\tau_n-}]$, $n \ge 0$, and we shall denote by $\mathcal{A}$ the set of all admissible make and take strategies $\alpha = (\alpha^{make}, \alpha^{take})$.

Remark 6.2.1 Let us define the process $V_t = X_t + Y_t P_t$, which represents at time t the marked-to-market value of the portfolio (or book value of the portfolio). From (6.2.4)-

(6.2.5)-(6.2.6)-(6.2.7), we see that its dynamics is governed by:

$$dV_t = \frac{\delta}{2}(L_t^b d\vartheta_t^b + L_t^a d\vartheta_t^a) + Y_{t^-} dP_t, \tag{6.2.8}$$

$$V_{\tau_n} - V_{\tau_n -} = -|\xi_n|(\frac{\delta}{2} + \varepsilon) - \varepsilon_0 1_{\xi_n \neq 0}, . \tag{6.2.9}$$

In equation (6.2.9), we notice that a trade at market will always diminish the marked to market value of our portfolio, due to the fact that we have to "cross the spread", hence trade at a least favorable price. On the other hand, in equation (6.2.8), the term $\int \frac{\delta}{2}(L_t^b d\vartheta_t^b + L_t^a d\vartheta_t^a)$ is always positive, and represents the profit obtained from a limit order execution, while the term $\int Y_{t^-} dP_t$ represents the portfolio value when holding shares in the stock, hence inducing an inventory risk, which one wants to reduce its variance.

6.3 Market making optimization procedure

6.3.1 Control problem formulation

The market model in the previous section is fully determined by the state variables (X, Y, P) controlled by the limit/market orders strategies $\alpha = (\alpha^{make}, \alpha^{take}) \in \mathcal{A}$. The market maker wants to optimize her profit over a finite time horizon T (typically short term), while keeping control of her inventory risk, and to get rid of any risky asset by time T. We choose a mean-variance optimization criterion, and the goal is to

$$\text{maximize} \quad \mathbf{E}\Big[X_T - \gamma \int_0^T Y_{t^-}^2 d < P >_t \Big] \text{ over all strategies } \alpha \in \mathcal{A}, \text{ s.t } Y_T = 0, \tag{6.3.1}$$

with the convention that $\infty - \infty = -\infty$, as usually done in expected utility maximization. The integral $\int_0^T Y_{t^-}^2 d < P >_t$ is a quadratic penalization term for holding a non zero inventory in the stock, and $\gamma > 0$ is a risk aversion parameter chosen by the investor. The penalty term $\gamma \mathbf{E}\big[\int_0^T Y_{t^-}^2 d < P >_t \big]$ can further be motivated by noting that the variance of the total value of the investor's inventory in the case where P is a martingale is by the Itô isometry:

$$\mathbf{Var}\Big(\int_0^T Y_{t^-} dP_t\Big) = \mathbf{E}\Big[\int_0^T Y_{t^-}^2 d < P >_t \Big],$$

which is our penalty term, up to the scale factor γ. As pointed out by Cartea and Jaimungal [17], this running penalty is much more effective than the terminal inventory constraint.

Let us now rewrite problem (6.3.1) in a more standard formulation where the constraint $Y_T = 0$ on the inventory control is removed. For this, let us introduce the liquidation function:

$$L(x, y, p) = x + yp - |y|\big(\frac{\delta}{2} + \varepsilon\big) - \varepsilon_0 1_{y \neq 0},$$

which represents the cash obtained after an immediate liquidation of the inventory via a (non zero) market order. Then, problem (6.3.1) is formulated equivalently as

$$\text{maximize} \quad \mathbf{E}\Big[L(X_T,Y_T,P_T) - \gamma \int_0^T Y_{t^-}^2 \, d < P >_t \Big] \quad \text{over all strategies } \alpha \in \mathcal{A}, \tag{6.3.2}$$

Indeed, the maximal value of problem (6.3.1) is clearly smaller than the one of problem (6.3.2) since for any $\alpha \in \mathcal{A}$ s.t. $Y_T = 0$, we have $L(X_T, Y_T, P_T, S_T) = X_T$. Conversely, given an arbitrary $\alpha \in \mathcal{A}$, let us consider the control $\tilde{\alpha} \in \mathcal{A}$, coinciding with α up to time T, and to which one add at the terminal date T the admissible market order consisting in liquidating all the inventory Y_T if it is nonzero. The associated state process $(\tilde{X}, \tilde{Y}, P, S)$ satisfies: $\tilde{X}_t = X_t$, $\tilde{Y}_t = Y_t$ for $t < T$, and $\tilde{X}_T = L(X_T, Y_T, P_T, S_T)$, $\tilde{Y}_T = 0$. This shows that the maximal value of problem (6.3.2) is smaller and then equal to the maximal value of problem (6.3.1).

Recalling (6.2.2), let us then define the value function for the problem (6.3.2):

$$v(t,x,y,p) \;=\; \sup_{\alpha \in \mathcal{A}} \mathbf{E}_{t,x,y,p}\Big[L(X_T,Y_T,P_T) - \gamma \int_t^T Y_s^2 \varrho(P_s) ds\Big], \tag{6.3.3}$$

for $t \in [0,T]$, $(x,y,p) \in \mathbb{R}^2 \times \mathbb{P}$. Here, given $\alpha \in \mathcal{A}$, $\mathbf{E}_{t,x,y,p}$ denotes the expectation operator under which the process (X,Y,P) solution to (6.2.4)-(6.2.5)-(6.2.6)-(6.2.7) with initial state $(X_{t^-}, Y_{t^-}, P_{t^-}) = (x,y,p)$, is taken. Problem (6.3.3) is a mixed impulse/regular control problem in Markov model with jumps that we shall study by dynamic programming methods.

First, we state some bounds on the value function, which shows in particular that the value function is finite and locally bounded.

Proposition 6.3.1 *There exists some constant K_P (depending only on the price process and γ) such that for all $(t,x,y,p) \in [0,T] \times \mathbb{R}^2 \times \mathbb{P}$,*

$$L(x,y,p) \;\le\; v(t,x,y,p) \;\le\; x + yp + \frac{\delta}{2}\big(\lambda^a \bar{\mu}^a + \lambda^b \bar{\mu}^b\big)(T-t) + K_P, \tag{6.3.4}$$

where $\bar{\mu}^a = \int_0^\infty z\mu^a(dz)$, $\bar{\mu}^b = \int_0^\infty z\mu^b(dz)$ are the mean of the distribution laws μ^a and μ^b.

Proof. The lower bound in (6.3.4) is derived easily by considering the particular strategy, which consists of liquidating immediately all the current inventory (if non zero) via a market order, and then doing nothing else until the final horizon. Let us now focus on the upper bound. Observe that in the definition of the value function in (6.3.3), we can restrict obviously to controls $\alpha \in \mathcal{A}$ s.t.

$$\mathbf{E}\Big[\int_0^T Y_{t^-}^2 \, d < P >_t \Big] \;<\; \infty. \tag{6.3.5}$$

For such strategies, we have:

$$\begin{aligned}
& \mathbf{E}_{t,x,y,p}\Big[L(X_T,Y_T,P_T) - \gamma\int_t^T Y_{s^-}^2 d<P>_s\Big] \\
\leq\ & \mathbf{E}_{t,x,y,p}\Big[V_T - \gamma\int_t^T Y_s^2 d<P>_s\Big] \\
\leq\ & x+yp+\mathbf{E}_{t,x,y,p}\Big[\frac{\delta}{2}\big(\vartheta^a_{T-t}+\vartheta^b_{T-t}\big)+\int_t^T Y_{s^-}dP_s - \gamma\int_t^T Y_{s^-}^2 d<P>_s\Big] \\
=\ & x+yp+\mathbf{E}_{t,x,y,p}\Big[\frac{\delta}{2}\big(\vartheta^a_{T-t}+\vartheta^b_{T-t}\big)+\int_t^T \big(Y_{s^-}\theta_s - \gamma Y_{s^-}^2\big)d<P>_s\Big].
\end{aligned}$$

Here, the second inequality follows from the relation (6.2.8), together with the fact that $L^a, L^b \leq 1$, ϑ^a, ϑ^b are increasing processes, and also that jumps of V are negative by (6.2.9). The last equality holds true by (6.2.1) and the fact that $\int Y_- dM$ is a square-integrable martingale from (6.3.5), where M is the martingale part of the semimartingale P. Since θ is bounded and $\gamma > 0$, this shows that for all strategies α satisfying (6.3.5), we have:

$$\begin{aligned}
& \mathbf{E}_{t,x,y,p}\Big[L(X_T,Y_T,P_T) - \gamma\int_t^T Y_{s^-}^2 d<P>_s\Big] \\
\leq\ & x+yp+\frac{\delta}{2}\mathbf{E}\big[\vartheta^a_{T-t}+\vartheta^b_{T-t}\big]+K\mathbf{E}[<P>_T],
\end{aligned}$$

for some positive constant K, which proves the required result by recalling the characteristics of the compound Poisson processes ϑ^a and ϑ^b, and since $<P>_T$ is assumed to be integrable. □

Remark 6.3.1 The terms of the upper bound in (6.3.4) has a financial interpretation. The term $x+yp$ represents the marked-to-market value of the portfolio evaluated at mid-price, whereas the term K_P stands for a bound on profit for any directional frictionless strategy on the fictive asset that is priced P. The term $\frac{\delta}{2}\big(\lambda^a\bar{\mu}^a+\lambda^b\bar{\mu}^b\big)(T-t)$, always positive, represents the upper bound on profit due to market making, i.e. the profit of the strategy participating in every trade, but with no costs associated to managing its inventory.

6.3.2 Dynamic programming equation

For any $(\ell^a,\ell^b) \in \{0,1\}^2$, we introduce the non-local operator associated with the limit order control:

$$\mathcal{L}^{\ell^a,\ell^b} \quad = \quad \mathcal{P} + \ell^a\Gamma^a + \ell^b\Gamma^b, \tag{6.3.6}$$

where

$$\begin{aligned}\Gamma^a\phi(t,x,y,p) &= \lambda^a\int_0^\infty \big[\phi\big(t,x+z(p+\frac{\delta}{2}),y-z,p\big)-\phi(t,x,y,p)\big]\mu^a(dz)\\ \Gamma^b\phi(t,x,y,p) &= \lambda^b\int_0^\infty \big[\phi\big(t,x-z(p-\frac{\delta}{2}),y+z,p\big)-\phi(t,x,y,p)\big]\mu^b(dz),\end{aligned}$$

for (t,x,y,p) $[0,T]\times\mathbb{R}\times\mathbb{R}\times\mathbb{P}$. Let us also consider the impulse operator associated with admissible market order controls, and defined by:

$$\mathcal{M}\phi(t,x,y,p) = \sup_{e\in[-M_Y-y,M_Y-y]}\phi\big(t,x-ep-|e|(\frac{\delta}{2}+\varepsilon)-\varepsilon_0 1_{e\neq 0},y+e,p\big).$$

The dynamic programming equation (DPE) associated to the control problem (6.3.3) is a quasi-variational inequality (QVI) in the form:

$$\min\Big[-\frac{\partial v}{\partial t}-\sup_{(\ell^a,\ell^b)\in\{0,1\}^2}\mathcal{L}^{\ell^a,\ell^b}v+\gamma g\ ,\ v-\mathcal{M}v\Big] = 0,\ \text{ on }\ [0,T)\times\mathbb{R}^2\times\mathbb{P}, \quad (6.3.7)$$

together with the terminal condition:

$$v(T,.) = L,\qquad \text{on }\ \mathbb{R}^2\times\mathbb{P}, \quad (6.3.8)$$

where we denoted by g the function: $g(y,p)=y^2\varrho(p)$. By standard methods of dynamic programming, one can show that the value function in (6.3.3) is the unique viscosity solution under growth conditions determined by (6.3.4) to the DPE (6.3.7)-(6.3.8) of dimension 3 (in addition to the time variable), see e.g. Chap. 9 in [57].

6.3.3 Dimension reduction in the Lévy case

We now consider a special case on the mid-price process where the market making control problem can be reduced to the resolution of a one-dimensional variational inequality involving only the inventory state variable. We shall suppose actually that P is a Lévy process so that

$$\mathcal{P}I_{\mathbb{P}} = c_{_P},\quad \text{and}\quad \varrho \text{ is a constant}, \quad (6.3.9)$$

where $I_{\mathbb{P}}$ is the identity function on $\mathbb{P}$, i.e. $I_{\mathbb{P}}(p)=p$, and $\varrho>0$, $c_{_P}$ are real constants depending on the characteristics triplet of P. Two practical examples are:

• **Martingale case**: The mid-price process P is a martingale, so that $\mathcal{P}I_{\mathbb{P}}=0$. This martingale assumption in a high-frequency context reflects the idea that the market maker has no information on the future direction of the stock price.

• **Trend information**: To remove the martingale assumption, one can introduce some knowledge about the price trend. A typical simple example is when P follows an arithmetic

Brownian motion (Bachelier model). A more relevant example is described by a pure jump process P valued in the discrete grid $\delta\mathbb{Z}$ with tick $\delta > 0$, and such that

$$\begin{aligned}\mathbf{P}\big(P_{t+h} - P_t = \delta \,|\mathcal{F}_t\big) &= \pi^+ h + o(h)\\ \mathbf{P}\big(P_{t+h} - P_t = -\delta \,|\mathcal{F}_t\big) &= \pi^- h + o(h)\\ \mathbf{P}\big(|P_{t+h} - P_t| > \delta \,|\mathcal{F}_t\big) &= o(h),\end{aligned}$$

where $\pi^+, \pi^- > 0$, and $o(h)$ is the usual notation meaning that $\lim_{h\to 0} o(h)/h = 0$. Relation (6.3.9) then holds with $c_{_P} = \varpi\delta$, where $\varpi = \pi^+ - \pi^-$ represents a constant information about price direction, and $\varrho = (\pi^+ + \pi^-)\delta^2$. In a high-frequency context, this model is of practical interest as it provides a way to include a (predictive) information about price direction. For example, work have been done in [21] to infer the future prices movements (at the scale of a few seconds) from the current state of the limit order book in a Poisson framework. In this work, as well as in our real data tests, the main quantities of interest are the volume offered at the best prices in the limit order book, also known as the imbalance.

In this Lévy context, we can decompose the value function v is decomposed into the form:

$$v(t,x,y,p) \;=\; L_0(x,y,p) + w(t,y), \tag{6.3.10}$$

where $L_0(x,y,p) = x + yp - |y|\big(\frac{\delta}{2} + \varepsilon\big) = L(x,y,p) + \varepsilon_0 1_{y\neq 0}$ is the liquidation function up to the fixed fee, and where w is solution to the integral variational inequality:

$$\min\Big[-\frac{\partial w}{\partial t} - yc_{_P} + \gamma\varrho y^2 - \mathcal{I}^a w - \mathcal{I}^b w\;,\; w - \tilde{\mathcal{M}} w\Big] \;=\; 0, \;\text{ on }\; [0,T)\times\mathbb{R}, \tag{6.3.11}$$

together with the terminal condition:

$$w(T,y) \;=\; -\varepsilon_0 1_{y\neq 0}, \quad \forall y\in\mathbb{R}, \tag{6.3.12}$$

with $\mathcal{I}^a$ and $\mathcal{I}^b$, the nonlocal integral operators:

$$\begin{aligned}\mathcal{I}^a w(t,y) &= \lambda^a\Big(\int_0^\infty \Big[w(t,y-z) - w(t,y) + z\frac{\delta}{2} + (\frac{\delta}{2}+\varepsilon)(|y| - |y-z|)\Big]\mu^a(dz)\Big)_+\\ \mathcal{I}^b w(t,y) &= \lambda^b\Big(\int_0^\infty \Big[w(t,y+z) - w(t,y) + z\frac{\delta}{2} + (\frac{\delta}{2}+\varepsilon)(|y| - |y+z|)\Big]\mu^b(dz)\Big)_+,\end{aligned}$$

and $\tilde{\mathcal{M}}$, the nonlocal operator:

$$\tilde{\mathcal{M}} w(t,y) \;=\; \sup_{e\in[-M_{_Y}-y, M_{_Y}-y]}\Big[w(t,y+e) - (\frac{\delta}{2}+\varepsilon)(|y+e| + |e| - |y|) - \varepsilon_0 1_{e\neq 0}\Big].$$

The interpretation of the decomposition (6.3.10) is the following. The term $L_0(x,y,p)$ represents the book value that the investor would obtain by liquidating immediately with

a market order (up to the fixed fee), and w is an additional correction term taking into account the illiquidity effects induced by the bid-ask spread and the fees, as well as the execution risk when submitting limit orders. Moreover, in the Lévy case, this correction function w depends only on time and inventory. From (6.3.4), we have the following bounds on the function w:

$$-\varepsilon_0 1_{y\neq 0} \;\leq\; w(t,y) \;\leq (\frac{\delta}{2}+\varepsilon)|y| + \frac{\delta}{2}\big(\lambda^a\bar{\mu}^a + \lambda^b\bar{\mu}^b\big)(T-t) + K_P, \quad \forall (t,y)\in[0,T]\times\mathbb{R}.$$

Actually, we have a sharper upper bound in the Lévy context.

Proposition 6.3.2 *Under* (6.3.9)*, we have:*

$$-\varepsilon_0 1_{y\neq 0} \;\leq\; w(t,y) \;\leq\; (T-t)\Big[\frac{c_P^2}{4\gamma\rho} + \lambda^a(\delta+\epsilon)\bar{\mu}^a + \lambda^b(\delta+\epsilon)\bar{\mu}^b\Big], \tag{6.3.13}$$

for all $(t,x,y,p)\in[0,T]\times\mathbb{R}^2\times\mathbb{P}$.

Proof. For any $(x,y,p)\in\mathbb{R}^2\times\mathbb{P}$, we notice that

$$\begin{aligned} & L_0(x,y,p) - \sup_{e\in[-M_Y-y,M_Y-y]} L_0(x-ep-|e|(\frac{\delta}{2}+\varepsilon)-\varepsilon_0 1_{e\neq 0}, y+e, p) \\ = & \inf_{e\in[-M_Y-y,M_Y-y]} \Big[(\frac{\delta}{2}+\varepsilon)\Big(|e|+|y+e|-|y|\Big) + \varepsilon_0 1_{e\neq 0}\Big] \;\geq\; 0. \end{aligned} \tag{6.3.14}$$

We also observe that for all $z\geq 0$:

$$\begin{aligned} L_0(x+z(p+\frac{\delta}{2}), y-z, p) - L_0(x,y,p) &= z\frac{\delta}{2} + (\frac{\delta}{2}+\varepsilon)\big(|y|-|y-z|\big) \\ &\leq (\delta+\varepsilon)z, \end{aligned} \tag{6.3.15}$$

and similarly:

$$L_0(x-z(p-\frac{\delta}{2}), y+z, p) - L_0(x,y,p) \;\leq\; (\delta+\varepsilon)z. \tag{6.3.16}$$

Let us then consider the function $\phi(t,x,y,p) = L_0(x,y,p)+(T-t)u$, for some real constant u to be determined later. Then, $\phi(T,.) = L_0$, and by (6.3.15)-(6.3.16), we easily check that:

$$\begin{aligned} & -\frac{\partial\phi}{\partial t} - \sup_{(\ell^a,\ell^b)\in\{0,1\}^2} \mathcal{L}^{\ell^a,\ell^b}\phi \;+\; \gamma g \\ \geq\; & u - \lambda^a(\delta+\varepsilon)\bar{\mu}^a - \lambda^b(\delta+\varepsilon)\bar{\mu}^b - yc_P + \gamma y^2\rho. \end{aligned}$$

The r.h.s. of this last inequality is a second order polynomial in y and therefore it is always nonnegative iff:

$$c_P^2 - 4\gamma\rho(u - \lambda^a(\delta+\varepsilon)\bar{\mu}^a - \lambda^b(\delta+\varepsilon)\bar{\mu}^b) \;\leq\; 0,$$

which is satisfied once the constant u is large enough, namely:

$$u \;\geq\; \hat{u} \;:=\; \frac{c_P^2}{4\gamma\rho} + \lambda^a(\delta+\epsilon)\bar{\mu}^a + \lambda^b(\delta+\epsilon)\bar{\mu}^b.$$

For such choice of $u = \hat{u}$, and denoting by $\hat{\phi}$ the associated function: $\hat{\phi}(t,x,y,p) = L_0(x,y,p) + (T-t)\hat{u}$ we have

$$-\frac{\partial\hat{\phi}}{\partial t} - \sup_{(\ell^a,\ell^b)\in\{0,1\}^2} \mathcal{L}^{\ell^a,\ell^b}\hat{\phi} + \gamma g \;\geq\; 0,$$

which shows, together with (6.3.14), that $\hat{\phi}$ is a supersolution of (6.3.7)-(6.3.8). From comparison principle for this variational inequality, we deduce that

$$v \;\leq\; \hat{\phi} \;\text{ on }\; [0,T]\times\mathbb{R}^2\times\mathbb{P},$$

which shows the required upper bound for $w = v - L_0$. □

Finally, from (6.3.11)-(6.3.12), and in the case where $\lambda^a = \lambda^b$, $\mu^a = \mu^b$, and by stressing the dependence of w in $c_{_P}$, we see that w satisfies the symmetry relation:

$$w(t,y,c_{_P}) \;=\; w(t,-y,-c_{_P}), \quad \forall (t,y)\in[0,T]\times\mathbb{R}. \tag{6.3.17}$$

6.4 Numerical resolution

In this section, we focus on the numerical resolution of the integral variational inequality (6.3.11)-(6.3.12), which characterizes the reduced value function of the market-making problem in the Lévy case.

6.4.1 Numerical scheme

We provide a computational scheme for the integral variational inequality (6.3.11). We first consider a time discretization of the interval $[0,T]$ with time step $h = T/N$ and a regular time grid $\mathbb{T}_N = \{t_k = kh\,,\, k = 0,\ldots,N\}$. Next, we discretize and localize the inventory state space on a finite regular grid: for any $M > 0$ (in practice we choose $M = M_Y$) and $N_Y \in \mathbb{N}$, and denoting by $\Delta_Y = \dfrac{M}{N_Y}$, we set:

$$\mathbb{Y}_M \;=\; \big\{y_i = i\Delta_Y,\; i = -N_Y,\ldots,N_Y\big\}.$$

We denote by $\mathrm{Proj}_M(y) := -M \vee (y \wedge M)$, and consider the discrete approximating distribution of μ^a and μ^b, defined by:

$$\hat{\mu}^a \;=\; \sum_{i\in\mathbb{Z}^+} \mu^a([i\Delta_Y;(i+1)\Delta_Y))\delta_{i\Delta_Y}\;, \qquad \hat{\mu}^b \;=\; \sum_{i\in\mathbb{Z}^+} \mu^b([i\Delta_Y;(i+1)\Delta_Y))\delta_{i\Delta_Y},$$

with δ_x the Dirac measure at x. We then introduce the operator associated to the explicit time-space discretization of the integral variational inequality (6.3.11): for any real-valued function φ on $[0,T]\times\mathbb{R}$, $t\in[0,T]$, and $y\in\mathbb{R}$, we define:

$$\mathcal{S}^{h,\Delta_Y,M}(t,y,\varphi) \;=\; \max\Big[\mathcal{T}^{h,\Delta_Y,M}(t,y,\varphi)\;;\;\tilde{\mathcal{M}}^{h,\Delta_Y,M}(t,y,\varphi)\Big],$$

where

$$\begin{aligned}\mathcal{T}^{h,\Delta_Y,M}(t,y,\varphi) \;=\;& \varphi(t,y)-h\gamma\varrho y^2+hyc_P\\ &+\lambda^a h\Big(\int_0^\infty\big[\varphi(t,\mathrm{Proj}_M(y-z))-\varphi(t,y)\big]\hat{\mu}^a(dz)\\ &\qquad+\int_0^\infty\big[\frac{\delta}{2}z+(\frac{\delta}{2}+\varepsilon)(|y|-|y-z|)\big]\mu^a(dz)\Big)_+\\ &+\lambda^b h\Big(\int_0^\infty\big[\varphi(t,\mathrm{Proj}_M(y+z))-\varphi(t,y)\big]\hat{\mu}^b(dz)\\ &\qquad+\int_0^\infty\big[\frac{\delta}{2}z+(\frac{\delta}{2}+\varepsilon)(|y|-|y+z|)\big]\mu^b(dz)\Big)_+,\end{aligned}$$

and

$$\begin{aligned}&\tilde{\mathcal{M}}^{h,\Delta_Y,M}(t,y,\varphi)\\ =\;&\sup_{e\in\mathbb{Y}_M\cap[-M_Y-y,M_Y-y]}\big[\varphi(t,\mathrm{Proj}_M(y+e))-(\frac{\delta}{2}+\varepsilon)(|y+e|+|e|-|y|)-\varepsilon_0 1_{e\neq 0}\big].\end{aligned}\tag{6.4.1}$$

By recalling that $x_+=\max_{\ell\in\{0,1\}}\ell x$, we see that the operator $\mathcal{T}^{h,\Delta_Y,M}$ may be written also as:

$$\begin{aligned}\mathcal{T}^{h,\Delta_Y,M}(t,y,\varphi) \;=\;& -h\gamma\varrho y^2+hyc_P+\max_{\ell^a,\ell^b\in\{0,1\}}\Big[\varphi(t,y)(1-\lambda^a h\ell^a-\lambda^b h\ell^b)\\ &+\lambda^a h\ell^a\Big(\int_0^\infty\varphi(t,\mathrm{Proj}_M(y-z))\hat{\mu}^a(dz)\\ &\qquad+\int_0^\infty\big[\frac{\delta}{2}z+(\frac{\delta}{2}+\varepsilon)(|y|-|y-z|)\big]\mu^a(dz)\Big)\\ &+\lambda^b h\ell^b\Big(\int_0^\infty\varphi(t,\mathrm{Proj}_M(y+z))\hat{\mu}^b(dz)\\ &\qquad+\int_0^\infty\big[\frac{\delta}{2}z+(\frac{\delta}{2}+\varepsilon)(|y|-|y+z|)\big]\mu^b(dz)\Big)\Big].\end{aligned}\tag{6.4.2}$$

Notice that on the boundary $y=M_Y$ (resp. $y=-M_Y$) the set of admissible market orders is $[-2y,0]$ (resp. $[0,-2y]$) which implies that we only allow sell (resp. buy) market orders. Limit orders controls can be of any type on the boundary, since we do not set a global constraint on the inventory.

We then approximate the solution w to (6.3.11)-(6.3.12) by the function $w^{h,\Delta_Y,M}$ on $\mathbb{T}_N \times \mathbb{Y}_M$ solution to the computational scheme:

$$w^{h,\Delta_Y,M}(t_N, y) = -\varepsilon_0 1_{y\neq 0}, \quad y \in \mathbb{Y}_M, \tag{6.4.3}$$

$$w^{h,\Delta_Y,M}(t_k, y) = \mathcal{S}^{h,\Delta_Y,M}(t_{k+1}, y, w^{h,\Delta_Y,M}) \ , \ k = 0, \ldots, N-1 \ , \ y \in \mathbb{Y}_M. \tag{6.4.4}$$

This algorithm is described explicitly in backward induction by the following pseudo-code:

- Timestep $t_N = T$: for each $y \in \mathbb{Y}_M$, set $w^{h,\Delta_Y,M}(t_n, y) := -\varepsilon_0 1_{y\neq 0}$
- For $k = N-1 \ldots 0$, from timestep t_{k+1} to timestep t_k, and for each $y \in \mathbb{Y}_M$:
  - Compute $\mathcal{T}^{h,\Delta_Y,M}(t_{k+1}, y, w^{h,\Delta_Y,M})$ from (6.4.2), and store $\ell^{a,\star}$, $\ell^{b,\star}$ the argmax
  - Compute $\tilde{\mathcal{M}}^{h,\Delta_Y,M}(t_{k+1}, y, w^{h,\Delta_Y,M})$ from (6.4.1), and store $e^\star$ the argmax
  - If $\mathcal{T}^{h,\Delta_Y,M}(t_{k+1}, y, w^{h,\Delta_Y,M}) \geq \tilde{\mathcal{M}}^{h,\Delta_Y,M}(t_{k+1}, y, w^{h,\Delta_Y,M})$ then set

$$w^{h,\Delta_Y,M}(t_k, y) := \mathcal{T}^{h,\Delta_Y,M}(t_{k+1}, y, w^{h,\Delta_Y,M})$$

    and the policy is make $(\ell^{a,\star}, \ell^{b,\star})$. Otherwise, set

$$w^{h,\Delta_Y,M}(t_k, y) := \tilde{\mathcal{M}}^{h,\Delta_Y,M}(t_{k+1}, y, w^{h,\Delta_Y,M})$$

    and the policy is take $e^\star$.

6.4.2 Convergence of the numerical scheme

In this section, we study the convergence of the numerical scheme (6.4.3)-(6.4.4) by showing the monotonicity, stability and consistency properties of this scheme. We denote by $C_b^1([0,T] \times \mathbb{R})$ the set of bounded continuously differentiable functions on $[0,T] \times \mathbb{R}$, with bounded derivatives.

Proposition 6.4.1 *(Monotonicity)*
For any $h > 0$ s.t. $h < \dfrac{1}{\lambda^a + \lambda^b}$ the operator $\mathcal{S}^{h,\Delta_Y,M}$ is non-decreasing in φ, i.e. for any $(t,y) \in [0,T] \times \mathbb{R}$ and any $\varphi, \psi \in C_b^1([0,T] \times \mathbb{R})$, s.t. $\varphi \leq \psi$:

$$\mathcal{S}^{h,\Delta_Y,M}(t, y, \varphi) \leq \mathcal{S}^{h,\Delta_Y,M}(t, y, \psi)$$

Proof. From the expression (6.4.2), it is clear that $\mathcal{T}^{h,\Delta_Y,M}(t, y, \varphi)$, and then also $\mathcal{S}^{h,\Delta_Y,M}(t, y, \varphi)$ is monotone in φ once $1 - \lambda^a h - \lambda^b h > 0$. □

Proposition 6.4.2 *(Stability)*
For any $h, \Delta_Y, M > 0$ there exists a unique solution $w^{h,\Delta_Y,M}$ to (6.4.3)-(6.4.4), and the sequence $(w^{h,\Delta_Y,M})$ is uniformly bounded: for any $(t,y) \in \mathbb{T}_N \times \mathbb{Y}_M$,

$$-\varepsilon_0 1_{y\neq 0} \;\leq\; w^{h,\Delta_Y,M}(t,y) \;\;\leq\;\; (T-t)\Big[\frac{c_P^2}{4\gamma\rho} + \lambda^a(\delta+\epsilon)\bar{\mu}^a + \lambda^b(\delta+\epsilon)\bar{\mu}^b\Big].$$

Proof. Existence and uniqueness of $w^{h,\Delta_Y,M}$ follows from the explicit backward scheme (6.4.3)-(6.4.4). Let us now prove the uniform bounds. We consider the function

$$\Psi^\star(t) \;\;=\;\; (T-t)\left[\frac{c_P^2}{4\gamma\rho} + \lambda^a(\delta+\epsilon)\bar{\mu}^a + \lambda^b(\delta+\epsilon)\bar{\mu}^b\right]$$

and notice that $\Psi^\star(t) \geq \mathcal{S}^{h,\Delta_Y,M}(t+h, y, \Psi^\star)$ by the same arguments as in Proposition 6.3.2. Moreover, we have, by definition, $w^{h,\Delta_Y,M}(T,y) = -\varepsilon_0 1_{y\neq 0} \leq \Psi^\star(T) = 0$, and therefore, a direct recurrence from (6.4.3)-(6.4.4) shows that $w^{h,\Delta_Y,M}(t,y) \leq \Psi^\star(t)$ for all $(t,y) \in \mathbb{T}_n \times \mathbb{Y}_M$, which is the required upper bound for $w^{h,\Delta_Y,M}$.

On the other hand, we notice that $\mathcal{S}^{h,\Delta_Y,M}(t,0,\varphi) \geq \varphi(t,0)$ for any function φ on $[0,T]\times\mathbb{R}$, and $t \in [0,T]$, by considering the"diffusive" part of the numerical scheme with the particular controls $\ell^a = \ell^b = 0$. Therefore, since $w^{h,\Delta_Y,M}(T,0) = 0$, we obtain by induction on (6.4.3)-(6.4.4) that $w^{h,\Delta_Y,M}(t,0) \geq 0$ for any $t \in \mathbb{T}_N$. Finally, considering the obstacle part of the numerical scheme, with the particular control $e = -y$, shows that $w^{h,\Delta_Y,M}(t,y) \geq w^{h,\Delta_Y,M}(t,0) - \varepsilon_0 1_{y\neq 0} \geq -\varepsilon_0 1_{y\neq 0}$ for any $(t,y) \in \mathbb{T}_N \times \mathbb{Y}_M$, which proves the required lower bound for $w^{h,\Delta_Y,M}$. □

Proposition 6.4.3 *(Consistency)*
For all $(t,y) \in [0,T)\times\mathbb{R}$ and $\varphi \in C_b^1([0,T]\times\mathbb{R})$, we have

$$\lim_{\substack{(h,\,\Delta_Y,\,M)\,\to\,(0,\,0,\,\infty)\\(t',\,y')\,\to\,(t,\,y)}} \frac{1}{h}\Big[\varphi(t',y') - \mathcal{T}^{h,\Delta_Y,M}(t'+h,y',\varphi)\Big] \tag{6.4.5}$$
$$= \;\; -\frac{\partial\varphi}{\partial t}(t,y) - yc_{_P} + \gamma\varrho y^2 - \mathcal{I}^a\varphi(t,y) - \mathcal{I}^b\varphi(t,y)$$

and

$$\lim_{\substack{(h,\,\Delta_Y,\,M)\,\to\,(0,\,0,\,\infty)\\(t',\,y')\,\to\,(t,\,y)}} \tilde{\mathcal{M}}^{h,\Delta_Y,M}(t'+h,y',\varphi) \;\;=\;\; \tilde{\mathcal{M}}\varphi(t,y) \tag{6.4.6}$$

Proof. The consistency relation (6.4.6) follows from the continuity of the function (t,y,e) $\to \varphi(t,y+e) - (\frac{\delta}{2}+\varepsilon)(|y+e|+|e|-|y|) - \varepsilon_0$. On the other hand, we have for all (t',y') $\in [0,T)\times\mathbb{R}$,

$$\begin{aligned}\frac{1}{h}\Big[\varphi(t',y') - \mathcal{T}^{h,\Delta_Y,M}(t'+h,y',\varphi)\Big] \;\;&=\;\; \frac{1}{h}\big[\varphi(t',y') - \varphi(t'+h,y')\big] - y'c_{_P} + \gamma\rho y'^2 \\ &\quad - \mathcal{I}_a^{h,\Delta_Y,M}(t'+h,y',\varphi) - \mathcal{I}_b^{h,\Delta_Y,M}(t'+h,y',\varphi),\end{aligned} \tag{6.4.7}$$

where

$$\begin{aligned}
\mathcal{I}_a^{h,\Delta_Y,M}(t,y,\varphi) &= \lambda^a\Big(\int_0^\infty \big[\varphi(t,\mathrm{Proj}_M(y-z))-\varphi(t,y)\big]\hat{\mu}^a(dz) \\
&\qquad + \int_0^\infty \big[\frac{\delta}{2}z+(\frac{\delta}{2}+\varepsilon)(|y|-|y-z|)\big]\mu^a(dz)\Big)_+ \\
\mathcal{I}_a^{h,\Delta_Y,M}(t,y,\varphi) &= \lambda^b\Big(\int_0^\infty \big[\varphi(t,\mathrm{Proj}_M(y+z))-\varphi(t,y)\big]\hat{\mu}^b(dz) \\
&\qquad + \int_0^\infty \big[\frac{\delta}{2}z+(\frac{\delta}{2}+\varepsilon)(|y|-|y+z|)\big]\mu^b(dz)\Big)_+.
\end{aligned}$$

The three first terms of (6.4.7) converge trivially to $-\dfrac{\partial\varphi}{\partial t}(t,y)-yc_{_P}+\gamma\varrho y^2$ as h goes to zero and (t',y') goes to (t,y). Hence, in order to get the consistency relation, it remains to prove the convergence of $\mathcal{I}_a^{h,\Delta_Y,M}(t'+h,y',\varphi)$ to $\mathcal{I}^a\varphi(t,y)$ as (h,Δ_Y,M) goes to $(0,0,\infty)$, and (t',y') goes to (t,y) (an identical argument holds for $\mathcal{I}_b^{h,\Delta_Y,M}(t'+h,y',\varphi)$). By writing that $|x_+-x'_+|\le|x-x'|$, we have

$$\begin{aligned}
&\Big|\mathcal{I}_a^{h,\Delta_Y,M}(t'+h,y',\varphi)-\mathcal{I}^a\varphi(t,y)\Big| \\
\le\ & \lambda^a\big|\varphi(t'+h,y')-\varphi(t,y)\big| \\
&+\lambda^a\Big|\int_0^\infty \varphi(t'+h,\mathrm{Proj}_M(y'-z))\hat{\mu}^a(dz)\ -\ \int_0^\infty \varphi(t,y-z)\mu^a(dz)\Big| \\
\le\ & \lambda^a\big|\varphi(t'+h,y')-\varphi(t,y)\big| \\
&+\lambda^a\Big|\int_0^{M+y'} \varphi(t'+h,y'-z)\hat{\mu}^a(dz)\ -\ \int_0^{M+y'} \varphi(t,y-z)\mu^a(dz)\Big| \\
&+\lambda^a\Big|\int_{M+y'}^\infty \varphi(t'+h,-M)\hat{\mu}^a(dz)\ -\ \int_{M+y'}^\infty \varphi(t,y-z)\mu^a(dz)\Big| \\
\le\ & \lambda^a\big|\varphi(t'+h,y')-\varphi(t,y)\big| \\
&+\lambda^a\int_0^\infty \big|\varphi(t'+h,y'-\kappa(z))\ -\ \varphi(t,y-z)\big|\mu^a(dz) \\
&+2\lambda^a\|\varphi\|_\infty\mu^a\big([M+y',\infty)\big),
\end{aligned}$$

where we denote by $\kappa(z)=\lfloor\dfrac{z}{\Delta_Y}\rfloor\Delta_Y$. Since the smooth function φ has bounded derivatives, say bounded by $\|\varphi^{(1)}\|_\infty$, it follows that

$$\begin{aligned}
\Big|\mathcal{I}_a^{h,\Delta_Y,M}(t'+h,y',\varphi)-\mathcal{I}^a\varphi(t,y)\Big| \ &\le\ \lambda^a\|\varphi^{(1)}\|_\infty\big(h+2|y'-y|+\Delta_Y\big) \\
&\quad +2\lambda^a\|\varphi\|_\infty\mu^a\big([M+y',\infty)\big),
\end{aligned}$$

which proves that

$$\lim_{\substack{(h,\,\Delta_Y,\,M)\,\to\,(0,\,0,\,\infty)\\(t',\,y')\,\to\,(t,\,y)}} \mathcal{I}_a^{h,\Delta_Y,M}(t'+h,y',\varphi) \ =\ \mathcal{I}^a\varphi(t,y),$$

hence completing the consistency relation (6.4.5). □

Theorem 6.4.1 *(Convergence)*
The solution $w^{h,\Delta_Y,M}$ to the numerical scheme ((6.4.3)-(6.4.4)) converges locally uniformly to w on $[0,T)\times\mathbb{R}$, as (h,Δ_Y,M) goes to $(0,0,\infty)$.

Proof. Given the above monotonicity, stability and consistency properties, the convergence of the sequence $(w^{h,\Delta_Y,M})$ towards w, which is the unique bounded viscosity solution to (6.3.11)-(6.3.12), follows from [8]. We report the arguments for sake of completeness. From the stability property, the semi-relaxed limits:

$$\begin{aligned} w_*(t,y) &= \liminf_{\substack{(h,\Delta_Y,M)\to(0,0,\infty)\\(t',y')\to(t,y)}} w^{h,\Delta_Y,M}(t',y'),\\ w^*(t,y) &= \limsup_{\substack{(h,\Delta_Y,M)\to(0,0,\infty)\\(t',y')\to(t,y)}} w^{h,\Delta_Y,M}(t',y'), \end{aligned}$$

are finite lower-semicontinuous and upper-semicontinuous functions on $[0,T]\times\mathbb{R}$, and inherit the boundedness of $(w^{h,\Delta_Y,M})$. We claim that w_* are w^* are respectively viscosity super and subsolution of (6.3.11)-(6.3.12). Assuming for the moment that this claim is true, we obtain by the strong comparison principle for (6.3.11)-(6.3.12) that $w^*\le w_*$. Since the converse inequality is obvious by the very definition of w_* and w^*, this shows that w_* $= w^* = w$ is the unique bounded continuous viscosity solution to (6.3.11)-(6.3.12), hence completing the proof of convergence.

In the sequel, we prove the viscosity supersolution property of w_* (a symmetric argument for the viscosity subsolution property of w^* holds true). Let $(\bar t,\bar y)\in[0,T)\times\mathbb{R}$ and φ a test function in $C_b^1([0,T]\times\mathbb{R})$ s.t. $(\bar t,\bar y)$ is a strict global minimimum point of $w_*-\varphi$. Then, one can find a sequence (t'_n,y'_n) in $[0,T)\times\mathbb{R}$, and a sequence (h_n,Δ_Y^n,M_n) such that:

$$\begin{aligned} (t'_n,y'_n)\ \to\ (\bar t,\bar y),\quad (h_n,\Delta_Y^n,M_n)\ \to\ (0,0,\infty),\quad w^{h_n,\Delta_Y^n,M_n}\ \to\ w_*(\bar t,\bar y),\\ (t'_n,y'_n)\ \text{ is a global minimum point of } w^{h_n,\Delta_Y^n,M_n}-\varphi. \end{aligned}$$

Denoting by $\zeta_n=(w^{h_n,\Delta_Y^n,M_n}-\varphi)(t'_n,y'_n)$, we have $w^{h_n,\Delta_Y^n,M_n}\ge\varphi+\zeta_n$. From the definition of the numerical scheme $\mathcal{S}^{h_n,\Delta_Y^n,M_n}$, and its monotonicity, we then get:

$$\begin{aligned} \zeta_n+\varphi(t'_n,y'_n) &= w^{h_n,\Delta_Y^n,M_n}(t'_n,y'_n)\\ &= \mathcal{S}^{h_n,\Delta_Y^n,M_n}(t'_n+h_n,y'_n,w^{h_n,\Delta_Y^n,M_n})\\ &\ge \mathcal{S}^{h_n,\Delta_Y^n,M_n}(t'_n+h_n,y'_n,\varphi+\zeta_n)\ =\ \mathcal{S}^{h_n,\Delta_Y^n,M_n}(t'_n+h_n,y'_n,\varphi)+\zeta_n\\ &= \max\Big[\mathcal{T}^{h_n,\Delta_Y^n,M_n}(t'_n+h_n,y'_n,\varphi)\ ,\ \tilde{\mathcal{M}}^{h_n,\Delta_Y^n,M_n}(t'_n+h_n,y'_n,\varphi)\Big]+\zeta_n, \end{aligned}$$

which implies

$$\min\Big[\frac{\varphi(t'_n,y'_n)-\mathcal{T}^{h_n,\Delta_Y^n,M_n}(t'_n+h_n,y'_n,\varphi)}{h_n}\ ,\ \varphi(t'_n,y'_n)-\tilde{\mathcal{M}}^{h_n,\Delta_Y^n,M_n}(t'_n+h_n,y'_n,\varphi)\Big]\ \geq\ 0.$$

By the consistency properties (6.4.5)-(6.4.6), and by sending n to infinity in the above inequality, we obtain the required viscosity supersolution property:

$$\min\Big[-\frac{\partial\varphi}{\partial t}(\bar{t},\bar{y})-\bar{y}c_{_P}+\gamma\varrho\bar{y}^2-\mathcal{I}^a\varphi(\bar{t},\bar{y})-\mathcal{I}^b\varphi(\bar{t},\bar{y})\ ,\ \varphi(\bar{t},\bar{y})-\tilde{\mathcal{M}}\varphi(\bar{t},\bar{y})\Big]\ \ \geq\ \ 0.$$

□

6.4.3 Numerical tests

In this section, we provide numerical results for the (reduced-form) value function and optimal policies in the martingale case and the trend information case, and a backtest on simulated data for the trend information case.

Within this section, we will denote by w^h the value function and by $\alpha^\star$ the make/take strategy associated with the backward numerical scheme (6.4.3)-(6.4.4). Given a generic controlled process Z and a control $\alpha\in\mathcal{A}$, we will denote Z^α the process controlled by α. Unless specified otherwise, such processes will be supposed to start at zero: typically, we assume that the investor starts from zero cash and zero inventory at date $t=0$ in the following numerical tests. Finally, we will write indifferently $w^h(t,y,c_{_P})$ or $w^h(t,y):=w^h(t,y,0)$ to either stress or omit the dependence in $c_{_P}$.

• **The martingale case**: in the martingale case, we performed the algorithm (6.4.3)-(6.4.4) with parameters shown in Table 6.1. This set of parameters are chosen to be consistent with calibration data on the front maturity for 3-months EURIBOR future, see for example [25].

Figure 6.3 displayed the reduced-form value function w^h on $[0,T]\times[-N_Y;N_Y]$. This result illustrates the linear bound (6.3.13) as noticed in proposition 6.3.2, and also the symmetry of w^h as pointed out in (6.3.17). We also observe the monotonicity over $\mathbb{R}_+$ and $\mathbb{R}_-$ of the value function $w^h(t,.)$.

In Figure 6.4, we display the optimal make and take policies. The optimal take policy (on the left side) is represented as the volume to buy or sell with a market order, as a function of the time and inventory $(t,y)\in[0,T]\times[-N_Y;N_Y]$. We notice that a market order only occurs when the inventory becomes to large, and therefore, the take policy can be interpreted as a "stop-loss" constraint, i.e. an emergency rebalancing of the portfolio when the inventory risk is too large.

The optimal make policy is represented as the regime of limit orders posting as a function of the time and inventory $(t,y)\in[0,T]\times[-N_Y;N_Y]$. For sake of simplicity, we represented the sum of ℓ^a and ℓ^b on the map. The meaning of this surface is as follows: 0 means that

Parameter	Value
δ	12.5 EUR/contract
ε	1.05 EUR/contract
ε_0	0
λ	$0.05s^{-1}$
μ	$\exp(1/\bar{\mu})$
$\bar{\mu}$	20 contracts
γ	$2.5.10^{-5}$
T	100 s

(a) Market and risk parameters

Parameter	Value
N_Y	100
N_T	500

(b) Discretization parameters

Table 6.1: Parameters for numerical results in the martingale case.

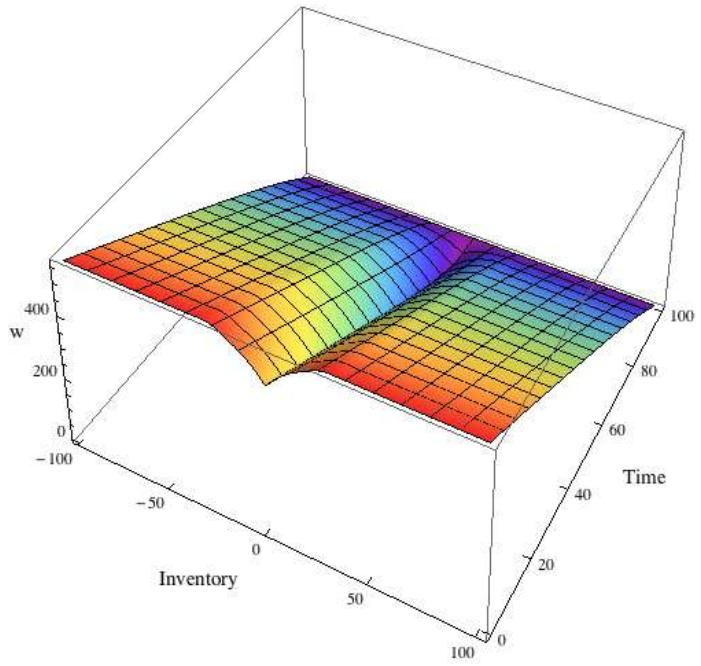

Figure 6.3: Reduced form value function w^h.

there is no active limit orders on either sides, 2 means that there is active limit orders on both bid and ask sides, and 1 means that there is only one active limit order either on the bid or the ask side, depending on the sign of y (if $y < 0$ only the bid side is active, and if $y > 0$ only the ask side is active). We notice that when close to maturity T, the optimal strategy tends to be more agressive, in the sense that it will seek to get rid of any positive or negative inventory, to match the terminal liquidation constraint. Moreover, we

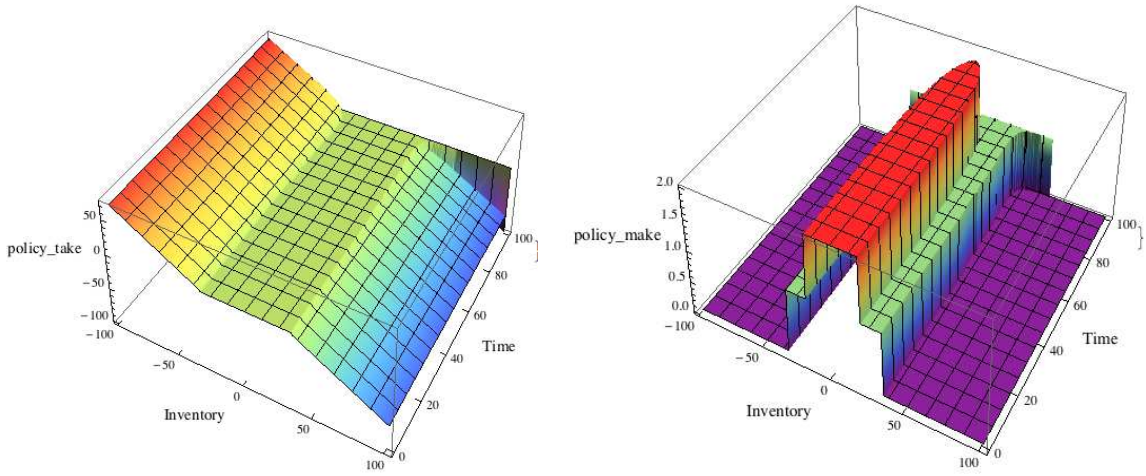

(a) Optimal take policy.

(b) Optimal make policy

Figure 6.4: Numerical results for the martingale case: representation of optimal make and take policies $\alpha^\star$. In the take policy, we represent the signed volume of the market order, in the make policy, 2 represent two-sided limit order posting, and 1 is one-sided order posting.

notice that close to date 0, the dependence in t seems to be negligible, which indicates that a"stationary regime" may be attained for large T. Figure 6.5 plots the cross-section of the optimal strategy when we are close to the initial date, i.e. far from the horizon T.

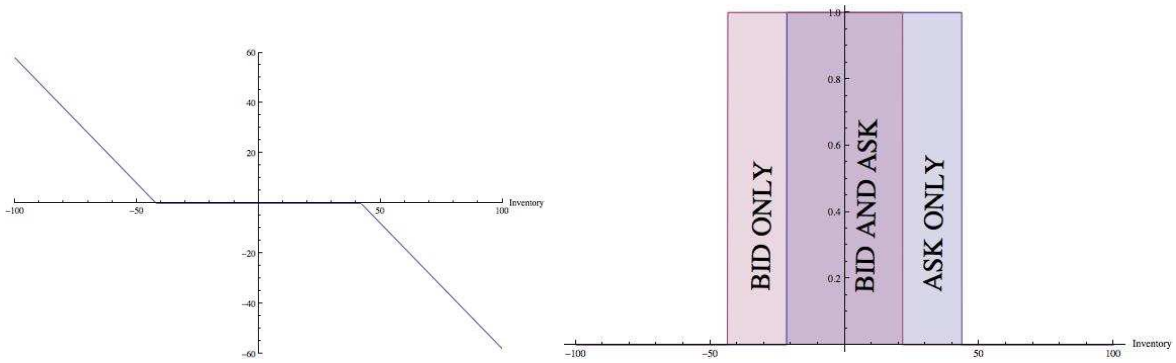

(a) Optimal take policy.

(b) Optimal make policy

Figure 6.5: Cross section of $\alpha^\star$ close to $t = 0$.

• **The trend information case:** in this case, we provide a backtest of the optimal strategy on simulated data in addition to the plot of the optimal policy $\alpha^\star$. We kept the

same parameters for execution intensity and volume, price characteristics and costs, but we choosed a wider time period in order to observe multiple trade event, see Table 6.2. With this set of parameters, we expect to observe about 100 trade events of average volume 20. Note that the execution intensity $\lambda = 0.05$, a value consistent with market activity of the front quaterly EURIBOR future, is independent in our model to the trend information ϖ that we will describe below.

Parameter	Value
δ	12.5 EUR/contract
ε	1.05 EUR/contract
ε_0	0
λ	$0.05s^{-1}$
μ	$\exp(1/\bar{\mu})$
$\bar{\mu}$	20 contracts
γ	$2.5.10^{-5}$
T	2000 s

(a) Market and risk parameters

Parameter	Value
N_Y	100
N_T	500
N_ϖ	50

(b) Discretization parameters

Table 6.2: Parameters for numerical results in the trend information case.

Figure 6.6 displays the optimal policy at date $t = 0$, in the plane (y, c_P). The policy has central symmetry properties as expected in (6.3.17), and should be read as follows: dark green zones represent situation where a market order to buy must be sent, light green means that a limit order is active only at bid, white means that limit orders are active on both sides, light red means that a limit order is active only at ask, and dark red means that a market order to sell must be sent. Let us provide a qualitative example: assume that after the high frequency trader acquired a positive inventory, the adverse selection effect implies that price should go down; therefore, using the fact that in this case we should have $c_P < 0$, the optimal strategy will be either to cancel the bid limit order (light red zone) and keep ask limit order active, or depending on the value of $|c_P|$, send a market order to get rid of our positive inventory (dark red zone).

We performed a benchmarked backtest on simulated data and a performance analysis in this case. The first benchmark strategy $\alpha^{WoMO} = (\alpha^{make\,,WoMO}, 0)$ correspond to the case where we do not allow the high-frequency trader to use market orders. It is computed using the backward numerical scheme (6.4.3)-(6.4.4), but without taking into account the obstacle part, which is equivalent to setting $\varepsilon_0 = \infty$. The second benchmark strategy is made of constant controls (a.k.a *symmetric* or *constant* strategy):

$$\begin{aligned} \alpha^{cst} &:= (\alpha^{make\,,cst}, 0) \\ \alpha^{make\,,cst} &:= (1, 1) \end{aligned}$$

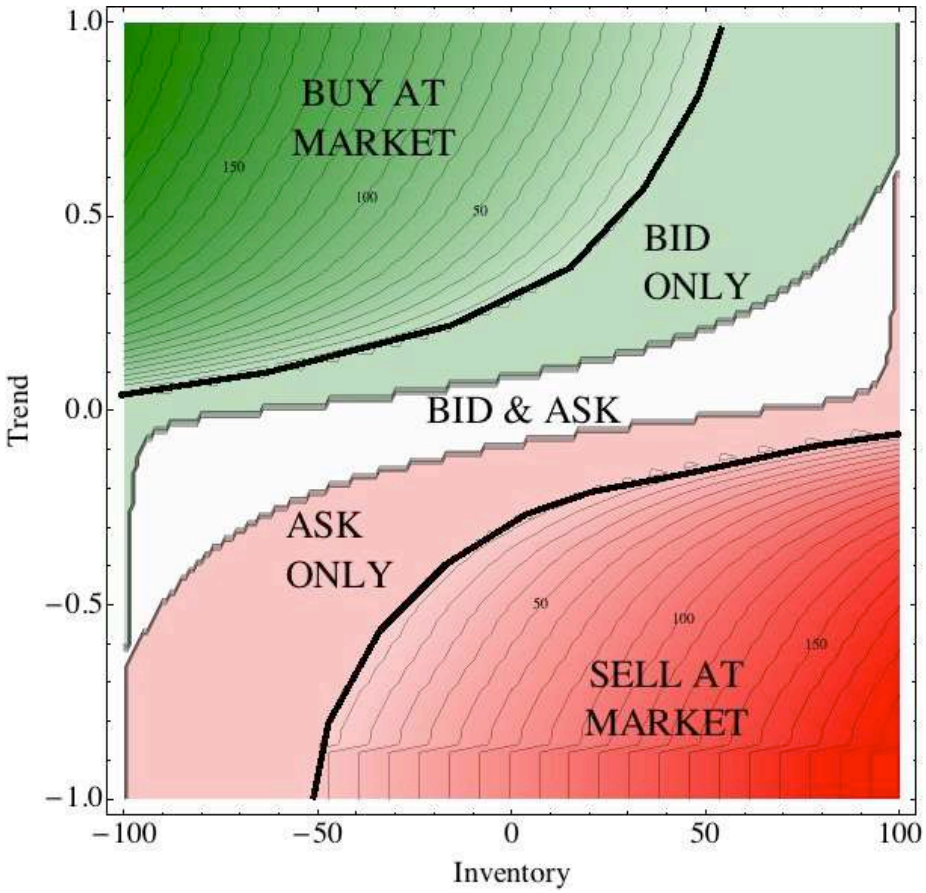

Figure 6.6: Optimal policy $\alpha^\star$ at date $t = 0$.

In order to make our simulated data backtest closer to the reality, we chosed to slightly deviate from the original price model, and use a varying price trend. We simulate a price process model given by

$$\hat{P}_t = \hat{P}_0 + \delta(N_t^+ - N_t^-),$$

where N^+ and N^- are the Euler scheme simulation of Cox processes of respective intensities π^+ and π^- defined as follows

$$\begin{aligned} \pi^+ + \pi^- &\equiv K = \varrho/\delta^2 \\ d\pi_t^+ - d\pi_t^- := d\varpi_t &= -\theta\varpi_t dt + \sigma dB_t \end{aligned}$$

where $K > 0$, $\theta > 0$ and $\sigma > 0$ are positive constants, and B is an independent Brownian motion. Note that we choosed the sum $\pi^+ + \pi^-$ to be the constant K, for simplicity sake: it means that, disregarding the direction of price variation, the mean number of price change

per second is assumed to be constant $\mathbf{P}\left(|P_{t+h} - P_t| = \delta\right) = Kh + o(h)$, which provides an easy way to calibrate the parameter K while reducing the dimension of the simulation. The interpretation of this simulation model is as follows: we add an exogenous risk factor B, which drives the price trend information ϖ as an Ornstein-Uhlenbeck process. Notice that this supplementary risk factor B is not taken into account in our optimization procedure and thus has a penalizing impact on the strategy's performance: therefore it does not spoil the backtest. This model choice for the process (ϖ_t) is an convenient way to simulate the real-world situation, where the high-frequency trader continuously updates her predictive information about short-term price movements, based e.g. on the current state of the limit order book.

Therefore, qualitatively speaking, our optimization procedure is consistent with this simulation model if we choose θ and σ s.t. the variation of the (reduced-form) value function w due to predictive information is very small compared to the variation of the value function due to other market events (e.g. an execution event).

This assumption is consistent with HFT practice since the HF trader is able to adapt very quickly to a modification of this predictive information. Backtest parameters involved in this simulation are shown in Table 6.3.

Parameter	Value
K	0.2
θ	$0.2s^{-1}$
σ	$0.2s^{-1}$
N_{MC}	50000

Table 6.3: Backtest parameters

The interpretation of the trend information parameters is the following: independently from the trade intensity $\lambda = 0.05$, we consider the price trend, which is interpreted as the expected return of the midprice over the next few milliseconds, and is directed by the state variable ϖ. In the stationnary regime, this variable ϖ has a marginal distribution $\mathcal{L}(\varpi_t)$ which is essentially a centered normal law of standard deviation $\sigma/\sqrt{2\theta} \simeq 0.32$ with this set of parameters. Qualitatively speaking, using the 2-sigma rule, this means that the process ϖ spends most of the time in the range -0.6 to 0.6. The value $\varpi_t = 0.6$ (resp. $\varpi_t = -0.6$) represents qualitatively a 60% probability of an uptick (resp. a downtick) in the next second. Such signal can be computed for example using the methods developed in [21]. Moreover, ϖ is a mean reverting process, of reversion speed $0.2s^{-1}$, which can be qualitatively interpreted as the timescale during which a prediction remains valid, in this case $\dfrac{1}{\theta_\varpi} = 5\text{s}$. This can be viewed as the timescale on which the high-frequency trader will update her prediction about the price trend. Note that in the case of STIR futures trading,

this choice of reversion speed is consistent with other market activity statistics: indeed, this reversion speed is greater than mid-price update intensity (of order 0.01 s^{-1}) and smaller than order book update intensity (of order 1 to 10 s^{-1}), see [25] for precise statistics.

Let us denote by $\hat{\vartheta}^a$ and $\hat{\vartheta}^b$ the Euler scheme simulation of the compound poisson processes ϑ^a and ϑ^b, with dynamics (6.2.3). Therefore, for $\alpha \in \{\alpha^\star, \alpha^{WoMO}, \alpha^{cst}\}$, we were able to compute the Euler scheme simulation $\hat{X}^\alpha$ (resp. $\hat{Y}^\alpha$) of X^α (resp. Y^α), starting at 0 at $t = 0$, by replacing ϑ^a (resp. ϑ^b) by $\hat{\vartheta}^a$ (resp. $\hat{\vartheta}^b$) in equation (6.2.4) (resp. (6.2.5)).

We performed N_{MC} simulation of the above processes. For each simulation $\omega \in [1...N_{MC}]$ and for $\alpha \in \{\alpha^\star, \alpha^{WoMO}, \alpha^{cst}\}$, we stored the following quantities: the terminal wealth after terminal liquidation $\hat{V}_T^\alpha(\omega) := L(\hat{X}^\alpha(\omega), \hat{Y}^\alpha(\omega), \hat{P}(\omega))$, called "performance" in what follows ; the total executed volume $\hat{Q}^{total,\alpha}(\omega) := \sum_{[0,T]} |\hat{Y}_t^\alpha(\omega) - \hat{Y}_{t-}^\alpha(\omega)|$; and the volume executed at market $\hat{Q}^{market,\alpha}(\omega) := \sum_{[0,T]} |\xi_n(\omega)^\alpha|$. Finally, we denote by $m(.)$ the empirical mean, by $\sigma(.)$ the empirical standard deviation, by skew(.) the empirical skewness, and by kurt(.) the empirical kurtosis, taken over $\omega \in [1...N_{MC}]$.

Quantity	Definition	$\alpha^\star$	α^{WoMO}	α^{cst}
Info ratio over T	$m(\hat{V}_T)/\sigma(\hat{V}_T)$	3.67	0.89	0.18
Profit per trade	$m(\hat{V}_T)/m(\hat{Q}^{total,.})$	8.06	16.31	5.57
Risk per trade	$\sigma(\hat{V}_T)/m(\hat{Q}^{total,.})$	2.19	18.31	29.56
Mean performance	$m(\hat{V}_T)$	31446.4	28246.3	21737.2
Standard deviation of perf	$\sigma(\hat{V}_T)$	8555.46	31701.2	115312
Skew of perf	skew$(\hat{V}_T)$	0.64	0.16	-0.007
Kurtosis of perf	kurt$(\hat{V}_T)$	3.82	3.31	7.02
Mean total executed volume	$m(\hat{Q}^{total,.})$	3900.68	1730.82	3900.61
Mean at market volume	$m(\hat{Q}^{market,.})$	1932.29	0	0
Ratio market over total exec	$m(\hat{Q}^{market,.})/m(\hat{Q}^{total,.})$	0.495	0	0

Table 6.4: Synthetis table for backtest. Categories are, from top to bottom: relative performance metrics, period-adjusted performance metrics, absolute performance metrics and absolute activity metrics.

Table 6.4 displayed a synthesis of descriptive statistics for this backtest. We first notice that the information ratio over T of $\alpha^\star$ is more than 4 times that of α^{WoMO}, which itself is about 4 times that of α^{cst}. Second, the per trade metrics can be compared to the half-spread $\frac{\delta}{2} = 6.25$ EUR/contract, and we see that although the mean profit per trade is smaller for the optimal strategy, the risk associated to each trade is dramatically reduced compared to the benchmark.

This is confirmed by the empirical distribution of performance, also shown in Figure 6.7, where the dark blue represents the performance distribution of the optimal strategy, the

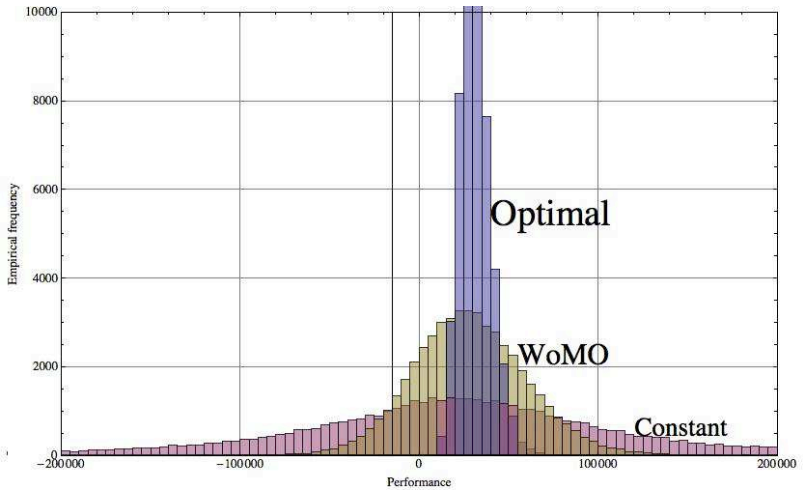

Figure 6.7: Empirical distribution of performance $\hat{V}_{\dot{T}}$. The graph shows the number of occurences for each bin on $N_{MC} = 50000$ simulations.

light yellow represents the performance distribution of the WoMO strategy and the light purple represents the performance distribution of the benchmark strategy. We see that not only benchmark has higher standard deviation, but also higher excess kurtosis and heavy tails: this is due to the fact that inventory can be very large for the constant strategy, and therefore it bears a non-negligible market risk (or inventory risk). Finally, we see that about 49% of the trades are done with market orders.

Our last numerical test is devoted to displaying the influence of the risk aversion parameter γ. All other parameters remaining the same, we tested several values of γ (as indicated in Table 6.5), and characterized the performance of the corresponding strategy by the pair $(\sigma(\hat{V}_{\dot{T}}), m(\hat{V}_{\dot{T}}))$, which gives the *efficient frontier* plot displayed in Figure 6.8. We measure the performance of each strategy empirically, both $\alpha^\star$ and α^{WoMO}, by running N_{MC} simulations of our market model, and therefore we can observe a slight measurement error on the points $(\sigma(\hat{V}_{\dot{T}}), m(\hat{V}_{\dot{T}}))$. As expected, a reduction of γ increases the standard deviation of the strategy: this is due to the fact that a small γ allows for large open position i.e. large inventory, and therefore the market risk is greater. For small γ, performance is also better since the investor can sustain large inventories, and therefore is less impatient to get rid of it: in particular, the proportion of volume executed at market is increasing in γ. In real trading condition, the value of γ should be tuned to attain the desired ratio of mean / volatility of PnL. The paper [17] observes similar behavior and determine these frontiers analytically when market-orders are absent. We also display in figure 6.8 the plot for α^{WoMO}, that clearly exhibits a larger risk, which indicates that the market orders in our optimal strategy are not only used to gain an extra performance, but also cut part of

γ	$\alpha^\star : \sigma(\hat{V}_{\dot{T}})$	$\alpha^\star : m(\hat{V}_{\dot{T}})$	$\alpha^{WoMO} : \sigma(\hat{V}_{\dot{T}})$	$\alpha^{WoMO} : m(\hat{V}_{\dot{T}})$
0,00000	8585,66	31619,86	32095,72	28161,49
0,00002	8556,70	31594,73	31861,96	28201,85
0,00004	8545,42	31593,76	31591,41	28214,79
0,00006	8534,72	31625,64	31404,14	28132,71
0,00008	8498,87	31559,85	30971,55	28123,92
0,00010	8427,40	31492,04	30613,55	27903,22
0,00012	8407,00	31547,71	30337,09	27919,02
0,00014	8331,14	31482,38	29934,89	27768,66
0,00016	8223,93	31399,91	29814,76	27731,23
0,00018	8193,31	31324,68	29380,58	27530,12
0,00020	8152,06	31185,16	29043,01	27319,45
0,00022	8139,15	31145,02	28595,55	27272,89
0,00024	8054,56	30856,74	28456,77	27482,97
0,00026	7965,48	30778,70	28082,26	27128,46
0,00028	7961,17	30658,60	27749,49	27274,58
0,00030	7880,11	30550,28	27322,27	26857,24
0,00032	7845,00	30463,95	27035,09	26905,98
0,00034	7748,16	30260,12	26839,15	26683,35
0,00036	7663,32	30049,39	26440,91	26606,25
0,00038	7661,55	29974,39	25993,18	26367,21
0,00040	7594,45	29783,59	25752,06	26287,87
0,00042	7551,23	29687,68	25338,63	26227,97
0,00044	7473,28	29479,14	24985,62	26072,39
0,00046	7416,75	29404,31	24743,03	25947,35
0,00048	7347,36	29128,09	24425,99	25807,40
0,00050	7252,69	29041,66	24300,78	25742,08

Table 6.5: Varying risk aversion parameter γ: data.

the risk of holding a non-zero position, especially when we expect the mid-price to move adversely.

6.5 Best execution problem and overtrading risk

In this section, we apply our market model framework to a best execution problem. The trading objective of the investor is to liquidate $Y_0 > 0$ assets over the finite time interval $[0, T]$. She is not allowed to purchase stock during the liquidation period, and may only

buy back the asset in case of short position. In this context, the investor posts continuously a limit sell order (with a volume much larger that the required quantity Y_0) at the best ask price, and also runs market (sell) orders strategy until she reaches either a negative inventory or the terminal date. By doing so, she hopes to trade as much as possible at the ask price, and therefore avoiding to *cross the spread*.

Mathematically, this means that the investor uses a subset $\mathcal{A}_\ell$ of strategies $\alpha = (\alpha^{make} = (L^a, L^b), \alpha^{take})$ in $\mathcal{A}$ such that:

$$\begin{aligned}(L^a_t, L^b_t) &= \begin{cases} (1,0) & \text{for } t \le \tau, \\ (0,0) & \text{for } t > \tau \end{cases} \\ \alpha^{take} &= (\tau_n, \zeta_n)_n \cup (\tau, -Y_\tau), \quad \text{with } \tau_n < \tau,\ \zeta_n < 0,\end{aligned}$$

where $\tau = \inf\{t \ge 0 : Y_t \le 0\} \wedge T$. The value function associated to this liquidation problem is then defined by

$$v_\ell(t,x,y,p) = \sup_{\alpha \in \mathcal{A}_\ell} \mathbf{E}_{t,x,y,p}\Big[L(X_T, Y_T, P_T) - \gamma \int_t^T Y_s^2 \varrho(P_s) ds\Big], \tag{6.5.1}$$

for $(t,x,y,p) \in [0,T] \times \mathbb{R}^2 \times \mathbb{P}$. With the notation in (6.3.6), the operator corresponding to the limit order in $\mathcal{A}_\ell$ is given by $\mathcal{L}^{1,0} = \mathcal{P} + \Gamma^a$, while the impulse operator associated to the market order in $\mathcal{A}_\ell$ is defined by:

$$\mathcal{M}_\ell\, \varphi(t,x,y,p) = \sup_{e \in [-M_Y - y, -(M_Y - y)_-]} \varphi\big(t, x - ep - |e|(\frac{\delta}{2} + \varepsilon) - \varepsilon_0 1_{e \ne 0}, y + e, p\big),$$

where $m_- = \max(-m, 0)$. The dynamic programming equation associated to (6.5.1) takes the form:

$$\min\big[-\frac{\partial v_\ell}{\partial t} - \mathcal{P} v_\ell - \Gamma^a v_\ell \ +\ \gamma g\ ,\ v_\ell - \mathcal{M}_\ell v_\ell\big] = 0, \text{ on } [0,T) \times \mathbb{R} \times (0,\infty) \times \mathbb{P},$$

together with the terminal and boundary conditions:

$$v_\ell = L, \quad \text{on } \big(\{T\} \times \mathbb{R} \times \mathbb{R} \times \mathbb{P}\big) \cup \big([0,T) \times \mathbb{R} \times \mathbb{R}_- \times \mathbb{P}\big).$$

The above boundary condition for nonpositive inventory is related to the *overtrading risk*, which is the risk that the investor sold too much assets via the (oversized) limit order at the best ask price. This risk occurs typically in execution problems on pro-rata limit order book, see [25].

Again, in the Lévy case (6.3.9), the value function v_ℓ is reduced into:

$$v_\ell(t,x,y,p) = L_0(x,y,p) + w_\ell(t,y),$$

where w_ℓ is solution to the integro-variational inequality:

$$\begin{aligned}
&\min\Big\{ -\frac{\partial w_\ell}{\partial t} - yc_{_P} + \gamma\varrho y^2 \\
&-\lambda^a \int_0^\infty \Big[w_\ell(t,y-z) - w_\ell(t,y) + z\frac{\delta}{2} + (\frac{\delta}{2}+\varepsilon)(|y|-|y-z|)\Big]\mu^a(dz)\ ; \\
&w_\ell(t,y) - \sup_{e\in[-M_{_Y}-y,-(M_{_Y}-y)_-]} \Big[w_\ell(t,y+e) - (\frac{\delta}{2}+\varepsilon)(|y+e|+|e|-|y|) - \varepsilon_0\Big]\Big\} \;=\; 0,
\end{aligned}$$

for $(t,y)\in[0,T)\times(0,\infty)$, together with the terminal and boundary conditions:

$$w_\ell(t,y) \;=\; -\varepsilon_0 1_{y\neq 0}, \quad \forall (t,y)\in\big(\{T\}\times\mathbb{R}\big)\cup\big([0,T)\times\mathbb{R}_-\big).$$

The associated numerical scheme reads now as follows:

$$\begin{aligned}
w_\ell^h(t_N,y) &= -\varepsilon_0 1_{y\neq 0}, \quad y\in\mathbb{R}, \\
w_\ell^h(t_k,y) &= 0, \quad k=0,\dots,N-1,\ y\le 0, \\
w_\ell^h(t_k,y) &= \max\Big[\mathcal{T}_\ell^{h,\Delta_Y,M}(t,y,\varphi)\ ;\ \mathcal{M}_\ell^{h,\Delta_Y,M}(t,y,\varphi)\Big],\ k=0,\dots,N-1\,,\ y\in\mathbb{Y}_M^+,
\end{aligned}$$

where $\mathbb{Y}_M^+ = \mathbb{Y}_M\cap\mathbb{R}_+$,

$$\begin{aligned}
\mathcal{T}_\ell^{h,\Delta_Y,M}(t,y,\varphi) &= \varphi(t,y) - h\gamma\varrho y^2 + hyc_P \\
&\quad + \lambda^a h\Big(\int_0^\infty \big[\varphi(t,\mathrm{Proj}_M(y-z)) - \varphi(t,y)\big]\hat{\mu}^a(dz) \\
&\quad + \int_0^\infty \big[\frac{\delta}{2}z + (\frac{\delta}{2}+\varepsilon)(|y|-|y-z|)\big]\mu^a(dz)\Big)
\end{aligned}$$

and

$$\begin{aligned}
&\mathcal{M}_\ell^{h,\Delta_Y,M}(t,y,\varphi) \\
&= \sup_{e\in\mathbb{Y}_M\cap[-M_{_Y}-y,-(M_{_Y}-y)_-]} \big[\varphi(t,\mathrm{Proj}_M(y+e)) - (\frac{\delta}{2}+\varepsilon)(|y+e|+|e|-|y|) - \varepsilon_0\big].
\end{aligned}$$

In this case, the optimal policy shown in Figure 6.9 is simple to describe. The state space is delimited in two zones: when the inventory is small, the HFT must wait for her limit sell order to be executed; and when the inventory is large, the HFT must send a market sell order to avoid the market risk related to holding a large position.

The frontier between the two zones (indicated in bold red in Figure 6.9) can be interpreted as an *optimal trading curve*, a concept that is extensively documented (see e.g. [35]) in the optimal execution literature. The optimal trading curve is the inventory that the investor should hold, seen as a function of time, in order to minimize overall trading costs. Therefore, in the typical setting, the execution strategy consists in trading via market orders to get as close as possible to the optimal trading curve. Similarly, in our case, we can

see on Figure 6.9 that the optimal strategy will behave similarly for large inventories (i.e. when above the trading curve): indeed, we observe that the quantities to sell are such that the market orders strategy would keep the inventory close to the optimal trading curve, if no limit orders were allowed. Now, in our case, we observe two specific features of the optimal strategy: 1) the optimal trading curve does not reach 0 at maturity, and therefore the HFT has to get rid of her inventory at market at final date to match the constraint $Y_T = 0$. This is due to the fact that a supplemental gain is always achievable when the limit order is executed. Therefore, this features leads to an execution strategy where the final trade is bigger than intermediary trades; 2) below the optimal trading curve, i.e. in the region where the HFT trades via limit orders only, the sell limit order is always active, and can lead to an execution. Therefore, the inventory is always below the optimal trading curve, and the distance between the current inventory and the optimal trading curve equals the volume executed via limit orders. This differs from classic pattern-based best execution strategies, for example the U-shaped execution strategy that consists in trading a large quantity at the beginning and at the end of the liquidation, and trade regularly small quantities in between. Indeed, the optimal strategy does not provide a fixed pattern for every execution, but provide the optimal action to take given the observation of the inventory that is still to be sold and the market characteristics as e.g. the mean traded volume at ask per second $\lambda^a \bar{\mu}^a$, or trades volume distributions at ask μ^a.

Finally, let us notice that this strategy can be interpreted as a convenient way to avoid the cost of crossing the spread during the liquidation of a portfolio, but we did not take into account the impact of the market order on the transaction price. In the case of a pro-rata microstructure, available volumes offered at best prices are usually about 200 times larger than the mean volume of market orders (see [25]), and therefore it is consistent to consider that there is no impact on the price for our market orders. Yet, the model can easily be modified by adding an impact component in the obstacle operator $\mathcal{M}_\ell$ to take care of this effect. We also did not model the possibility that the intensities λ^a and λ^b of execution processes may vary, and postpone this investigation for future research.

6.6 Conclusion

In this paper, we investigate a framework to build up mixed high-frequency trading strategies in an exotic microstructure, the pro-rata microstructure. This microstructure can be encountered for example on short-term interest rates futures. We consider the situation of an investor willing to maximize her terminal profit over a finite time horizon, and able to trade with limit and market orders. We adopt the perspective of inventory management, which means that the investor primary objective is to keep her position on the risky asset close to zero at all times, in order to avoid being exposed to market risk.

We provide a tractable market model that mimics the major features of our target microstructure, while being parsimonious enough to fit a large range of products. We detail the optimization procedure, by means of stochastic control, as well as the numerical scheme used to solve the resulting HJB equation. Dimension reduction techniques as well as interpretable decomposition of the profit's dynamics are described. We also discuss the practical implementation of such strategy.

In this particular microstructure, we are able to define and address two specific types of risk: the overtrading risk, which is the risk of brutal variations in the investor inventory, due to the fact that she does not control the quantity traded at limit; and the adverse selection risk, which is the risk of market reacting unfavorably to the investor quotes.

For this last purpose, we introduce a new state variable, that we interpret as a predictive price indicator, that allows us to balance our position before the price changes. This last feature also provides an extra performance on our empirical tests. We provide several examples of application of our framework, including a mixed limit/market strategy when no information is available on price, a mixed strategy with superior information on price, and a liquidation strategy without information on price. Moreover, we point out the advantages of using market orders in this setup by benchmarking the performance of our strategy against a pure limit order strategy, and we find that the mean/volatility ratio is much smaller in this last case.

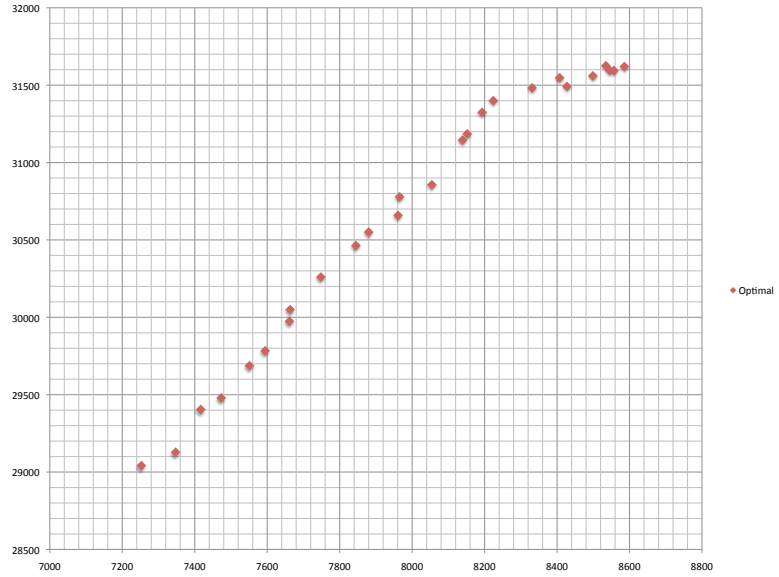

(a) Optimal

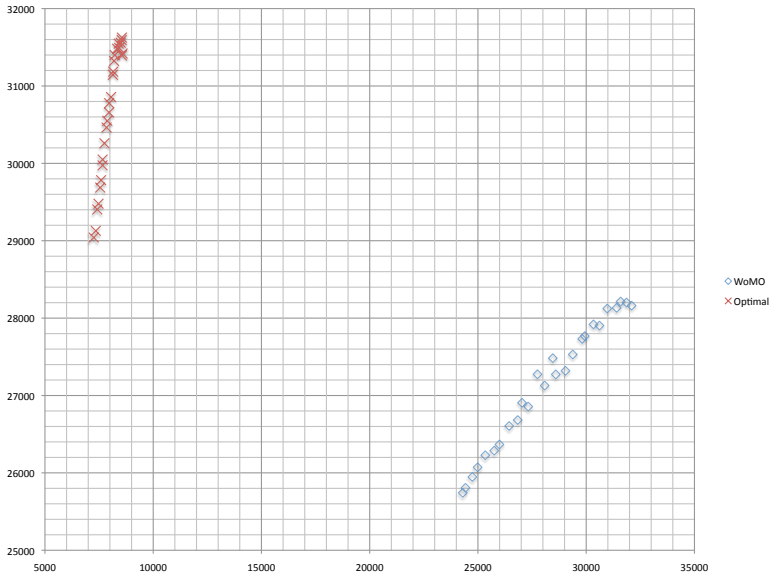

(b) Optimal and WoMO

Figure 6.8: Varying parameter γ for $\alpha^\star$ and α^{WoMO}. The X-axis represent the standard deviation of performance $\sigma(\hat{V}_T)$ and the Y-axis the average performance $m(\hat{V}_T)$ estimated on $N_{MC} = 50000$ simulations of the model.

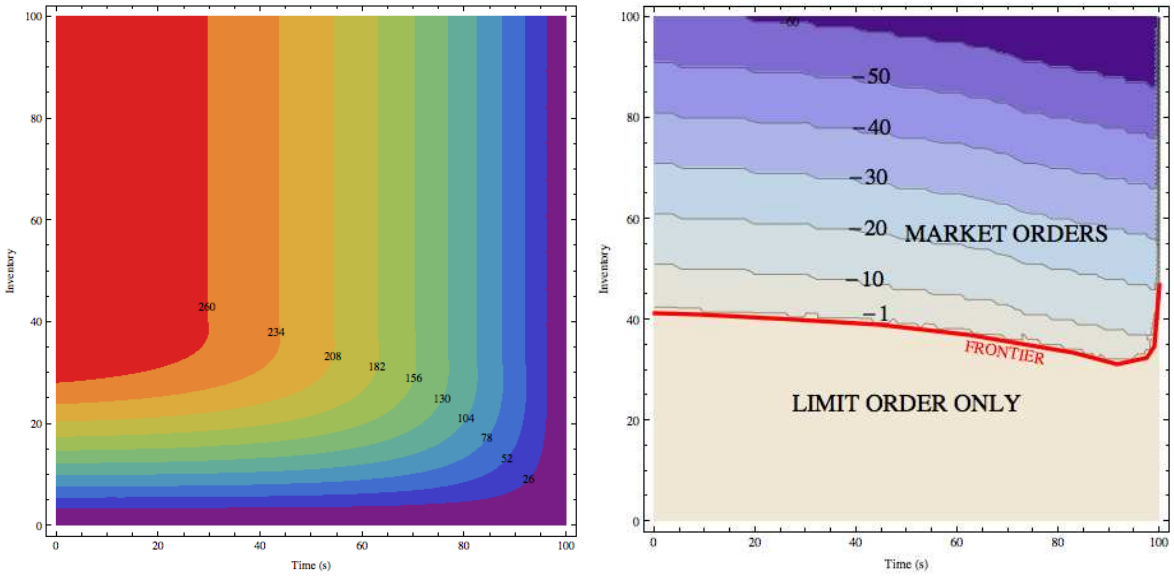

(a) Value Function w_ℓ

(b) Optimal policy (take)

Figure 6.9: Numerical results for the simple liquidation problem (for $c_P = 0$). On the left side, level lines are indicated for the value function w_ℓ. On the right side, numbers indicated on the figure represent the quantity to sell in the optimal market order control.

Bibliography

[1] Aikin, S. (2006): "Trading STIR Futures : An Introduction to Short-Term Interest Rate Futures", Harriman House Publishing, ISBN: 978-1897597811

[2] Alfonsi A., Schied A. and A. Slynko (2009): "Order book resilience, price manipulation, and the positive portfolio problem", preprint

[3] Almgren R. and N. Chriss (2001): "Optimal execution of portfolio transactions", *Journal of Risk*, **3**, 5-39.

[4] Almgren R., Thum C., Hauptmann E. and H. Li (2005): "Equity market impact", *Risk*, July 2005, 58-62.

[5] Amihud, Y. and Mendelson, H. (1980): "Dealership market: Market-making with inventory." *Journal of Financial Economics* 8, 31–53.

[6] M. Avellaneda M. (2010): "Statistical Arbitrage in the US equities market", *Quantitative Finance*, 10(7), 761-782.

[7] Avellaneda M. and S. Stoikov (2008): "High frequency trading in a limit order book", *Quantitative Finance*, 8(3), 217-224.

[8] Barles G. and P. Souganidis (1991): "Convergence of approximation schemes for fully nonlinear second order equations", *Asymptotic Analysis*, 4, 271-283.

[9] Bayraktar E. and M. Ludkovski (2011): "Liquidation in limit order books with controlled intensity", available at: arXiv:1105.0247

[10] Bertsimas D. and A. Lo (1998): "Optimal control of execution costs", *Journal of Financial Markets*, **1**, 1-50.

[11] Biais, B., Glosten, L. and C. Spatt, (2005): "Market microstructure: a survey of microfoundations, empirical results and policy implications." *J. Financ. Markets*, 8, 217–264.

[12] Bouchard B., Dang N.M., and C.A. Lehalle (2010): "Optimal control of trading algorithms: a general impulse control approach", preprint, University Paris Dauphine.

[13] C.Bowsher (2007): "Modelling Security Market Events in Continuous Time: Intensity Based, Multivariate Point Process Models", Journal of Econometrics

[14] Bozdog D., Florescu I., Khashanah K., and J. Wang (2012): "A Study of Persistence of Price Movement using High Frequency Financial Data" preprint

[15] Cartea, A and S. Jaimungal (2010): "Modeling asset prices for algorithmic and high frequency trading". SSRN eLibrary, http://ssrn.com/paper=1722202.

[16] Cartea A, Jaimungal S. and J. Ricci (2011): "Buy low sell high: a high frequency trading perspective", preprint

[17] Cartea, A. and Jaimungal, S, (2012): "Risk Metrics and Fine Tuning of High Frequency Trading Strategies", to appear in *Mathematical Finance*.

[18] Chakraborti A., Muni Toke I., Patriarca M. and F. Abergel (2011): "Econophysics: empirical facts and agent-based models", *Quantitative Finance*, 11, 991-1041.

[19] Chancelier J.P., Oksendal B. and A. Sulem (2002): "Combined stochastic control and optimal stopping, and application to numerical approximation of combined stochastic and impulse control", *Tr. Mat. Inst. Steklova, 237(Stokhast. Finans. Mat.)*, 149?172.

[20] Chen Z. and and P.A. Forsyth (2008): "A numerical scheme for the impulse control formulation for pricing variable annuities with a guaranteed minimum withdrawal benefit (gmwb)", *Numerische Mathematik*, **109**, 535-569.

[21] Cont R. and A. de Larrard (2011): "Price dynamics in a Markovian limit order market", preprint LPMA, available at: arXiv:1104.4596

[22] Cont R. and A. de Larrard (2012): "Order Book Dynamics in Liquid Markets: Limit Theorems and Diffusion Approximations" preprint LPMA

[23] Crandall M., Ishii H. and P.L. Lions (1992): "User's guide to viscosity solutions of second order partial differential equations", *Bull. Amer. Math. Society*, 27, 1-67.

[24] Al Dayri K.A. (2011): "Market Microstructure and Modeling of the Trading Flow" Ph.D. Thesis, Ecole Polytechnique

[25] Field J. and J. Large (2008): "Pro-Rata Matching and One-Tick Futures Markets", preprint

[26] Florescu I., Mariani M.C. and F. Viens (Eds.) (2012): "Handbook of High Frequency Data modeling" Wiley, ISBN-13 978-0470876886

[27] Fodra P. and M.Labadie (2012) :"High-frequency market-making with inventory constraints and directional bets" preprint

[28] Forsyth P. (2009): "A Hamilton-Jacobi-Bellman approach to trade execution", preprint, University of Waterloo.

[29] Frey S. and Grammig (2008): "Liquidity supply and adverse selection in a pure limit order book market" *High Frequency Financial Econometrics, Physica-Verlag HD*, 83-109, J. Bauwens L., Pohlmeier W. and D. Veredas (Eds.)

[30] Gatheral J. (2010): "No-dynamic-arbitrage and market impact", *Quantitative Finance*, **10**, 749 - 759

[31] Gatheral J., Schied A. and A. Slynko (2010): "Transient linear price impact and Fredholm integral equations", preprint

[32] Guéant O., Fernandez Tapia J. and C.-A. Lehalle (2011): "Dealing with inventory risk", preprint.

[33] Gould M.D., Porter M.A, Williams S., McDonald M., Fenn D.J. and S.D. Howison (2010): "The limit order book: a survey", available at: arXiv:1012.0349

[34] Grillet-Aubert L. (2010): "Négociation d'actions: une revue de la littérature à l'usage des régulateurs de marché", *AMF*, available at: http://www.amf-france.org/documents/general/9530_1.pdf.

[35] Guéant O., Fernandez Tapia J. and C.-A. Lehalle (2011): "Dealing with inventory risk", available at: arXiv:1105.3115

[36] Guilbaud F., Mnif M. and H. Pham (2010): "Numerical methods for an optimal order execution problem", to appear in *Journal of Computational Finance*.

[37] Guilbaud F. and H. Pham (2011): "Optimal high frequency trading with limit and market orders", to appear in *Quantitative Finance*.

[38] Guilbaud F. and H. Pham (2012): "Optimal high frequency trading in a pro-rata microstructure with predictive information", preprint.

[39] He H. and H. Mamaysky (2005): "Dynamic trading policies with price impact", *Journal of Economic Dynamics and Control*, **29**, 891-930.

[40] Hendershott T., Jones C.M. and A.J. Menkveld (2010): "Does algorithmic trading improve liquidity?", *Journal of Finance*, to appear.

[41] P.Hewlett (2006): "Clustering of order arrivals, price impact and trade path optimisation", preprint

[42] Ho, T. and H. Stoll (1981), "Optimal dealer pricing under transactions and return uncertainty". *J. Financ. Econ.*, 9, 47–73.

[43] Karel Jane·cek, Martin Kabrhel (2007): "Matching Algorithms of International Exchanges", working paper.

[44] Kharroubi I. and H. Pham (2010): "Optimal portfolio liquidation and execution risk", *SIAM J. Financial Mathematics*, 1, 897-931.

[45] Kühn C. and M. Stroh (2010): "Optimal portfolios of a small investor in a limit order market: a shadow price approach", *Mathematics and Financial Economics*, 3(2), 45-72.

[46] Kühn C. and M. Stroh (2012): "Continuous time trading of a small investor in a limit order market", preprint

[47] Lapeyre B., Sulem A. and D. Talay (2007): Understanding Numerical Analysis for Option Pricing, *Cambridge University Press.*

[48] Laruelle S., Lehalle C.-A. and G.Pagès (2010): "Optimal split of orders across liquidity pools: a stochastic algorithm approach" preprint

[49] Laruelle S., Lehalle C.-A. and G.Pagès (2010): "Optimal posting distance of limit orders: a stochastic algorithm approach" preprint

[50] Lillo F., Farmer J. and R. Mantagna (2003): "Master curve for price impact function", *Nature*, **421**, 129-130.

[51] Ly Vath V., Mnif M. and H. Pham (2007): "A model of optimal portfolio selection under liquidity risk and price impact", *Finance and Stochastics*, **11**, 51-90.

[52] Maroso S. (2006): Analyse numérique de problèmes de contrôle stochastique, PhD thesis, University Paris 6.

[53] A.J. Menkveld (2011): "High frequency trading and the new-market makers", preprint.

[54] I. Muni Toke (2010): "Market making" in an order book model and its impact on the spread", preprint

[55] Obizhaeva A. and J. Wang (2005): "Optimal trading strategy and supply/demand dynamics", to appear in *Journal of Financial Markets*.

[56] M. O'Hara (1995), "Market Microstructure Theory", Blackwell, Oxford, ISBN 1-55786-443-8.

[57] Øksendal B. and A. Sulem (2007): Applied stochastic control of jump diffusions, Springer, Universitext.

[58] Pagès G., Pham H and J. Printems (2004): "Optimal quantization methods and applications to numerical problems in finance", Handbook of computational and numerical methods in finance, ed. Z. Rachev, Birkhauser.

[59] Pham H. (2009): Continuous-time stochastic control and optimization with financial applications, Springer Series SMAP.

[60] Potters M. and J.P. Bouchaud (2003): "More statistical properties of order books and price impact", *Physica A*, **324**, 133-140.

[61] Predoiu S., Shaikhet G. and S.Shreve (2010): "Optimal Execution in a General One-Sided Limit-Order Book", Preprint.

[62] Protter P. (2005): Stochastic Integration and Differential Equations, Second Edition, Version 2.1, Springer-Verlag, Heidelberg.

[63] Rogers L.C.G. and S. Singh (2008): "The cost of illiquidity and its effects on hedging", to appear in *Mathematical Finance*.

[64] Schied A. and T. Schöneborn (2009): "Risk aversion and the dynamics of optimal liquidation strategies in illiquid markets", *Finance and Stochastics*, **13**, 181-204.

[65] Stoikov S. and M. Saglam (2009): "Option market making under inventory risk", *Review of Derivatives Research*, 12(1), 55-79.

[66] Ulibarri C. A. and P. C. Anselmo (2012):"A Market Microstructure Model of Ultra High Frequency Trading" preprint

[67] Veraart L.A.M. (2010): "Optimal Market Making in the Foreign Exchange Market", *Applied Mathematical Finance*, 17, 359-372.

[68] Veraart L.A.M. (2011): "Optimal Investment in the Foreign Exchange Market with Proportional Transaction Costs", *Quantitative Finance*, 11(4): 631-640. 2011.