Probabilistic numerical methods for high-dimensional stochastic control and valuation problems on electricity markets

Nicolas Langrené

Université Paris Diderot (Paris 7) Sorbonne Paris Cité

École doctorale ED386 Sciences Mathématiques de Paris Centre

Laboratoire de Probabilités et Modèles Aléatoires

THÈSE DE DOCTORAT

Discipline : Mathématiques Appliquées

Présentée par

NICOLAS LANGRENÉ

Méthodes numériques probabilistes en grande dimension pour le contrôle stochastique et problèmes de valorisation sur les marchés d'électricité

Probabilistic numerical methods for high-dimensional stochastic control and valuation problems on electricity markets

---

Sous la direction de **Huyên PHAM**
et la codirection de **Luciano CAMPI**

Soutenue publiquement le 5 mars 2014 devant le jury composé de:

René Aïd	*Électricité de France R&D*	Examinateur
Luciano Campi	*London School of Economics*	Codirecteur
René Carmona	*Princeton University*	Rapporteur
Romuald Élie	*Université Marne-la-Vallée*	Examinateur
Huyên Pham	*Université Paris Diderot*	Directeur
Denis Talay	*INRIA Sophia Antipolis*	Rapporteur
Peter Tankov	*Université Paris Diderot*	Président du Jury
Xavier Warin	*Électricité de France R&D*	Examinateur

Remerciements

Suivant en cela une convention séculaire, je commencerai par célébrer la dimension collective de l'odyssée doctorale au moyen d'un inventaire précis de remerciements.

Tout d'abord, je souhaite exprimer ma gratitude envers mon directeur de thèse, Huyên Pham, pour son soutien, sa patience et ses conseils, sans lesquels mes efforts eurent maintes fois achoppé.

Ensuite, j'adresse mes remerciements à mon codirecteur de thèse, Luciano Campi, pour sa disponibilité, sa diligence et son aide précieuse face à plusieurs preuves probabilistes à la technicité rebutante.

Enfin, j'adresse mes remerciements à René Aïd, pour son énergie, son enthousiasme inébranlable et la confiance qu'il m'accorde depuis mes premiers travaux de stage à la R&D d'EDF.

Je remercie très chaleureusement Denis Talay et René Carmona, dont j'apprécie les travaux, d'avoir accepté de rapporter les miens. Je leur en suis reconnaissant. Je remercie également Romuald Élie, Peter Tankov et Xavier Warin d'avoir accepté de faire partie de mon jury.

Je remercie tous les membres du LPMA , du laboratoire FiME, et du département OSIRIS de la R&D d'EDF que j'ai eu la chance de cotoyer, ainsi que les autres chercheurs, notamment Idris Kharroubi, avec qui j'ai eu la chance de travailler.

Je souhaite remercier les secrétaires du laboratoire, de l'école doctorale et d'EDF, en particulier Nathalie, Valérie et Agnès, pour leur disponibilité, leur efficacité et leur bonne humeur.

Je remercie tous mes camarades doctorants que j'ai eu la chance de cotoyer, à Chevaleret d'abord, puis à Sophie Germain[1], ainsi qu'à d'EDF ; par ordre d'apparition : Thomas, Marie, Adrien, Victor, Christophe, Ennio, Paul, Laurent, Jordan, Oriane, Lorick, Maud, Marc-Antoine, Sébastien, Guillaume, Sandrine, Aser, Pierre, Jiatu, Pietro, Shanqiu, Huy, Sophie, Vu-Lan, Hugo, Anna, Arturo et Clément.

Je souhaite aussi explicitement remercier de manière générale tous les chercheurs m'ayant précédé sur mes thèmes actuels, car comme le relativisait Montaigne[2], "nous échelons ainsi de degré en degré, et advient de là que le plus haut monté a souvent plus d'honneur que de mérite, car il n'est monté que d'un grain sur les épaules du pénultième."

Symétriquement, je remercie les hypothétiques lecteurs futurs, et félicite ceux qui auraient la témérité de découdre, entre deux interludes graphiques, les lemmes les plus techniques de ces présents travaux.

Ayant réalisé l'essentiel des présents travaux dans la capitale française, je souhaite remercier Paris, notamment ses bibliothèques et son système dense de transports en commun, pour avoir contribué (en moyenne) à la bonne marche quotidienne au cours de ces 3,31 années de doctorat.

Je remercie pour finir mes amies Stéphanie, Nourimane et Gohar pour leur soutien salvateur, ainsi que toutes les personnes non précédemment susnommées ayant directement ou indirectement concouru à la complétion effective de ce manuscrit.

Paris, 21 février 2014

[1] Qui, avant de devenir un bâtiment, fut d'abord une insolite mathématicienne de chair et d'os.
[2] Essais, Livre III, Chapitre XIII.

Contents

1 Introduction (en français)

Ce présent travail est consacré pour l'essentiel à la résolution numérique de problèmes d'évaluation et de couverture d'actifs, et plus généralement de problèmes de contrôle stochastique, liés aux marchés d'énergies, en particulier au marché de l'électricité. Ayant pour optique la résolution de problèmes de dimension élevée, tels que des problèmes d'investissements en centrales électriques, nous avons privilégié la construction et l'analyse de schémas numériques de type probabilistes.

Dans une première partie, nous avons commencé par nous intéresser à la modélisation proprement dite du prix de l'électricité. Motivés par des considérations empiriques sur la formation du prix d'électricité, nous avons commencé par proposer et étudier un nouveau modèle structurel de prix d'électricité, lequel présente l'avantage crucial de pouvoir rendre compte de la dépendence fine entre le prix de l'électricité et les prix des autres énergies (pétrole, charbon, gaz,...). Dans cette première partie, nous avons fait usage de ce modèle pour des problèmes de valorisation et de couverture d'actifs contingents faisant intervenir le prix d'électricité. En particulier, nous avons cherché à améliorer les stratégies de couverture possibles pour ces produits en autorisant l'usage d'actifs de couverture basés sur les autres énergies, et en exploitant leur dépendance fine avec l'électricité. Divers tests numériques viennent illustrer l'effet de cet apport, lesquels nécessitèrent la mise en place d'algorithmes efficaces de calcul de stratégies de couverture.

Dans une deuxième partie, nous avons utilisé ce nouveau modèle d'électricité pour un nouveau problème, à savoir la valorisation de centrales électriques et, par extension, la détermination de stratégies d'investissements économiquement optimales en nouvelles centrales dans le futur. D'un point de vue mathématique, ce genre de problème peut s'écrire sous la forme d'un problème de commutation optimale, qui correspond à un problème de contrôle stochastique dont toute modification de la stratégie, impulsionnelle, génère un coût associé. Le nombre de dimensions du problème étant élevé, nous avons donc proposé un schéma probabiliste pour le résoudre, qui combine programmation dynamique et régressions. Ce schéma s'inspire de ceux existant pour la résolution du problème plus simple de valorisation d'option américaine ([31]). Nous avons analysé mathématiquement la convergence de ce schéma, ainsi que sa vitesse pour un choix spécifique de base de régression (fonctions constantes par morceaux). Une partie numérique vient illustrer la faisabilité de ce schéma, et viennent démontrer que les techniques et améliorations mises en œuvre ici rendent, de manière plus générale, envisageable l'idée d'une résolution probabiliste directe de problèmes de contrôle stochastique multidimensionnels difficiles.

Enfin, dans une troisième partie nous avons étendu le schéma probabiliste précédent à des problèmes plus généraux, à savoir des problèmes de contrôle stochastique dont la dérive et la volatilité du processus d'état peuvent dépendre du contrôle. Jusqu'à présent, les schémas de type probabilistes ne pouvaient pas traîter ce genre de problèmes, puisque l'initiation de ceux-ci nécessite la simulation de trajectoires du processus d'état sous-jacent, ce qui n'est pas possible ici puisque le contrôle optimal est initialement inconnu. Pour surmonter ce problème nous avons fait usage d'une représentation du problème par Équations Différentielles Stochastiques Rétrogrades (EDSRs) à sauts contraints, qui est une généralisation de la classe des EDSRs introduite dans [72] qui permet d'englober les problèmes de contrôle stochastique qui nous intéressent ici, et plus généralement les équations de Hamilton-Jacobi-Bellman (HJB) totalement non-linéaires. Nous avons donc proposé un schéma numérique probabiliste implémentable pour ces EDSRs contraintes, basé sur une randomisation initiale du contrôle puis sur l'élimination ultérieure de

ce bruit additionnel par l'ajout d'un calcul de supremum dans le schéma rétrograde. Nous avons analysé mathématiquement ce schéma, en particulier l'étape de discrétisation de ces EDSRs contraintes. Enfin, nous avons illustré numériquement les capacités de ce schéma sur le problème de sur-réplication d'option dans un modèle à volatilité incertaine ([5]).

1.1 Un modèle structurel risque-neutre pour la valorisation et la couverture de produits dérivés sur l'électricité

Cette première partie de la thèse a pour objet un modèle structurel de prix d'électricité adapté à la valorisation et la couverture partielle de produits dérivés sur ce marché. Elle a donné lieu à la publication de l'article [1].

En mathématiques financières, l'approche classique pour étudier un produit dérivé donné est de commencer par proposer un modèle pour la dynamique de l'actif sous-jacent, à partir de laquelle la valorisation et la stratégie de couverture du produit dérivé peuvent être déduites. C'est l'approche dite par *forme réduite*.

Ainsi, dans le cadre de cette approche, plusieurs auteurs ont cherché à proposer des modèles prenant en compte les caractéristiques propres au prix de l'électricité, tout en permettant de manière souple la valorisation et la couverture de produits dérivés standards (cf. [15] par exemple).

Cependant, de par sa nature, d'autres types de modélisation du prix de l'électricité sont possibles. En effet, l'électricité présente des caractéristiques uniques qui la différencient significativement de toutes les autres classes d'actifs. La plus importante est que l'électricité n'est pas stockable[1], ce qui implique que l'électricité soit produite en temps réel en quantité adéquate pour répondre exactement à la demande.

Différents types de centrales électriques permettent cette production, notamment à base d'énergies renouvelables (hydraulique, éolien, ...) ou d'énergies fossiles (charbon, gaz,...). Par conséquent, il est clair que la nature du parc électrique et le mécanisme de formation du prix vont significativement influer sur le prix de l'électricité. Ainsi, à l'opposé des modèles à forme réduite, les modèles d'*empilement* cherchent à exploiter ce mécanisme. Ils consistent à modéliser de manière fine la demande d'électricité, les prix de combustibles, ainsi que toutes les centrales électriques disponibles. Le prix d'électricité résulte alors d'une optimisation globale, qui permet de répondre à la demande d'électricité en utilisant les actifs de production disponibles de manière optimale (cf. [61] par exemple).

Si cette approche par optimisation globale permet une grande précision de modélisation (il est possible de prendre en compte nombre de détails, comme par exemple les contraintes dynamiques de production des centrales thermiques, cf. [80]), elle a l'inconvénient d'être très lourde à mettre en œuvre, et d'être très peu adaptée à l'étude de produits dérivés.

À mi-chemin entre ces deux extrêmes, la classe des *modèles structurels* cherche à tenir compte de manière simplifiée du mécanisme de formation du prix d'électricité, tout en utilisant les outils mathématiques classiques de l'approche par forme réduite. En particulier, cette approche est particulièrement bien adaptée à la valorisation d'option multi-sous-jacents faisant intervenir l'électricité et d'autres énergies comme le gaz par exemple, car elle permet de prendre en compte la structure fine de dépendance entre ces variables. Un état de l'art de la littérature sur les modèles structurels est disponible dans [30].

[1] Hormis via la production d'origine hydraulique, les barrages permettant un stockage indirect.

Le point de départ de ce chapitre est le modèle structurel par coût marginal de [3], dont voici la description. Considérons un marché d'électricité donné, sur lequel existe n types de centrales électriques. Pour chaque type $i = 1, \dots, n$ soit:

- S_t^i le prix du combustible utilisé par ce type de centrale (s'il s'agit d'une énergie renouvelable alors $S_t^i \equiv 0$).
- h_i son rendement énergétique, supposé constant (de sorte que $h_i S_t^i$ soit exprimé en €/MWh).
- C_t^i la capacité de production d'électricité à partir de ce type de centrale (en MW).

De plus, soit D_t la demande d'électricité (en MW). Sans perte de généralité, supposons que les technologies $i = 1, \dots, n$ sont classées par ordre croissant de coût de production : $h_1 S_t^1 \leq \dots \leq h_n S_t^n$. Alors, dans [3], le prix d'électricité P_t est modélisé directement par le *coût marginal* de production, défini par:

$$CM_t = \sum_{i=1}^{n} h_i S_t^i \mathbf{1} \left\{ \sum_{k=1}^{i-1} C_t^k \leq D_t \leq \sum_{k=1}^{i} C_t^k \right\} . \tag{1.1.1}$$

L'équation (1.1.1) signifie que, si la demande D_t dépasse la capacité totale de production des $i-1$ technologies les moins chères ($\sum_{k=1}^{i-1} C_t^k$) mais n'excède pas celle des i technologies les moins chères ($\sum_{k=1}^{i} C_t^k$), alors le coût marginal de production d'électricité est fixé par la i-ème technologie, et est donné par $h_i S_t^i$.

L'inconvénient du modèle $P_t = CM_t$ proposé dans [3] est que, si le coût marginal est effectivement un indicateur important du niveau de prix d'électricité, les deux peuvent être différents en pratique, du fait de tous les phénomènes secondaires négligés par cette modélisation simple[2]. Cette différence peut parfois être substantielle, comme par exemple lorsqu'un pic de prix inattendu se produit. En effet, le modèle $P_t = CM_t$ ne permet pas de générer de pic de prix, puisque par définition $CM_t \leq h_n S_t^n$ (et qu'une modélisation adéquate du prix S_t^n du combustible le plus onéreux ne comporte pas de pics).

La première étape de ce chapitre est d'enrichir le modèle (1.1.1) pour au moins pouvoir rendre compte des pics de prix d'électricité observables sur les marchés. En confrontant le coût marginal réalisé au prix réalisé (Figure 3.2.1a) et en préservant les quelques pics du jeu de données en choisissant de modéliser le lien entre les quantiles des variables (Proposition 3.2.2), le modèle suivant apparaît comme empiriquement pertinent (Figure 3.2.1b) :

$$P_t = g\left(\sum_{k=1}^{n} C_t^k - D_t \right) CM_t \, , \tag{1.1.2}$$

où la fonction g, que nous appelons *fonction de rareté*, est définie par :

$$g(x) = \min\left(M, \frac{\gamma}{x^\nu} \right) \mathbf{1}\{x > 0\} + M \mathbf{1}\{x \leq 0\} \, . \tag{1.1.3}$$

Ainsi, nous corrigeons le coût marginal par un terme multiplicatif qui est d'autant plus grand (en γ / x^ν) que la capacité résiduelle disponible du système ($\sum_{k=1}^{n} C_t^k - D_t$) est faible. En particulier, ce modèle permet de générer des pics de prix précisément lorsque le système est tendu (faible capacité résiduelle, proche de la défaillance). Ce terme correctif est plafonné à une valeur maximale M. L'amélioration entre les modèles (1.1.1) et (1.1.2) est illustrée par la Figure 3.2.2.

[2] Citons la non homogénéité des centrales de même type, les contraintes dynamiques de production, les réserves de production pour la stabilité réseau, les effet des imports-exports, les comportements stratégiques des producteurs,...

La seconde étape de ce chapitre est d'utiliser le nouveau modèle (1.1.2) pour valoriser et couvrir des produits dérivés sur l'électricité. Comme le prix d'électricité dans ce modèle est influencé à la fois par des sources de risques couvrables (les prix de combustibles S_t^i) et des sources de risques non couvrables (la demande D_t et les capacités de production C_t^i), le marché est donc incomplet. Autrement dit, la technique classique de couverture dynamique ne permet pas ici de supprimer totalement le risque lié à un produit dérivé donné. Néanmoins, sans le supprimer totalement, il reste possible de diminuer ce risque, en cherchant à minimiser un critère de risque préalablement choisi.

Dans le cadre des marchés incomplets, il existe toute une gamme de critères de risques potentiels pour construire des couvertures dynamiques partielles. Ici, nous choisissons le critère de *minimisation du risque local*, introduit par [52], pour sa simplicité, et pour le fait qu'il permet de séparer naturellement la partie couvrable de la partie non couvrable de flux basés sur un sous-jacent de type (1.1.2). Notons que le travail ci-dessous a par la suite été étendu au critère d'indifférence d'utilité dans [13].

À ce stade, nous avons besoin de poser un modèle sur les dynamiques des facteurs C_t, D_t et S_t. Nous choisissons des modèles très simples:

- Les écarts $Y_t^i := h_i S_t^i - h_{i-1} S_t^{i-1}$ sont modélisés par des browniens géométriques.
- La demande D_t et les capacités C_t^i sont modélisées par des diffusions générales (équations (3.3.2) et (3.3.3)).

Le point important est que les browniens liés aux prix des combustibles sont supposés indépendants des browniens liés à la demande et aux capacités (Hypothèse 1), hypothèse qui est raisonnable en pratique.

Après avoir identifié l'ensemble des mesures martingales équivalentes dans notre modèle (Proposition 3.3.1), dont la *mesure martingale minimale* $\widehat{\mathbb{Q}}$ liée au critère de minimisation du risque local, nous cherchons à valoriser quelques actifs dérivés significatifs.

Le premier exemple que nous étudions est celui d'un contrat à terme d'électricité $F_t^e(T)$ avec période de livraison instantanée T. Il s'agit bien sûr d'un actif fictif, étant donné que les contrat à terme d'électricité effectivement échangés sur les marchés comportent une période de livraison $[T_1, T_2]$ avec $T_1 < T_2$. Néanmoins, il est toujours possible de s'appuyer sur la brique intermédiaire $F_t^e(T)$ pour reconstruire le prix d'un vrai contrat à terme $F_t^e(T_1, T_2)$.

Dans notre modèle, le prix à la date t du contrat à terme d'électricité avec date de livraison T est donné par

$$F_t^e(T) = \sum_{i=1}^{n} h_i G_i^T(t, C_t, D_t) F_t^i(T) , \tag{1.1.4}$$

où, pour $i = 1, \dots, n$, $F_t^i(T)$ correspond au prix à la date t du contrat à terme sur le combustible i avec date de livraison T, et où la quantité $G_i^T(t, C_t, D_t)$, appelée fonction *Espérance Conditionnelle de la Rareté (ECR)* est définie par

$$G_i^T(t, C_t, D_t) := \mathbb{E}\left[g\left(\sum_{k=1}^{n} C_T^k - D_T \right) \mathbf{1}\left\{ \sum_{k=1}^{i-1} C_T^k \leq D_T \leq \sum_{k=1}^{i} C_T^k \right\} | \mathcal{F}_t^{D,C} \right] . \tag{1.1.5}$$

Ainsi, l'équation (1.1.4) nous indique que le prix à terme d'électricité peut s'écrire comme une combinaison linéaire des prix à terme sur combustibles. Les poids de cette combinaison linéaire sont donnés par le rendement énergétique h_i multiplié par la fonction ECR (1.1.5). S'il n'y avait pas la pénalité de rareté ($g \equiv 1$), alors la fonction ECR du combustible i correspondrait simplement à la probabilité que le combustible i soit marginal à la date T (comme dans [3]).

Ici, la présence de la fonction de rareté entraîne que la probabilité de marginalité est pénalisée de la façon décrite dans l'équation (1.1.5).

L'équation (1.1.4) permet de distinguer clairement l'influence des différentes sources de risques. Ainsi, les sources de risques couvrables S_t^i influencent les prix à terme de combustibles $F_t^i(T)$, tandis que l'influence des sources de risques non couvrables C_t^i et D_t se fait sentir sur les poids $G_i^T(t, C_t, D_t)$.

Outre la valeur de $F_t^e(T)$, il est possible, à partir de l'équation (1.1.4), de déduire sa dynamique, laquelle fait intervenir les dérivées partielle de la fonction ECR (cf. équations (3.3.13) et (3.3.14)).

Remarquons que, en utilisant des arguments similaires à ceux menant à l'équation (1.1.4), il est possible d'écrire la prime de risque d'électricité comme une combinaison linéaire des primes de risques sur combustibles (Proposition 3.3.2).

Nous passons ensuite à l'étude d'autres actifs. Outre les contrats à terme sur électricité, nous explicitons, dans le cas $n = 2$ combustibles, la valorisation d'une option d'achat sur la différence $P_T - h_i S_T^i$ (Proposition 3.3.4), ainsi que celle d'une option d'achat sur un contrat à terme d'électricité (Proposition 3.3.6).

Enfin, nous nous intéressons à la stratégie de couverture des ces produits dérivés. Considérons une option générale dont le résultat à échéance peut dépendre de la demande, des capacités, et de prix à terme de combustibles et d'électricité. Considérons des portefeuilles de couverture pouvant contenir des contrats à terme sur combustibles et électricité. Alors, et cela constitue le résultat principal de ce chapitre, nous sommes en mesure d'expliciter les quantités optimales à investir dans chacun des actifs de couverture (Proposition 3.3.7). Ces poids dépendent de la forme de l'option, et des différents termes intervenant dans la dynamique des prix à terme d'électricité (autrement dit de la fonction ECR et ses dérivées partielles). De plus, nous sommes aussi en mesure de quantifier le risque résiduelle de la couverture, lequel résulte des sources de risque non-couvrables (demande et capacités).

Le reste du chapitre consiste en des applications numériques de ces résultats.

Jusqu'ici nous avions modélisé la demande D_t et les capacités C_t^i par des diffusions générales. Pour les besoins de la partie numérique, il nous faut expliciter un choix de modèle pour ces variables. Nous choisissons des les modéliser par la combinaison d'une partie déterministe (permettant de prendre en compte les différents niveaux de saisonnalité) et d'une partie stochastique modélisée par un simple processus d'Ornstein-Uhlenbeck. Nous calibrons ces modèles sur les données de la zone France (cf. Images 3.4.1 et 3.4.2). Quoique très simple, ces modèles s'avèrent relativement fidèles aux observations.

Avec ces modèles, nous cherchons ensuite à calculer en pratique le prix d'un contrat à terme sur électricité. Comme l'indique l'équation (1.1.4), le point crucial est d'être capable de calculer efficacement la fonction ECR (1.1.5) et ses dérivées partielles.

Dans la littérature sur les modèles de prix d'électricité, que ce soit les modèles à forme réduite ou les modèles structurels, la capacité à obtenir des formules explicites pour les prix à terme d'électricité est généralement une caractéristique recherchée lors de la construction du modèle. Par exemple, dans le cadre des modèles structurels les auteurs peuvent préférer modéliser la courbe d'offre (ou ses variantes) par une simple fonction exponentielle (cf. [30]). Au prix d'une moindre précision, dont l'absence de pics de prix extrêmes, ce choix garantit des calculs simples.

Dans ce chapitre, motivé par nos observations empiriques, nous avons fait le choix d'utiliser une fonction puissance pour prendre en compte les pics de prix (équation (1.1.3)). Ce choix ne permet malheureusement pas de calculer explicitement les poids (1.1.5). Néanmoins, conscients de leur importance, nous avons cherché à rendre leur calcul le plus efficace possible numériquement.

Tout d'abord, nous avons mis en évidence le fait que le calcul de cette fonction ECR se résume à la capacité de calculer les quantités suivantes :

$$\widetilde{\mathcal{G}}(x,y;\nu) := \int_x^\infty \frac{1}{(y+z)^\nu} e^{-z^2} dz \tag{1.1.6}$$

$$\widetilde{\mathcal{H}}(m_1,m_2,\sigma_1,\sigma_2;\nu) := \int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^\infty \widetilde{\mathcal{G}}\left(\frac{\widetilde{x}-m_1-m_2-\sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}, \frac{m_1+m_2+\sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}; \nu\right) e^{-u^2} du \tag{1.1.7}$$

(cf. Propositions 3.4.3, 3.4.4 et 3.4.9). Considérons tout d'abord la quantité $\widetilde{\mathcal{G}}(x,y;\nu)$. Comme dans le cas $\nu = 1$ cette intégrale correspond à une fonction spéciale appelée intégrale de Goodwin-Staton incomplète (cf. [45]), nous avons baptisé la nouvelle fonction (1.1.6) *intégrale de Goodwin-Staton incomplète étendue*. Cette fonction peut s'interpréter de manière probabiliste, modulo renormalisations, comme densité de la somme d'une variable aléatoire gaussienne et d'une variable aléatoire de Pareto indépendantes (cf. Proposition B.1).

Nous avons établi le développement en série suivant:

$$\widetilde{\mathcal{G}}(x,y;\nu) = \frac{1}{2} e^{-y^2} \sum_{n=0}^{\infty} \Gamma\left(\frac{1-\nu}{2} + \frac{n}{2}, (x+y)^2\right) \frac{(2y)^n}{n!},$$

où $\Gamma(\alpha,x)$ correspond à la fonction Gamma incomplète. Combiné à des relations de récurrence et des développements asymptotiques, ce développement permet de calculer très efficacement $\widetilde{\mathcal{G}}(x,y;\nu)$. De la même manière, nous avons établi un développement en série pour la fonction $\widetilde{\mathcal{H}}(m_1,m_2,\sigma_1,\sigma_2;\nu)$, ainsi que des relations de récurrence et des développements asymptotiques appropriés, rendant son estimation numérique très efficace. Finalement, nous avons explicité comment calculer les dérivées partielle de la fonction ECR (1.1.5) à l'aide des deux fonctions $\widetilde{\mathcal{G}}$ et $\widetilde{\mathcal{H}}$.

Finalement, avec ces outils de calcul, nous avons pu tester numériquement sur deux exemples la valorisation et la couverture partielle d'actifs dérivés sur l'électricité, à l'aide des formules générales établies auparavant.

Le premier exemple numérique consiste en la valorisation d'un contrat à terme d'électricité de maturité $T = 3\text{mois}$, avec test de la couverture partielle à base de contrats à terme sur combustibles. Le comportement temporel de la couverture est très intéressant, et est résumé par la Figure 3.4.4. Deux phases temporelles bien distinctes apparaissent:

- Loin de la maturité, pour $t \in [0, T-\Delta]$ avec $\Delta \sim 2\text{semaines}$, la couverture partielle est pratiquement parfaite. En effet, dans cette phase, ce sont les risques sur les prix de combustibles qui dominent. Or, il s'agit précisément des risques couvrables, que la stratégie de couverture permet donc d'éliminer. Dans cette phase, le contrat à terme sur électricité se comporte pratiquement comme un panier de contrats à terme sur combustibles.

- Proche de la maturité, pour $t \in [T-\Delta, T]$, la couverture partielle est totalement inefficace. En effet, dans cette phase finale, ce sont les risques sur le niveau de demande et de capacités qui dominent. Comme il s'agit des sources de risques non couvrables, la réduction de risque apportée par la stratégie de couverture devient négligeable.

Cette observation empirique importante est très générale. Elle n'est ni liée à la forme spécifique de l'actif contingent choisi, ni à celle du critère de couverture, mais au modèle de prix (1.1.2) et au fait que les sources de risque couvrables sont de type martingale, tandis que les sources de risque non couvrables sont de type retour à la moyenne (de sorte que les lois de C_T et D_T

ne dépendent guère des valeurs courantes C_t et D_t lorsque $T-t$ est grand, et que la valeur de l'instant de changement de phase $T-\Delta$ dépend essentiellement des valeurs des coefficients de retour à la moyenne). De manière plus générale, cette observation fixe les limites de la capacité des produits sur combustibles à réduire le risque de positions sur l'électricité.

Enfin, le deuxième exemple numérique consiste en la valorisation d'options sur différence $P_T - h_i S_T^i$. Cette fois, il est nécessaire de recourir à des méthodes d'intégration numérique. Les résultats sont donnés par la Figure 3.4.5. L'intérêt pratique de cet exemple est d'illustrer certains avantages pratiques d'une modélisation structurelle, dont la prise en compte sur les prix des effets de saisonnalité, mais aussi de modifications structurelles du marché, telle que l'installation anticipée de nouvelles centrales électriques.

1.2 Un algorithme probabiliste pour la résolution de problèmes de commutation optimale en grande dimension

La seconde partie de la thèse a pour objet la résolution pratique de problèmes de contrôle stochastique en grande dimension, plus précisément des problèmes appartenant à la classe des problèmes de commutation optimale. Elle a donné lieu à l'article [2], actuellement en cours de révision pour publication.

Commençons par définir cette classe de problèmes. Nous considérons les éléments suivants :

- $X = (X_t)_{t\geq 0}$ un processus stochastique à valeurs dans $\mathbb{R}^d$, partant de $x_0 \in \mathbb{R}^d$ à l'instant initial $t = 0$.
- $I^\alpha = (I_t^\alpha)_{t\geq 0}$ un processus constant par morceaux à valeurs dans $\mathbb{R}^{d'}$, partant de $i_0 \in \mathbb{R}^{d'}$ à l'instant initial $t = 0$. Plus précisément, on suppose que I^α prend ses valeurs dans un sous-ensemble fini $\mathbb{I}_q = \{i_1, \dots, i_q\}$ de $\mathbb{R}^{d'}$. Ce processus I^α est commandé au cours du temps par une stratégie α.
- α est une stratégie impulsionnelle définie par une séquence $(\tau_n, \iota_n)_{n\in\mathbb{N}}$ de temps d'arrêts croissants $\tau_n \geq 0$ et de variables aléatoires $\mathcal{F}_{\tau_n}-$mesurables à valeurs dans $\mathbb{I}_q$. Le processus commandé I^α se déduit de cette séquence comme suit :

$$I_t^\alpha = \iota_n \text{ quand } t \in [\tau_n, \tau_{n+1}[\ .$$

- Parmi l'ensemble des stratégies α possibles, nous considérons seulement celles appartenant à la classe des stratégies admissibles $\mathcal{A}$. Essentiellement, cela revient à considérer seulement les stratégies telles que $\tau_n \to +\infty$ p.s. quand $n \to \infty$ (on exclut les points d'accumulation).
- $f : \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}^{d'} \to \mathbb{R}$ et $k : \mathbb{R} \times \mathbb{R}^{d'} \times \mathbb{R}^{d'} \to \mathbb{R}$ deux fonctions mesurables.

Alors, le problème de contrôle stochastique que nous considérons est le suivant :

$$v(0, x_0, i_0) = \sup_{\alpha\in\mathcal{A}} \mathbb{E}\left[\int_0^\infty f(t, X_t, I_t^\alpha)\, dt - \sum_{\tau_n \geq 0} k(\tau_n, \iota_{n-1}, \iota_n)\right]. \tag{1.2.1}$$

Ainsi, l'objectif est de maximiser le gain apporté par la fonction f au cours du temps. Celui-ci dépend d'une part de la variable d'état non contrôlée X, et de la variable contrôlée I^α. Le moyen de maximiser le gain est donc d'ajuster de la manière la plus avantageuse possible ce contrôle au cours du temps. Cependant, le fait de modifier l'état du contrôle génère un coût défini par la fonction k. Ce coût dépend des valeurs du contrôle avant et après l'action.

Le problème (1.2.1) doit son appellation de commutation optimale au fait que le processus commandé I^α prend ses valeurs dans un ensemble fini.

De manière plus précise, pour s'assurer que le problème (1.2.1) est bien posé, il convient de faire quelques hypothèses ordinaires de régularité sur les divers éléments qui interviennent (cf. Section 4.2.2).

Tout d'abord, expliquons pourquoi nous nous sommes intéressés à ce problème. Dans le premier chapitre, nous avons construit un modèle structurel de prix d'électricité qui permettait de modéliser finement la dépendance temporelle entre le prix d'électricité et le prix des autres énergies. En particulier, ce modèle permet de valoriser de manière adéquate des options sur différence de prix (entre électricité et une autre énergie), et donc, par extension, de quantifier, par une approche d'option réelle, la valeur d'une centrale électrique donnée. Avec un tel outil à disposition, nous avons voulu savoir s'il était possible de déterminer les meilleurs investissements possibles en centrales électriques au cours du temps. (Quel type de centrale contruire ? Combien ? Et quand ?) Nous verrons plus loin que ce problème d'investissement peut effectivement s'écrire sous la forme d'un problème de commutation optimale de type (1.2.1) en dimension élevée ($d+d' \gg 3$). C'est pour cela que nous avons cherché à construire une méthode numérique capable de résoudre en pratique le problème (1.2.1) en dimension élevée.

Maintenant, détaillons la façon dont nous allons procéder pour résoudre le problème (1.2.1). Si l'on n'autorise les modifications de stratégie que sur une grille temporelle fixe $\Pi = \{t_0 = 0 < t_1 < \ldots < t_N = T\}$, alors, en utilisant le principe de programmation dynamique, la fonction valeur discrétisée v^Π va vérifier la relation d'induction rétrograde suivante:

$$v^\Pi(t_n, x, i) = \max_{j \in \mathbb{I}_q} \{E(t_n, x, j) - k(t_n, i, j)\}, \tag{1.2.2}$$

où:

$$E(T, x, i) := \mathbb{E}\left[\int_T^\infty f(s, X_s, i)\, dt \,|X_T = x\right]$$

$$E(t_n, x, i) := \mathbb{E}\left[\int_{t_n}^{t_{n+1}} f(s, X_s, i)\, dt \,|X_{t_n} = x\right] + \mathbb{E}\left[v^\Pi(t_{n+1}, X_{t_{n+1}}, i) \,|X_{t_n} = x\right], n = N-1, \ldots, 0.$$

En pratique, au delà des discrétisations de processus et du calcul des valeurs terminales, la difficulté principale dans la mise en œuvre pratique du schéma (1.2.2) réside dans le calcul de l'espérance conditionnelle $\mathbb{E}\left[v^\Pi(t_{n+1}, X_{t_{n+1}}, i) \,|X_{t_n} = x\right]$. En effet, cette quantité ne peut généralement pas être calculée de manière explicite, ce qui oblige à recourir à des méthodes d'approximations.

Dans la littérature sur les options américaines (qui est un problème plus simple que (1.2.1)), le même problème se pose, et diverses méthodologies ont été proposées pour y faire face. On peut notamment citer l'approche par régression linéaire ([83, 102]), l'approche par régression non paramétrique ([101, 74]), l'approche par quantification ([8]) et l'approche par calcul de Malliavin ([6]).

Dans le cadre des problèmes de contrôle stochastique sans coût de commutation ($k \equiv 0$), l'approche par calcul de Malliavin a été testée dans [88], et l'approche par régression dans [12]. Enfin, pour le problème de commutation optimale qui nous intéresse ici, l'approche par quantification est utilisée dans [55], et l'approche par régression linéaire dans [32]. Cependant, la méthodologie employée dans [32] n'utilise pas directement le principe de programmation dynamique comme dans l'équation (1.2.2) mais a recours à la représentation des problèmes de commutation optimale comme couches successives de problèmes d'arrêts optimaux.

Dans ce chapitre, eu égard aux conclusions du comparatif numérique [28] (qui a pour cadre les options américaines), nous avons décidé d'utiliser l'approche par régression locale pour résoudre notre problème de commutation optimale. Celle-ci apparaît en effet comme la plus apte aujourd'hui à manier des problèmes de dimension élevée.

La première partie de ce chapitre est consacrée à l'analyse de l'erreur entre la quantité initiale (1.2.1) que l'on veut estimer et le résultat du schéma numérique qui l'approxime. Entre les deux, il est nécessaire d'effectuer une série d'approximations. Nous avons déjà évoqué la discrétisation temporelle pour l'équation intermédiaire (1.2.2), mais d'autres couches d'approximations sont requises. En voici la liste :

- *[Horizon temporel fini]* La première est la gestion de l'horizon de temps. Il est nécessaire en effet de limiter l'horizon de temps des stratégies à une date $T > 0$ grande mais finie.

- *[Discrétisation temporelle]* Ensuite, on a besoin de discrétiser les processus X et I^α. Pour simplifier, on peut choisir une grille de temps $\Pi = \{t_0 = 0 < t_1 < \ldots < t_N = T\}$ avec un pas de temps fixe $h > 0$.

- *[Localisation spatiale]* Nous aurons aussi besoin de projeter le processus X (qui peut a priori prendre ses valeurs dans $\mathbb{R}^d$ tout entier) dans un sous-domaine $\mathcal{D}_\varepsilon$ vaste mais borné. La raison du besoin de cette étape sera donnée un peu plus loin.

- *[Approximation des espérances conditionnelles]* Enfin, la dernière étape consiste à approximer les espérances conditionnelles qui apparaissent lorsqu'on utilise le principe de programmation dynamique sur le problème discrétisé. Nous choisissons de les approximer par régression linéaire à partir d'un faisceau de trajectoires de Monte Carlo. Nous reviendrons sur le choix de la base de régression.

Pour chacune de ces étapes, nous avons analysé l'erreur commise par l'approximation en question.

Le passage à un horizon de temps fini est analysé dans la Proposition 4.3.1. L'élément clé est l'inclusion d'une actualisation exponentielle dans la définition de f et k (cf les hypothèses listés dans la Section 4.2.2). L'erreur se comporte en $e^{-\rho T}$.

La discrétisation temporelle est analysé dans les Propositions 4.3.2 et 4.3.3. L'essentiel de cette analyse provient de [55]. L'erreur se comporte en $\sqrt{h}$ (ou en $\sqrt{h}\sqrt{\log\left(\frac{2T}{h}\right)}$ lorsque la fonction de coût k dépend aussi de la variable d'état X).

Pour la localisation spatiale, le domaine $\mathcal{D}_\varepsilon$ est directement choisi tel que l'erreur commise se comporte en ε (Proposition 4.3.4).

Finalement, nous en arrivons à l'étape la plus complexe, qui est l'analyse de l'erreur de régression. À ce stade, la programmation dynamique appliquée à la fonction valeur approximée peut s'écrire de la manière simple suivante :

$$\begin{aligned}
&\bar{v}_\Pi(T,x,i) = g(T,x,i)\\
&\bar{v}_\Pi(t_n,x,i) = \max_{j\in\mathbb{I}_q}\left\{hf(t_n,x,j) - k(t_n,i,j) + \mathbb{E}\left[\bar{v}_\Pi\left(t_{n+1},\bar{X}_{t_{n+1}},j\right)\middle|\bar{X}_{t_n}=x\right]\right\}, n = N-1,\ldots,0.
\end{aligned} \tag{1.2.3}$$

Comme annoncé précédemment, nous choisissons d'approximer les espérances conditionnelles apparaissant dans l'équation (1.2.3) par régression. Afin de permettre l'obtention d'une borne d'erreur explicite, nous avons dû restreindre notre analyse à une base de régression bien spécifique, à savoir une base de fonctions indicatrices sur un ensemble d'hypercubes $(B_k)_{k=1,\ldots,K}$ partitionnant le domaine $\mathcal{D}_\varepsilon$ (cf. Hypothèse 6).

Pour la clarté et la simplicité, nous avons réalisé l'analyse d'erreur en deux étapes successives (même si les réaliser simultanément peut s'avérer plus précis, cf. [82] par exemple) :

- *[Régression théorique]* Nous remplaçons tout d'abord, à chaque pas de temps t_n, les espérances conditionnelles par une projection sur l'ensemble des variables aléatoires $\mathcal{F}_{t_n}$-mesurables générables à partir de la base de régression choisie. Cette erreur se comporte en δ/h, où δ est la taille maximale des côtés des hypercubes.

- *[Régression empirique]* Puis, nous remplaçons la projection ci-dessus par la projection empirique calculée à partir d'un échantillon de M réalisations indépendantes de la variable d'état X_{t_n}. Nous avons montré que la L_p-norme de la différence entre les variables régressées théoriques et empiriques se comporte essentiellement en $C_p/\left(\sqrt{M}\times P^{1-\frac{1}{p\vee 2}}\right)$, où $C_p>0$ et $P=\min_{t\in\Pi}\min_{B_k\subset\mathcal{D}_\varepsilon}\mathbb{P}\left(\bar{X}_t\in B_k\right)$ (cf. Proposition 4.3.6).

Cette dernière estimation d'erreur généralise à la norme L_p un résultat démontré dans [100] dans le cas de la norme L_2. La démonstration fait usage du résultat pratique suivant, démontré en Annexe 4.7.1 à l'aide des inégalités de Jensen et de Marcinkiewicz-Zygmund:

$$\left\|\frac{1}{M}\sum_{m=1}^{M}X_m\right\|_{L_p}\leq\frac{C_p}{\sqrt{M}}\left\|X_1\right\|_{L_{p\vee 2}},$$

où $X_1,\dots,X_M$ est un échantillon i.i.d. de variables aléatoires réelles centrées, dont le moment d'ordre $p\vee 2$ existe.

Nous pouvons maintenant expliquer l'intérêt de l'étape de localisation spatiale: elle permet de garantir que la variable P au dénominateur de l'erreur de régression reste strictement positive quel que soit M, ce qui assure la convergence de l'algorithme. Bien sûr, cette étape est requise sur le plan théorique seulement, car en pratique, pour M fixé, un faisceau de M trajectoires est déjà contenu dans un espace fini, et une restriction supplémentaire de l'espace serait donc redondante.

Finalement, le résultat théorique principal du chapitre est le Théorème 4.3.1, qui fournit une analyse complète, par combinaison de toutes les étapes décrites jusqu'ici, de l'erreur de convergence entre le problème de contrôle initial (1.2.1) et le résultat de l'algorithme que nous proposons.

La seconde partie de ce chapitre est consacrée à la complexité de l'algorithme, à la fois sur le plan du nombre d'opérations et de l'espace mémoire requis.

De manière générale nous montrons que la complexité algorithmique du schéma que nous proposons se comporte en $\mathcal{O}\left(q^2\cdot N\cdot M\right)$, où q est le nombre d'éléments de l'ensemble de commutation $\mathbb{I}_q$, N est le nombre de pas de temps sur la grille Π, et M est le nombre de trajectoires de Monte Carlo. Sous certaines conditions, explicitées à la Section 4.4.1.1, cette complexité peut être améliorée, pour atteindre $\mathcal{O}\left(q\cdot N\cdot M\right)$. Cette complexité améliorée est très satisfaisante, et est, selon nous, sans doute la meilleure que puisse atteindre un algorithme de résolution de problème de commutation optimale par Monte Carlo (le simple fait de parcourir chaque commutation possible pour chaque trajectoire et chaque pas de temps coûtant déjà $\mathcal{O}\left(q\cdot N\cdot M\right)$).

Comme le schéma d'Euler est progressif et que la programmation dynamique est rétrograde, la première étape du schéma est de simuler, en partant de x_0, des trajectoires du processus $\bar{X}$ jusqu'au temps terminal T, pour pouvoir ensuite initier la programmation dynamique. Comme, à chaque instant $t_n\in\Pi$, l'étape de régression nécessite d'avoir accès à l'échantillon $\left(\bar{X}_{t_n}^m\right)^{1\leq m\leq M}$,

la solution la plus simple, communément mise en œuvre, est de stocker entièrement les trajectoires $\left(\bar{X}_{t_n}^m\right)_{1\leq i\leq N}^{1\leq m\leq M}$ pendant le déroulement de la résolution. Cela nécessite un espace mémoire en $\mathcal{O}\left(N\cdot M\right)$.

Si cela ne pose pas de problème lorsque l'horizon de temps est proche (N petit) ou que la dimension du problème est peu élevée (M petit), cela peut rapidement devenir redhibitoire en dimension élevée avec horizon de temps lointain (une combinaison que vérifie précisément l'application numérique qui nous intéresse). Pour surmonter cette limitation, nous avons généralisé une méthode de réduction de mémoire, introduite dans [36] dans le cas d'un mouvement brownien géométrique, afin qu'elle puisse fonctionner pour n'importe quel processus discrétisé par schéma d'Euler. Celle-ci permet, au prix d'un doublement du nombre d'opérations, de diminuer la taille mémoire requise à $\mathcal{O}\left(N+M\right)$, ce qui est un gain considérable qui supprime toute contrainte mémoire.

L'idée, détaillée, analysée et illustrée à la Section 4.4.2, est la suivante. Lors de la simulation initiale du processus sous-jacent $\bar{X}$, à chaque instant $t_n\in\Pi$, au lieu de stocker tout l'échatillon $\left(\bar{X}_{t_n}^m\right)^{1\leq m\leq M}$, il suffit de stocker simplement la valeur de la graine du générateur aléatoire juste avant la simulation de l'échantillon courant. À la fin de ce premier passage, la mémoire contient un échantillon de $\bar{X}_T$ (taille M) et les graines du générateur aléatoire (taille N). Ensuite, lors de l'étape d'induction rétrograde, à chaque instant $t_n\in\Pi$, il est possible, à partir de l'échantillon $\left(\bar{X}_{t_{n+1}}^m\right)^{1\leq m\leq M}$ et de la graine de l'instant t_n, de reconstruire l'échantillon $\left(\bar{X}_{t_n}^m\right)^{1\leq m\leq M}$, en inversant la formule du schéma d'Euler et en resimulant, à partir de la graine stockée, les aléas ayant permis de passer précédemment de $\left(\bar{X}_{t_n}^m\right)^{1\leq m\leq M}$ à $\left(\bar{X}_{t_{n+1}}^m\right)^{1\leq m\leq M}$.

Finalement, la dernière partie du chapitre est consacrée à une application numérique de notre schéma de résolution au problème d'investissement en centrales électriques.

En adaptant le modèle de prix d'électricité développé au premier chapitre, il est possible d'exprimer le problème d'investissement en centrales électriques sous la forme d'un problème de commutation optimale (équation (4.5.9)). Celui-ci est néanmoins lourd et imposant, car il combine à la fois un horizon de temps très lointain (plusieurs décennies) et une variable d'état de grande dimension (demande, capacités, combustibles,...). Mais, comme nous l'avons vu lors de l'analyse de complexité, notre algorithme est précisément conçu pour pouvoir faire face, dans une certaine mesure, à de tels problèmes.

Nous avons appliqué avec succès notre algorithme sur un exemple numérique avec deux technologies. La Figure 4.5.1 illustre l'estimation de la distribution de nouvelles centrales à la date terminale, et les Figures 4.5.2 et 4.5.3 illustrent l'impact des nouvelles centrales sur l'évolution du prix de l'électricité. On y constate l'effet prévisible des nouvelles centrales (prix d'électricité plus faibles et moins volatils) ainsi que la moindre attractivité, relativement aux centrales de base, des centrales de pointe lorsque les prix des combustibles sont élevés.

Terminons par quelques remarques:

- Afin d'évaluer empiriquement si les paramètres numériques choisis étaient suffisants pour garantir une convergence convenable des résultats obtenus, nous avons adapté une méthode d'obtention d'intervalles de confiance empiriques décrite dans [28] (cf. Annexe 4.7.4).

- Pour garantir un certain réalisme du modèle de prix de combustibles sur des horizons de temps très grands, nous avons ajouté de la cointégration au modèle de browniens géométriques du premier chapitre:

$$dS_t = \Xi S_t dt + \mathrm{diag}\left(S_t\right)\Sigma dW_t\ ,\ S_0>0\,,$$

où Ξ est la matrice de cointégration et Σ la matrice de covariance. En Annexe 4.7.2, nous démontrons une condition nécessaire et suffisante pour garantir la positivité de S_t (il faut et il suffit que les termes non-diagonaux de la matrice de cointégration soient positifs).

- Pour finir, afin de représenter d'une manière graphique intuitive l'évolution temporelle d'un processus stochastique potentiellement multimodal (comme par exemple le prix d'électricité généré par notre modèle), nous avons généralisé la construction d'intervalles interquantiles à des boréliens (potentiellement non connexes) ajustés aux lignes de niveau de la distribution estimée (cf. Annexe 4.7.5). Cette construction générale pourrait s'avérer utile dans une large gamme de domaines.

1.3 Un algorithme numérique pour la résolution des équations de HJB totalement non-linéaires via EDSRs à sauts négatifs

La troisième et dernière partie de la thèse a de nouveau pour objet la résolution de problèmes de contrôle stochastique, mais cette fois nous nous intéressons à des problèmes dont la dynamique du processus d'état sous-jacent peuvent être influencées par le contrôle. Le contenu de cette partie synthétise le contenu deux articles [71] et [70].

Commençons par l'exemple introductif suivant :

$$\begin{aligned} v(t,x) &= \sup_{\alpha\in\mathcal{A}} \mathbb{E}\left[\int_t^T f\left(X_s^\alpha,\alpha_s\right)dt + g\left(X_T^\alpha\right)|X_t^\alpha = x\right] && (1.3.1)\\ dX_s^\alpha &= b\left(X_s^\alpha,\alpha_s\right)ds + \sigma\left(X_s^\alpha\right)dW_s\,, && (1.3.2)\end{aligned}$$

où la diffusion X^α prend ses valeurs dans $\mathbb{R}^d$, f et g sont deux fonctions mesurables, et les processus $\alpha = (\alpha_s)_{t\le s\le T}$ sont des stratégies choisies parmi un ensemble $\mathcal{A}$ de stratégies admissibles à valeurs dans un sous-domaine compact $A \subset \mathbb{R}^q$.

Sur cet exemple (1.3.1), les stratégies peuvent effectivement venir modifier la dynamique du processus sous-jacent X^α, mais seulement au travers de la dérive b. Dans ce cas précis, le problème peut se résoudre par le biais d'une Équation Différentielle Stochastique Rétrograde (EDSR).

En effet, le problème de contrôle stochastique (1.3.1) est relié, tout d'abord, à l'équation de Hamilton-Jacobi-Bellman (HJB) suivante :

$$\begin{aligned} &\frac{\partial v}{\partial t} + \sup_{a\in A}\left\{b(x,a).D_x v + \frac{1}{2}\mathrm{tr}\left(\sigma\sigma^\top(x)\,D_x^2 v\right) + f(x,a)\right\} = 0\ ,\ (t,x)\in[0,T)\times\mathbb{R}^d\\ &v(T,x) = g(x)\quad ,\ x\in\mathbb{R}^d\,. \end{aligned} \tag{1.3.3}$$

Il est possible d'essayer de résoudre cette EDP numériquement, mais tout comme au chapitre précédent, nous sommes ici à la recherche d'un algorithme de type probabiliste, afin d'être capable de traiter des problèmes de dimension élevée.

Dans le cas où σ est de rang plein, l'équation (1.3.3) peut se réécrire sous la forme de l'EDP semi-linéaire suivante :

$$\begin{aligned} &\frac{\partial v}{\partial t} + F\left(x,\sigma^\top D_x v\right) + \frac{1}{2}\mathrm{tr}\left(\sigma\sigma^\top(x)\,D_x^2 v\right) = 0\ ,\ (t,x)\in[0,T)\times\mathbb{R}^d\\ &v(T,x) = g(x)\quad ,\ x\in\mathbb{R}^d\,, \end{aligned} \tag{1.3.4}$$

où la transformée F est définie par $F(x,z) := \sup_{a\in A}\{\theta(x,a).z + f(x,a)\}$, où θ est tel que $\sigma(x)\theta(x,a) = b(x,a)$. Cette nouvelle EDP (1.3.4) est elle-même associée à l'EDSR suivante :

$$\begin{aligned} dX_s^0 &= \sigma\left(X_s^0\right) dW_s \\ Y_t &= g\left(X_T^0\right) + \int_t^T F\left(X_s^0, Z_s\right) ds - \int_t^T Z_s dW_s \,. \end{aligned}$$

Comme divers schémas numériques probabilistes sont disponibles pour résoudre une EDSR donnée ([27, 82, 58]), il est donc possible de résoudre effectivement le problème de contrôle stochastique (1.3.1) par un schéma numérique probabiliste.

Maintenant, intéressons nous au problème plus complexe suivant :

$$\begin{aligned} v(t,x) &= \sup_{\alpha\in\mathcal{A}} \mathbb{E}\left[\int_t^T f(X_s^\alpha, \alpha_s)\, dt + g(X_T^\alpha) \,|X_t^\alpha = x\right] \\ dX_s^\alpha &= b(X_s^\alpha, \alpha_s)\, ds + \sigma(X_s^\alpha, \alpha_s)\, dW_s \,, \end{aligned} \tag{1.3.5}$$

lequel est associé à l'équation de HJB totalement non-linéaire suivante:

$$\begin{aligned} &\frac{\partial v}{\partial t} + \sup_{a\in A}\left\{ b(x,a).D_x v + \frac{1}{2}\mathrm{tr}\left(\sigma\sigma^\top(x,a) D_x^2 v\right) + f(x,a)\right\} = 0 \;,\; (t,x)\in[0,T)\times\mathbb{R}^d \\ &v(T,x) = g(x) \quad ,\; x\in\mathbb{R}^d \,. \end{aligned} \tag{1.3.6}$$

Cette fois-ci, la stratégie α peut influer à la fois sur la dérive b et sur la volatilité σ du processus X^α. De fait, il s'agit d'un problème bien plus complexe que le précédent exemple (1.3.1), qui ne peut pas se réduire à la résolution d'une simple EDSR.

Ainsi, l'objet de ce chapitre est précisément de construire un algorithme probabiliste capable de résoudre des problèmes de contrôle stochastique de type (1.3.5). L'idée est de s'inspirer de la résolution décrite ci-dessus du problème plus simple (1.3.1), et de l'étendre de manière adéquate pour pouvoir englober les problèmes de type (1.3.5).

La première étape est d'identifier l'objet probabiliste adéquat pour représenter la solution du problème (1.3.5). Cet objet est défini et étudié dans [72], et consiste en une EDSR à saut négatifs. Commençons par rappeler cette construction, et son lien avec le problème (1.3.6).

Soit I un processus de sauts pur, indépendant de W, à valeurs dans A. Sa dynamique est définie comme suit :

$$dI_s = \int_A (a - I_{s^-})\mu_A(ds, da)\,, \tag{1.3.7}$$

où $\mu_A(dt, da)$ est une mesure aléatoire de Poisson sur $\mathbb{R}_+ \times A$, dont la mesure d'intensité $\lambda_A(da)\, dt$ est finie.

Ensuite, considérons la diffusion markovienne à basculement de régime (X, I) dont la dynamique de I est donnée par (1.3.7) et la dynamique de X par :

$$dX_s = b(X_s, I_s)\, ds + \sigma(X_s, I_s)\, dW_s \,. \tag{1.3.8}$$

Autrement dit, nous remplaçons le contrôle α dans la dynamique (1.3.5) de X^α par un nouveau processus stochastique I indépendant du reste. L'espérance $\mathbb{E}\left[\int_t^T f(X_s, I_s)\, dt + g(X_T)\,|X_t = x\right]$ est alors associée à la solution de l'EDSR à sauts suivante :

$$Y_t = g(X_T) + \int_t^T f(X_s, I_s)\, ds - \int_t^T Z_s dW_s - \int_t^T \int_A U_s(a)\, \tilde{\mu}_A(ds, da)\;, \tag{1.3.9}$$

où $\tilde{\mu}_A$ est la mesure compensée de μ_A.

Enfin, afin de retomber sur la solution de l'équation de HJB (1.3.6), il faut et il suffit d'ajouter la contrainte suivante sur le terme $U_t(a)$ de l'équation (1.3.9) :

$$U_t(a) \leq 0, \ d\mathbb{P} \otimes dt \otimes \lambda(da) \ p.s. \ .$$

Sous les hypothèses suivantes, que nous supposerons vérifiée dans toute la suite du chapitre :

- b et σ lipschitziennes (H1),
- f et g lipschitziennes (H2),

il est démontré dans [72] que la solution minimale (Y, Z, U, K) de l'EDSR contrainte

$$\begin{aligned} Y_t &= g(X_T) + \int_t^T f(X_s, I_s)\, ds - \int_t^T Z_s dW_s \\ &\quad + K_T - K_t - \int_t^T \int_A U_s(a)\, \tilde{\mu}_A(ds, da) \ , \ 0 \leq t \leq T \ , \ p.s. \\ U_t(a) &\leq 0 \ , \ d\mathbb{P} \otimes dt \otimes \lambda(da) \ \ p.s. \text{ sur } \Omega \times [0,T] \times A \end{aligned} \tag{1.3.10}$$

est reliée à la solution de l'équation de HJB (1.3.6). En effet, il y est démontré qu'il est possible d'écrire le Y de la solution minimale de (1.3.10) comme une fonction de t et de X_t seulement :

$$Y_t = v(t, X_t) \ , \quad t \in [0,T] \ , \tag{1.3.11}$$

et que cette fonction v est l'unique solution de viscosité continue à croissance linéaire de l'équation de HJB (1.3.6). Ce résultat reste valide si f dépend aussi de Y et Z (ce qui permet d'englober des problèmes plus généraux que le problème de contrôle stochastique (1.3.5)). De plus, soulignons que ce résultat (ainsi que tous les résultats de ce chapitre) ne nécessitent aucune condition d'ellipticité uniforme sur σ.

Une conséquence de ce théorème important est que pour trouver une solution à l'équation de HJB totalement non-linéaire (1.3.6) (et donc au problème de contrôle stochastique (1.3.5)), il suffit de trouver une solution à l'EDSR contrainte (1.3.10). Aussi, l'objet de ce chapitre sera donc de proposer, analyser et tester un schéma numérique probabiliste pour l'EDSR contrainte (1.3.10).

La première étape consiste à discrétiser cette EDSR contrainte. Soit $\pi := \{0 = t_0 < \ldots < t_N = T\}$ une grille de temps déterministe entre les instants 0 et T, de pas $|\pi| = \max_{0 \leq k \leq n-1} \{t_{k+1} - t_k\}$. Alors, en s'inspirant des schémas pour les EDSRs classiques, et en interprétant de manière adéquate la contrainte sur $U_t(a)$, il apparait que le schéma discret suivant approxime l'EDSR contrainte (1.3.10) (dans le cas général où f dépend de y et z) :

$$\begin{cases} Y_N & = g(X_N) \\ \Delta_i \mathcal{Z}_i & = \mathbb{E}_i \left[Y_{i+1} \Delta W_i^\top \right] \\ \mathcal{Y}_i & = \mathbb{E}_i \left[Y_{i+1} + f(X_i, I_i, Y_{i+1}, \mathcal{Z}_i)\, \Delta_i \right] \\ Y_i & = \operatorname{ess\,sup}_{a \in A} \mathbb{E}_{i,a} \left[\mathcal{Y}_i \right] \ , \end{cases} \tag{1.3.12}$$

où $(X_i, I_i)_{i=1,\ldots,N}$ correspond à la discrétisation du processus $(X_t, I_t)_{0 \leq t \leq T}$ par schéma d'Euler, $\mathbb{E}_i[.] := \mathbb{E}[. \,|\mathcal{F}_{t_i}]$, et $\mathbb{E}_{i,a}[.] := \mathbb{E}[. \,|X_i, I_i, I_i = a] = \mathbb{E}[. \,|X_i, I_i = a]$. Pour bien comprendre ce schéma, il faut remarquer que, par la propriété de Markov, les quantités intermédiaires $\mathcal{Y}_i$ et $\mathcal{Z}_i$ peuvent s'écrire comme fonction de X_i et I_i :

$$(\mathcal{Y}_i, \mathcal{Z}_i) = (\tilde{y}_i(X_i, I_i), \tilde{z}_i(X_i, I_i)) \ .$$

On retrouve cette même étape dans les schémas de discrétisation des EDSRs classiques. Ici vient s'y ajouter la prise en compte de la contrainte sur $U_t(a)$, qui permet de retomber sur la vraie approximation $(Y_i, Z_i) = (y_i(X_i), z_i(X_i))$, qui ne dépend pas de I_i. D'après la définition de $U_t(a)$, la contrainte peut se réécrire

$$\tilde{y}_i(X_i, a) - y_i(X_i) = U_t(a) \leq 0 \ p.s. \ \forall a \in A,$$

ce qui implique que le terme Y_i de la solution minimale doit satisfaire

$$Y_i = y_i(X_i) = \operatorname*{ess\,sup}_{a \in A} \tilde{y}_i(X_i, a) = \operatorname*{ess\,sup}_{a \in A} \mathbb{E}_{i,a}[\mathcal{Y}_i] .$$

Si nécessaire, il est possible d'extraire le terme Z_i du schéma. En effet, en notant $a_i^* = \operatorname{argess\,sup}_{a \in A} \mathbb{E}_{i,a}[\mathcal{Y}_i]$, alors $Z_i = z_i(X_i) = \tilde{z}_i(X_i, a_i^*)$.

Une première partie de ce chapitre est consacrée à l'analyse théorique de cette discrétisation. Nous procédons en deux étapes. Nous commençons d'abord par discrétiser la contrainte sans modifier le processus X, puis nous remplaçons X par son schéma d'Euler.

Ainsi, nous commençons par définir la séquence intermédiaire d'EDSRs $(Y^\pi, \mathcal{Y}^\pi, \mathcal{Z}^\pi, \mathcal{U}^\pi)$ (équations (5.2.13)-(5.2.14)-(5.2.15)), qui correspond à l'EDSR contrainte initiale (1.3.10) avec discrétisation de la contrainte sur U. Nous définissons aussi les fonctions valeurs v^π et ϑ^π associées à Y^π et $\mathcal{Y}^\pi$ (équations (5.2.16)-(5.2.17)-(5.2.18)).

Cette construction d'une séquence d'EDSRs contraintes n'étant pas standard, nous commençons par démontrer que tous ces éléments sont bien définis (Propositions 5.2.1-5.2.2-5.2.3). Ensuite, nous démontrons que la fonction v^π associée au schéma discret converge bel et bien vers la fonction v de l'EDSR contrainte initiale lorsque $|\pi| \to 0$ (Corollaire 5.2.2). Pour finir, nous nous intéressons à la vitesse de convergence de v^π vers v. Pour cette étape, nous avons besoin des hypothèses supplémentaires suivantes :

- b et σ uniformément bornées (H1'),
- f ne dépend pas de z, $f(.,.,0)$ et g uniformément bornées, et $y \mapsto f(x,a,y)$ convexe (H2').

Avec ces hypothèses supplémentaires, en adaptant la méthode de perturbation de coefficients de Krylov ([78], cf. aussi [49]) aux diffusions à sauts, nous parvenons à obtenir la vitesse de convergence suivante (Théorème 5.2.1) :

$$0 \leq v(t,x) - v^\pi(t,x) \leq C|\pi|^{\frac{1}{10}}$$

pour tout $(t,x) \in [0,T] \times \mathbb{R}^d$ et toute grille π. De plus, dans le cas où $f = f(x,a)$ ne dépend pas de y, la représentation par problème de contrôle stochastique combinée aux résultats de [77] permet d'améliorer cette vitesse en $C|\pi|^{\frac{1}{6}}$. Autrement dit, la solution de l'EDSR contrainte discrétisée converge vers la solution de l'EDSR contrainte continue à la vitesse $\mathcal{O}\left(|\pi|^{\frac{1}{10}}\right)$ (ou $\mathcal{O}\left(|\pi|^{\frac{1}{6}}\right)$ quand f ne dépend pas de y), cf. Corollaire 5.2.3.

Ensuite, la seconde partie de l'étape de discrétisation consiste à remplacer la diffusion X par son schéma d'Euler, aboutissant au schéma discret (1.3.12). En adaptant les propriétés du schémas d'Euler (Lemme 5.2.1) et des résultats de régularité sur les EDSRs (Proposition 5.2.3), nous parvenons à obtenir pour cette étape une estimation d'erreur en $\mathcal{O}\left(|\pi|^{\frac{1}{2}}\right)$ (Théorème 5.2.2). Par conséquent, en combinant les résultats des deux étapes, nous obtenons, pour la discrétisation temporelle complète, une vitesse de convergence en $\mathcal{O}\left(|\pi|^{\frac{1}{10}}\right)$ ($\mathcal{O}\left(|\pi|^{\frac{1}{6}}\right)$ quand f ne dépend pas

de y), ou, plus précisément, une erreur négative en $\mathcal{O}\left(|\pi|^{\frac{1}{2}}\right)$ et une erreur positive en $\mathcal{O}\left(|\pi|^{\frac{1}{10}}\right)$ ($\mathcal{O}\left(|\pi|^{\frac{1}{6}}\right)$ quand f ne dépend pas de y), cf. Corollaire 5.2.4.

Ce résultat vient améliorer les résultats connus jusqu'alors (cf. Remarque 5.2.3). Par exemple, dans le cas où f ne dépend pas de y, nos bornes $|\pi|^{\frac{1}{2}}$ et $|\pi|^{\frac{1}{6}}$ améliorent respectivement les bornes $|\pi|^{\frac{1}{4}}$ et $|\pi|^{\frac{1}{10}}$ obtenues par EDSR du second ordre (cf. [49]). Ensuite, et c'est un point important, nos résultats sont obtenus sans aucune hypothèse d'ellipticité uniforme sur σ. Enfin, et cela est très utile en pratique, le schéma permet d'obtenir une approximation du contrôle optimal (cf. Section 5.2.4.2). Tous ces éléments nous indiquent que les EDSRs contraintes semblent un bon outil pour l'étude des équations de HJB totalement non-linéaires.

La seconde partie de ce chapitre est consacrée à l'approximation des espérances conditionnelles du schéma discret (1.3.12). En effet, tout comme au chapitre précédent, nous avons affaire à un schéma numérique rétrograde faisant apparaître des espérances conditionnelles qui ne peuvent généralement pas être calculées de manière explicite. Une nouvelle fois, nous avons choisi des les approximer par méthode de simulation et régression. Cela a ici plusieurs avantages. Tout d'abord, la simulation par Monte Carlo du processus progressif (X, I) permet de s'affranchir dans une certaine mesure des contraintes de dimension. Ensuite, ce choix permet de réduire le calcul de supremum du schéma discrétisé (1.3.12) à une simple maximisation d'une fonction paramétrique donnée. Enfin, cela permet d'obtenir un estimateur paramétrique du contrôle optimal.

Au chapitre précédent, dans le cadre des problèmes de commutation optimale, nous n'étions parvenus à analyser complètement l'erreur commise par cette approximation que dans le cas spécifique d'une base de fonctions constantes par morceaux. Dans ce chapitre, nous avons voulu analyser cette erreur sans faire d'hypothèse spécifique sur la base de fonctions de régression (à l'image des travaux de [82] et [58] dans le cadre des EDSRs standards). Malheureusement, dans le cadre des EDSRs contraintes, l'étape additionnelle de supremum complexifie significativement l'analyse de l'erreur, et nous n'avons pu obtenir que des résultats partiels pour cette étape.

L'étape préalable de localisation spatiale se passe sans encombres (Lemme 5.3.1 et Proposition 5.3.1). Considérons maintenant l'étape de régression proprement dite. Soit $\left\{p_b : \mathbb{R}^d \times \mathbb{R}^q \to \mathbb{R}\right\}$, $b = 1, \dots, B$, une base de fonctions de régressions. Étant donnée une variable aléatoire $\mathcal{F}_T$-mesurable U, l'idée est de remplacer l'espérance conditionnelle $\mathbb{E}_i[U]$ par la projection $\mathcal{P}_i(U)$ définie comme suit :

$$\begin{aligned} \hat{\lambda}_i(U) &:= \arg \inf_{\lambda \in \mathbb{R}^B} \mathbb{E}\left[\left(\lambda.p(X_i, I_i) - U\right)^2\right] \\ \mathcal{P}_i(U) &:= \hat{\lambda}_i(U).p(X_i, I_i) \ . \end{aligned} \tag{1.3.13}$$

Il s'agit ici d'une régression théorique (l'étape finale de régression empirique aboutissant à un schéma implémentable consiste à remplacer lors de la minimisation de moindres carrés ci-dessus l'espérance $\mathbb{E}$ par une espérance empirique $\frac{1}{M}\sum_{m=1}^{M}$). En effectuant cette substitution dans le schéma discret (1.3.12), on obtient le nouveau schéma suivant :

$$\begin{aligned} \hat{Y}_N &= g(X_N) \\ \Delta_i \hat{\mathcal{Z}}_i &= \mathcal{P}_i\left(\hat{Y}_{i+1} \Delta W_i^\top\right) \\ \hat{\mathcal{Y}}_i &= \mathcal{P}_i\left(\hat{Y}_{i+1} + f_i\left(X_i, I_i, \hat{Y}_{i+1}, \hat{\mathcal{Z}}_i\right)\Delta_i\right) =: \hat{\lambda}_i^Y.p_i(X_i, I_i) \\ \hat{Y}_i &= \operatorname*{ess\,sup}_{a \in \mathcal{A}_i} \mathbb{E}_{i,a}\left[\hat{\mathcal{Y}}_i\right] = \operatorname*{ess\,sup}_{a \in \mathcal{A}_i} \hat{\lambda}_i^Y.p_i(X_i, a) \ , \end{aligned} \tag{1.3.14}$$

où $\mathcal{A}_i$ est l'ensemble des variables aléatoires $\sigma(X_i)$-mesurables à valeurs dans A. (Pour simplifier, on omet ici l'étape de localisation spatiale). Même à ce stade théorique, si l'on veut étudier

l'erreur entre le schéma discret (1.3.12) et son approximation (1.3.14), la difficulté suivante apparaît : à cause de l'étape finale de maximisation, la variable $\hat{Y}_i$, contrairement à la variable intermédiaire $\hat{\mathcal{Y}}_i$, ne constitue pas en soi la projection d'une variable aléatoire donnée. Par conséquent, il n'est pas possible d'utiliser les outils usuels (cf. Lemme 5.3.3) qui permettraient de se défaire de l'opérateur $\mathcal{P}_i$ lors de l'analyse de la différence entre (1.3.12) et (1.3.14).

Il est intéressant de noter que si l'on modifie la régression comme suit :

$$\begin{aligned} \hat{\lambda}_{i,a}(U) &:= \arg \inf_{\lambda \in \mathbb{R}^B} \mathbb{E}\left[\left(\lambda.p(X_i, a) - U_a\right)^2\right] \\ \mathcal{P}_{i,a}(U) &:= \hat{\lambda}_i(U).p(X_i, a) \ , \end{aligned} \tag{1.3.15}$$

où $U_a := U|a$, $a \in \mathcal{A}_i$, alors il devient possible d'analyser l'erreur de régression (Proposition 5.3.2). La différence vient du fait que la maximisation et les régressions sont cette fois effectuées simultanément, et non itérativement comme pour le schéma (1.3.14). Cependant, contrairement à la version empirique du schéma (1.3.14), un schéma empirique basé sur la régression (1.3.15) serait lourd et malaisé à mettre en œuvre. C'est pourquoi nous choisissons d'implémenter malgré tout un schéma basé sur la régression (1.3.13), tout en devant abandonner l'obtention d'une analyse d'erreur exhaustive à des travaux ultérieurs.

Finalement, la dernière partie de ce chapitre est consacrée à des applications numériques de la version empirique du schéma (1.3.14).

Tout d'abord, nous commençons par un exemple de problème de contrôle stochastique de type linéaire-quadratique (équations (5.4.1) et (5.4.2)). La solution explicite de ce type de problèmes étant connue (cf. [107]), cet exemple permet de fournir un première vérification de l'exactitude du schéma. Les résultats numériques indiquent, comme l'illustre la Figure 5.4.1b, que le schéma parvient à approximer le contrôle optimal avec une grande précision.

Après cette vérification préliminaire, nous avons consacré le reste de nos applications numériques au problème de sur-réplication d'options dans des modèles à volatilités et/ou corrélations incertaines. Nous commençons avec un modèle à deux actifs S_1 et S_2 dont les dynamiques sont données par :

$$dS_i(t) = \sigma_i S_i(t)\, dW_i(t) \ , \ i = 1, 2 \tag{1.3.16}$$

$$\langle dW_1(t), dW_2(t) \rangle = \rho dt \, , \tag{1.3.17}$$

où le processus ρ est incertain ; on suppose seulement qu'il ne peut prendre ses valeurs que dans un intervalle donné $[\rho_{\min}, \rho_{\max}] \subseteq [-1, 1]$. Les prix de sur-réplication de l'option de valeur finale

$$\Phi = \left(S_1(T) - S_2(T) - K_1\right)^+ - \left(S_1(T) - S_2(T) - K_2\right)^+ \ , \tag{1.3.18}$$

(non triviale car ni convexe ni concave en $S_1(T) - S_2(T)$) obtenus par notre algorithme sont illustrés sur la Figure 5.4.3. On retrouve en particulier le fait que le modèle à corrélation incertaine est plus à même d'estimer le risque associé au prix de l'option qu'une simple variation d'une corrélation constante.

Après cet exemple (1.3.18) qui permet d'illustrer concrètement l'intérêt des modèles avec paramètres incertains, nous avons testé notre algorithme sur plusieurs autres options, dans l'optique cette fois d'évaluer la qualité de l'algorithme même. Ainsi, nous avons comparé les résultats de notre algorithme sur les exemples numériques proposés dans [62]. Cet article présente en effet l'intérêt de fournir des estimations de prix de sur-réplication d'options par méthode d'EDP, ainsi que par un algorithme Monte Carlo alternatif basé sur les EDSRs de second-ordre (cf. [49]).

Sur tous ces exemples, et comme au chapitre précédent, nous avons calculé deux prix complémentaires formant un intervalle de confiance empirique asymptotique du prix (équations (5.4.8)

et (5.4.9)), l'étendue de cet intervalle permettant d'évaluer la qualité de la base de régression choisie.

Les Figures 5.4.4 à 5.4.9 illustrent nos résultats. Nous y faisons varier le nombre de trajectoires de Monte Carlo, ainsi que la taille du pas de temps de discrétisation. Les résultats obtenus sont ceux attendus, avec les deux prix encadrant le vrai prix (tel qu'estimé par la méthode d'EDP), et ceux-ci étant d'autant plus précis que le pas de discrétisation est petit et que le nombre de trajectoires de Monte Carlo est grand.

Si notre algorithme et celui proposé dans [62] présentent des résultats numériques comparables, le nôtre possède l'avantage crucial de comporter un processus progressif (X, I) aisément simulable, ce qui, nous l'espérons, devrait permettre de l'étendre facilement à des problèmes plus généraux tels que les équations de HJB-Isaacs et les jeux à champs moyens.

2 Introduction (in English)

The present work is devoted for the most part to the numerical solution of problems of pricing and hedging of contingent claims, and more generally of stochastic control problems, related to energy markets, in particular to electricity markets. Having high dimensional problems in prospect, like investment problems in power plants, we opted for the construction and analysis of probabilistic numerical schemes.

In the first part, we started by looking for a suitable model for the price of electricity. In the light of empirical considerations on the price formation mechanism, we proposed and studied a new structural model for the price of electricity, which has the crucial advantage of being able to reproduce the fine dependence structure between the price of electricity and the prices of other enegies (oil, coal, gas,...). In this first part, we made use of this model to price and hedge electricity derivatives. In particular, we tried to improve the hedging strategies for these assets by allowing, as hedging assets, the use of other energy derivatives, taking advantage of their fine dependence with electricity. The effect of this addition is displayed on numerical tests, the completion of which required the construction of efficient algorithms to compute the hedging strategies.

In the second part, we used one again this new structural model, but on another application, namely the problem of valuing power plants, and, by extension, the computation of economically optimal investment strategies in new power plants in the future. From a mathematical point of view, this kind of problem can be expressed as an optimal switching problem, which is a stochastic control problem for which every policy modification generates a corresponding cost. This problem is high-dimensional, therefore we proposed a probabilistic scheme to solve it, which combines dynamic programming with least-squares regressions. This scheme is inspired by those existing for the simpler problem of American option valuation ([31]). We performed a mathematical analysis of the convergence of this scheme, and retrieved its convergence rate for a specific choice of regression basis (piecewise constant functions). Numerical applications illustrate the feasibility of this algorithm, and demonstrate more generally that the techniques and improvements implemented here make it conceivable to try to tackle difficult, multi-dimensional stochastic control problems by a direct probabilistic approach.

Finally, we extended in the third part our previous probabilistic scheme to more general problems, namely stochastic control problems for which the drift and the volatility of the underlying state variable can depend upon the strategy. So far, probabilistic schemes could not handle this kind of problems, as their initiation necessitates the simulation of trajectories of the underlying state variable, which is not possible here as the optimal control is unknown at first. To overcome this difficulty, we made use of a probabilistic representation of the problem using Backward Stochastic Differential Equations (BSDEs) with constrained jumps, which is a generalization, introduced in [72], of the class of BSDEs. This generalization encompasses the kind of stochastic control problems we are interested in here, and, more generally, encompasses fully nonlinear Hamilton-Jacobi-Bellman (HJB) equations. Therefore, we proposed an implementable probabilistic numerical scheme for these constrained BSDEs, based on an initial randomization of the control, combined with a subsequent removal of this additional randomness by the adjunction of the computation of a supremum within the numerical backward scheme. We performed a mathematical analysis of this scheme, including the time discretization of these constrained

BSDEs. Finally, we illustrated numerically the possibilities of the scheme on the problem of superreplication of options under uncertain volatility ([5]).

2.1 A structural risk neutral model for pricing and hedging electricity derivatives

This first part of the thesis deals with an electricity structural model suitable for pricing and partially hedging power derivatives. It led to the publication of the article [1].

The classical approach in financial mathematics in order to study a given derivative product is to first propose a model for the dynamics of the underlying asset, from which the pricing and the hedging strategy can be deduced. This is the so-called *reduced-form* approach.

In this way, several authors tried to propose models that aim at taking into account the idiosyncrasies of the price of electricity, while remaining practical for pricing and hedging purposes (cf. [15] for example). One of these distinguishing features is that electricity is non storable [1], which implies that electricity production has to be adjusted in real time to the exact level of electricity demand.

There exist several types of power plants, based either on renewable energies (hydro, wind, ...) or fossil energies (coal, gas, ...). Consequently, it is clear that the composition of the electric fleet as well as the price formation mechanism will significantly impact the price of electricity. Against this backdrop, on the opposite side from reduced-form models, the *stacking* models try to take into account this mechanism. These models require a detailed modeling of the electricity demand, of the prices of fuels, and of all the available power plants. A global optimization then leads to the price of electricity, which is the smallest price that allows to satisfy the level of demand, making an optimal use of the available production assets (cf. [61] for example).

This approach allows for a great modeling precision (it is possible to take into account many details, like the dynamic production constraints of thermal plants for example, cf. [80]). However, its main drawback is that it is very heavy and clumsy to implement, and is ill-suited to the study of derivative products.

Halfway between these two extremes, the class of structural models tries to take account in a simplified manner of the price formation mechanism, all the while making use of the classical mathematical finance tools from the reduced-form approach. In particular, this intermediary approach is very well suited to the pricing of multi-asset options that include the electricity as well as other energies like gas for example, as it allows to take into account the fine dependence structure between these variables. A survey on this class of structural models is available in [30].

The starting point of this chapter is the structural model by marginal cost developed [3], that we recall below. Consider a given power market, which includes n types of power plants. For each $i = 1, \dots, n$ let:

- S_t^i be the price of the fuel used by this type of plant (if it is based on a renewable energy, then $S_t^i \equiv 0$).
- h_i be its heat rate, that we assume to be constant (it is such that $h_i S_t^i$ is expressed in €/MWh).
- C_t^i be the power generation capacity from this type of plants (in MW).

[1] Apart from hydroelectric production, as dams provide an indirect power storage

Moreover, let D_t be the electricity demand (in MW). Without loss of generality, suppose that the technologies $i = 1, \dots, n$ are sorted in increasing production cost order: $h_1 S_t^1 \leq \dots \leq h_n S_t^n$. Then, in [3], the power price P_t is directly modeled by the *marginal cost* of production, defined by:

$$CM_t = \sum_{i=1}^{n} h_i S_t^i \mathbf{1} \left\{ \sum_{k=1}^{i-1} C_t^k \leq D_t \leq \sum_{k=1}^{i} C_t^k \right\} . \tag{2.1.1}$$

Equation (2.1.1) means that, if the demand D_t exceeds the total production capacity of the $i-1$ cheapest technologies ($\sum_{k=1}^{i-1} C_t^k$) but does not exceed that of the i cheapest technologies ($\sum_{k=1}^{i} C_t^k$), then the marginal cost of producing electricity is fixed by the i-th technology, and is given by $h_i S_t^i$.

The drawback of the model $P_t = CM_t$ proposed in [3] is that, if the marginal cost is indeed an important indication of the price level of electricity, these two quantities can be different in practice, because of all the secondary phenomenons neglected by this simple model [2]. This difference can sometimes be substantial, like for instance when an unexpected price spike occurs. Indeed, the simple model $P_t = CM_t$ cannot generate price spikes, as by definition $CM_t \leq h_n S_t^n$ (and an adequate modeling of the price S_t^n of the most expensive fuel does not require spikes).

The first part of this chapter is devoted to the improvement of the model (2.1.1) in order to be at least able to produce the kind of price spikes observable on power markets. After a comparison between the realized marginal cost and the realized power price (Figure 3.2.1a), and a careful preservation of the few spikes of the dataset by opting for a dependence model between the quantiles of the variables at hand (Proposition 3.2.2), the following model appears to be empirically satisfactory (Figure 3.2.1b) :

$$P_t = g\left(\sum_{k=1}^{n} C_t^k - D_t \right) CM_t \, , \tag{2.1.2}$$

where the function g, that we called *scarcity function*, is defined by :

$$g(x) = \min\left(M, \frac{\gamma}{x^\nu} \right) \mathbf{1}\{x > 0\} + M \mathbf{1}\{x \leq 0\} \, . \tag{2.1.3}$$

Therefore, we correct the marginal cost by a multiplicative term which is all the more so big ($\mathrm{in}\gamma/x^\nu$) as the available residual capacity of the system ($\sum_{k=1}^{n} C_t^k - D_t$) is low. In particular, this models generates price spikes precisely when the system is tensed (low residual capacity, close to a blackout). This correction term is capped to a maximum value M. The improvement from the model (2.1.1) to the model (2.1.2) is illustrated on Figure 3.2.2.

The second part of this chapter is devoted the use of the new model (2.1.2) for pricing and hedging power derivatives. As the price of electricity in this model depends on hedgeable risks (the fuel prices S_t^i) as well as non-hedgeable risks (the demand D_t and the production capacities C_t^i), the market is therefore incomplete. In other words, the classical dynamic hedging procedure cannot suppress all the risk stemming from a given derivative. Nevertheless, without suppressing all the risk, it is still possible to tame it, by minimizing a predefined risk criterion.

In the context of incomplete markets, there exists a whole range of possible risk criterions in order to build partial dynamic hedging strategies. Here, we choose the *local risk minimization* criterion, introduced in [52], for its simplicity and the fact that it enables a natural split between the hedgeable part and the unhedgeable part of any derivative product based on an underlying

[2] Like for instance the non homogeneity between plants of the same type, the dynamic constraints of production, the mandatory production reserves for the stability of the network, the impact of imports and exports, the strategic behavior of the producers, ...

modeled by (2.1.2). Remark that our present work was subsequently extended to the utility indifference criterion in [13].

At this stage, we need to assume specific models for the dynamics of the factors C_t, D_t and S_t. We choose very simple models:

- The spreads $Y_t^i := h_i S_t^i - h_{i-1} S_t^{i-1}$ are modeled by geometric Brownian motions.
- The demand D_t and the capacities C_t^i are modeled by general diffusions (equations (3.3.2) and (3.3.3)).

The key point is that the Brownian motions used for the fuel prices are assumed independent from the Brownian motions used for demand and capacities (Assumption 1). This assumption is reasonable in practice.

We identified the whole range of equivalent martingale measures in our model (Proposition 3.3.1), including the *minimal martingale measure* $\widehat{\mathbb{Q}}$ related to the local risk minimization criterion. Then, we tried to price a few significant power derivatives.

First, we studied a forward contract on electricity $F_t^e(T)$ with instantaneous delivery period T. It is of course a fictitious asset, as real forward contracts on electricity involve a real delivery period $[T_1, T_2]$ with $T_1 < T_2$. Nevertheless, it is always possible to make use of the intermediary block $F_t^e(T)$ to reconstruct the price of a true forward contract $F_t^e(T_1, T_2)$.

In our model, the price at time t of a forward contract on electricity with delivery at time T is given by

$$F_t^e(T) = \sum_{i=1}^{n} h_i G_i^T(t, C_t, D_t) F_t^i(T) , \tag{2.1.4}$$

where, for $i = 1, \ldots, n$, $F_t^i(T)$ corresponds to the price at time t of the forward contract on the fuel i with delivery at time T, and where the quantity $G_i^T(t, C_t, D_t)$, called *Conditional Expectation of Scarcity (CES)* function, is defined by

$$G_i^T(t, C_t, D_t) := \mathbb{E}\left[g\left(\sum_{k=1}^{n} C_T^k - D_T \right) \mathbf{1}\left\{ \sum_{k=1}^{i-1} C_T^k \leq D_T \leq \sum_{k=1}^{i} C_T^k \right\} | \mathcal{F}_t^{D,C} \right] . \tag{2.1.5}$$

Thus, equation (2.1.4) tells us that the forward price of electricity can be expressed as a linear combination of the forward fuel prices. The weights in this linear combination are given by the heat rate h_i multiplied by the CES function (2.1.5). Without the scarcity penalization ($g \equiv 1$), the CES function of the i-th fuel would simply correspond to the probability for the i-th fuel to be marginal at time T (as in [3]). Here, the presence of the scarcity function means that the probability of marginality is penalized in the manner described by equation (2.1.5).

Equation (2.1.4) provides a clear separation between the impact of the different sources of risk: the hedgeable risks S_t^i impact the forward fuel prices $F_t^i(T)$, while the non-hedgeable risks C_t^i and D_t impact the stochastic weights $G_i^T(t, C_t, D_t)$.

Aside from the value of $F_t^e(T)$, it is possible, from equation (2.1.4), to deduce its dynamics, which makes use of the partial derivatives of the CES function (cf. equations (3.3.13) and (3.3.14)).

Remark that, using arguments similar to those leading to equation (2.1.4), it is possible to express the electricity risk premium as a linear combination of the risk premiums on fuels (Proposition 3.3.2).

After this analysis of power futures, we then study other derivatives. In the case with $n = 2$ fuels, we provide an explicit pricing formula for a call option on the spread $P_T - h_i S_T^i$ (Proposition 3.3.4) and for a call option on a power futures (Proposition 3.3.6).

Finally, we study the hedging strategy of these derivatives. Consider a general option whose final payoff can can depend on demand, capacities, as well as power and fuel prices. Consider hedging porfolios that can contain forward contracts on fuel and electricity. Then, and this is the main result of this chapter, we are able to explicit the optimal weights of the hedging porfolio (Proposition 3.3.7). These weights depend on the shape of the payoff, as well as on the quantities involved in the dynamics of the power futures (namely the CES function and its partial derivatives). Moreover, we are also able to assess the residual risk of the hedge, which stems from the non-hedgeable sources of risk (demand and capacities).

The rest of this chapter is made of numerical applications of these results.

So far, we modeled the demand D_t and the capacities C_t^i by general diffusions. In order to perform numerical tests, we now need to assume a more explicit model for these variables. We choose to model them by a combination of a deterministic part (which accounts for the various seasonalities) and stochastic part, modeled by a simple Ornstein-Uhlenbeck process. We calibrated these models on the French dataset (cf. Pictures 3.4.1 and 3.4.2). Though very simple, these models fit the data quite well.

Using these models, we then try to compute in practice the price of a forward contract on electricity. As shown by equation (2.1.4), the crucial step is to be able to compute efficiently the CES function (2.1.5) and its partial derivatives.

In the literature on power price modeling, for both reduced-form models as well as structural models, being able to obtain closed-form formulas for the forward price of electricity is a much sought-after ability when it comes to building and proposing a new model. For example, in the context of structural model, people may prefer to model the stack curve (and its variants) by a simple exponential function (cf. [30]). Though such a choice sacrifices some features, including the possibility of extreme price spikes, it ensures that computations remain very simple.

In this chapter, motivated by empirical observations, we chose to use a power function to model the price spikes (equation (2.1.3)). Unfortunately, this choice does not allow for explicit expressions for the weights (2.1.5). However, being aware of their importance, we tried to make their numerical computation as efficient as possible.

First of all, we established that the computation of this CES function boils down to the computation of the following expressions:

$$\widetilde{\mathcal{G}}(x,y;\nu) := \int_x^\infty \frac{1}{(y+z)^\nu} e^{-z^2} dz \tag{2.1.6}$$

$$\widetilde{\mathcal{H}}(m_1,m_2,\sigma_1,\sigma_2;\nu) := \int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^\infty \widetilde{\mathcal{G}}\left(\frac{\widetilde{x}-m_1-m_2-\sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}, \frac{m_1+m_2+\sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}};\nu\right) e^{-u^2} du \tag{2.1.7}$$

(cf. Propositions 3.4.3, 3.4.4 and 3.4.9). Consider first the function $\widetilde{\mathcal{G}}(x,y;\nu)$. As in the case $\nu = 1$ this integral corresponds to the special function called incomplete Goodwin-Staton integral (cf. [45]), we called the new function (2.1.6) *extended incomplete Goodwin-Staton integral*. This function can be interpreted probabilistically, up to some coefficient adjustments, as the density of the sum of two independent Gaussian and Pareto random variables (cf. Proposition B.1).

We established the following series expansion:

$$\widetilde{\mathcal{G}}(x,y;\nu) = \frac{1}{2} e^{-y^2} \sum_{n=0}^\infty \Gamma\left(\frac{1-\nu}{2} + \frac{n}{2}, (x+y)^2\right) \frac{(2y)^n}{n!},$$

where $\Gamma(\alpha,x)$ corresponds to the incomplete Gamma function. Combined with recurrence relations and asymptotic expansions, this series expansion allows to compute $\widetilde{\mathcal{G}}(x,y;\nu)$ very efficiently. Similarly, we established a series expansion for $\widetilde{\mathcal{H}}(m_1,m_2,\sigma_1,\sigma_2;\nu)$, as well as adequate

recurrence relations and asymptotic expansions, making its numerical estimation very efficient. Finally, we explicited how to compute the partial derivatives of the CES function (2.1.5) using the two functions $\widetilde{\mathcal{G}}$ and $\widetilde{\mathcal{H}}$.

Finally, using these numerical tools, we were able to test numerically the pricing and partial hedging of power derivatives on two examples, using the previously established general formulas.

First, we tested power futures with maturity $T = 3$months, with a test of a partial hedging strategy based on fuel futures. We observed an interesting time behavior for this hedge, which is displayed on Figure 3.4.4. Two very distinct temporal phases emerge:

- Far from maturity, for $t \in [0, T - \Delta]$ with $\Delta \sim 2$weeks, the partial hedge is almost perfect. Indeed, in this phase, the prices of fuels are the dominant sources of risk. Now, these risks correspond precisely to the hedgeable factors, that the hedging strategy is able to eliminate. In this phase, the power futures behaves almost like a basket of fuel futures.

- Close to maturity, for $t \in [T - \Delta, T]$, the partial hedge is completely inefficient. Indeed, in this phase, the levels of demand and capacities become the dominant sources of risk. As they correspond to the non-hedgeable factors, the mitigation provided by the hedging strategy becomes negligible.

This important empirical observation is very general. It is neither due to the specific shape of the payoff, nor to the hedging criterion, but to the power price model (2.1.2) and the fact that the hedgeable factors are of a martingale type, while the non-hedgeable factors are of a mean-reversion type (meaning that the laws of C_T and D_T barely depend on the current values C_t and D_t when $T - t$, and that the value of time threshold $T - \Delta$ between the two phases depends essentially on the mean-reverting parameters. More generally, this observation shows the limits of the ability of fuel derivatives to mitigate the risks from power derivatives.

Then, we tested the pricing of a spread option, more precisely a call option on the spread $P_T - h_i S^i_T$. This time, it is necessary to resort to numerical integration methods. Our results are displayed on Figure 3.4.5. The practical interest of this example is to show that a structural model can easily accomodate seasonality effects, as well as structural changes in the market, like the expected setting up of new power plants.

2.2 A probabilistic numerical method for optimal multiple switching problem in high dimension

The second part of this thesis deals with the numerical solution of stochastic control problems in high dimension, more precisely of optimal switching problems. It led to the article [2], currently under review for publication.

First, here is a definition of this specific class of problems. We consider the following elements:

- $X = (X_t)_{t \geq 0}$ is a stochastic process taking values in $\mathbb{R}^d$, starting from $x_0 \in \mathbb{R}^d$ at time $t = 0$.

- $I^\alpha = (I^\alpha_t)_{t \geq 0}$ is a piecewise constant process taking values in $\mathbb{R}^{d'}$, starting from $i_0 \in \mathbb{R}^{d'}$ at time $t = 0$. More precisely, I^α is supposed to take its values into a finite subset $\mathbb{I}_q = \{i_1, \ldots, i_q\}$ of $\mathbb{R}^{d'}$. This process I^α is controlled over time by a strategy α.

- α is an impulse control defined by a sequence $(\tau_n, \iota_n)_{n \in \mathbb{N}}$ of increasing stopping times $\tau_n \geq 0$ and $\mathbb{I}_q$-valued $\mathcal{F}_{\tau_n}-$measurable random variables. The controlled processed I^α can be deduced

from this sequence as follows:

$$I_t^\alpha = \iota_n \text{ when } t \in [\tau_n, \tau_{n+1}[\ .$$

- Among the possible strategies α, we only consider those that belong to an admissible set $\mathcal{A}$. Broadly speaking, it means that we only consider the strategies such that $\tau_n \to +\infty$ a.s. when $n \to \infty$ (i.e. accumulation points are excluded).
- $f : \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}^{d'} \to \mathbb{R}$ and $k : \mathbb{R} \times \mathbb{R}^{d'} \times \mathbb{R}^{d'} \to \mathbb{R}$ two measurable functions.

Then, the stochastic problem that we consider is the following:

$$v(0, x_0, i_0) = \sup_{\alpha \in \mathcal{A}} \mathbb{E}\left[\int_0^\infty f(t, X_t, I_t^\alpha)\, dt - \sum_{\tau_n \geq 0} k(\tau_n, \iota_{n-1}, \iota_n) \right] . \tag{2.2.1}$$

The goal is to maximize the gains brought by the function f over time. These gains depend uncontrolled state variable X, as well as upon the controlled variable I^α. The goal is thus to adjust adequately the strategy over time. However, every modification of the strategy generates a cost given by the function k. This cost depends on the values of the control immediately before and after the move.

The problem (2.2.1) is called optimal switching because th controlled process I^α takes its values within a finite discrete set.

To be more precise, some usual regularity assumptions are required in order for the problem (2.2.1) to be well defined (cf. Section 4.2.2).

First of all, here is why we were interested in the study of this specific problem. In the first chapter, we built a structural model for the price of electricity, which is able to model properly the time dependence between the prices of electricity and other energies. In particular, this model allows to price spread options properly (between electricity and another energy), and thus, by extension, to value, in a real option framework, a given power plant. Using this tool, we wanted to know if it was possible to detect the best possible investments in power plants over time. (Which type of plant to build? How much? And when?) We will see later that this kind of investment problem can indeed be expressed as an optimal switching problem of the form (2.2.1) in high dimension ($d + d' \gg 3$). This is why we tried to build a numerical method able to solve in practice the problem (2.2.1) in high dimension.

Now, let us detail how to proceed to solve the problem (2.2.1). If the strategy is allowed to change only on a fixed time grid $\Pi = \{t_0 = 0 < t_1 < \ldots < t_N = T\}$, then, using the dynamic programming principle, the discretized value function v^Π satisfies the following backward induction:

$$v^\Pi(t_n, x, i) = \max_{j \in \mathbb{I}_q} \{E(t_n, x, j) - k(t_n, i, j)\} , \tag{2.2.2}$$

where

$$E(T, x, i) := \mathbb{E}\left[\int_T^\infty f(s, X_s, i)\, dt \,|X_T = x \right]$$

$$E(t_n, x, i) := \mathbb{E}\left[\int_{t_n}^{t_{n+1}} f(s, X_s, i)\, dt \,|X_{t_n} = x \right] + \mathbb{E}\left[v^\Pi(t_{n+1}, X_{t_{n+1}}, i) \,|X_{t_n} = x \right], n = N-1, \ldots, 0.$$

In practice, aside from time discretizations and the computation of final values, the major difficulty to implement the scheme (2.2.2) consists in computing the conditional expectations $\mathbb{E}\left[v^\Pi(t_{n+1}, X_{t_{n+1}}, i) \,|X_{t_n} = x \right]$. Indeed, this expression cannot in general be computed explicitly, therefore one has to resort to approximations.

In the literature on American options (which is a simpler problem than (2.2.1)), the same problem arise, and several methodologies have been proposed to cope with it. Among them are the linear regression approach ([83, 102]), the non-parametric regression approach ([101, 74]), the quatization approach ([8]) and the Malliavin calculus approach ([6]).

In the context of stochastic control problems without switching cost ($k \equiv 0$), the Malliavin approach has been used in [88], and the regression approach in [12]. Finally, for the optimal switching problem we are interested in here, the quantization has been used in [55], and the linear regression approach in [32]. However, the methodology used in [32] does not make a direct use of the dynamic programming principle as in equation (2.2.2). Instead, they resort to the representation of optimal switching problems as successive layers of optimal stopping problems.

In this chapter, in light of the conclusions of the numerical comparison tests in [28] (in the American options framework), we choose to use the local regression approach to solve our optimal switching problem. This approach appears indeed to be currently to most capable of handling high dimensional problems.

The first part of this chapter is devoted to the error analysis between the initial quantity (2.2.1) that we want to estimate and the result of the approximating numerical scheme. Between them, several approximations are required, listed below:

- *[Finite time horizon]* The first is concerned with the management of the time horizon. It is indeed necessary to limit the time horizon of the strategies to a large but fixed date $T > 0$.
- *[Time discretization]* Then, it is necessary to discretize the processes X and I^α. To keep things simple, one can choose a time grid $\Pi = \{t_0 = 0 < t_1 < \ldots < t_N = T\}$ with a fixed time step $h > 0$.
- *[Space localization]* We will also need to project the process X (which can a priori take its values anywhere in $\mathbb{R}^d$) in a large but bounded subset $\mathcal{D}_\varepsilon$. The reason for this approximation will be given later.
- *[Approximation of conditional expectations]* Finally, the last step consists in approximating the conditional expectations arising when using the dynamic programming principle. We approximate them by empirical least-squares regression on a bundle of Monte Carlo trajectories. We will come back later on the choice of regression basis.

For each approximation step, we computed a bound on the corresponding error.

The transition to a finite time horizon is analyzed in Proposition 4.3.1. The key element is the inclusion of an exponential discounting in the definition of f and k (cf. the assumptions listed in Section 4.2.2). The error behaves like $e^{-\rho T}$.

The time discretization is analyzed in Propositions 4.3.2 and 4.3.3. Most of the analysis comes from [55]. The error behaves like $\sqrt{h}$ (or $\sqrt{h}\sqrt{\log\left(\frac{2T}{h}\right)}$ when the cost fuction k depends also from the state variable X).

For the state localization, the bounded domain $\mathcal{D}_\varepsilon$ is directly chosen such that the corresponding error behaves like ε (Proposition 4.3.4).

Finally, we come to the most complex step, namely the analysis of the regression error. So far, the dynamic programming principle applied to the approximated value function can be simply

expressed as follows:

$$\begin{aligned}\bar{v}_{\Pi}(T,x,i) &= g(T,x,i)\\ \bar{v}_{\Pi}(t_n,x,i) &= \max_{j\in\mathbb{I}_q}\left\{hf(t_n,x,j) - k(t_n,i,j) + \mathbb{E}\left[\bar{v}_{\Pi}\left(t_{n+1},\bar{X}_{t_{n+1}},j\right)\middle|\bar{X}_{t_n}=x\right]\right\}, n = N-1,\dots,0.\end{aligned} \tag{2.2.3}$$

As previously stated, we choose to approximate the conditional expectations arising in (2.2.3) by regression. In order to be able to obtain an explicit error bound, we had to restrict our analysis to a specific regression basis, namely a basis of piecewise constant functions on a set of hypercubes $(B_k)_{k=1,\dots,K}$ which partition the domain $\mathcal{D}_\varepsilon$ (cf. Assumption 6).

For clarity and simplicity, we proceeded in two successive step (even though performing them simultaneously may improve the resulting bound, cf. [82] for example):

- *[Theoretical regression] First, we replace, at each time step t_n*, the conditional expectations by a projection of the set of $\mathcal{F}_{t_n}$-mesurable random variables generated by the chosen regression basis. The error behaves like δ/h, where δ is the maximum length of an edge of a hypercube.

- *[Empirical regression]* Then, we replace the previous projections by empirical projections computed from a sample of M independent draws of the state variable X_{t_n}. We established that the L_p-norm between the theoretical and empirical regressions behaves essentially like $C_p/\left(\sqrt{M}\times P^{1-\frac{1}{p\vee 2}}\right)$, where $C_p>0$ and $P=\min_{t\in\Pi}\min_{B_k\subset\mathcal{D}_\varepsilon}\mathbb{P}\left(\bar{X}_t\in B_k\right)$ (cf. Proposition 4.3.6).

This estimate generalizes to the L_p-norm a result established in [100] for the L_2-norm. The proof makes use of the following useful result, established in Appendix 4.7.1, which combines Jensen's inequality with Marcinkiewicz-Zygmund's inequality:

$$\left\|\frac{1}{M}\sum_{m=1}^{M}X_m\right\|_{L_p} \leq \frac{C_p}{\sqrt{M}}\left\|X_1\right\|_{L_{p\vee 2}},$$

where $X_1,\dots,X_M$ is an i.i.d. sample of real-valued random variable with zero mean, with a finite $p\vee 2$-order moment.

We now can explain the interest of the spatial localization step: it ensures that the variable P in the denominator of the regression error remains strictly positive for every M, ensuring the convergence of the algorithm. Of course, this step is only a theoretical requirement, as in practice, for fixed M, a sample of M trajectories would already be included in a finite set, therefore an additional space restriction would be redundant.

Finally, the main theoretical result of the chapter is the Theorem 4.3.1, which provides a comprehensive analysis, combining all the steps previously discussed, of the convergence error between the initial control problem (2.2.1) and the result of our algorithm.

The second part of the chapter is devoted to the complexity of the algorithm, in terms of computation as well as memory.

We establish that the general computational complexity of the algorithm behaves like $\mathcal{O}(q^2\cdot N\cdot M)$, where q is the number of elements in the switching set $\mathbb{I}_q$, N is the number of time steps on the grid Π, and M is the number of Monte Carlo trajectories. Under some conditions, specified in Section 4.4.1.1, this complexity can be improved to $\mathcal{O}(q\cdot N\cdot M)$. This improved complexity is very satisfactory, and is, in our view, probably the best attainable by a Monte Carlo algorithm for optimal switching (as the mere access to every possible switch for every trajectory at each time step already costs $\mathcal{O}(q\cdot N\cdot M)$).

As the Euler scheme is a forward scheme, while dynamic programming is a backward scheme, the first stage of the scheme is to simulate, starting from x_0, trajectories of the process $\bar{X}$ up to time T, so as to be able to initiate the dynamic programming. As, at each time $t_n \in \Pi$, the regression step requires to have access to the sample $\left(\bar{X}_{t_n}^m\right)^{1\leq m\leq M}$, the simplest implementation consists in storing the whole samples $\left(\bar{X}_{t_n}^m\right)_{1\leq i\leq N}^{1\leq m\leq M}$ beforehand. This requires a memory space in $\mathcal{O}\left(N \cdot M\right)$.

Such a memory space remains manageable when the time horizon is close (small N) or when the dimension of the problem is low (small M), but it can become a major limitation in case of both a distant time horizon and high dimension (two features of the numerical applications we are interested in). To overcome this limitation, we generalized a memory reduction method, introduced in [36] in the case of a geometric Brownian motion, to any stochastic processed discretized using the Euler scheme. This technique allows, at the price of doubling the number of operations, to decrease the required memory space to $\mathcal{O}\left(N + M\right)$, a considerable improvement that solves the memory issue.

The idea, detailed, analyzed and illustrated in Section 4.4.2, is the following. During the intial sumulation of the underlying process $\bar{X}$, at each time step $t_n \in \Pi$, instead of storing the whole sample $\left(\bar{X}_{t_n}^m\right)^{1\leq m\leq M}$, it is enough to store only the current state of the random generator immediately before the simulation of the current sample. After this first passage, the memory contains a sample of the final $\bar{X}_T$ (size M) as well as the seeds of the random generator (size N). Then, during the backward induction, at each time step $t_n \in \Pi$, it is possible, from the sample $\left(\bar{X}_{t_{n+1}}^m\right)^{1\leq m\leq M}$ and the seed from time t_n, to rebuild the sample $\left(\bar{X}_{t_n}^m\right)^{1\leq m\leq M}$, using the inverse of the Euler scheme combined with the resimulation, using the stored seed, of the sample that was used previously to go from $\left(\bar{X}_{t_n}^m\right)^{1\leq m\leq M}$ to $\left(\bar{X}_{t_{n+1}}^m\right)^{1\leq m\leq M}$.

Finally, the last part of the chapter is devoted to a numerical application of our scheme on an investment problem in power plants.

Using the power price model from the previous chapter, it is possible to express the power plant investment problem as an optimal switching problem (equation (4.5.9)). This is however a heavy problem, as it combines a very distant time horizon (several decades) with a high-dimensional state variable (demand, capacities, fuels,...). Nevertheless, as shown by the complexity analysis, our algorithm is precisely designed to cope, to some extent, with such problems.

We successfully implemented our algorithm on a numerical example with two power plant technologies. Figure 4.5.1 displays the estimated distribution of new power plants at time T, and Figures 4.5.2 et 4.5.3 display the impact of these new power plants on the time evolution of the price of electricity. We observe the expected impact of the new plants (lower and less volatile prices) along with the decreased attractivity, compared with base power plants, of peak power plants when energy prices are high.

Finally, here are a few closing remarks:

- In order to obtain an empirical assessment of the distance to convergence of our numerical results, we adapted a technique of empirical confidence intervals borrowed from [28] (cf. Appendix 4.7.4).

- To ensure sensible fuel price trajectories on very long time horizons, we added cointegration to the geometric Brownian motions modelling from the previous chapter:

$$dS_t = \Xi S_t dt + \mathrm{diag}\left(S_t\right)\Sigma dW_t \ , \ S_0 > 0 \, ,$$

where Ξ is the cointegration matrix and Σ is the covariance matrix. In Appendix 4.7.2 is established a necessary and sufficient condition to ensure that S_t remains non-negative over time (the non-diagonal elements of the cointegration matrix must be non-negative).

- Finally, in order to provide an intuitive visual display of the time evolution of a possibly multimodal stochastic process (like the price of electricity generated with our model), we generalized the construction of interquantile sets to (possibly disconnected) Borel sets adjusted to the level sets of the estimated distribution (cf. Appendix 4.7.5). This general construction could prove to be useful in a wide variety of contexts.

2.3 A numerical algorithm for fully nonlinear HJB equations: an approach by control randomization

The third and last part of the thesis deals once again with the numerical solution of stochastic control problems, but this time we consider problems such that the dynamics of the underlying state process can be modified by the control. The content of this part sums up the content of the two articles [71] et [70].

Consider first the following introductory example:

$$v(t,x) = \sup_{\alpha\in\mathcal{A}} \mathbb{E}\left[\int_t^T f(X_s^\alpha,\alpha_s)\,dt + g(X_T^\alpha)\,|X_t^\alpha = x\right] \tag{2.3.1}$$

$$dX_s^\alpha = b(X_s^\alpha,\alpha_s)\,ds + \sigma(X_s^\alpha)\,dW_s\,, \tag{2.3.2}$$

where the diffusion X^α takes values in $\mathbb{R}^d$, f and g are two measurable functions, and the processes $\alpha = (\alpha_s)_{t\le s\le T}$ are strategies chosen among a set $\mathcal{A}$ of admissible strategies taking values in a subset $A \subset \mathbb{R}^q$.

On this example (2.3.1), the dynamics of the underlying process X^α can be modified by the strategy α, but only through the drift b. In this very specific case, the problem can be solved using a Backward Stochastic Differential Equation (BSDE).

Indeed, the stochastic control problem (2.3.1) is, first of all, related to the following Hamilton-Jacobi-Bellman (HJB) equation:

$$\begin{aligned}&\frac{\partial v}{\partial t} + \sup_{a\in A}\left\{b(x,a).D_x v + \frac{1}{2}\mathrm{tr}\left(\sigma\sigma^\top(x)\,D_x^2 v\right) + f(x,a)\right\} = 0\ ,\ (t,x)\in[0,T)\times\mathbb{R}^d\\ &v(T,x) = g(x)\quad,\ x\in\mathbb{R}^d\,.\end{aligned} \tag{2.3.3}$$

A possibility is to try direct numerical methods on this PDE, but, as in the previous chapter, we look for a probabilistic algorithm, in order to be able to handle high-dimensional problems.

If σ is of full rank, then equation (2.3.3) can be transformed into the following semilinear PDE:

$$\begin{aligned}&\frac{\partial v}{\partial t} + F\left(x,\sigma^\top D_x v\right) + \frac{1}{2}\mathrm{tr}\left(\sigma\sigma^\top(x)\,D_x^2 v\right) = 0\ ,\ (t,x)\in[0,T)\times\mathbb{R}^d\\ &v(T,x) = g(x)\quad,\ x\in\mathbb{R}^d\,,\end{aligned} \tag{2.3.4}$$

where F is defined by $F(x,z) := \sup_{a\in A}\{\theta(x,a).z + f(x,a)\}$, where θ is such that $\sigma(x)\,\theta(x,a) = b(x,a)$. This new PDE (2.3.4) is itself associated with the following BSDE:

$$\begin{aligned}dX_s^0 &= \sigma\left(X_s^0\right)dW_s\\ Y_t &= g\left(X_T^0\right) + \int_t^T F\left(X_s^0,Z_s\right)ds - \int_t^T Z_s dW_s\,.\end{aligned}$$

As several probabilistic numerical schemes are available to solve a standard BSDE ([27, 82, 58]), it is therefore possible to solve the stochastic control problem (2.3.1) using a probabilistic numerical scheme.

Now, consider the following, more difficult problem:

$$\begin{aligned} v(t,x) &= \sup_{\alpha\in\mathcal{A}} \mathbb{E}\left[\int_t^T f(X_s^\alpha,\alpha_s)\,dt + g(X_T^\alpha)\,|X_t^\alpha = x\right] \\ dX_s^\alpha &= b(X_s^\alpha,\alpha_s)\,ds + \sigma(X_s^\alpha,\alpha_s)\,dW_s\,, \end{aligned} \tag{2.3.5}$$

which is associated with the following fully-nonlinear HJB equation:

$$\begin{aligned} &\frac{\partial v}{\partial t} + \sup_{a\in A}\left\{b(x,a).D_x v + \frac{1}{2}\mathrm{tr}\left(\sigma\sigma^\top(x,a)\,D_x^2 v\right) + f(x,a)\right\} = 0\ ,\ (t,x)\in[0,T)\times\mathbb{R}^d \\ &v(T,x) = g(x)\quad ,\ x\in\mathbb{R}^d\,. \end{aligned} \tag{2.3.6}$$

Now, the strategy α can impact both the drift b and the volatility σ of the process X^α. Such a problem is much more difficult than the previous example (2.3.1). In particular, it cannot be reduced to a simple standard BSDE.

Therefore, the aim of this chapter is to build a probabilistic algorithm able to solve stochastic control problems of the type (2.3.5). The idea is to extend the methodology used for the simpler problem (2.3.1), so as to be able to handle problems of the type (2.3.5).

The first step is to identify the adequate probabilistic representation for the problem (2.3.5). The object we need is a BSDE with nonpositive jumps, which is defined and studied in [72]. Let us recall this construction, and its link with the problem (2.3.6).

Let I be a pure jump process, independent of W, taking values in A. Its dynamics is defined as follows:

$$dI_s = \int_A (a - I_{s^-})\mu_A(ds,da)\,, \tag{2.3.7}$$

where $\mu_A(dt,da)$ is a Poisson random measure on $\mathbb{R}_+\times A$ with finite intensity measure $\lambda_A(da)\,dt$.

Then, consider the forward Markov regime-switching diffusion process (X,I) such that the dynamics of I is given by (2.3.7) and the dynamics of X by:

$$dX_s = b(X_s,I_s)\,ds + \sigma(X_s,I_s)\,dW_s\,. \tag{2.3.8}$$

In other words, we replace the control α in the dynamics (2.3.5) of X^α by a new stochastic process I independent of everything else. The expectation $\mathbb{E}\left[\int_t^T f(X_s,I_s)\,dt + g(X_T)\,|X_t = x\right]$ is then associated with the solution of the following BSDE with jumps:

$$Y_t = g(X_T) + \int_t^T f(X_s,I_s)\,ds - \int_t^T Z_s dW_s - \int_t^T\int_A U_s(a)\,\tilde{\mu}_A(ds,da)\ , \tag{2.3.9}$$

where $\tilde{\mu}_A$ is the compensated measure of μ_A.

Finally, in order to retrieve the solution of the HJB equation (2.3.6), it suffices to constrain the term $U_t(a)$ in equation (2.3.9) as follows:

$$U_t(a) \le 0,\ d\mathbb{P}\otimes dt\otimes\lambda(da)\ a.s.\ .$$

Under the following standing assumptions:

- b et σ Lipschitz (H1),

- f et g Lipschitz (H2),

it is proved in [72] that the minimal solution (Y, Z, U, K) of the constrained BSDE

$$\begin{aligned} Y_t &= g(X_T) + \int_t^T f(X_s, I_s)\, ds - \int_t^T Z_s dW_s \\ &\quad + K_T - K_t - \int_t^T \int_A U_s(a)\, \tilde{\mu}_A(ds, da)\ ,\ 0 \leq t \leq T\ ,\ p.s. \\ U_t(a) &\leq 0\ ,\ d\mathbb{P} \otimes dt \otimes \lambda(da)\ \ p.s.\ \text{sur}\ \Omega \times [0,T] \times A \end{aligned} \tag{2.3.10}$$

is related to the solution of the HJB equation (2.3.6). Indeed, it is shown that it is possible to express the Y component of the minimal solution of (2.3.10) as a function of t and X_t only (and not I_t):

$$Y_t = v(t, X_t)\ ,\quad t \in [0,T]\ , \tag{2.3.11}$$

and that this function v is the unique viscosity solution with linear growth of the HJB equation (2.3.6). This result remains valid if f also depends on Y and Z (which allows to consider problems more general than the stochastic control problem (2.3.5)). Moreover, we stress that this result (as well as all the results from this chapter) does not require any uniform ellipticity assumption on σ.

A consequence of this important theorem is that in order to solve the fully nonlinear HJB equation (2.3.6) (and therefore the stochastic control problem (2.3.5)), one can try to find a solution to the constrained BSDE (2.3.10) instead. Therefore, the focus of this chapter will be to propose, analyze and test a probabilistic numerical scheme for the constrained BSDE (2.3.10).

First of all, one needs to discretize this constrained BSDE. Let $\pi := \{0 = t_0 < \ldots < t_N = T\}$ be a deterministic time grid between 0 and T, with mesh $|\pi| = \max_{0 \leq k \leq n-1} \{t_{k+1} - t_k\}$. Then, drawing inspiration from the schemes for classical BSDEs, and taking adequately the constraint on $U_t(a)$ into account, it holds that the following discrete scheme approximates the constrained BSDE (2.3.10) (in the general case when f depends also on y and z):

$$\begin{cases} Y_N &= g(X_N) \\ \Delta_i \mathcal{Z}_i &= \mathbb{E}_i\left[Y_{i+1} \Delta W_i^\top\right] \\ \mathcal{Y}_i &= \mathbb{E}_i\left[Y_{i+1} + f(X_i, I_i, Y_{i+1}, \mathcal{Z}_i)\, \Delta_i\right] \\ Y_i &= \operatorname{ess\,sup}_{a \in A} \mathbb{E}_{i,a}\left[\mathcal{Y}_i\right]\ , \end{cases} \tag{2.3.12}$$

where $(X_i, I_i)_{i=1,\ldots,N}$ corresponds to the Euler discretization of the process $(X_t, I_t)_{0 \leq t \leq T}$, $\mathbb{E}_i[.] := \mathbb{E}[.\,|\mathcal{F}_{t_i}]$, and $\mathbb{E}_{i,a}[.] := \mathbb{E}[.\,|X_i, I_i, I_i = a] = \mathbb{E}[.\,|X_i, I_i = a]$. In order to understand this scheme, notice that, by Markov property, the intermediary quantities $\mathcal{Y}_i$ and $\mathcal{Z}_i$ can be expressed as a function of X_i and I_i:

$$(\mathcal{Y}_i, \mathcal{Z}_i) = (\tilde{y}_i(X_i, I_i), \tilde{z}_i(X_i, I_i))\ .$$

The steps involving $\mathcal{Y}_i$ and $\mathcal{Z}_i$ are similar to those from classical BSDE schemes. Here, a new operation is required in order to take the constraint on $U_t(a)$ into account, and to retrieve the true approximation $(Y_i, Z_i) = (y_i(X_i), z_i(X_i))$, which does not depend on I_i. From the definition of $U_t(a)$, this constraint can be expressed as follows:

$$\tilde{y}_i(X_i, a) - y_i(X_i) = U_t(a) \leq 0\ p.s.\ \forall a \in A\,,$$

which implies that the Y_i component of the minimal solution has to satisfy:

$$Y_i = y_i(X_i) = \operatorname*{ess\,sup}_{a \in A} \tilde{y}_i(X_i, a) = \operatorname*{ess\,sup}_{a \in A} \mathbb{E}_{i,a}\left[\mathcal{Y}_i\right]\ .$$

If necessary, it is possible to retrieve the component Z_i from the scheme indeed, defining $a_i^* = \text{argess sup}_{a\in A} \mathbb{E}_{i,a}[\mathcal{Y}_i]$, then $Z_i = z_i(X_i) = \tilde{z}_i(X_i, a_i^*)$.

A first part of this chapter is devoted to the theoretical analysis of this discretization. We proceed in two steps. We first discretize the constraint without changing the process X, and then we replace X by its Euler scheme.

To do so, we start by defining the intermediary sequence of BSDEs $(Y^\pi, \mathcal{Y}^\pi, \mathcal{Z}^\pi, \mathcal{U}^\pi)$ (equations (5.2.13)-(5.2.14)-(5.2.15)), which corresponds to the initial constrained BSDE (2.3.10) with discretization of the constraint on U. We also define the value functions v^π and ϑ^π associated with Y^π and $\mathcal{Y}^\pi$ (equations (5.2.16)-(5.2.17)-(5.2.18)).

As this construction of a sequence of constrained BSDEs is not standard, we start by proving that all these elements are well defined (Propositions 5.2.1-5.2.2-5.2.3). Then, we prove that the function v^π associated with the discrete scheme does converge to the function v from the initial constrained BSDE when $|\pi| \to 0$ (Corollary 5.2.2). Finally, we establish a convergence rate of v^π towards v. To do so, we need the following additional assumptions:

- b et σ uniformly bounded (H1'),
- f does not depend on z, $f(.,.,0)$ and g uniformly bounded, and $y \mapsto f(x,a,y)$ convex (H2').

With these additional assumptions, adapting the shaking coefficients method from Krylov ([78], cf. also [49]) to the jump-diffusion framework, we manage to establish the following convergence rate (Theorem 5.2.1) :

$$0 \leq v(t,x) - v^\pi(t,x) \leq C|\pi|^{\frac{1}{10}}$$

for every $(t,x) \in [0,T] \times \mathbb{R}^d$ and every grid π. Moreover, in the case when $f = f(x,a)$ does not depend on y, the stochastic control representation combined with the results from [77] allows to improve this rate to $C|\pi|^{\frac{1}{6}}$. In other words, the solution of the discretely constrained BSDE converges to the solution of the continuously constrained BSDE at a rate $\mathcal{O}\left(|\pi|^{\frac{1}{10}}\right)$ (or $\mathcal{O}\left(|\pi|^{\frac{1}{6}}\right)$ when f does not depend on y), cf. Corollary 5.2.3.

Then, the second part of the discretization consists in replacing the diffusion X by its Euler scheme, ending up with the discrete scheme (2.3.12). Adapting the properties of the Euler scheme (Lemma 5.2.1) and regularity results on BSDEs (Proposition 5.2.3), we obtain for this second step an error estimate in $\mathcal{O}\left(|\pi|^{\frac{1}{2}}\right)$ (Theorem 5.2.2). Consequently, combining the results from both steps, we obtain, for a full time-discretization, a convergence rate in $\mathcal{O}\left(|\pi|^{\frac{1}{10}}\right)$ ($\mathcal{O}\left(|\pi|^{\frac{1}{6}}\right)$ when f does not depend on y), or, more precisely, a negative error in $\mathcal{O}\left(|\pi|^{\frac{1}{2}}\right)$ and a positive error in $\mathcal{O}\left(|\pi|^{\frac{1}{10}}\right)$ ($\mathcal{O}\left(|\pi|^{\frac{1}{6}}\right)$ when f does not depend on y), cf. Corollary 5.2.4.

This convergence rate improves the existing results (cf. Remark 5.2.3). For example, when f does not depend on y, our rates $|\pi|^{\frac{1}{2}}$ and $|\pi|^{\frac{1}{6}}$ improve respectively the bounds $|\pi|^{\frac{1}{4}}$ an $|\pi|^{\frac{1}{10}}$ obtained via second-order BSDEs (cf. [49]). Then, and it is a major feature of our approach, our rates are obtained without any uniform ellipticity condition on σ. Finally, a very useful practical feature is that our scheme provides an estimate of the optimal control (cf. Section 5.2.4.2). All these observations indicate that constrained BSDEs seem to be a suitable tool for the study of fully-nonlinear HJB equations.

The second part of this chapter is devoted to the approximation of the conditional expectations from the discrete scheme (2.3.12). Indeed, as in the previous chapter, we are dealing with a backward numerical scheme that requires the computation of conditional expectations, for which no explicit expressions are available in general. One again, we choose to approximate them by

least-squares regression. This choice has several advantages. First the Monte Carlo simulation of the forward process (X, I) allows to get rid, to some extent, of dimensional constraints. Then, this choice reduces the computation of supremum from the discrete scheme (2.3.12) to a simple maximization of a given parametric function. Finally, it provides a parametric estimate of the optimal control.

In the previous chapter, in the context of optimal switching problems, we only managed to assess the regression error for a very specific basis of piecewise constant functions. In this chapter, we wanted to establish a general error bound without specific assumption on the chosen regression basis (in the spirit of [82] and [58] in the context of standard BSDEs). Unfortunately, in the context of constrained BSDEs, the additional supremum step makes the analysis substantially more difficult, and we could only manage to obtain partial results for this error analysis.

The preliminary space localization step does not generate any major difficulty (Lemma 5.3.1 and Proposition 5.3.1). Consider now the actual regression step. Let $\left\{p_b : \mathbb{R}^d \times \mathbb{R}^q \to \mathbb{R}\right\}$, $b = 1, \dots, B$, be a basis of regression functions. An $\mathcal{F}_T$-measurable random variable U being given, the idea is to replace the conditional expectation $\mathbb{E}_i[U]$ by the projection $\mathcal{P}_i(U)$ defined as follows:

$$\begin{aligned} \hat{\lambda}_i(U) &:= \arg \inf_{\lambda \in \mathbb{R}^B} \mathbb{E}\left[\left(\lambda.p(X_i, I_i) - U\right)^2\right] \\ \mathcal{P}_i(U) &:= \hat{\lambda}_i(U).p(X_i, I_i) \ . \end{aligned} \tag{2.3.13}$$

This constitutes a theoretical regression (the final empirical regression step, in order to get an implementable scheme, consists in replacing the expectation $\mathbb{E}$ by an empirical expectation $\frac{1}{M}\sum_{m=1}^{M}$ in the above least-squares minimization). Performing this substitution within the discrete scheme (2.3.12), one obtains the following new scheme:

$$\begin{aligned} \hat{Y}_N &= g(X_N) \\ \Delta_i \hat{\mathcal{Z}}_i &= \mathcal{P}_i\left(\hat{Y}_{i+1} \Delta W_i^\top\right) \\ \hat{\mathcal{Y}}_i &= \mathcal{P}_i\left(\hat{Y}_{i+1} + f_i\left(X_i, I_i, \hat{Y}_{i+1}, \hat{\mathcal{Z}}_i\right)\Delta_i\right) =: \hat{\lambda}_i^Y.p_i(X_i, I_i) \\ \hat{Y}_i &= \operatorname*{ess\,sup}_{a \in \mathcal{A}_i} \mathbb{E}_{i,a}\left[\hat{\mathcal{Y}}_i\right] = \operatorname*{ess\,sup}_{a \in \mathcal{A}_i} \hat{\lambda}_i^Y.p_i(X_i, a) \ , \end{aligned} \tag{2.3.14}$$

where $\mathcal{A}_i$ is the set of $\sigma(X_i)$-measurable random variables taking values in A. (For readability purposes, we omit here the space localization step). Even at this theoretical stage, the analysis between the discrete scheme (2.3.12) and its approximation (2.3.14) raises the following difficulty: because of the final maximization step, the random variable $\hat{Y}_i$, contrary to the intermediary variable $\hat{\mathcal{Y}}_i$, does not constitute in itself the projection of a given random variable. Consequently, it is not possible to use the usual tools (cf. Lemma 5.3.3) that would allow to get rid of the operator $\mathcal{P}_i$ when performing the analysis of the difference between (2.3.12) and (2.3.14).

It is interesting to point out that, should the regression be modified as follows:

$$\begin{aligned} \hat{\lambda}_{i,a}(U) &:= \arg \inf_{\lambda \in \mathbb{R}^B} \mathbb{E}\left[\left(\lambda.p(X_i, a) - U_a\right)^2\right] \\ \mathcal{P}_{i,a}(U) &:= \hat{\lambda}_i(U).p(X_i, a) \ , \end{aligned} \tag{2.3.15}$$

where $U_a := U|a$, $a \in \mathcal{A}_i$, then it would become possible to analyze the regression error (Proposition 5.3.2). The difference here lies in the fact that regression and maximization are now performed simultaneously, and not iteratively as with the scheme (2.3.14). However, unlike the empirical version of the scheme (2.3.14), an empirical scheme based on the regression (2.3.15) would be tricky to implement. Therefore, we choose to stick after all with a scheme based on

the regression (2.3.13), the comprehensive theoretical analysis of which being left for further research.

Finally, the last part of this chapter is devoted to numerical applications of the empirical version of the scheme (2.3.14).

First, we implement our scheme on a linear quadratic stochastic control problem (equations (5.4.1) and (5.4.2)). The explicit solution to these problems being known (cf. [107]), this example allows to check the accuracy of the scheme. As shown by our numerical results (cf. Figure 5.4.1b), our scheme manage the approximate the optimal control with great accuracy.

After this preliminary check, we devoted our numerical tests to the problem of superreplication of options under uncertain volatility (and/or correlation) models. We start with an example with two assets S_1 et S_2 whose dynamics are given by:

$$dS_i(t) = \sigma_i S_i(t)\, dW_i(t) \ , \ i = 1, 2 \tag{2.3.16}$$

$$\langle dW_1(t), dW_2(t) \rangle = \rho dt\,, \tag{2.3.17}$$

where the process ρ is deemed uncertain; the only assumption is that it can only take values within a given interval $[\rho_{\min}, \rho_{\max}] \subseteq [-1, 1]$. The superreplication prices of the option with final payoff

$$\Phi = (S_1(T) - S_2(T) - K_1)^+ - (S_1(T) - S_2(T) - K_2)^+ \ , \tag{2.3.18}$$

(non trivial because neither convex nor concave in $S_1(T) - S_2(T)$) estimated by our algorithm are displayed on Figure 5.4.3. We retrieve the fact that uncertain correlation models are more suited for risk analysis and model risk purposes than the mere variation of constant correlation.

After this example (2.3.18) which illustrates the practical interest of models with uncertain parameters, we tested our algorithm on several other payoffs, in order to assess the quality of the algorithm itself. We implemented our scheme on the numerical examples provided in [62], as their paper provides superreplication price estimates obtained by PDE methods and by an alternative Monte Carlo algorithm based on second-order BSDE (cf. [49]).

On all these examples, and as in the previous chapter, we computed two complementary prices that form an asymptotic empirical confidence interval for the price (equations (5.4.8) and (5.4.9)). The extent of this interval provides an empirical assessment of the suitability of the chosen regression basis.

Figures 5.4.4 to 5.4.9 illustrate our results. We vary the number of Monte Carlo trajectories, as well as the size of the time step. Our results behave as expected: both prices surround the true price (as estimated by PDE methods), and those prices are all the more so accurate as the time step is small and the number of Monte Carlo trajectories is high.

Even though our numerical results and those from [62] are relatively similar, the strong advantage of our scheme is that our forward process (X, I) is easy to simulate, a fact which, we hope, should allow to easily extend it to more general problems like HJB-Isaacs equations and mean-field games.

3 A structural risk-neutral model for pricing and hedging power derivatives

In this chapter, we develop a structural risk-neutral model for energy market modifying along several directions the approach introduced in [3]. In particular a scarcity function is introduced to allow important deviations of the spot price from the marginal fuel price, producing price spikes. We focus on pricing and hedging electricity derivatives. The hedging instruments are forward contract on fuels and electricity. The presence of production capacities and electricity demand makes such a market incomplete. We follow a local risk minimization approach to price and hedge energy derivatives. Despite the richness of information included in the spot model, we obtain closed-form formulae for futures prices and semi-explicit formulae for spread options and European options on electricity forward contracts. An analysis of the electricity price risk premium is provided showing the contribution of demand and capacity to the futures prices. We show that when far from delivery, electricity futures behave like a basket of futures on fuels.

3.1 Introduction

This chapter is a contribution to the development of electricity price models that can provide explicit or semi-explicit formulae for European derivatives on electricity markets. Since the beginning of the liberalization process of electricity markets in the 90s in Europe and in the USA, there has been an important research effort devoted to electricity price modeling for pricing derivatives. Due to the non-storable nature of electricity, it was — and still is — a challenge to reach to a completely satisfying methodology that would suit the needs of trading desks: a realistic and robust model, computational tractability of prices and Greeks, consistency with market data. Two main standard approaches have usually been used to tackle this problem. The first one consists in directly modeling the forward curve dynamics and to deduce the spot price as a futures with immediate delivery. Belonging to this approach are, e.g., [42] and [18]. This approach is pragmatic in the sense that it models the prices of the available hedging instruments. However, it makes it difficult to capture the right dependencies between fuels and electricity prices (without cointegration). The second approach starts from a spot price model to deduce futures price as the expectation of the spot under a risk-neutral probability. The main benefit of this approach is that it provides a consistent framework for all possible derivatives. This approach has been successfully applied to commodities in [96] seminal work. Its main drawback is that it generally leads to complex computations for the prices of electricity derivatives. Within this approach, most of the authors use an exogenous dynamics for the electricity spot price [47, 16, 29, 75, 34, 17, 19, 60] and only a few try to deduce futures and option prices inside an equilibrium model or a model including a price formation mechanism [35, 93, 84, 3].

The main contribution of this work is to provide analytical formulae for electricity futures and semi-explicit expressions for European options in an electricity spot price model that includes demand and capacities as well as fuel dynamics. Being able to model the dependencies between fuels and electricity is of great importance for spread options evaluation. To our knowledge, this is the first attempt performed in that direction.

Concerning the use of an equilibrium model or a price mechanism for pricing electricity derivatives, the closest works to ours can be found in [93, 35, 84]. It has been recognized that the mechanism leading to the electricity spot price was too complex to allow for a complete modeling that would fit the constraints of derivatives pricing. The simplest one is maybe [11]'s model where the price is determined by the matching of a simple parametric offer curve and a random demand. Many authors have then derived reduced equilibrium models for electricity prices in this spirit [68, 43]. In [93]'s work, electricity dependency on fuel prices is taken into account by modeling directly the dynamics of the marginal fuel. The authors manage to provide the partial differential equation and its boundary conditions for the price of a European derivative. The approach followed by [35] and [84] is quite similar. Therein, the price is modeled as an exponential of a linear combination of demand and capacity. In general, it is difficult to introduce in the same framework the dependency of the electricity spot price on fuels and at the same time its dependency on demand and capacity. Dependencies among fuels are generally captured by simple correlations between Ornstein-Uhlenbeck processes as in [54]'s paper or by cointegration method as in [14]'s work.

Here, we start from the marginal price model developed in [3] and enrich it substantially to take into account the effect of the margin capacity uncertainty on futures prices. In order to include the biggest price spikes in our model, we introduce a multiplying factor allowing the electricity spot price to deviate from the marginal fuel price when demand gets closer to the capacity limit. Since electricity is a non-storable commodity, this factor accounts directly for the scarcity of production capacity. Although such an additional feature complexifies the model, we can still provide closed form formulae for futures prices. Under this model, any electricity futures contract behaves almost as a portfolio of futures contracts on fuels as long as the product is far from delivery. In contrast, near delivery, electricity futures prices are determined by the scarcity rent, ie. demand and capacity uncertainties.

The term *scarcity* is in line with the concept introduced by [56] since the multiplying factor is the inverse of the excess capacity which can be considered the analog of the inventory level for storable commodities. It is also consistent with Working's seminal observations on the relation between inventory and time spread price for agricultural commodities [106]. Moreover, the idea that electricity spot price spikes are related to a lack of capacity is some sort of common knowledge in the power system literature (see for instance [22]).

This chapter is structured as follows. We first present the spot price model in Section 3.2, where we perform the estimation of the scaling factor allowing the spot price to deviate from the marginal fuel price and comment on the production capacity scarcity effect. Then, in Section 3.3, we apply the spot model for pricing and hedging derivatives. We first choose realistic as well as tractable models for the dynamics of demand, capacities and fuel prices (Section 3.3.1). Then, since we work in an incomplete market setting, we discuss in Section 3.3.2 our choice for [52]'s Local Risk Minimization hedge criterion. Using such a criterion, a closed-form expression for futures prices is provided (Section 3.3.3) and semi-explicit formulae are given for spread options and for European options on futures (Section 3.3.4). Finally, an important part of this work is devoted to numerical simulations and backtesting presented in Section 3.4. Despite the apparent complexity of the model, the numerical computations essentially involve integration of functionals against Gaussian kernels. This part requires simple but long and tedious calculations. Thus, for the sake of readability, they have been relegated in Section 3.6.2 of the Appendix.

3.2 Electricity spot market model

3.2.1 Spot model

We denote by P_t the electricity spot price at time t. At any time t, the electricity producer can choose, among n possible fuels, the most convenient to produce electricity at that particular moment, called the marginal fuel. We will define the electricity spot price as a proportion of the spot price of the marginal fuel, corrected by a scarcity factor, ie. a factor depending on the current difference between available capacity and demand.

Denote as $(S_t^1, \dots, S_t^n)$ the fuel prices at time t, and $(h_1, \dots, h_n)$ the corresponding heat rates, assumed to be constant. That means that $h_i S_t^i$ corresponds to the t-price of the quantity of i-th fuel necessary to produce 1MWh of electricity. Unlike the model proposed in [3], we make the further assumption that the production costs are ordered among fuels, ie. that $h_1 S_t^1 \leq \cdots \leq h_n S_t^n$. This ranking is supposed fixed and known. This assumption is realistic, at least when considering not too long maturities. Now, how does the electricity producer choose the most convenient fuel to use? For each $i = 1, \dots, n$, let $C_t^i > 0$ denotes the capacity of the i-th technology for electricity production at time t. Denote as $\overline{C}_t^i$ the total capacity at time t of the first i fuels, i.e. $\overline{C}_t^i := \sum_{j \leq i} C_t^j$. For the maximal capacity $\overline{C}_t^n$ we will use, alternatively, the notation C_t^{max}. We define the following production intervals:

$$I_t^1 := (-\infty, C_t^1), \quad I_t^i := \left[\overline{C}_t^{i-1}, \overline{C}_t^i\right), \, 2 \leq i \leq n-1, \quad I_t^n := \left[\overline{C}_t^{n-1}, +\infty\right), \tag{3.2.1}$$

with the conventions that when $n = 1$ there is only one interval $I_t^1 := \mathbb{R}$, and when $n = 2$ there are only two intervals, the extreme ones, i.e. $I_t^1 := (-\infty, C_t^1)$ and $I_t^2 := [C_t^1, +\infty)$. Thus, if the market demand at time t for electricity D_t belongs to the interval I_t^i, the last (marginal) unit of electricity is produced using the i-th fuel, when the corresponding plant is available. Otherwise, it is produced with the next fuel, more expensive, in the ranking. Having said that, we model the electricity spot price P_t by the following relation:

$$P_t = g\left(C_t^{max} - D_t\right) \sum_{i=1}^{n} h_i S_t^i \mathbf{1}_{\left\{D_t \in I_t^i\right\}}, \quad t \geq 0 \tag{3.2.2}$$

where recall that $C_t^{max} = \overline{C}_t^n = \sum_{i=1}^n C_t^i$ is the maximal capacity of the whole system at time t, and g is a bounded real-valued function given by:

$$g(x) = \min\left(M, \frac{\gamma}{x^\nu}\right) \mathbf{1}_{\{x>0\}} + M \mathbf{1}_{\{x \leq 0\}} \tag{3.2.3}$$

where γ, M and ν are positive and constant parameters.

The term $C_t^{max} - D_t$ is the *margin capacity* of the system. It is an indicator of the tension in the system due to scarcity, since electricity is non-storable, and the term $g\left(C_t^{max} - D_t\right)$ models the effect of this tension on prices. This is why we will call g the *scarcity function* of electricity prices. The margin capacity is better at capturing electricity price spikes than only demand as in [11] and [68]. This remark has already been pointed out by several authors as, e.g., [35, 43] and [84]. [43] provide the most complex framework, which is based on a model for the bidding curves. Unfortunately their model does not allow for analytical expressions of futures prices as soon as there are more than two fuels. On the contrary, in [35] and [84]'s works, drastic simplifications of the electricity market rules allow them to provide analytical expressions for futures prices, and, in [84]'s model, for European option prices as well. In their models, the electricity spot price is expressed as an exponential of the margin capacity, which simplifies considerably the computation of derivative prices and of the risk premium. However,

the exponential fonction does not enable price spikes as sharp as those observed on electricity spot data. On the contrary, as will be shown below, our choice (3.2.3) of a power law of margin capacity can accurately reproduce such a behaviour, even for smooth and rather simple dynamics of demand and capacity processes.

3.2.2 Estimation and backtesting

Before making further use of our model (3.2.2), we aim in this section at assessing its accuracy on historical data. In particular, a methodology to estimate the scarcity function (3.2.3) will be proposed.

3.2.2.1 Data set

We choose to test the model on the French deregulated power market. We retrieve the required data from the following sources:

- PowerNext for the hourly power spot price P_t.
- RTE, the French transmission system operator, for the hourly power demand D_t and capacity C_t^i for each fuel.
- TFS (Tradition Financial Services) for the daily coal price (API#2).
- IPE (International Petroleum Exchange) for the daily oil price (Brent).
- ECB (the European Central Bank) for the daily USD/EUR exchange rate, to convert the coal and oil prices, which are denominated in US dollars, into euros.
- ECX (European Climate Exchange) for the daily CO_2 price.
- EDF, French power utility, for heat rates h_i and CO_2 emission rates.

We focus our analysis on one particular hour of the day, namely the 19^{th}, which usually bears the highest demand level of a day. As a consequence, during these peak hours, the electricity price is almost always fixed by one of the two most expensive technologies that are coal and oil. Consequently such a choice simplifies the model since it makes possible to work with only $n = 2$ technologies.

Thus, S_t^1 corresponds to the daily coal price, converted to EUR/MWh using the USD/EUR exchange rate and the coal heat rate h_1. We include the price of CO_2 emissions in S_t^1, using the daily CO_2 price and the emission rate of French coal plants. Similarly, S_t^2 corresponds to the daily overall oil price.

For the simplification to $n = 2$ technologies to hold, it suffices to define D_t as the residual demand corresponding to coal and oil. This quantity can be extracted from RTE demand, production and capacity data. More details are given in Appendix 3.6.1.

Our dataset covers the period going from November 13^{th}, 2006 to April 30^{th}, 2010. The beginning of the period was fixed by the availability of the production capacity data. On this period, during the 19^{th} hours, the average electricity spot price was 74.5€/MWh (see Figure 3.2.2a), the average coal price (including heat rate and CO_2) was 47.4€/MWh, and the average oil price (including heat rate and CO_2) was 102.0€/MWh.

3.2.2.2 Parameters estimation

So far, we retrieved all the necessary data to test equation (3.2.2). The last remaining step is to estimate the parameters M, γ and ν that characterize the scarcity function (3.2.3) recalled below:

$$g(x) = \min\left(M, \frac{\gamma}{x^\nu}\right)\mathbf{1}_{\{x>0\}} + M\mathbf{1}_{\{x\leq 0\}}$$

First, we fix M so as to roughly match the high cap on electricity spot price, defined by the market as 3000€/MWh[1]. Our estimate is $M = 30$.

Now, we turn our attention to the parameters γ and ν. Remark first that γ, unlike M and ν, depends on the unit in which D_t and C_t^i, $1 \leq i \leq n$, are denominated. This is a consequence of the following result:

Proposition 3.2.1. *(Change of Unit) Let $N > 0$ be a constant. The following holds:*

$$P_t(C_t, D_t, S_t, M, \gamma, \nu) = P_t\left(\frac{C_t}{N}, \frac{D_t}{N}, S_t, M, \frac{\gamma}{N^\nu}, \nu\right)$$

where we have explicited the parameters involved in the definition of P_t.

Proof. First, one can easily check that $g(x; M, \gamma, \nu) = g\left(\frac{x}{N}; M, \frac{\gamma}{N^\nu}, \nu\right)$. The result follows then from equation (3.2.2) and the fact that $\left\{\frac{D_t}{N} \in \frac{I_t^i}{N}\right\} = \{D_t \in I_t^i\}$ a.s., $1 \leq i \leq n$. □

This useful change of unit formula indicates that we are free to choose the desired unit for D_t and C_t. In our case, we choose to convert these data, which are provided in MWh, into GWh ($N = 1000$), and to estimate the corresponding "GWh-γ". This renormalisation will prove to be numerically convenient in Section 3.4.2.

Going back to the estimation of ν and γ, Figure 3.2.1a depicts the quantity $y_t := \frac{P_t}{\sum_{i=1}^n h_i S_t^i \mathbf{1}_{\{D_t \in I_t^i\}}}$ as a function of $x_t := C_t^{max} - D_t$ for the period we consider, i.e. November 13$^{\text{th}}$, 2006 to April 30$^{\text{th}}$, 2010. The quantity y_t corresponds to the price of electricity corrected from the effect of the marginal fuel price, while x_t corresponds to the remaining available production capacity (margin capacity). Recalling equation (3.2.2), the relation between x_t and y_t is to be captured by g. Observe that the highest prices (high y_t) occur mostly when there is not much available capacity left (low x_t), which simply translates the law of supply and demand. It suggests a decreasing causality relationship between x_t and y_t.

First, note that as the high cap price is not reached in our time period, the parameter M will not intervene in the following. Now, remark that such classical tools as least squares or maximum likelihood are not adapted to the estimation of g. Indeed, both x_t and y_t are random, and the slope of the relation between x_t and y_t becomes extremely steep as x_t approaches zero. These tools, that measure the error in one dimension only (y_t) are bound to miss much of the slope part. To overcome this difficulty, our idea is intuitively the following: if a strictly decreasing deterministic relation between x_t and y_t was to be inforced, then a similar relation would link their quantiles, from which the estimation would be easier. This is the content of the following proposition:

Proposition 3.2.2. *Let X and Y be two real random variables on the probability space $(\mathbb{R}, \mathcal{B}(\mathbb{R}), \mathbb{P})$, where $\mathcal{B}(\mathbb{R})$ denotes the corresponding Borel σ-field. Assume that both X and Y have continuous*

[1] see http://www.epexspot.com/en/product-info/auction/france

and strictly increasing cumulative distribution functions. Thus, their quantile functions q_X *and* q_Y *exist, are unique, and defined for all* $p \in [0,1]$ *by:*

$$\mathbb{P}\left(X \leq q_X\left(p\right)\right) = p \quad , \qquad \mathbb{P}\left(Y \leq q_Y\left(p\right)\right) = p$$

Suppose that there exists a strictly decreasing function h *such that the relation* $Y = h\left(X\right)$ *holds* $\mathbb{P}$*-as. Then for all* $p \in [0,1]$*:*

$$q_Y\left(1-p\right) = h\left(q_X\left(p\right)\right)$$

Proof. For every $p \in [0,1]$:

$$1 - p = 1 - \mathbb{P}\left(X \leq q_X\left(p\right)\right) = 1 - \mathbb{P}\left(h\left(X\right) \geq h\left(q_X\left(p\right)\right)\right) = \mathbb{P}\left(Y \leq h\left(q_X\left(p\right)\right)\right)$$
$$1 - p = \mathbb{P}\left(Y \leq q_Y\left(1-p\right)\right)$$

Thus $\mathbb{P}\left(Y \leq q_Y\left(1-p\right)\right) = \mathbb{P}\left(Y \leq h\left(q_X\left(p\right)\right)\right)$, and the unicity of the quantile function yields $q_Y\left(1-p\right) = h\left(q_X\left(p\right)\right)$. □

Consequently, if one assumes a strictly decreasing deterministic relation h between x_t and y_t, then h can be estimated from the quantiles of x_t and y_t. This will prove to be very simple and much more robust and accurate than working directly with the realizations of x_t and y_t.

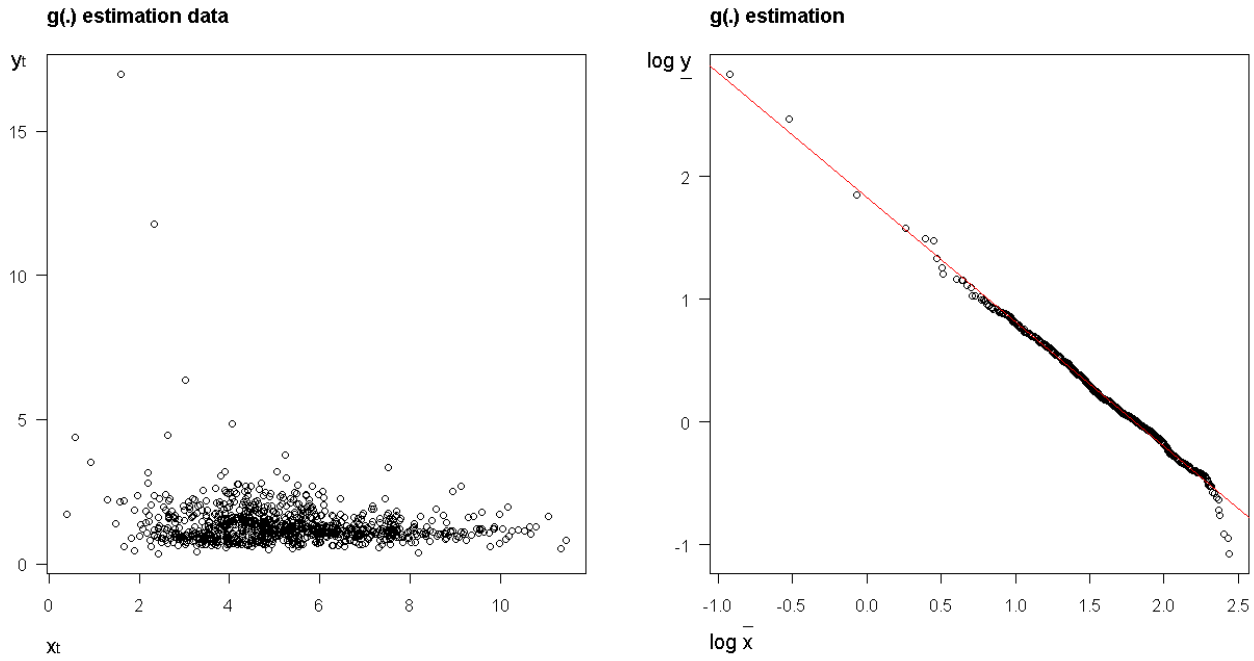

(a) Data

(b) Regression of $\log\left(\underline{y}\right)$ versus $\log\left(\overline{x}\right)$

Figure 3.2.1: g estimation

Denote as m the sample size, and $\left(x_1, y_1\right), \ldots, \left(x_m, y_m\right)$ the realizations of $\left(x_t, y_t\right)$. Define the sorted vectors $\overline{x} := \left(\overline{x}_1, \ldots, \overline{x}_m\right)$ and $\underline{y} := \left(\underline{y}_1, \ldots, \underline{y}_m\right)$ in such a way that $\overline{x}_1 \leq \ldots \leq \overline{x}_m$ (increasing order) and $\underline{y}_1 \geq \ldots \geq \underline{y}_m$ (decreasing order). The last step is to choose some $p \in [0,1]$ and to estimate the quantiles $q_Y\left(p\right)$ and $q_X\left(1-p\right)$. A simple and natural choice is to select $p_i := \frac{i}{m}$, $1 \leq i \leq m$, and to choose as quantile estimator for $q_X\left(p_i\right)$ and $q_Y\left(1-p_i\right)$ the inverse of the empirical cumulative distribution function, which yield in this case the estimates

$\overline{x}_{\lceil mp_i \rceil} = \overline{x}_i$ and $\underline{y}_{\lceil mp_i \rceil} = \underline{y}_i$ ($\lceil x \rceil$ stands for the upper integer part of any real number x). Figure 3.2.1b represents the quantity $\log\left(\underline{y}\right)$ as a function of $\log\left(\overline{x}\right)$.

A strikingly accurate linear relation appears. In particular it retranscribes very well the highest peak prices, which advocates a power-law description of electricity price spikes. The very lowest prices seem not to fit the linear relation. However a specific analysis of these few points reveals that they correspond to the few holidays (hence low demand) of the period considered for which the marginal fuel was gaz during the 19^{th} hour of the day. As we neglected on purpose this possibilily to work with only two fuels (coal and oil), the consequence is that the residual capacity is undervalued at these specific dates, breaking the smooth relation. Otherwise, the linear relation appears very plausible.

Denote as $-\nu$ and $\log(\gamma)$ the coefficients of the regression, ie. $\log\left(\underline{y}\right) = -\nu \log\left(\overline{x}\right) + \log(\gamma)$. Taking the exponential yields $\underline{y} = \frac{\gamma}{\overline{x}^{\nu}}$. Thus our estimation of the relation between the quantiles $q_{y_t}(1-p)$ and $q_{x_t}(p)$ is given by $h(x) = \frac{\gamma}{x^{\nu}}$. Numerically, we found the estimates $\nu = 1.022$ and $\gamma = 6.241$ with the corresponding 95% confidence intervals of $[1.017, 1.028]$ and $[6.186, 6.296]$ respectively, and an adjusted R-squared equal to 0.9935 (see Figure 3.2.1b).

Now, using Proposition 3.2.2, we can trace back to the relation between $y_t = \frac{P_t}{\sum_{i=1}^{n} h_i S_t^i \mathbf{1}_{\{D_t \in I_t^i\}}}$ and $x_t = C_t^{max} - D_t$. Indeed, the consequence of Proposition 3.2.2 combined with the empirical power law relation between $\overline{x}$ and $\underline{y}$ is the following: if one assumes a strictly decreasing deterministic relation between x_t and y_t, then it ought be a power law relation, defined by $h(x) = \frac{\gamma}{x^{\nu}}$. This, combined with the modeling of the market cap price, leads to the relation (3.2.3) and to our model (3.2.2). Therefore, our estimation of the parameters ν and γ provides also empirical evidence for the relevance of our spot model.

3.2.2.3 Backtest

Finally, with our data and our estimated parameters, we can backtest our model, i.e. we can compute the quantity $g(C_t^{max} - D_t) \sum_{i=1}^{n} h_i S_t^i \mathbf{1}_{\{D_t \in I_t^i\}}$ for each date t of our dataset, and comparing it to the realized electricity spot prices for the same dates. Figure 3.2.2a illustrates this comparison.

It can be seen that the base prices are described rather well by the model (thanks to the marginal price $\sum_{i=1}^{n} h_i S_t^i \mathbf{1}_{\{D_t \in I_t^i\}}$) and that the model is able to produce price spikes of relatively good size (thanks to the multiplication to the scarcity function g and its power law shape (3.2.3)) and timing (thanks to the choice of the margin capacity $C_t^{max} - D_t$ as a state variable). It is these adequate prisms $C_t^{max} - D_t$ and $\sum_{i=1}^{n} h_i S_t^i \mathbf{1}_{\{D_t \in I_t^i\}}$ that enable g to be stationary over time, making the model robust.

Figure 3.2.2b compares the model (3.2.2) to the simpler model $\sum_{i=1}^{n} h_i S_t^i \mathbf{1}_{\{D_t \in I_t^i\}}$ (i.e. when $g \equiv 1$), with the price scale limited to 600€/MWh for readability. As one can see, both models behave similarly as long as $C_t^{max} - D_t$ is large, but the simpler model is unable to produce large price spikes in periods of tension (i.e. when $C_t^{max} - D_t$ is low). This lack is corrected by the present model (3.2.2).

Remark 3.2.1. The estimation problem studied in this Section (see Figure 3.2.1a) is related to the problem of estimating the electricity supply curve. The usual approach is to propose a parametric modeling of this curve (or of a variant of this curve) and to estimate the corresponding parameters via such methods as maximum likelihood ([11, 35, 43]), or non-linear regression ([84, 68]). In [93], a semi-parametric estimation is performed using local quadratic regression. In our case, the supply curve considered is corrected from the marginal fuel price (as in [93]),

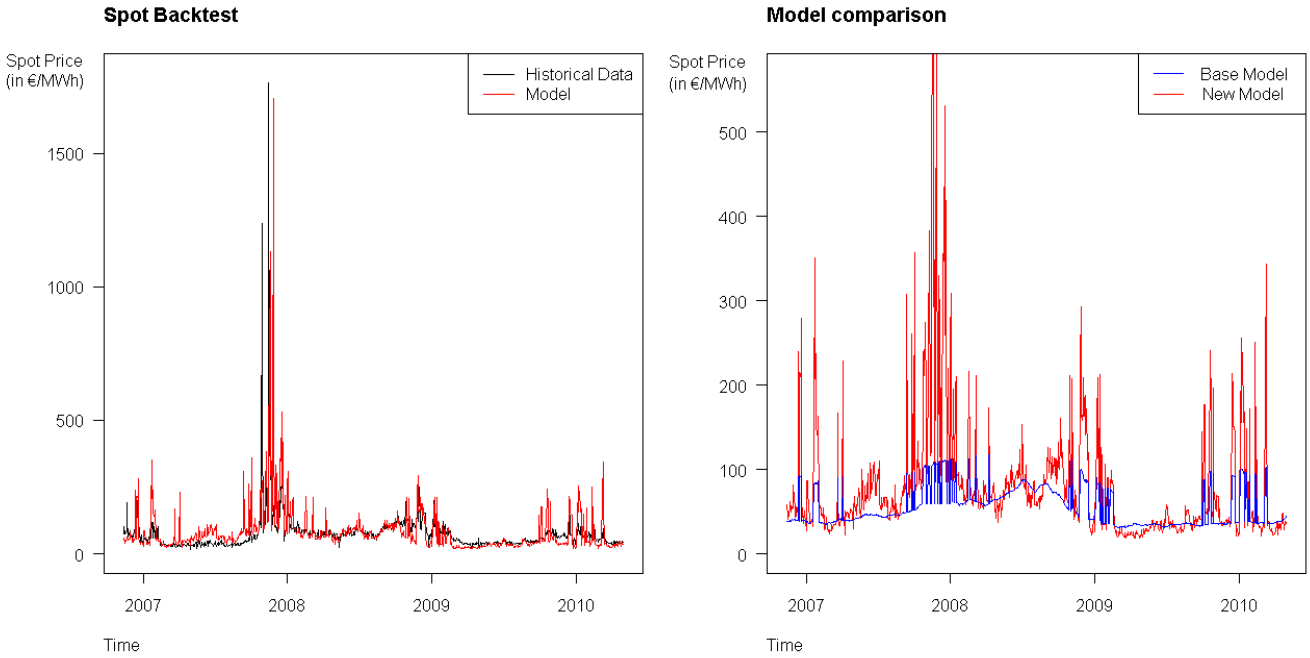

(a) Spot model vs. Historical prices

(b) Spot model vs. Marginal fuel price model

Figure 3.2.2: Backtests

the input of the correction is the difference $C_t - D_t$ (unlike the more common choices of D_t or D_t/C_t), and the distinctive feature of our estimation procedure is that it is performed on the quantiles of the inputs, which preserves the information provided by the peaks in the data, and allow for a simple estimation using one simple linear regression. Consequently, we did not investigate for other estimation techniques, as we did not need it. However, it cannot be excluded that more involved estimation techniques may be required, for instance, to apply our model to other specific electricity markets. For the time being, we leave these empirical investigations and extensions for further research.

3.3 Pricing and hedging

In this section, we use the model (3.2.2) to derive pricing and hedging formulae for power derivatives, including forward contracts. We first model the different processes involved in the equation (3.2.2), namely the capacity processes C_t^i, the electricity demand D_t and the fuel prices S_t^i, $1 \leq i \leq n$. The fuels S^i are clearly tradable, but demand and capacities D and C^i are not, which means that we are going to work in an incomplete market setting. Consequently, a perfect hedging will not be possible, and, equivalently, the market will have infinitely many Equivalent (local) Martingale Measures (henceforth EMM). Many criteria for the choice of a 'good' EMM are available in the literature. In this chapter, we use the local risk minimization criterion introduced by [52], which is based on a financially meaningful decomposition of contingent claims into hedgeable and non-hedgeable parts, and gives rise to explicit price formulae as we will see next. Investigation of other criteria such as, e.g., utility indifference pricing or mean-variance hedging, and their comparison are left for future research (one can refer to [26] and [60] for applications of these other criteria to electricity markets).

3.3.1 Model for capacity, demand and fuel prices

Let $(\Omega, \mathbb{P}, \mathcal{F})$ be a given probability space, where $\mathbb{P}$ is the historical (or statistical) probability measure. $\mathbb{E}$ will denote the expectation operator taken with respect to $\mathbb{P}$. All the subsequent processes, namely C, D and S, will be defined on this probability space. The market filtration $\mathcal{F}_t$ will be the natural filtration generated by all Brownian motions driving the dynamics of all such processes. We assume from the beginning that the spot interest rate r is a positive constant. For the sake of simplicity, we set storage cost and convenience yield of every fuel equal to zero.

Market of fuels. We recall from Section (3.2) that for $1 \leq i \leq n$, S_t^i denotes the price of fuel i at time t, that h_i denotes the corresponding heat rate, such that $h_i S_t^i$ is expressed in units of currency per MWh, and that we assume a fixed order $h_1 S_t^1 \leq \cdots \leq h_n S_t^n$ between fuels. In order to enforce this assumption, we model the dynamics of the fuels spreads $Y_t^i := h_i S_t^i - h_{i-1} S_t^{i-1}$ (with the convention $S^0 \equiv 0$) rather than the dynamics of the fuels directly. We choose $Y = (Y^1, \ldots, Y^n)$ to be a vector of independent geometric Brownian motions under the statistical measure $\mathbb{P}$, i.e. for $1 \leq i \leq n$:

$$dY_t^i = \mu_i Y_t^i dt + \sigma_i Y_t^i dW_t^i, \quad Y_0^i > 0,$$

where $W = (W^1, \ldots W^n)$ is a n-dimensional Brownian motion for the measure $\mathbb{P}$ and μ_i, σ_i $(1 \leq i \leq n)$ are fixed real numbers with $\sigma_i > 0$. $\mathcal{F}^W = (\mathcal{F}_t^W)$ denotes the natural (and $\mathbb{P}$-saturated) filtration generated by W. Note that $h_i S^i = \sum_{j \leq i} Y^j$ and, as a consequence of the positivity of geometric Brownian motions, the condition $h_i S^1 < \cdots < h_n S^n$ is satisfied.

The $\mathbb{P}$-dynamics of S^i, $1 \leq i \leq n$, is then given by:

$$dS_t^i = \frac{1}{h_i} \sum_{j \leq i} (\mu_j Y_t^j dt + \sigma_j Y_t^j dW_t^j). \tag{3.3.1}$$

Since each S_t^i is a.s. positive, we can rewrite its $\mathbb{P}$-dynamics in the following way:

$$\frac{dS_t^i}{S_t^i} = \mu_t^{S,i} dt + \sigma_t^{S,i} dW_t^{S,i}$$

where the drift $\mu_t^{S,i}$, the volatility $\sigma_t^{S,i}$ and the Brownian motion $W^{S,i}$ are given by:

$$\mu_t^{S,i} = \sum_{j \leq i} \frac{Y_t^j}{h_i S_t^i} \mu_j, \qquad \sigma_t^{S,i} = \sqrt{\sum_{j \leq i} \left(\frac{Y_t^j}{h_i S_t^i} \right)^2 \sigma_j^2}, \qquad dW_t^{S,i} = \frac{1}{\sigma_t^{S,i}} \sum_{j \leq i} \frac{Y_t^j}{h_i S_t^i} \sigma_j dW_t^j$$

To verify that $W^{S,i}$ is a Brownian motion it suffices to check that $\left\langle W^{S,i} \right\rangle_t = t$ and use Lévy's characterization theorem to conclude. Notice that $\mathcal{F}^W = \mathcal{F}^{W^S}$, where $\mathcal{F}^{W^S}$ denotes the natural filtration generated by the n-dimensional process $W^S := \left(W^{S,1}, \ldots, W^{S,n} \right)$.

Electricity demand and production capacities. The demand for electricity is modeled by a process D, adapted to the natural filtration $\mathcal{F}_t^D$ generated by a Brownian motion W^D under $\mathbb{P}$. Similarly, the production capacities from each fuel are modeled by processes C^i, $1 \leq i \leq n$, adapted to the natural filtration $\mathcal{F}_t^C$ generated by an n-dimensional Brownian motion $W^C = (W^{C,1}, \ldots, W^{C,n})$ under $\mathbb{P}$. We assume the following dynamics:

$$dD_t = a\left(t, D_t\right) dt + b\left(t, D_t\right) dW_t^D \tag{3.3.2}$$

$$dC_t^i = \alpha_i \left(t, C_t^i\right) dt + \beta_i \left(t, C_t^i\right) dW_t^{C,i} \tag{3.3.3}$$

where $a, b, \alpha_i, \beta_i : \mathbb{R}_+ \times \mathbb{R} \mapsto \mathbb{R}$ are measurable functions such that the SDE (3.3.2) and (3.3.3) admit unique strong solutions on, respectively, $\mathcal{F}^D$ and on the natural filtration generated by $W^{C,i}$.[1] We make the following standing assumption.

Assumption 1. *We assume that the Brownian motions W, W^C and W^D are mutually independent under the true probability $\mathbb{P}$. Moreover, the market filtration, denoted $(\mathcal{F}_t)$, is the natural filtration — satisfying usual conditions — generated by all these Brownian motions, i.e. $\mathcal{F}_t = \mathcal{F}_t^S \vee \mathcal{F}_t^C \vee \mathcal{F}_t^D$.*

3.3.2 Choice of pricing measure

3.3.2.1 Some preliminaries on local risk minimization

We recall some basic facts on local risk minimization (henceforth LRM) approach for pricing and hedging in incomplete markets. This approach has been introduced by [52]. We will essentially follow the two survey papers by [91] and [97]. All the processes we will introduce in this section refer to a given filtration $(\mathcal{F}_t)$ satisfying usual conditions and representing the market's information flow.

Let X be a discounted continuous price process, i.e. X is an adapted continuous $\mathbb{R}^d$-valued semimartingale with decomposition $X = X_0 + M + A$, where X_0 is a constant in $\mathbb{R}^d$, M is a local martingale and A a finite variation process such that $M_0 = A_0 = 0$. We assume that there exists a square integrable EMM $\mathbb{Q}$ for X, i.e. X is a local martingale under $\mathbb{Q}$ with $d\mathbb{Q}/d\mathbb{P} \in L^2(\mathbb{P})$. It is well known that under such an assumption, the finite variation part A is absolutely continuous with respect to M's quadratic variation, i.e. it satisfies

$$A_t = \int_0^t d\langle M \rangle_s \lambda_s, \quad t \in [0, T]$$

for some predictable $\mathbb{R}^n$-valued process λ. A portfolio strategy is a pair $\varphi = (V, \vartheta)$ where V is a real-valued adapted process such that $V_T \in L^2(\mathbb{P})$ and ϑ is a predictable, $\mathbb{R}^d$-valued, X-integrable process such that $\int_0^T \vartheta_t dX_t \in L^2(\mathbb{P})$ and $\int \vartheta dX$ is a Q-martingale for all $Q \in \mathcal{M}_2^e$, the set of all $\mathbb{P}$-equivalent probability measures with square integrable derivative and making X a local martingale. The set of all such strategies ϑ will be denoted by Θ.

We now associate to each portfolio strategy $\varphi = (V, \vartheta)$ a process, which will be very useful in the sequel in describing the main features of the LRM approach: the *cost process* $\mathrm{Cost}(\varphi)$. The cost process of a portfolio strategy $\varphi = (V, \vartheta)$ is defined by:

$$\mathrm{Cost}_t(\varphi) = V_t - \int_0^t \vartheta_u dX_u, \quad t \in [0, T]$$

A portfolio strategy φ is called *self-financing* if its cost process $\mathrm{Cost}(\varphi)$ is constant $\mathbb{P}$ a.s.. It is called *mean self-financing* if $\mathrm{Cost}(\varphi)$ is a martingale under $\mathbb{P}$. Let H be a square-integrable, $\mathcal{F}_T$-measurable contingent claim. We say that a portfolio strategy $\varphi = (V, \vartheta)$ is *H-admissible* if $V_T = H$, $\mathbb{P}$ a.s.. Therefore, an H-admissible portfolio strategy φ is called *locally risk minimizing* (henceforth LRM-strategy) if the corresponding cost process $\mathrm{Cost}(\varphi)$ belongs to $\mathbb{H}^2(\mathbb{P})$ (the Banach space of all $\mathbb{P}$-martingales bounded in $L^2(\mathbb{P})$, equipped with the norm $\|M\|^2 = \mathbb{E}[\sup_t |M_t|]^2$) and is orthogonal to X under $\mathbb{P}$. There exists a LRM-strategy if and only if H admits a decomposition:

$$H = H_0 + \int_0^T \vartheta_t^H dX_t + L_T^H, \quad \mathbb{P} \text{ a.s.}, \tag{3.3.4}$$

[1] See, e.g., [69] or [94] for standard assumptions ensuring such properties.

where H_0 is a constant, $\vartheta^H \in \Theta$ and $L^H \in \mathbb{H}^2(\mathbb{P})$ is orthogonal to X. Such a decomposition is called the *Föllmer-Schweizer decomposition* of H under $\mathbb{P}$, and the portfolio strategy $\varphi = (V, \vartheta^H)$ with

$$V_t = H_0 + \int_0^t \vartheta_s^H dX_s + L_t^H, \quad \mathbb{P} \text{ a.s.}, \quad t \in [0,T].$$

is a LRM-strategy for X. There is a very useful characterization of the LRM-strategy by means of the Galtchouk-Kunita-Watanabe decomposition (henceforth GKW-decomposition) of H under a suitable EMM, namely the *minimal EMM* introduced by [52]. We recall now some basic facts about this measure and its deep relation with the LRM approach. We denote by $\widehat{Z}$ the *minimal martingale density:*

$$\widehat{Z}_t = \mathcal{E}\left(-\int \lambda dM\right)_t, \quad t \in [0,T]. \tag{3.3.5}$$

Since [46] we know that the existence of a Föllmer-Schweizer decomposition (and so of a unique LRM-strategy) for every $H \in L^2(\mathbb{P}, \mathcal{F}_t)$, for any $t \in [0,T]$ is equivalent to assuming an additional integrability condition on $\widehat{Z}$, which is usually called $R_2\,(\mathbb{P})$ (see [46] for details). Such a condition will be verified in our model. Moreover, under such a condition, we can define on $\mathcal{F}_t$, for all $t \in [0,T]$, an EMM $\widehat{\mathbb{Q}}$ for X, given by:

$$\left.\frac{d\widehat{\mathbb{Q}}}{d\mathbb{P}}\right|_{\mathcal{F}_t} = \widehat{Z}_t$$

which is called *minimal EMM* for X. We will denote $\widehat{\mathbb{E}}$ the expectation operator under $\widehat{\mathbb{Q}}$. We now quote without proof (for which we refer to [52], Theorem 3.14, p. 403) the following fundamental result relating the minimal EMM and the LRM-strategy:

Theorem 3.3.1. *Let H be a contingent claim in $L^2(\mathbb{P}, \mathcal{F}_T)$. The LRM-strategy $\widehat{\varphi}$, hence also the corresponding Föllmer-Schweizer decomposition (3.3.4), is uniquely determined. It can be computed in terms of the minimal EMM $\widehat{\mathbb{Q}}$: if $\widehat{V}_t^H$, $t \in [0,T]$, denotes a right-continuous version of the $\widehat{\mathbb{Q}}$-martingale $\widehat{\mathbb{E}}[H|\mathcal{F}_t]$, $t \in [0,T]$, with GKW-decomposition:*

$$\widehat{V}_t^H = \widehat{V}_0^H + \int_0^t \widehat{\vartheta}_s^H dX_s + \widehat{L}_t^H, \quad t \in [0,T],$$

then the portfolio strategy $\widehat{\varphi}^H = (\widehat{V}^H, \widehat{\vartheta}^H)$ is the LRM-strategy for H and its cost process is given by $\mathrm{Cost}(\widehat{\varphi}) = \widehat{\mathbb{E}}[H] + \widehat{L}^H$.

This theorem gives us a practical way for computing the LRM-strategy of a given contingent claim H. Indeed, in order to find its LRM-strategy and the associated cost process, we need to compute only its GKW-decomposition and identify the integral part and the orthogonal part. The integral part represents the hedgeable part of the claim H, while the orthogonal part L^H represents the unhedgeable part or residual risk. The expectation of H under the minimal EMM $\widehat{\mathbb{Q}}$, $\widehat{\mathbb{E}}[H]$, is clearly one of the infinitely many no-arbitrage prices of H and it can be also viewed as the initial wealth allowing to hedge the hedgeable part of H. Moreover, in [66], it is shown that in a large class of diffusion market models with non-tradable assets, the expectation of a contingent claim H under the minimal EMM, i.e. $\widehat{\mathbb{E}}[H]$, is an upper bound for bid utility indifference prices of H. In the next sections, we will use this approach to price and (partially) hedge power derivatives.

3.3.2.2 LRM in our energy market model

We have already noticed that our market is incomplete. Thus it admits infinitely many EMMs. Here is a complete description of such measures.

Consider first the submarket composed of fuels only. Since each Y^i admits a unique EMM $\mathbb{Q}^i$ and their corresponding Radon-Nikodym derivatives are independent, it is easy to see that the measure $\widehat{\mathbb{Q}}$ defined as the product of all the $\mathbb{Q}^i$'s is an EMM for S. More precisely,

$$\frac{d\widehat{\mathbb{Q}}}{d\mathbb{P}} = \prod_{i=1}^{n} \frac{d\mathbb{Q}^i}{d\mathbb{P}} = \prod_{i=1}^{n} e^{-\lambda_i W_T^i - \frac{\lambda_i^2}{2}T}, \quad \text{on } \mathcal{F}_T \tag{3.3.6}$$

where $\lambda_i = \frac{\mu_i - r}{\sigma_i}$ is the market price of risk of the i-th spread Y_t^i, i.e. the spread between fuels $i-1$ and i. Thus, the assumption of non arbitrage is satisfied in the market of fuels (indeed, one can easily prove that each $\widetilde{Y}^i$ is a $\mathbb{Q}^i$-martingale if and only if each $\widetilde{S}^i$ is a $\widehat{\mathbb{Q}}$-martingale, where $\widetilde{\cdot}$ denotes discounting). Let $\widehat{W} = (\widehat{W}^1, \dots, \widehat{W}^n)$ be the $\widehat{\mathbb{Q}}$-BM defined via Girsanov's theorem as:

$$\widehat{W}_t^i = W_t^i + \lambda_i t. \tag{3.3.7}$$

Finally, one can show that the $\widehat{\mathbb{Q}}$-dynamics of the i-th fuel, $1 \le i \le n$, is given by:

$$dS_t^i = \frac{1}{h_i} \sum_{j \le i} Y_t^j (r dt + \sigma_j d\widehat{W}_t^j) = r S_t^i dt + \frac{1}{h_i} \sum_{j \le i} Y_t^j \sigma_j d\widehat{W}_t^j \tag{3.3.8}$$

and that of its forward price $F_t^i(T)$ with any maturity $T > 0$ is:

$$dF_t^i(T) = e^{r(T-t)} \frac{1}{h_i} \sum_{j \le i} Y_t^j \sigma_j d\widehat{W}_t^j$$

In order to generate the whole family of EMMs, we need to consider also possible changes of measure for the demand D (equivalently its driving Brownian motion W^D) and all the capacities C^i (equivalently their driving multivariate Brownian motion $W^{C,i}$). Thanks to the mutual independence between fuels, capacities and electricity demand, it is not difficult to obtain the following result. Let $T > 0$ be a given finite horizon, e.g. the maturity of a forward contract on electricity. The next proposition gives a full characterization of the (square-integrable) EMMs of our model.

Proposition 3.3.1. *The set of all EMMs $\mathcal{M}_e$ of our model over the time horizon $[0, T]$ is given by all $\mathcal{F}_T$-measurable random variables Z_T such that there exists adapted processes $\eta = (\eta^C, \eta^D)$ verifying:*

$$Z_T = \left(\prod_{i=1}^{n} \frac{d\mathbb{Q}^i}{d\mathbb{P}} \right) \mathcal{E}_T \left(\int_0^{\cdot} \eta_s^C dW_s^C \right) \mathcal{E}_T \left(\int_0^{\cdot} \eta_s^D dW_s^D \right) \tag{3.3.9}$$

and such that $\mathbb{E}_{\mathbb{P}}[Z_T] = 1$, where $\mathcal{E}_T(\cdot)$ denotes the stochastic exponential at time T and $\mathbb{Q}^i$ is the unique EMM for the i-th spread Y^i, for $i = 1, \dots, n$, as in (3.3.6).

Proof. Use the mutual independence of (W, W^C, W^D) and the representation theorem for Brownian martingales. □

In view of (3.3.5) and (3.3.9), it is easy to see that in our model the Föllmer and Schweizer minimal EMM corresponds to the case when $\eta^C \equiv \eta^D \equiv 0$ in (3.3.9).[2] Thus, consistently with our notation, the minimal EMM is exactly the measure $\widehat{\mathbb{Q}}$ previously introduced.

[2] Indeed, in our model, the Radon-Nikodym derivative $\widehat{Z}_t = \mathbb{E}\left[d\widehat{Q}/d\mathbb{P} \,|\mathcal{F}_t \right]$ clearly satisfies the so-called integrability condition $R_2(\mathbb{P})$, so that Theorem C in [46] can be applied.

3.3.3 Electricity futures

As a first important application of the LRM approach, we derive the price dynamics of a forward contract on electricity with instantaneous delivery at a given time T under our model (3.2.2). Such a dynamics will be of a great importance to obtain hedging strategies for energy derivatives. To do so, we will apply LRM approach for pricing and hedging a futures contract on electricity via trading on fuels. Then, from these results, we will be able to study the risk premium of electricity. Forward contracts on electricity will be used as additional hedging instruments in the next subsections on hedging and pricing of more complex energy derivatives.

3.3.3.1 Price and dynamics

We now derive the price and dynamics of electricity forward contracts. Using the notation of the previous section, the contingent claim to hedge has terminal payoff $H = P_T$ and the hedge is performed by trading in the fuel process S, i.e. $X = S$, or equivalently in the fuel futures.[3] Recall from Section 3.3.1 that we assume a constant spot interest rate r (futures and forwards are thus identical, and we use both terms interchangeably), as well as independence between fuels, demand and capacities. In this setting, and recalling our spot model (3.2.2), it is not difficult to see that the formula relating the electricity forward price $F_t^e(T)$ and the forward prices of fuels $F_t^i(T)$, both with instantaneous delivery period, is given by:

$$F_t^e(T) = \sum_{i=1}^{n} h_i F_t^i(T) \mathbb{E}\left[g(C_T^{max} - D_T)\mathbf{1}_{\{D_T \in I_T^i\}} | \mathcal{F}_t^{D,C}\right] \tag{3.3.10}$$

for $t \in [0, T]$, where $\mathcal{F}_t^{D,C} := \mathcal{F}_t^D \vee \mathcal{F}_t^C$ is the natural filtration generated by both W^D and W^C. Recall that, by the definition of minimal EMM, one has $\widehat{\mathbb{Q}} = \mathbb{P}$ on $\mathcal{F}^{D,C}$.

Remark 3.3.1. The previous formula is key in our approach. A similar formula is obtained in the previous paper [3] for a slightly different model, the arguments used to prove it being the same. The idea behind it is that, unlike usual approaches in energy market models, in our model the use of a risk-neutral measure is motivated by embedding the energy market into the larger market including the possibility of trading in fuels. Then, since in the latter market one may in principle trade on fuels, the risk-neutral approach that proved to be successful in stock markets can be applied. The final step is that, since a forward contract on energy can be viewed as an option whose payoff is exactly the spot price at maturity and the latter is a function of fuels (via the relation (3.2.2)), such a forward contract can be priced taking expectation under a risk-neutral measure for fuels. For more details, we refer to [3]. The price to pay is, in some sense, that the production function linking fuels and energy contains other non-tradable factors as well, e.g. demand and capacities.

Remark 3.3.2. The same kind of factorization as in formula (3.3.10) can be obtained under any EMM $\mathbb{Q}$ such that fuels S are independent from demand and capacities (D, C). For instance, when $\mathbb{Q}$ is an EMM with deterministic (time-dependent) market prices of fuel, demand and capacity risks, i.e. deterministic η^D, η^C as in Proposition 3.3.1, the same computations lead to the same formula. Now, if one take an arbitrary EMM $\mathbb{Q}$ (under which fuels, demand and capacities are not necessarily independent), one may obtain an analogue formula for electricity futures price via an additional probability change. Indeed, set $d\widehat{\mathbb{Q}}^i/d\widehat{\mathbb{Q}} = e^{-rT} S_T^i / S_0^i$ on $\mathcal{F}_T$ for all $i = 1, \ldots, n$. Thus we get:

$$F_t^e(T) = \widehat{\mathbb{E}}[P_T | \mathcal{F}_t] = \sum_{i=1}^{n} h_i F_t^i(T) \widehat{\mathbb{E}}^i \left[g(C_T^{max} - D_T)\mathbf{1}_{\{D_T \in I_T^i\}} | \mathcal{F}_t^{D,C}\right]$$

[3] Indeed, one may easily switch from $F_t^i(T)$ to S_t^i via the well-known formula $F_t^i(T) = e^{r(T-t)} S_t^i$. Recall that we assumed constant interest rate r and zero convenience yield and storage cost, for any fuel i.

where $\widehat{\mathbb{E}}^i$ denotes expectation under $\widehat{\mathbb{Q}}^i$. Of course, here the difficulty would be to compute the weights multiplying the futures on fuels. The laws of D and C under each $\widehat{\mathbb{Q}}^i$ can of course be obtained through Girsanov's theorem, but their parameters might depend of S, C and D in a complicated way, and obtaining closed-form formulae may be very difficult.

The next step consists in evaluating the conditional expectation appearing in formula (3.3.10). Since the process (C, D) is Markov, we have:

$$\mathbb{E}\left[g\left(C_T^{max} - D_T\right)\mathbf{1}_{\{D_T \in I_T^i\}} | \mathcal{F}_t^{D,C}\right] = G_i^T\left(t, C_t, D_t\right)$$

for some real-valued measurable function $G_i^T(t, c, d)$ defined on $[0, T] \times \mathbb{R}^n \times \mathbb{R}$. We will call this function the *Conditional Expectation of Scarcity Function* (henceforth CES function). Under specific dynamics for C_t and D_t, the CES function, as well as its partial derivatives, can be computed explicitly. This will be the purpose of Section 3.4.2. Our key relation (3.3.10) between electricity forwards and fuel forwards now simply reads:

$$F_t^e(T) = \sum_{i=1}^n h_i G_i^T\left(t, C_t, D_t\right) F_t^i(T) \tag{3.3.11}$$

meaning that an electricity forward can be seen as a basket of fuels forwards, with stochastic weights given by the CES function, driven by electricity demand and production capacities. Note that this relation does not depend on the specific model chosen for fuels in Section 3.3.1, except for the assumption that fuels are independent of capacities and demand.

We now derive the dynamics of electricity forwards. Assume that $G_i^T \in \mathcal{C}^{1,2,2}([0, T] \times \mathbb{R}^n \times \mathbb{R})$.[4] Then Itō's lemma provides the dynamics of $G_i^T(t, C_t, D_t)$ as follows:

$$dG_i^T(t, C_t, D_t) = \sum_{k=1}^n \frac{\partial G_i^T}{\partial c_k}(t, C_t, D_t)\beta_k(t, C_t^k)dW_t^{C,k} + \frac{\partial G_i^T}{\partial d}(t, C_t, D_t)b(t, D_t)dW_t^D. \tag{3.3.12}$$

A simple application of integration-by-parts formula gives the following $\widehat{\mathbb{Q}}$-dynamics for electricity T-forward prices :

$$\begin{aligned} dF_t^e(T) &= \sum_{i=1}^n h_i\left[G_i^T(t, C_t, D_t)dF_t^i(T) + F_t^i(T)dG_i^T(t, C_t, D_t)\right] \\ &= e^{r(T-t)}\sum_{i=1}^n\left(\sum_{k=i}^n G_k^T(t, C_t, D_t)\right)\sigma_i Y_t^i d\widehat{W}_t^i + \sum_{i=1}^n h_i F_t^i(T)\frac{\partial G_i^T}{\partial d}(t, C_t, D_t)b(t, D_t)dW_t^D \\ &\quad + \sum_{i=1}^n h_i F_t^i(T)\sum_{k=1}^n \frac{\partial G_i^T}{\partial c_k}(t, C_t, D_t)\beta_k(t, C_t^k)dW_t^{C,k} \end{aligned} \tag{3.3.13}$$

where recall from Section 3.3.2.2 that $\widehat{W}_t^i$ is a $\widehat{\mathbb{Q}}$-BM. Notice that the quadratic covariation between $F_t^i(T)$ and $G_i(t, C_t, D_t)$ vanishes, due once more to the independence between S^i and (C, D). From equation (3.3.7), one can deduce the $\mathbb{P}$-dynamics of $F^e(T)$:

$$\begin{aligned} dF_t^e(T) &= e^{r(T-t)}\sum_{i=1}^n\left(\sum_{k=i}^n G_k^T(t, C_t, D_t)\right)\sigma_i Y_t^i\left(dW_t^i + \lambda_i dt\right) \\ &\quad + \sum_{i=1}^n h_i F_t^i(T)\frac{\partial G_i^T}{\partial d}(t, C_t, D_t)b(t, D_t)dW_t^D \\ &\quad + \sum_{i=1}^n h_i F_t^i(T)\sum_{k=1}^n \frac{\partial G_i^T}{\partial c_k}(t, C_t, D_t)\beta_k(t, C_t^k)dW_t^{C,k} \end{aligned} \tag{3.3.14}$$

[4] Whether such an assumption is verified or not will depend on the regularity of the coefficients of the SDEs governing the dynamics of D and C. Notice that the dynamics that we will use in Section 3.4 are such that the CES functions $G_i^T(t, c, z)$ are smooth enough to apply Itō's formula (see also the explicit formulae for the derivatives of CES functions in 3.4.2.7)

3.3.3.2 Risk premium

Using the previous results, we are able to study the electricity risk premium $\pi^e(t,T)$, defined as:

$$\pi^e(t,T) := F_t^e(T) - \mathbb{E}_{\mathbb{P}}[P_T|\mathcal{F}_t], \quad t \leq T. \tag{3.3.15}$$

Just like the relation (3.3.11) between futures prices, the electricity risk premium can be expressed as a weighted linear combination of fuel risk premiums. Let $\pi^i(t,T)$ denote the risk premium of the i-th fuel, i.e. $\pi^i(t,T) := F_t^i(T) - \mathbb{E}_{\mathbb{P}}[S_T^i|\mathcal{F}_t]$ for $t \leq T$. Recall that forward prices are computed under the minimal EMM $\widehat{\mathbb{Q}}$.

Proposition 3.3.2. *Under our model assumptions, one has:*

$$\pi^e(t,T) = \sum_{i=1}^{n} h_i G_i(t, C_t, D_t)\pi_t^i(t,T), \quad t \in [0,T]. \tag{3.3.16}$$

Proof. It follows from formulae (3.2.2) and (3.3.11), and the fact that the law of the processes (C,D) under $\widehat{\mathbb{Q}}$ is the same as under $\mathbb{P}$. □

Remark 3.3.3. An easy consequence of the previous equality is that if all the fuels are in normal backwardation (or in contango), then it also holds for electricity. To our knowledge, no studies have been performed to assess this particular point. Nevertheless, in [81] a systematic analysis of commodity term structures is done. The authors show that energy prices (crude, heating oil, gasoil, gas) are at least 40% of the time in contango. Hence, it is not unreasonable to imagine that situations where they are all in contango at the same time may arise.

3.3.3.3 Electricity futures as hedging instruments

In the next sections, we will use these electricity forward contracts as tradable hedging instruments to improve the hedging of more complex derivatives on electricity spot and forwards. In other terms, we will consider an enlarged market $(S, F^e(T^*))$ where agents can trade on fuels as well as on a forward contract with a given maturity T^*. While the minimal EMM for the market of fuels S is $\widehat{\mathbb{Q}}$, it is not garanteed a priori that the minimal EMM for the richer market $(S, F^e(T^*))$ is still $\widehat{\mathbb{Q}}$. It will depend on the $\mathbb{P}$-dynamics of the forward contract $F^e(T^*)$. We conclude this part of the chapter by showing that if the $\mathbb{P}$-dynamics of $F^e(T^*)$ is given exactly by (3.3.14), then the minimal EMM for $(S, F^e(T^*))$ is given by $\widehat{\mathbb{Q}}$.

Proposition 3.3.3. *Let T^* be any positive finite maturity. Assume that the $\mathbb{P}$-dynamics of S and $F^e(T^*)$ are given by, respectively, (3.3.1) and (3.3.14). Then, the minimal EMM for $(S, F^e(T^*))$ is given by $\widehat{\mathbb{Q}}$.*

Proof. By the definition of the minimal EMM (see Definition (3.2) in [52]), we have to verify that any square-integrable $\mathbb{P}$-martingale M that is orthogonal to both S and $F^e(T^*)$, must be a $\widehat{\mathbb{Q}}$-martingale. By the representation theorem of Brownian martingales and since $\mathcal{F}_t^W = \mathcal{F}_t^{W^S}$ for any t, such a $\mathbb{P}$-martingale M satisfies

$$M_t = M_0 + \int_0^t \alpha_s dW_s^S + \int_0^t \beta_s dW_s^C + \int_0^t \gamma_s dW_s^D$$

for some predictable processes α, β, γ. Being M orthogonal to S, i.e. $\langle M, W^S \rangle = 0$, we have $\alpha \equiv 0$. Moreover, M is also orthogonal to $F^e(T^*)$, which implies that

$$\langle M, F^e(T^*) \rangle_t = \int_0^t \beta_s \theta_s^C ds + \int_0^t \gamma_s \theta_s^D ds = 0$$

for all t, where $\theta^{C,k}$ (for $k = 1, \dots, n$) and θ^D are the integrads in, respectively, the $dW^{C,k}$ part and the dW^D part appearing in the dynamics (3.3.14) of forward prices. Thus, $\beta_t\theta_t^C + \gamma_t\theta_t^D \equiv 0$ for all t. As a consequence, since W^C and W^D are Brownian motions under $\widehat{\mathbb{Q}}$, M is a $\widehat{\mathbb{Q}}$-martingale. By uniqueness of the minimal EMM, we can conclude. □

3.3.4 Pricing formulae

In this section we compute the price of energy derivatives via LRM approach. It consists in computing the expectation of the terminal pay-off under the minimal EMM $\widehat{\mathbb{Q}}$. Such an expectation represents the initial wealth allowing to approximately replicate a given option in a local risk minimization sense, as explained in Section 3.3.2.1.

In what follows, we will focus on options on spreads (between electricity and fuels) as well as on options on electricity forward contracts. We will see in particular that any European options on an electricity forward contract can be viewed as a basket option on fuels with random (but independent) coefficients. Thus, numerical methods developed to price basket options on securities can be applied to evaluate energy options as well. Finally, we will show how to obtain explicit formulae in the case of two fuels, i.e. $n = 2$.

We will use the notation $BS_t(\sigma, K)$ for the Black-Scholes formula of the t-price of a European call option with volatility σ and strike K. As the other parameters (like maturity or interest rate) are fixed, they will not appear in the notation. Finally, $f_X(\cdot)$ (resp. $\hat{f}_X(\cdot)$) will denote the density at time T of a process X under the statistical measure $\mathbb{P}$ (resp. under the minimal EMM $\widehat{\mathbb{Q}}$).

3.3.4.1 Spread options

Let us consider a spread option with maturity T between electricity and a fuel j chosen among the n fuels used to produce electricity. Then, the corresponding pay-off is given by:

$$H := \varphi(P_T - h_j S_T^j), \tag{3.3.17}$$

where φ is a real-valued function such that $H \in L^2(\widehat{\mathbb{Q}})$, e.g. $\varphi(x) = (x - K)^+$ for $K \geq 0$. For the sake of simplicity, we compute the price of this option at time $t = 0$. The price at any time t can be easily deduced from that case by using the Markov property of the price processes. Using equation (3.2.2), one obtain:

$$\pi_0 := e^{-rT}\widehat{\mathbb{E}}[\varphi(P_T - h_j S_T^j)] = e^{-rT}\sum_{i=1}^{n}\widehat{\mathbb{E}}\left[\varphi\left(g(C_T^{max} - D_T)h_i S_T^i - h_j S_T^j\right)\mathbf{1}_{\{D_T \in I_T^i\}}\right].$$

Now, consider the case of two fuels ($n = 2$) and let $j = 1$. The other case $j = 2$ can be treated similarly. Recall from equation (3.2.1) that in the two fuels case, the intervals I_T^i are $I_T^1 = (-\infty, C_T^1)$ and $I_T^2 = [C_T^1, +\infty)$. Using the independence between fuels S, demand D and capacities C, and the fact that $\widehat{\mathbb{Q}} = \mathbb{P}$ on $\mathcal{F}_T^{C,D}$, we have:

$$\pi_0 = \int_{\mathbb{R}^2} f_{C_T^1 - D_T}(z) f_{C_T^2}(c)\left\{\phi_1(c,z)\mathbf{1}_{\{z>0\}} + \phi_2(c,z)\mathbf{1}_{\{z\leq 0\}}\right\} dcdz, \tag{3.3.18}$$

where $f_{C_T^1 - D_T}(.)$ and $f_{C_T^2}(.)$ are the $\mathbb{P}$-densities of, respectively, $C_T^1 - D_T$ and C_T^2, while $\phi_1(c,z)$ and $\phi_2(c,z)$ are given by:

$$\begin{aligned}
\phi_1(c,z) &:= e^{-rT}\ \widehat{\mathbb{E}}\left[\varphi\left(\left(g(c+z) - 1\right)Y_T^1\right)\right] \\
\phi_2(c,z) &:= e^{-rT}\ \widehat{\mathbb{E}}\left[\varphi\left(g(c+z)h_2 S_T^2 - h_1 S_T^1\right)\right] \\
&= e^{-rT}\ \widehat{\mathbb{E}}\left[\varphi\left(\left(g(c+z) - 1\right)Y_T^1 + g(c+z)Y_T^2\right)\right].
\end{aligned}$$

We used the fact that $h_1 S^1 = Y^1$ and $h_2 S^2 = Y^1 + Y^2$ (see Section 3.3.1). Recall that, by assumption, Y^1 and Y^2 are independent geometric Brownian motions. We need to compute both terms $\phi_i(c,z)$. Since c, z are fixed, we can simplify the notation by dropping (c,z) in $\phi_i(c,z)$ $(i=1,2)$ and $g(c+z)$. We will simply write ϕ_i and g instead.

To push further our computations, let us consider a call spread option, i.e. $\varphi(x) = (x-K)^+$ for a given strike $K > 0$. In this case, the quantities ϕ_i can be computed more explicitly. Indeed, if $g \leq 1$ one has $\phi_1 = 0$, while on the event $\{g > 1\}$ the quantity ϕ_1 is just $(g-1)$ times the Black-Scholes formula for a European call option with strike $K/(g-1)$ and underlying Y^1, a geometric Brownian motion with volatility σ_1, i.e.

$$\phi_1 = (g-1) BS_0(\sigma_1, \frac{K}{g-1}) \mathbf{1}_{\{g>1\}}.$$

On the other hand, we show that ϕ_2 is a mixture of Black-Scholes formulae with respect to strikes. Recall that $\hat{f}_{Y_T^i}(y_i)$ denotes the log-normal density of Y_T^i under $\widehat{\mathbb{Q}}$ for $i = 1, 2$. We have:

$$\phi_2 = e^{-rT} \widehat{\mathbb{E}} \left[\Big((g-1) Y_T^1 + g Y_T^2 - K \Big)^+ \right]$$

On the event $\{g \leq 1\}$, we have:

$$\phi_2 = g e^{-rT} \widehat{\mathbb{E}} \left[\left(Y_T^2 - \frac{K + (1-g) Y_T^1}{g} \right)^+ \right] = g \int_0^\infty \hat{f}_{Y_T^1}(y) BS_0 \left(\sigma_2, \frac{K + (1-g)y}{g} \right) dy.$$

On the opposite event $\{g > 1\}$, we have:

$$\begin{aligned}
\phi_2 &= g e^{-rT} \widehat{\mathbb{E}} \left[\left(Y_T^2 - \frac{K - (g-1) Y_T^1}{g} \right)^+ \right] \\
&= g \int_0^\infty \hat{f}_{Y_T^1}(y) e^{-rT} \widehat{\mathbb{E}} \left[\left(Y_T^2 - \frac{K - (g-1)y}{g} \right)^+ \right] dy \\
&= g \int_0^{\frac{K}{g-1}} \hat{f}_{Y_T^1}(y) BS_0 \left(\sigma_2, \frac{K - (g-1)y}{g} \right) dy + g \int_{\frac{K}{g-1}}^\infty \hat{f}_{Y_T^1}(y) e^{-rT} \widehat{\mathbb{E}} \left[Y_T^2 - \frac{K - (g-1)y}{g} \right] dy \\
&= g \int_0^{\frac{K}{g-1}} \hat{f}_{Y_T^1}(y) BS_0 \left(\sigma_2, \frac{K - (g-1)y}{g} \right) dy + \left(g Y_0^2 - e^{-rT} K \right) \widehat{\mathbb{Q}} \left(Y_T^1 \geq \frac{K}{g-1} \right) \\
&\quad + (g-1)\, e^{-rT} \widehat{\mathbb{E}} \left[Y_T^1 \mathbf{1}_{\{Y_T^1 > \frac{K}{g-1}\}} \right].
\end{aligned}$$

Observe that:

$$e^{-rT} \widehat{\mathbb{E}} \left[Y_T^1 \mathbf{1}_{\{Y_T^1 > \frac{K}{g-1}\}} \right] = BS_0 \left(\sigma_1, \frac{K}{g-1} \right) + e^{-rT} \frac{K}{g-1} \widehat{\mathbb{Q}} \left[Y_T^1 > \frac{K}{g-1} \right]$$

so that on $\{g > 1\}$ we have:

$$\phi_2 = g \int_0^{\frac{K}{g-1}} \hat{f}_{Y_T^1}(y) BS_0 \left(\sigma_2, \frac{K - (g-1)y}{g} \right) dy + g Y_0^2 \widehat{\mathbb{Q}} \left(Y_T^1 \geq \frac{K}{g-1} \right) + (g-1)\, BS_0 \left(\sigma_1, \frac{K}{g-1} \right).$$

Since, from Sections 3.3.1 and 3.3.2.2, Y^1 is a geometric Brownian motion under $\widehat{\mathbb{Q}}$, with volatility σ_1 and drift r, we have

$$\widehat{\mathbb{Q}} \left(Y_T^1 \geq \frac{K}{g-1} \right) = \mathcal{N} \left(\frac{\left(r - \frac{\sigma_1^2}{2} \right) T - \ln \left(\frac{K}{(g-1) Y_0^1} \right)}{\sigma_1 \sqrt{T}} \right),$$

where $\mathcal{N}$ is the cumulative distribution function of a standard normal random variable. We summarize these results in the following proposition:

Proposition 3.3.4. *Let $n = 2$, i.e. electricity is produced out of two fuels. The price π_0 at time $t = 0$ of a call spread option with pay-off $H = (P_T - h_1 S_T^1 - K)^+$, $K > 0$, is given by the following formula:*

$$\pi_0 = \int_{\mathbb{R}^2} f_{C_T^1 - D_T}(z) f_{C_T^2}(c) \left\{ \phi_1(c,z) \mathbf{1}_{\{z>0\}} + \phi_2(c,z) \mathbf{1}_{\{z \leq 0\}} \right\} dcdz, \tag{3.3.19}$$

where the quantities $\phi_i = \phi_i(c,z)$, $i = 1, 2$, are given by:

$$\begin{aligned}
\phi_1 &= (g-1) BS_0(\sigma_1, \frac{K}{g-1}) \mathbf{1}_{\{g>1\}} \\
\phi_2 &= g \int_0^\infty \hat{f}_{Y_T^1}(y) BS_0 \left(\sigma_2, \frac{K + (1-g)y}{g} \right) \left(\mathbf{1}_{\{g \leq 1\}} + \mathbf{1}_{\{g>1\}} \mathbf{1}_{\{y < \frac{K}{g-1}\}} \right) dy \\
&\quad + \left(g Y_0^2 \mathcal{N} \left(\frac{\left(r - \frac{\sigma_1^2}{2} \right) T - \ln \left(\frac{K}{(g-1) Y_0^1} \right)}{\sigma_1 \sqrt{T}} \right) + (g-1)\, BS_0 \left(\sigma_1, \frac{K}{g-1} \right) \right) \mathbf{1}_{\{g>1\}}
\end{aligned}$$

where we have set $g := g(c+z)$.

3.3.4.2 Options on electricity forwards

Let us consider a contingent claim H with maturity T whose payoff is given by a function $\varphi : \mathbb{R} \mapsto \mathbb{R}$ of a forward contract on electricity with instantaneous delivery period at $T^* > T$, i.e.

$$H := \varphi(F_T^e(T^*))$$

We assume that φ is such that $H \in L^2(\widehat{\mathbb{Q}})$. The next proposition gives a pricing formula for such a contingent claim.

Proposition 3.3.5. *Under the above assumptions, the price at time $t \leq T$ of the contingent claim H is given by:*

$$\widehat{\mathbb{E}}[e^{-r(T-t)} H | \mathcal{F}_t] = \widehat{\mathbb{E}} \left[\psi(t, F_t(T^*), C_T, D_T) | \mathcal{F}_t^{D,C} \right] \tag{3.3.20}$$

where:

$$\psi(t, F_t(T^*), C_T, D_T) := e^{-r(T-t)} \widehat{\mathbb{E}} \left[\varphi \left(\sum_{i=1}^n h_i G_i^{T^*}(T, C_T, D_T) F_T^i(T^*) \right) \Big| \mathcal{F}_t^W \right]. \tag{3.3.21}$$

Proof. It follows from equation (3.3.11), independence between W and (W^D, W^C) and the properties of conditional expectations. □

Remark 3.3.4. The previous pricing formula (3.3.20) provides an easy way of computing prices using basket options pricing algorithms. Indeed, the formula suggests the following procedure:

1. Evaluate first the expectation (3.3.21) with respect to $\mathcal{F}_t^W$, i.e. the function ψ, taking the weights $G_i^{T^*}(T, C_T, D_T)$ as fixed, using a basket options evaluation procedure.

2. Finally, take the average with respect to the weights $G_i^{T^*}(T, C_T, D_T)$.

These two steps can be performed separately thanks to independence between W and (W^D, W^C).

As for spread options, we now look for explicit formulae for European call on electricity forward in the two fuels case, i.e. $n = 2$, at time $t = 0$. We compute the function ψ from equation (3.3.21) in this case. To simplify notation, we set $w_i := e^{-r(T^*-T)}G_i^{T^*}(T, C_T, D_T)$. Recalling that, regarding futures on fuels, $F_T^i(T^*) = e^{r(T^*-T)}S_T^i$ for $i = 1, 2$, we have:

$$\begin{aligned}\psi(0) &= e^{-rT}\widehat{\mathbb{E}}\left[\left(w_1h_1S_T^1 + w_2h_2S_T^2 - K\right)^+\right] = e^{-rT}\widehat{\mathbb{E}}\left[\left((w_1+w_2)Y_T^1 - (K - w_2Y_T^2)\right)^+\right] \\ &= e^{-rT}\widehat{\mathbb{E}}\left[\left((w_1+w_2)Y_T^1 - (K - w_2Y_T^2)\right)^+ \mathbf{1}_{\{Y_T^2 \le K/w_2\}}\right] \\ &\quad + e^{-rT}\widehat{\mathbb{E}}\left[\left((w_1+w_2)Y_T^1 - (K - w_2Y_T^2)\right)^+ \mathbf{1}_{\{Y_T^2 > K/w_2\}}\right] =: A_1 + A_2.\end{aligned}$$

Let us compute separately A_1 and A_2. For A_1, we obtain:

$$\begin{aligned}A_1 &= (w_1+w_2)\int_0^{K/w_2} \hat{f}_{Y_T^2}(y)\widehat{\mathbb{E}}\left[e^{-rT}\left(Y_T^1 - \frac{K - w_2y}{w_1+w_2}\right)^+\right]dy \\ &= (w_1+w_2)\int_0^{K/w_2} \hat{f}_{Y_T^2}(y)BS_0\left(\sigma_1, \frac{K - w_2y}{w_1+w_2}\right)dy.\end{aligned}$$

For A_2, we have:

$$\begin{aligned}A_2 &= e^{-rT}\widehat{\mathbb{E}}\left[\left((w_1+w_2)Y_T^1 - (K - w_2Y_T^2)\right)\mathbf{1}_{\{Y_T^2 > K/w_2\}}\right] \\ &= (w_1+w_2)\widehat{\mathbb{E}}\left[e^{-rT}Y_T^1\mathbf{1}_{\{Y_T^2 > K/w_2\}}\right] + w_2\widehat{\mathbb{E}}\left[e^{-rT}\left(Y_T^2 - \frac{K}{w_2}\right)\mathbf{1}_{\{Y_T^2 > K/w_2\}}\right] \\ &= (w_1+w_2)Y_0^1\widehat{\mathbb{Q}}(Y_T^2 > K/w_2) + w_2BS_0\left(\sigma_2, \frac{K}{w_2}\right)\end{aligned}$$

We summarize our findings in the following proposition.

Proposition 3.3.6. *Consider the two fuels case, i.e. $n = 2$. The price π_0^F at time $t = 0$ of a European call option with maturity T on a T^*-forward contract on electricity with $T^* > T$ is given by the following formula:*

$$\pi_0^F = \int_{\mathbb{R}} f_{D_T}(z)\int_{\mathbb{R}^2} f_{C_T^1}(c_1)f_{C_T^2}(c_2)\psi_0(c_1, c_2, z)dzdc_1dc_2,$$

where the function $\psi_0(c_1, c_2, z)$ is given by:

$$\begin{aligned}\psi_0(c_1, c_2, z) = (w_1+w_2)&\left\{\int_0^{K/w_2} \hat{f}_{Y_T^2}(y)BS_0\left(\sigma_1, \frac{K - w_2y}{w_1+w_2}\right)dy + Y_0^1\mathcal{N}\left(\frac{\left(r - \frac{\sigma_2^2}{2}\right)T - \ln\left(\frac{K}{w_2}\right)}{\sigma_2\sqrt{T}}\right)\right\} \\ &+ w_2BS_0\left(\sigma_2, \frac{K}{w_2}\right)\end{aligned}$$

where $w_i := e^{-r(T^*-T)}G_i^{T^*}(T, c_1, c_2, z)$, $i = 1, 2$.

To make full use of this result, the weights w_i, and thus the CES function $G_i^{T^*}(T, c_1, c_2, z)$, must be computed explicitly. This will be done in Section 3.4.2 under more specific assumptions on dynamics of capacities C and demand D.

3.3.5 Hedging derivatives

Now we turn to hedging. In this subsection, we will identify the hedgeable and the unhedgeable part of any contingent claim written on electricity as well as fuels.

As hedging instruments, we consider forward contracts on electricity and forward contracts on fuels (or, equivalently, spot fuels). We consider a $\widehat{\mathbb{Q}}$-square integrable European-type contingent claim H written on a forward contract on electricity and fuels as well as on capacities and energy demand, i.e.

$$H = \varphi(F_T^e(T^*), F_T(T^*), C_T, D_T)$$

such that $H \in L^2(\widehat{\mathbb{Q}})$. Notice that spread options (Section 3.3.4.1) and options on electricity forwards (Section 3.3.4.2) are of this type, as well as any option on electricity spot price.

Since $\widehat{\mathbb{Q}}$ is the minimal EMM for the market of fuels S (see Section 3.3.2.2) as well as for the larger market $(S, F^e(T^*))$ of fuels and electricity T^*-forward contract (see Proposition 3.3.3), we have to find, according to Theorem 3.3.1, an explicit expression for the GKW decomposition of such an H under $\widehat{\mathbb{Q}}$. More precisely, being $F^e(T^*)$ and $F^i(T^*)$, $1 \leq i \leq n$ our hedging instruments, we look for self-financing strategies ξ^e and $\xi = (\xi^1, \ldots, \xi^n)$ such that:

$$H = \widehat{\mathbb{E}}[H] + \int_0^T \xi_t \cdot dF_t(T^*) + \int_0^T \xi_t^e dF_t^e(T^*) + L_T^H \tag{3.3.22}$$

where L_T^H is the terminal value of a $\widehat{\mathbb{Q}}$-martingale L^H orthogonal to $F^e(T^*)$ and $F(T^*)$, representing the unhedgeable risk related to the contingent claim H. We are going to compute explicitly these strategies as well the unhedgeable risk L_T^H. We adapt methods developed in, e.g., [65], for stochastic volatility diffusion models.

First, let $\theta = (\theta^S, \theta^C, \theta^D)$ denote the integrands in the $\widehat{\mathbb{Q}}$-dynamics for the forward price $F^e(T^*)$ given by equation (3.3.13), i.e.:

$$dF_t^e = \theta_t^S . d\widehat{W}_t + \theta_t^C . dW_t^C + \theta_t^D dW_t^D \tag{3.3.23}$$

By the Markov property of the vector-valued process $(F(T^*), C, D)$ and by the fact that $F_t^e(T^*)$ is a function of $(t, F_t(T^*), C_t, D_t)$ (see formula (3.3.11)), we have

$$V_t^H := \widehat{\mathbb{E}}[H|\mathcal{F}_t] = \phi(t, F_t(T^*), C_t, D_t)$$

for some measurable function $\phi : [0,T] \times \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R} \mapsto \mathbb{R}$. Moreover, under some regularity assumptions on the coefficients of the underlying processes, which are satisfied by the model considered in, e.g., Section 3.4, such a function is of class $\mathcal{C}^{1,2,2,2}([0,T] \times \mathbb{R}^n \times \mathbb{R}^n \times \mathbb{R})$.[5] From now on we drop, for the sake of simplicity, the dependence from (t, F_t, C_t, D_t) from the function ϕ and its derivatives. In the next proposition we will use the notation $||\theta_t^C, \theta_t^D||$ for the norm of the vector (θ_t^C, θ_t^D).

Proposition 3.3.7. *Let $H = \varphi(F_T^e(T^*), F_T(T^*), C_T, D_T) \in L^2(\widehat{\mathbb{Q}})$ be a T-contingent claim with $T \leq T^*$. The local risk minimizing strategy (ξ^e, ξ) is given by:*

$$\xi_t^e = \frac{1}{||\theta_t^C, \theta_t^D||^2} \left\{ \sum_{i=1}^n \theta_t^{C,i} \frac{\partial \phi}{\partial c_i} \beta_i(t, C_t^i) + \theta_t^D \frac{\partial \phi}{\partial d} b(t, D_t) \right\} \tag{3.3.24}$$

$$\xi_t^i = \frac{\partial \phi}{\partial y_i} + \frac{h_i G_i^{T^*}(t, C_t, D_t)}{\|\theta_t^C, \theta_t^D\|^2} \left\{ \sum_{i=1}^n \theta_t^{C,i} \frac{\partial \phi}{\partial c_i} \beta_i(t, C_t^i) + \theta_t^D \frac{\partial \phi}{\partial d} b(t, D_t) \right\} \tag{3.3.25}$$

[5] See, e.g., Theorem 5.3 in [53] for such regularity assumptions

while the residual risk L^H satisfies:

$$dL_t^H = \sum_{i=1}^{n}\left(\frac{\partial\phi}{\partial c_i}\beta_i(t,C_t^i) - \frac{\sum_{k=1}^n \theta_t^{C,k}\frac{\partial\phi}{\partial c_k}\beta_k(t,C_t^k) + \theta_t^D\frac{\partial\phi}{\partial d}b(t,D_t)}{\|\theta_t^C,\theta_t^D\|^2}\theta_t^{C,i}\mathbf{1}_{\{\|\theta_t^C,\theta_t^D\|>0\}}\right)dW_t^{C,i} - \left(\frac{\partial\phi}{\partial d}b(t,D_t) - \frac{\theta_t^D\frac{\partial\phi}{\partial d}b(t,D_t)}{\|\theta_t^C,\theta_t^D\|^2}\mathbf{1}_{\{\|\theta_t^C,\theta_t^D\|>0\}}\right)dW_t^D. \tag{3.3.26}$$

Proof. The maturity T^* of the forward contracts being fixed, we drop it from the notation, so that we now write F_t^e, F_t^i and G_i instead of, respectively, $F_t^e(T^*)$, $F_t^i(T^*)$ and $G_i^{T^*}$. We also assume w.l.o.g. that $\widehat{\mathbb{E}}[H]=0$. Since $V_t^H = \widehat{\mathbb{E}}[H|\mathcal{F}_t] = \phi(t,F_t,C_t,D_t)$ with ϕ regular enough to apply Itō's lemma, we have that

$$V_t^H = \int_0^t \sum_i \frac{\partial\phi}{\partial y_i}dF_s^i + \int_0^t\sum_i\frac{\partial\phi}{\partial c_i}\beta_i(s,C_s^i)dW_s^{C,i} + \int_0^t\frac{\partial\phi}{\partial d}b(s,D_s)dW_s^D.$$

Notice that there is no dt-term in the expression above since we already know, by definition, that V_t^H is a martingale under $\widehat{\mathbb{Q}}$. Now, recall from equation (3.3.23) that:

$$dF_t^e = \theta_t^S.d\widehat{W}_t + \theta_t^C.dW_t^C + \theta_t^D dW_t^D$$

where explicit expressions for the integrands $\theta=(\theta^S,\theta^C,\theta^D)$ are provided by equation (3.3.13). We consider only the non-redundant part of F^e, i.e. the part that cannot be hedged using fuels, which is given by:

$$dF_t^{C,D} := \theta_t^C dW_t^C + \theta_t^D dW_t^D.$$

This process can be rewritten in terms of a suitable Brownian motion $W^{C,D}$ (use Lévy's characterization theorem to prove that $W^{C,D}$ is a Brownian motion) defined as:

$$W_t^{C,D} := \int_0^t \frac{\theta_s^C\cdot dW_s^C + \theta_s^D dW_s^D}{||\theta_s^C,\theta_s^D||}$$

From equation (3.3.13), one can check that$\|\theta_t^C,\theta_t^D\|>0$ for all t. Therefore:

$$dF_t^{C,D} = ||\theta_t^C,\theta_t^D||dW_t^{C,D}, \quad t\in[0,T^*].$$

Analogously, it holds that:

$$\sum_i\frac{\partial\phi}{\partial c_i}\beta_i(t,C_t^i)dW_t^{C,i} + \frac{\partial\phi}{\partial d}b(t,D_t)dW_t^D = \zeta_t d\overline{W}_t^{C,D}$$

where:

$$\zeta_t^2 := \sum_i\left(\frac{\partial\phi}{\partial c_i}\beta_i(t,C_t^i)\right)^2 + \left(\frac{\partial\phi}{\partial d}b(t,D_t)\right)^2$$

and $\overline{W}^{C,D}$ is a standard Brownian motion (use Lévy's criterion once more) defined as:

$$\overline{W}_t^{C,D} := \int_0^t \frac{\sum_i\frac{\partial\phi}{\partial c_i}\beta_i(s,C_s^i)dW_s^{C,i} + \frac{\partial\phi}{\partial d}b(s,D_s)dW_s^D}{\zeta_s}.$$

We require that $\zeta_t>0$ for all t, which basically means that the contingent claim H does depends on C and D. Observe that $W^{C,D}$ and $\overline{W}^{C,D}$ are correlated Brownian motions under both probabilities $\mathbb{P}$ and $\widehat{\mathbb{Q}}$, with quadratic covariation ρ_t given by:

$$\rho_t = \frac{\sum_i\theta_t^{C,i}\frac{\partial\phi}{\partial c_i}\beta_i(t,C_t^i) + \theta_t^D\frac{\partial\phi}{\partial d}b(t,D_t)}{||\theta_t^C,\theta_t^D||\zeta_t},$$

i.e. $d\langle W^{C,D}, \overline{W}^{C,D}\rangle_t = \rho_t dt$. We can define a new standard Brownian motion $W^\perp$ (under both $\mathbb{P}$ and $\widehat{\mathbb{Q}}$), independent of $W^{C,D}$, such that:

$$\overline{W}_t^{C,D} = \int_0^t \rho_s dW_s^{C,D} + \int_0^t \sqrt{1-\rho_s^2} dW_s^\perp,$$

i.e. $W_t^\perp = \int_0^t \frac{d\overline{W}_s^{C,D} - \rho_s dW_s^{C,D}}{\sqrt{1-\rho_s^2}}$. Therefore, we can write:

$$\sum_i \frac{\partial \phi}{\partial c^i}\beta_i(t, C_t^i) dW_t^{C,i} + \frac{\partial \phi}{\partial d} b(t, D_t) dW_t^D = \zeta_t \left(\rho_t dW_t^{C,D} + \sqrt{1-\rho_t^2} dW_t^\perp\right).$$

Finally, since $dW_t^{C,D} = \frac{dF_t^e - \theta_t^S d\widehat{W}_t}{||\theta_t^C, \theta_t^D||}$, we have:

$$\sum_i \frac{\partial \phi}{\partial c_i}\beta_i(t, C_t^i) dW_t^{C,i} + \frac{\partial \phi}{\partial d} b(t, D_t) dW_t^D = \zeta_t \left(\frac{\rho_t}{||\theta_t^C, \theta_t^D||}\left(dF_t^e - \theta_t^S d\widehat{W}_t\right) + \sqrt{1-\rho_t^2} dW_t^\perp\right)$$

and, using the fact that $\theta_t^{S,i} d\widehat{W}_t^i = h_i G_i(t, C_t, D_t) dF_t^i$ for all i, we also have:

$$V_t^H = \int_0^t \frac{\zeta_s \rho_s}{||\theta_s^C, \theta_s^D||} dF_s^e + \int_0^t \sum_i \left(\frac{\partial \phi}{\partial y_i} - \frac{\zeta_s \rho_s}{\|\theta_s^C, \theta_s^D\|} h_i G_i(s, C_s, D_s)\right) dF_s^i + \int_0^t \zeta_s \sqrt{1-\rho_s^2} dW_s^\perp,$$

which implies that:

$$\xi_t^e = \frac{\zeta_t \rho_t}{||\theta_t^C, \theta_t^D||}, \qquad \xi_t^i = \frac{\partial \phi}{\partial y_i} + \frac{\zeta_t \rho_t}{\|\theta_t^C, \theta_t^D\|} h_i G_i(t, C_t, D_t), \qquad L_t^H = \int_0^t \zeta_s \sqrt{1-\rho_s^2} dW_s^\perp,$$

are the good candidates for the local risk-minimizing strategy and the residual risk process. Thus, to conclude it suffices to verify that the proposed strategy and residual risk process provide the GKW decomposition under the minimal EMM $\widehat{\mathbb{Q}}$ and apply Theorem 3.3.1 to get the result. This verification being straightforward, the details are therefore omitted. □

To complete our characterization of LRM-strategy for $H = \varphi\left(F_T^e(T^*), F_T(T^*), C_T, D_T\right)$, we have to compute the pricing function ϕ and its derivatives, that appear in the formulae for ξ^e, ξ and L^H above. In Section 3.3.4, we have done so for the pricing function of some specific options. More generally, one can use standard PDE's techniques as follows. Notice that this part is rather formal. Let us consider:

$$V_t^H := \widehat{\mathbb{E}}[H|\mathcal{F}_t] = \phi(t, F_t(T^*), C_t, D_t)$$

where $F_t(T^*)$ denotes the vector $(F_t^1(T^*), \ldots, F_t^n(T^*))$. Under some regularity assumptions and using Itō's formula, one can prove that the function $\phi(t, y, c, d)$ is the solution to the following PDE:

$$\begin{aligned} 0 \;=\; & \frac{\partial \phi}{\partial t} + \sum_{i=1}^n \frac{\partial \phi}{\partial c_i}\alpha_i(t, c_i) + \frac{\partial \phi}{\partial d} a(t, d) + \frac{1}{2}\sum_{i=1}^n \frac{\partial^2 \phi}{\partial y_i^2}\left[\sum_{j\le i}(y_j - y_{j-1})^2 \sigma_j^2\right] \\ & + \frac{1}{2}\sum_{i=1}^n \frac{\partial^2 \phi}{\partial c_i^2}\beta_i(t, c_i)^2 + \frac{1}{2}\frac{\partial^2 \phi}{\partial d^2} b(t, d)^2 \end{aligned} \tag{3.3.27}$$

with boundary condition:

$$\phi(T, y, c, d) = \varphi(y, c, d), \quad \text{for all } (y, c, d) \in \mathbb{R}_+^n \times \mathbb{R} \times \mathbb{R}.$$

Notice that the hedging formulae from Proposition 3.3.7 (through θ_t^C and θ_t^D, see equation (3.3.11)), as well as the previous pricing PDE, contain the derivatives of the CES function G_i. These derivatives will be computed explicitly in Section 3.4.2.

3.4 Numerical results

3.4.1 Explicit model for capacities and demand

3.4.1.1 Model choice

So far, we have worked with the general diffusion models from equations (3.3.2) and (3.3.3) for the demand and capacities processes. Now, in order to push further the computations, we are going to choose and estimate more specific models. We decide to model the demand and capacities processes as follows:

$$\begin{aligned} D_t &= f_D(t) + Z_D(t) \\ C_t^i &= f_i(t) + Z_i(t) \end{aligned} \tag{3.4.1}$$

where f_D and f_i are deterministic functions, and Z_D and Z_i are independent Ornstein-Uhlenbeck (henceforth OU) processes under $\mathbb{P}$:

$$\begin{aligned} dZ_D(t) &= -\alpha_D Z_D(t)\,dt + b dW_t^D \\ dZ_i(t) &= -\alpha_i Z_i(t)\,dt + \beta_i dW_t^i \end{aligned} \tag{3.4.2}$$

where $\alpha_D, b, \alpha_i, \beta_i$, $1 \leq i \leq n$ are real constants. In other words, we choose the following functions as coefficients in (3.3.2) and (3.3.3):

$$\begin{aligned} a(t,d) &= \alpha_D \left(f_D(t) + \frac{f_D'(t)}{\alpha_D} - d \right) & b(t,d) &= b \\ \alpha_i(t,c) &= \alpha_i \left(f_i(t) + \frac{f_i'(t)}{\alpha_i} - c \right) & \beta_i(t,d) &= \beta_i \end{aligned} \tag{3.4.3}$$

We will see that this simple choice combines both satisfactory empirical accuracy and tractability. The ideas behind the definition (3.4.1) are the following:

- We decompose the demand between a deterministic part, that takes into account in a simple way the yearly and weekly seasonalities, and a stochastic part modeling the randomness of the demand.
- We use the same decomposition for the capacities, except that only a yearly seasonality is considered (as no statistically significant weekly seasonality appears), and in addition the deterministic part takes into account also the evolution of the installed capacity on the reference fleet. Indeed, should some plants be planned to be added or removed in the near future, the resulting shift in capacity has to be considered.

In equation (3.4.1), the deterministic functions f_D and f_i are defined as follows:

$$\begin{aligned} f_D(t) &= d_1 + d_2 \cos\left(2\pi \frac{t - d_3}{l_1}\right) + d_4 \cos\left(2\pi \frac{t - d_5}{l_2}\right) \\ f_i(t) &= c_1^i + c_2^i \cos\left(2\pi \frac{t - c_3^i}{l_1}\right) + f_i^{\text{evo}}(t) \end{aligned} \tag{3.4.4}$$

where d_j, $1 \leq j \leq 5$ and c_k^i, $1 \leq k \leq 3$, $1 \leq i \leq n$ are constants, and, assuming that t is expressed in years, $l_1 = 1$ (yearly seasonality) and $l_2 \simeq \frac{1}{52}$ (weekly seasonality), and f_i^{evo}, $1 \leq i \leq n$ are the deterministic installed capacities evolutions.

3.4.1.2 Model Estimation

Working on the dataset described in Section 3.2.2.1, we first estimate, using non-linear least squares, the deterministic functions from equation (3.4.4) (without the f_i^{evo} functions, which are provided by RTE[1]). Figure 3.4.1 illustrates these estimates for the demand process and the coal capacity process.

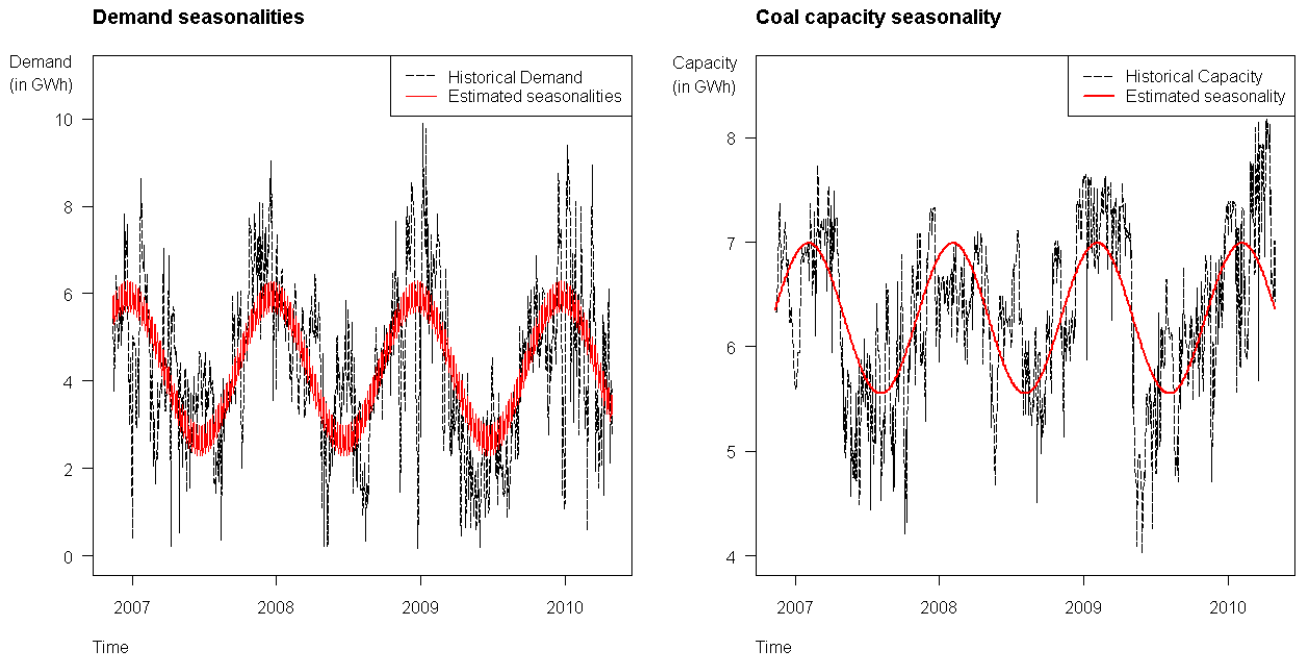

(a) Demand

(b) Coal capacity

Figure 3.4.1: Seasonalities estimates

All parameters are statistically significant. After doing that, one can subtract the estimated seasonalities from the data and work with the resulting series, that correspond to realizations of the processes $Z_D(t)$ and $Z_i(t)$ as in equation (3.4.1). We estimate the mean-reversion parameters α_D and α_i, $1 \leq i \leq n$, by exploiting the link between continuous OU processes and discrete AR(1) processes. Indeed, applying the Euler scheme to SDEs (3.4.2), denoting Δt as the time stepsize, yields:

$$Z_D(t+\Delta t) = Z_D(t)(1-\alpha_D \Delta t) + b\sqrt{\Delta t} N_D(t)$$
$$Z_i(t+\Delta t) = Z_i(t)(1-\alpha_i \Delta t) + \beta_i \sqrt{\Delta t} N_i(t)$$

for each time step t, where $N_D(t)$ and $N_i(t)$ are independent standard Gaussian white noises. Consequently, simple linear regressions will yield estimates for α_D, b, α_i and β_i, $1 \leq i \leq n$. From the different hypothesis of the model, a statistical analysis of the residuals reveals that the least accurate is the Gaussian hypothesis, as slight but statistically significant excess kurtosis appears. This phenomenon is illustrated by Figure 3.4.2, where the residual densities of Z_D and Z_1, estimated by the kernel method, are compared to Gaussian distributions with the same mean and the same variance. Nevertheless, given the moderate excess kurtosis estimates (around 2.0, see Table 3.4.1) and the tractability provided by the presence of Brownian motions, we choose

[1]The French transmission system operator.

to stick to the model built so far. Possible extensions of the model may later accomodate these empirical deviations.

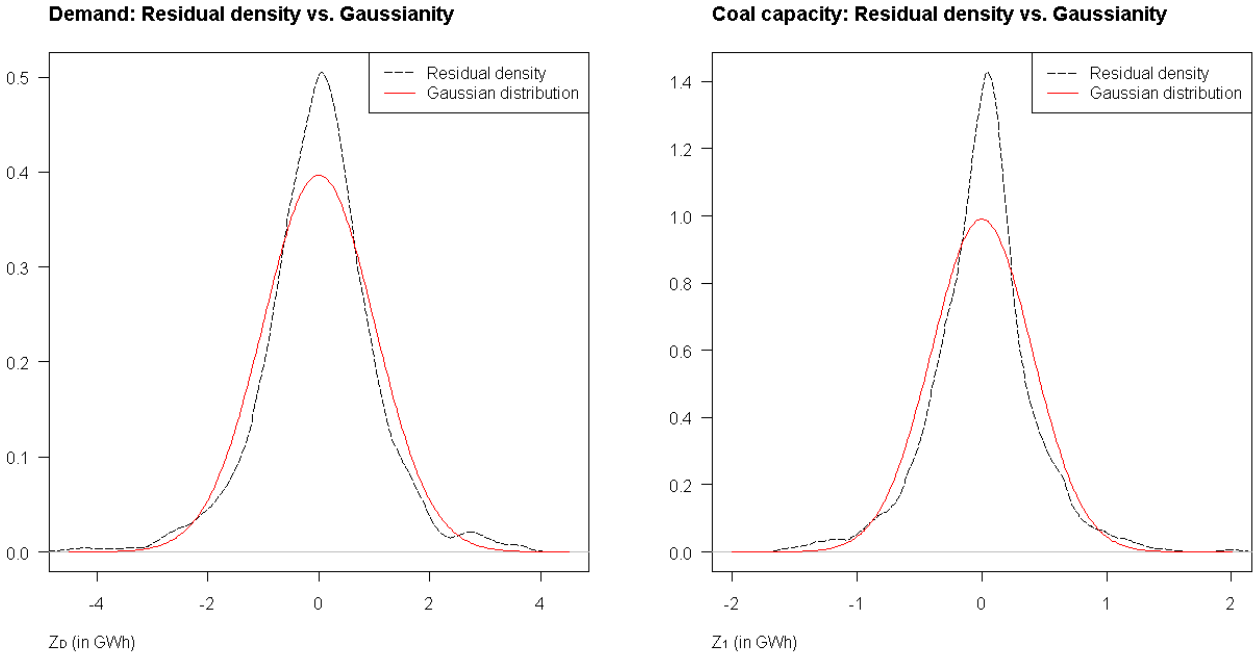

(a) Demand

(b) Coal capacity

Figure 3.4.2: Residuals estimated densities

residuals	mean	st.dev.	skewness	(excess) kurtosis
demand	0.00	1.01	−0.13	1.92
coal	0.00	0.40	−0.04	2.08

Table 3.4.1: Moments Estimates

3.4.2 Computing the Conditional Expectation of Scarcity Function

Using the dynamics of demand and capacities set in Section 3.4.1, we now study in detail the conditional expectation of scarcity function $G_i^T(t, C_t, D_t)$ and its partial derivatives, which are key quantities for pricing and hedging electricity derivatives (see Section 3.3). We will see how they can be mathematically and numerically computed.

3.4.2.1 Definition of the auxiliary function $G(m, \sigma)$

We recall from Section 3.3.3.1 the definition of the CES function:

$$G_i^T(t, C_t, D_t) = \mathbb{E}\left[g\left(C_T^{max} - D_T\right)\mathbf{1}_{\{D_T \in I_T^i\}} | \mathcal{F}_t^{D,C}\right], \quad 1 \leq i \leq n \tag{3.4.5}$$

Recall that the change of measure from $\mathbb{P}$ to the minimal EMM $\hat{\mathbb{Q}}$ does not impact the processes D and C^i, $1 \leq i \leq n$, meaning that the two measures $\hat{\mathbb{Q}}$ and $\mathbb{P}$ coincide on $\mathcal{F}_t^{D,C}$ for any t. Therefore, these processes and their dynamics are still defined by equations (3.4.1) and (3.4.2), which are easily solved to deduce the laws of $D_T|\mathcal{F}_t^{D,C}$ and $C_T^i|\mathcal{F}_t^{D,C}$:

Proposition 3.4.1. *Conditionally on $\mathcal{F}_t^{D,C}$, the random variable D_T is Gaussian, i.e. $D_T|\mathcal{F}_t^{D,C} \sim \mathcal{N}\left(m_{t,T}^D, \sigma_{t,T}^D\right)$ where:*

$$\begin{aligned} m_{t,T}^D &= f_D(T) + e^{-\alpha_D(T-t)}\left(D_t - f_D(t)\right) \\ \left(\sigma_{t,T}^D\right)^2 &= \frac{b^2}{2\alpha_D}\left[1 - e^{-2\alpha_D(T-t)}\right] \end{aligned} \tag{3.4.6}$$

Similarly, $C_T^i|\mathcal{F}_t^{D,C} \sim \mathcal{N}\left(m_{t,T}^i, \sigma_{t,T}^i\right)$ where:

$$\begin{aligned} m_{t,T}^i &= f_i(T) + e^{-\alpha_i(T-t)}\left(C_t^i - f_i(t)\right) \\ \left(\sigma_{t,T}^i\right)^2 &= \frac{\beta_i^2}{2\alpha_i}\left[1 - e^{-2\alpha_i(T-t)}\right] \end{aligned} \tag{3.4.7}$$

Proof. See Appendix 3.6.2.1. □

In view of equation (3.4.5), Proposition 3.4.1 indicates that the quantity of interest for computing the CES function is given by:

$$\mathcal{G}(m,\sigma) := \mathbb{E}\left[g(X)\right] \tag{3.4.8}$$

where $X \sim \mathcal{N}(m,\sigma)$ is a Gaussian random variable under the measure considered.

3.4.2.2 Expressing the CES function via $\mathcal{G}(m,\sigma)$

Indeed, one can express $G_i^T(t, C_t, D_t)$ using the auxiliary function $\mathcal{G}(m,\sigma)$. This is the purpose of this section. First, we consider the particular case with only one fuel ($n = 1$). The index i can be dropped, and the result is the following:

Proposition 3.4.2. *When $n = 1$, we have is given by:*

$$G^T(t, C_t, D_t) = \mathcal{G}(\bar{m}, \bar{\sigma}), \quad \bar{m} = m_{t,T} - m_{t,T}^D, \quad \bar{\sigma}^2 = (\sigma_{t,T})^2 + \left(\sigma_{t,T}^D\right)^2$$

Proof. See Appendix 3.6.2.2. □

In this particular case $n = 1$, the link between the two functions is simple. We now turn to the general case with n fuels. The link is now given by the following proposition:

Proposition 3.4.3. *For $2 \leq i \leq n-1$, we have:*

$$G_i^T(t, C_t, D_t) = H\left(\bar{m}_{i+1}^n, \bar{m}_1^{i,D}, \bar{\sigma}_{i+1}^n, \bar{\sigma}_1^{i,D}\right) - H\left(\bar{m}_i^n, \bar{m}_1^{i-1,D}, \bar{\sigma}_i^n, \bar{\sigma}_1^{i-1,D}\right) \tag{3.4.9}$$

$$H(m_1, m_2, \sigma_1, \sigma_2) := \int_0^\infty \mathcal{G}(x + m_1, \sigma_1)\, \Phi_N(x; m_2, \sigma_2)\, dx \tag{3.4.10}$$

where $\Phi_N(x; m, \sigma)$ is the probability density function of a normal random variable with mean m and variance $\sigma^2 > 0$ and where:

$$\begin{array}{rclrcl} \bar{m}_i^n &=& \sum_{k=i}^n m_{t,T}^k & (\bar{\sigma}_i^n)^2 &=& \sum_{k=i}^n \left(\sigma_{t,T}^k\right)^2 \\ \bar{m}_1^{i,D} &=& \sum_{k=1}^i m_{t,T}^k - m_{t,T}^D & \left(\bar{\sigma}_1^{i,D}\right)^2 &=& \sum_{k=1}^i \left(\sigma_{t,T}^k\right)^2 + \left(\sigma_{t,T}^D\right)^2 \end{array} \tag{3.4.11}$$

For $i = 1$ or $i = n$, $G_i^T(t, C_t, D_t)$ is given by:

$$\begin{aligned} G_1^T(t, C_t, D_t) &= H\left(\bar{m}_2^n, \bar{m}_1^{1,D}, \bar{\sigma}_2^n, \bar{\sigma}_1^{1,D}\right) \\ G_n^T(t, C_t, D_t) &= \mathcal{G}\left(\bar{m}_1^{n,D}, \bar{\sigma}_1^{n,D}\right) - H\left(\bar{m}_n^n, \bar{m}_1^{n-1,D}, \bar{\sigma}_n^n, \bar{\sigma}_1^{n-1,D}\right) \end{aligned} \tag{3.4.12}$$

Proof. See Appendix 3.6.2.3. □

Therefore, equations (3.4.9) and (3.4.12) indicate that computing the CES function reduces to compute the functions $\mathcal{G}(m,\sigma)$ and $H(m_1,m_2,\sigma_1,\sigma_2)$. This is why we devote the next sections to the evaluation of these quantities. We first consider $\mathcal{G}(m,\sigma)$, as computing $H(m_1,m_2,\sigma_1,\sigma_2)$ will involve results concerning $\mathcal{G}(m,\sigma)$.

3.4.2.3 Computation of $\mathcal{G}(m,\sigma)$

The following proposition corresponds to the first step in the calculation of $\mathcal{G}(m,\sigma)$.

Proposition 3.4.4. *Let $\mathcal{G}(m,\sigma)=\mathbb{E}[g(X)]$ where $X\sim\mathcal{N}(m,\sigma)$ and $g(x)=\min\left\{M,\frac{\gamma}{x^\nu}\right\}\mathbf{1}_{\{x>0\}}+M\mathbf{1}_{\{x\leq 0\}}$ with $M>0$, $\gamma>0$, $\nu>0$ and $\sigma>0$. Then $\mathcal{G}(m,\sigma)$ is given by the following expression:*

$$\mathcal{G}(m,\sigma)=M\mathcal{N}\left(\frac{\widetilde{x}-m}{\sigma}\right)+\frac{\gamma}{\left(\sigma\sqrt{2}\right)^{\nu}\sqrt{\pi}}\widetilde{\mathcal{G}}\left(\frac{\widetilde{x}-m}{\sigma\sqrt{2}},\frac{m}{\sigma\sqrt{2}};\nu\right) \tag{3.4.13}$$

where $\widetilde{x}=\left(\frac{\gamma}{M}\right)^{\frac{1}{\nu}}$, $\mathcal{N}$ is the cumulative distribution function of a standard normal random variable, and $\widetilde{\mathcal{G}}$ is defined by:

$$\widetilde{\mathcal{G}}(x,y;\nu):=\int_x^{\infty}\frac{1}{(y+z)^{\nu}}e^{-z^2}dz \tag{3.4.14}$$

where $x+y\geq 0$ ($x+y>0$ if $\nu\geq 1$)

Proof. See Appendix 3.6.2.4. □

The next step is to compute the function $\widetilde{\mathcal{G}}$ defined in equation (3.4.14).

3.4.2.4 On the Extended Incomplete Goodwin-Staton integral

Definition. We recall the definition of $\widetilde{\mathcal{G}}$ (equation (3.4.14)):

$$\widetilde{\mathcal{G}}(x,y;\nu)=\int_x^{\infty}\frac{1}{(y+z)^{\nu}}e^{-z^2}dz$$

where $(x,y,\nu)\in\mathbb{R}^3$ with $x+y>0$ if $\nu>1$, and $x+y\geq 0$ otherwise. These constraints ensure that $\widetilde{G}$ is well-defined. In the particular case $\nu=1$, $\widetilde{\mathcal{G}}$ corresponds to the incomplete Goodwin-Staton integral (cf. [45]). This is why we call $\widetilde{\mathcal{G}}$ the *extended incomplete Goodwin-Staton integral* (henceforth *EIGS* integral). Note that this extension of the Goodwin-Staton integral is different from the usual Generalized Goodwin-Staton integral (as defined in [85] for instance). Proposition 3.6.1 in Appendix 3.6.2.5 provides a probabilistic interpretation of the EIGS integral using the density of the sum of independent Pareto and Gaussian random variables.

Properties. We establish two useful properties verified by the EIGS integral.

Proposition 3.4.5. *Whenever $\nu\neq 1$, the following recurrence relation holds:*

$$\widetilde{\mathcal{G}}(x,y;\nu)=\frac{1}{1-\nu}\left(2\widetilde{\mathcal{G}}(x,y;\nu-2)-2y\widetilde{\mathcal{G}}(x,y;\nu-1)-(x+y)^{1-\nu}e^{-x^2}\right) \tag{3.4.15}$$

Proof. See Appendix 3.6.2.6. □

This recursive formula can be used, for instance, to compute $\widetilde{\mathcal{G}}(x,y;-n)$, $n \in \mathbb{N}$. Indeed, the initial values to start the recursion, $\widetilde{\mathcal{G}}(x,y;0)$ and $\widetilde{\mathcal{G}}(x,y;-1)$ are given by:

$$\widetilde{\mathcal{G}}(x,y;0) = \sqrt{\pi}\left[1-\mathcal{N}\left(x\sqrt{2}\right)\right] \quad , \qquad \widetilde{\mathcal{G}}(x,y;-1) = y\widetilde{\mathcal{G}}(x,y;0) + \tfrac{1}{2}e^{-x^2}$$

We will get advantage of this useful application later, in Section 3.4.2.6. Now, we establish the key property of the EIGS integral, which provides the ground for our evaluation algorithm.

Proposition 3.4.6. *The following identity holds:*

$$\widetilde{\mathcal{G}}(x,y;\nu) = \frac{1}{2}e^{-y^2}\Gamma\left(\frac{1-\nu}{2},(x+y)^2;-2y;-\frac{1}{2}\right) \tag{3.4.16}$$

where Γ *is the* extended incomplete gamma function *(cf. [40] p.266; henceforth* EIG *function):*

$$\Gamma(\alpha,x;b;\beta) := \int_x^\infty t^{\alpha-1}\exp\left(-t-bt^{-\beta}\right)dt \tag{3.4.17}$$

where $x \geq 0$ *and* $(\alpha,b,\beta) \in \mathbb{R}^3$ *are such that* $\Gamma(\alpha,x;b;\beta)$ *exists.*

Proof. See Appendix 3.6.2.7. □

3.4.2.5 Numerical computation of the EIGS integral

As indicated by equation (3.4.13),$G(m,\sigma)$ is composed of two parts. The first part involves the cumulative distribution function $\mathcal{N}$, for which efficient evaluation algorithms already exist. The second part involves the EIGS integral (3.4.14), which is the quantity that remains to be computed. Therefore, we provide in this section an algorithm to compute the EIGS integral.

Series expansion of the EIG function. Our idea to compute the EIGS integral is based on equation (3.4.16): if one can compute the EIG function (3.4.17), then one can compute the EIGS integral (3.4.14). In fact, there exist efficient algorithms to compute the (upper) incomplete Gamma function (henceforth IG function) defined by:

$$\Gamma(\alpha,x) = \int_x^\infty t^{\alpha-1}\exp(-t)\,dt \tag{3.4.18}$$

where $x \geq 0$ ($x > 0$ if $\alpha \leq 0$). However, this is not the case for the less standard EIG function. We therefore propose below a simple algorithm to compute efficiently this function. It is based on the following property (see [40] p.273):

Proposition 3.4.7. *The EIG function has the following series expansion :*

$$\Gamma(\alpha,x;b;\beta) = \sum_{n=0}^{\infty}\Gamma(\alpha-n\beta,x)\frac{(-b)^n}{n!} \tag{3.4.19}$$

Proof. Replace $\exp\left(-bt^{-\beta}\right)$ in (3.4.17) by its MacLaurin series expansion, and recall the definition of the IG function (equation (3.4.18)). □

Consequently, a simple way to numerically evaluate $\Gamma(\alpha,x;b;\beta)$ is to compute the first terms of the sum (3.4.19) up to numerical convergence, using an existing algorithm to compute $\Gamma(\alpha-n\beta,x)$. However, note that we only need to compute $\Gamma(\alpha,x;b;\beta)$ for $\beta = -\frac{1}{2}$ (see

equation (3.4.16) ; other values for β may require preliminary steps before exploiting equation (3.4.19), see Appendix 3.6.2.8). In other words, in the sum (3.4.19), the terms $\Gamma\left(\alpha+\frac{n}{2},x\right)$ are to be computed. To do so, the following recurrence relation of the IG function can be used:

$$\Gamma(\alpha+1,x)=\alpha\Gamma(\alpha,x)+x^{\alpha}e^{-x} \tag{3.4.20}$$

Thus, to compute the terms of the sum (3.4.19), the use of an IG evaluation algorithm will be necessary only for the two first IG terms $\Gamma(\alpha,x)$ and $\Gamma\left(\alpha+\frac{1}{2},x\right)$, as the next IG terms can be computed using (3.4.20). Such an implementation is provided in Appendix 3.6.3.1.

Approximation for large positive y. Combining equations (3.4.16) and (3.4.19), we obtain the following identity:

$$\widetilde{\mathcal{G}}(x,y;\nu)=\frac{1}{2}e^{-y^2}\sum_{n=0}^{\infty}\Gamma\left(\frac{1-\nu}{2}+\frac{n}{2},(x+y)^2\right)\frac{(2y)^n}{n!} \tag{3.4.21}$$

As a consequence, the smaller $|y|$ is, the more efficient the computation of the EIGS integral is. Therefore, finding accurate approximations when $|y|\gg 0$ sounds useful. In fact, when $y>0$, such an approximation will not only be useful but also necessary, because in our case where $x+y\geq 0$ is constant (see equation (3.4.13)) while y can take large positive values (see equation (3.4.9)), the decomposition (3.4.6) can involve the product of the very small term $\frac{1}{2}e^{-y^2}$ with the very large term $\Gamma\left(\frac{1-\nu}{2},(x+y)^2;-2y;-\frac{1}{2}\right)$. For large y, it is possible that $\frac{1}{2}e^{-y^2}$ reaches numerically zero, leading to the wrong estimate $\widetilde{\mathcal{G}}(x,y;\nu)=0$. This phenomenon is illustrated in Figure 3.4.3a, where the parameters have been chosen to be realistic in regard to the estimates in Section 3.4.1 and to the equations (3.4.6), (3.4.7) and (3.4.11) and the subsequent equations up to $\widetilde{G}$ and $\widetilde{H}$. For large positive y, we propose the following approximation:

$$\widetilde{\mathcal{G}}(x,y;\nu)\simeq\frac{1}{y^{\nu}}\exp\left(\frac{\nu^2}{4y^2}\right)\sqrt{\pi}\left[1-\mathcal{N}\left(\sqrt{2}x+\frac{\nu}{\sqrt{2}y}\right)\right] \tag{3.4.22}$$

Its derivation is detailed in Appendix 3.6.2.9. Figure 3.4.3a illustrates this approximating function on an example, and compares it with the series expansion (3.4.21). The accuracy of (3.4.22) for large y can be appreciated. Approximation (3.4.22) appears suitable to overcome the numerical problem described above.

Approximation for large negative y. When y takes large negative values, the previous numerical difficulty does not appear, as both terms $\frac{1}{2}e^{-y^2}$ and $\Gamma\left(\frac{1-\nu}{2},(x+y)^2;-2y;-\frac{1}{2}\right)$ become small. In fact, $\widetilde{\mathcal{G}}(x,y;\nu)$ can readily reach zero numerically, and detecting when this occurs can save much computation time. A simple way to check it is to compute the upper bound provided below:

Proposition 3.4.8. *For any $y<0$ (and as usual $x+y\geq 0$):*

$$0\leq\widetilde{\mathcal{G}}(x,y;\nu)\leq e^{-y^2-(x+y)^2}(-2y)^{\nu-1}\Gamma(1-\nu,(-2y)(x+y))$$

Proof. See Appendix 3.6.2.10. □

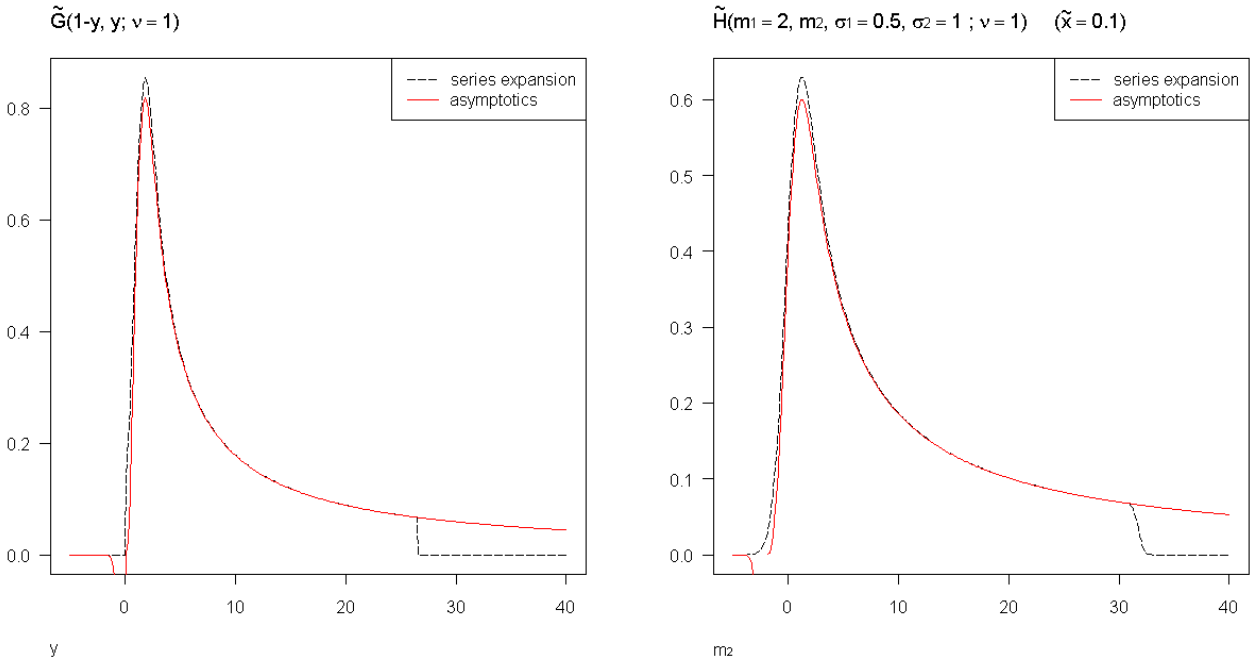

(a) Computing the EIGS integral

(b) Computing $\widetilde{H}$

Figure 3.4.3: Numerical approximations

3.4.2.6 Numerical computation of the CES function

So far, we have been giving a method for computing $\mathcal{G}(m, \sigma)$. We have shown in Section 3.4.2.2, more precisely with equations (3.4.9) and (3.4.12), that the evaluation of the CES function $G_i^T(t, C_t, D_t)$ involves $\mathcal{G}(m, \sigma)$ through the quantity $H(m_1, m_2, \sigma_1, \sigma_2)$ defined in equation (3.4.10), recalled below:

$$H(m_1, m_2, \sigma_1, \sigma_2) = \int_0^\infty \mathcal{G}(x + m_1, \sigma_1)\, \Phi_N(x; m_2, \sigma_2)\, dx$$

where $(m_1, m_2, \sigma_1, \sigma_2) \in \mathbb{R}^4$ with $\sigma_1 > 0$, $\sigma_2 > 0$. The following propositions are devoted to the computation of such a quantity.

Proposition 3.4.9. *The following holds:*

$$\begin{aligned} H(m_1, m_2, \sigma_1, \sigma_2) &= M\mathcal{N}\left(\frac{m_2}{\sigma_2}, \frac{\widetilde{x} - m_1 - m_2}{\sigma_1}; \begin{bmatrix} 0 & 0 \end{bmatrix}, \begin{bmatrix} 1 & -\frac{\sigma_2}{\sigma_1} \\ -\frac{\sigma_2}{\sigma_1} & 1 + \frac{\sigma_2^2}{\sigma_1^2} \end{bmatrix}\right) \\ &\quad + \frac{\gamma}{\pi\left(\sigma_1\sqrt{2}\right)^\nu} \widetilde{H}(m_1, m_2, \sigma_1, \sigma_2; \nu) \end{aligned} \tag{3.4.23}$$

where $\mathcal{N}\left(x, y; \begin{bmatrix} \mu_1 & \mu_2 \end{bmatrix}, \Sigma\right)$ *is the cumulative distribution function of the bivariate Gaussian random vector with mean* $\begin{bmatrix} \mu_1 & \mu_2 \end{bmatrix}$ *and covariance matrix* Σ*, and* $\widetilde{H}$ *is defined by:*

$$\widetilde{H}(m_1, m_2, \sigma_1, \sigma_2; \nu) := \int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^\infty \widetilde{\mathcal{G}}\left(\frac{\widetilde{x} - m_1 - m_2 - \sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}, \frac{m_1 + m_2 + \sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}; \nu\right) e^{-u^2} du \tag{3.4.24}$$

Proof. See Appendix 3.6.2.11. □

The cumulative distribution function $\mathcal{N}\left(x, y; \begin{bmatrix} \mu_1 & \mu_2 \end{bmatrix}, \Sigma\right)$ is efficiently computable. The task is now to compute $\widetilde{H}\left(m_1, m_2, \sigma_1, \sigma_2; \nu\right)$. This is done below.

Series expansion.

Proposition 3.4.10. $\widetilde{H}\left(m_1, m_2, \sigma_1, \sigma_2; \nu\right)$ has the following series expansion:

$$\widetilde{H}\left(m_1, m_2, \sigma_1, \sigma_2; \nu\right) = \frac{\sigma_1}{2\bar{\sigma}} e^{-\frac{\bar{m}^2}{2\bar{\sigma}^2}} \sum_{n=0}^{\infty} \Gamma\left(\frac{1-\nu}{2} + \frac{n}{2}, \frac{\widetilde{x}^2}{2\sigma_1^2}\right) \widetilde{\mathcal{G}}\left(\frac{m_1\sigma_2^2 - m_2\sigma_1^2}{\sigma_1\sigma_2\bar{\sigma}\sqrt{2}}, \frac{\bar{m}\sigma_1}{\sigma_2\bar{\sigma}\sqrt{2}}, -n\right) \frac{\left(\frac{2\sigma_2}{\bar{\sigma}}\right)^n}{n!} \tag{3.4.25}$$

where $\bar{m} = m_1 + m_2$ *and* $\bar{\sigma} = \sqrt{\sigma_1^2 + \sigma_2^2}$.

Proof. See Appendix 3.6.2.12. □

The terms $\Gamma\left(\frac{1-\nu}{2} + \frac{n}{2}, \frac{\widetilde{x}^2}{2\sigma_1^2}\right)$ can be computed using the recurrence relation (3.4.20) (see the discussion following Proposition 3.6.2), and the terms $\widetilde{\mathcal{G}}\left(\frac{m_1\sigma_2^2 - m_2\sigma_1^2}{\sigma_1\sigma_2\bar{\sigma}\sqrt{2}}, \frac{\bar{m}\sigma_1}{\sigma_2\bar{\sigma}\sqrt{2}}, -n\right)$ can be computed using the recurrence relation (3.4.15) (see the discussion following proposition 3.4.5). This provides an efficient way to compute $\widetilde{H}\left(m_1, m_2, \sigma_1, \sigma_2; \nu\right)$, and hence the CES function $G_i^T\left(t, C_t, D_t\right)$ (via equations (3.4.9), (3.4.12) and (3.4.23)). Such an implementation is provided in Appendix 3.6.3.3.

Approximation for large positive $\bar{m}$. Equation (3.4.25) involves the coefficient $\exp\left(-\frac{\bar{m}^2}{2\bar{\sigma}^2}\right)$. Similarly to Section 3.4.2.5, this can prevent $\widetilde{H}$ from being correctly computed when $\bar{m}$ takes large positive values. In that case, the following approximation can be used:

$$\widetilde{H}\left(m_1, m_2, \sigma_1, \sigma_2; \nu\right) \simeq \frac{\pi\left(\sigma_1\sqrt{2}\right)^{\nu}}{\bar{m}^{\nu}} \exp\left(\frac{\nu^2\bar{\sigma}^2}{2\bar{m}^2}\right) \mathcal{N}\left(\frac{m_2}{\sigma_2}, \frac{\bar{m} - \widetilde{x}}{\sigma_1}; \begin{bmatrix} \frac{\nu\sigma_2}{\bar{m}} & \frac{\nu\bar{\sigma}^2}{\sigma_1\bar{m}} \end{bmatrix}, \begin{bmatrix} 1 & \frac{\sigma_2}{\sigma_1} \\ \frac{\sigma_2}{\sigma_1} & 1 + \frac{\sigma_2^2}{\sigma_1^2} \end{bmatrix}\right) \tag{3.4.26}$$

The derivation of this expression is detailed in Appendix 3.6.2.13. Figure 3.4.3b illustrates the approximation (3.4.26) on an example, where it is compared to the series expansion (3.4.25). Again, the approximation is quickly very accurate, and enables to consider large positive $\bar{m}$.

Approximation for large negative $\bar{m}$. Similarly to Section 3.4.2.5, we provide an upper bound useful for large negative $\bar{m}$.

Proposition 3.4.11. *Suppose* $\nu \geq 0$. *Then the following holds:*

$$0 \leq \widetilde{H}\left(m_1, m_2, \sigma_1, \sigma_2; \nu\right) \leq \pi\left(\frac{\sigma_1\sqrt{2}}{\widetilde{x}}\right)^{\nu} \mathcal{N}\left(\frac{m_2}{\sigma_2}; -\frac{\widetilde{x}}{\sigma_1} + \frac{\bar{m}}{\sigma_1}; \begin{bmatrix} 0 & 0 \end{bmatrix}; \begin{bmatrix} 1 & \frac{\sigma_2}{\sigma_1} \\ \frac{\sigma_2}{\sigma_1} & 1 + \frac{\sigma_2^2}{\sigma_1^2} \end{bmatrix}\right)$$

Proof. See Appendix 3.6.2.14. □

3.4.2.7 Derivatives of the CES function

Thanks to the various results established in the previous sections, we are now able to compute the partial derivatives of the CESF function. They appear in the dynamics of the electricity forward prices (3.3.13) as well as in the hedging strategies of electricity derivatives (see Proposition 3.3.7).

Proposition 3.4.12. *The following derivatives hold:*

$$\begin{aligned}
\frac{\partial \mathcal{G}}{\partial m}(m,\sigma) &= -\frac{\nu\gamma}{\sqrt{\pi}\left(\sigma\sqrt{2}\right)^{\nu+1}}\widetilde{\mathcal{G}}\left(\frac{\widetilde{x}-m}{\sigma\sqrt{2}},\frac{m}{\sigma\sqrt{2}};\nu+1\right)\\
\frac{\partial H}{\partial m_1}(m_1,m_2,\sigma_1,\sigma_2) &= -\frac{\nu\gamma}{\pi\left(\sigma_1\sqrt{2}\right)^{\nu+1}}\widetilde{H}(m_1,m_2,\sigma_1,\sigma_2;\nu+1)\\
\frac{\partial H}{\partial m_2}(m_1,m_2,\sigma_1,\sigma_2) &= \mathcal{G}(m_1,\sigma_1)\,\Phi_N(0;m_2,\sigma_2) - \frac{\nu\gamma}{\pi\left(\sigma_1\sqrt{2}\right)^{\nu+1}}\widetilde{H}(m_1,m_2,\sigma_1,\sigma_2;\nu+1)\\
&= \mathcal{G}(m_1,\sigma_1)\,\Phi_N(0;m_2,\sigma_2) + \frac{\partial H}{\partial m_1}(m_1,m_2,\sigma_1,\sigma_2)
\end{aligned}$$

These formulae allow to compute the derivatives of $G_i^T(t,x,y)$:

For $2 \le i \le n-1$ *and* $1 \le k \le n$*:*

$$\begin{aligned}
\frac{\partial G_i^T}{\partial y}(t,C_t,D_t) &= -e^{-\alpha_D(T-t)}\left(\frac{\partial H}{\partial m_2}\left(\bar{m}_{i+1}^n,\bar{m}_1^{i,D},\bar{\sigma}_{i+1}^n,\bar{\sigma}_1^{i,D}\right) - \frac{\partial H}{\partial m_2}\left(\bar{m}_i^n,\bar{m}_1^{i-1,D},\bar{\sigma}_i^n,\bar{\sigma}_1^{i-1,D}\right)\right)\\
\frac{\partial G_i^T}{\partial x_k}(t,C_t,D_t) &= e^{-\alpha_k(T-t)}\Bigg(\frac{\partial H}{\partial m_1}\left(\bar{m}_{i+1}^n,\bar{m}_1^{i,D},\bar{\sigma}_{i+1}^n,\bar{\sigma}_1^{i,D}\right) + \mathcal{G}\left(\bar{m}_{i+1}^n,\bar{\sigma}_{i+1}^n\right)\Phi_N\left(0;\bar{m}_1^{i,D},\bar{\sigma}_1^{i,D}\right)\mathbf{1}_{\{k\le i\}}\\
&\quad - \frac{\partial H}{\partial m_1}\left(\bar{m}_i^n,\bar{m}_1^{i-1,D},\bar{\sigma}_i^n,\bar{\sigma}_1^{i-1,D}\right) - \mathcal{G}\left(\bar{m}_i^n,\bar{\sigma}_i^n\right)\Phi_N\left(0;\bar{m}_1^{i-1,D},\bar{\sigma}_1^{i-1,D}\right)\mathbf{1}_{\{k<i\}}\Bigg)
\end{aligned}$$

For $i = 1$ *or* $i = n$*:*

$$\begin{aligned}
\frac{\partial G_1^T}{\partial y}(t,C_t,D_t) &= -e^{-\alpha_D(T-t)}\frac{\partial H}{\partial m_2}\left(\bar{m}_2^n,\bar{m}_1^{1,D},\bar{\sigma}_2^n,\bar{\sigma}_1^{1,D}\right)\\
\frac{\partial G_1^T}{\partial x_k}(t,C_t,D_t) &= e^{-\alpha_k(T-t)}\left(\frac{\partial H}{\partial m_1}\left(\bar{m}_2^n,\bar{m}_1^{1,D},\bar{\sigma}_2^n,\bar{\sigma}_1^{1,D}\right) + \mathcal{G}\left(\bar{m}_2^n,\bar{\sigma}_2^n\right)\Phi_N\left(0;\bar{m}_1^{1,D},\bar{\sigma}_1^{1,D}\right)\mathbf{1}_{\{k=1\}}\right)\\
\frac{\partial G_n^T}{\partial y}(t,C_t,D_t) &= -e^{-\alpha_D(T-t)}\left(\frac{\partial \mathcal{G}}{\partial m}\left(\bar{m}_1^{n,D},\bar{\sigma}_1^{n,D}\right) - \frac{\partial H}{\partial m_2}\left(\bar{m}_n^n,\bar{m}_1^{n-1,D},\bar{\sigma}_n^n,\bar{\sigma}_1^{n-1,D}\right)\right)\\
\frac{\partial G_n^T}{\partial x_k}(t,C_t,D_t) &= e^{-\alpha_k(T-t)}\Bigg(\frac{\partial \mathcal{G}}{\partial m}\left(\bar{m}_1^{n,D},\bar{\sigma}_1^{n,D}\right) - \frac{\partial H}{\partial m_1}\left(\bar{m}_n^n,\bar{m}_1^{n-1,D},\bar{\sigma}_n^n,\bar{\sigma}_1^{n-1,D}\right)\\
&\quad - \mathcal{G}\left(\bar{m}_n^n,\bar{\sigma}_n^n\right)\Phi_N\left(0;\bar{m}_1^{n-1,D},\bar{\sigma}_1^{n-1,D}\right)\mathbf{1}_{\{k<n\}}\Bigg)
\end{aligned}$$

Proof. See Appendix 3.6.2.15. □

This ends our derivation of results and algorithms to compute $G_i^T(t,C_t,D_t)$ and its derivatives.

3.4.3 Pricing and Hedging

Finally, using the algorithms from Section 3.4.2 and our dataset, we briefly test and discuss the pricing and hedging of some of the derivatives analysed in Section 3.3.

3.4.3.1 Electricity futures

The simplest test to perform is the partial hedge of electricity futures using fuels futures. Recalling equations (3.3.11) and (3.3.13), one can see that the algorithms from Section 3.4.2 allow us to compute $F_t^e(T)$ and $dF_t^e(T)$, which are needed to simulate the partial hedge portfolio. We consider a 3-month electricity futures with a delivery period of one single hour. We use a constant rebalancing stepsize of one day, and neglect transaction costs. Figure 3.4.4a depicts the distribution of the final hedging error, estimated from a sample of 1000 realizations. For comparison, the distribution of the hedging error before maturity is represented as well, for different remaining maturities.

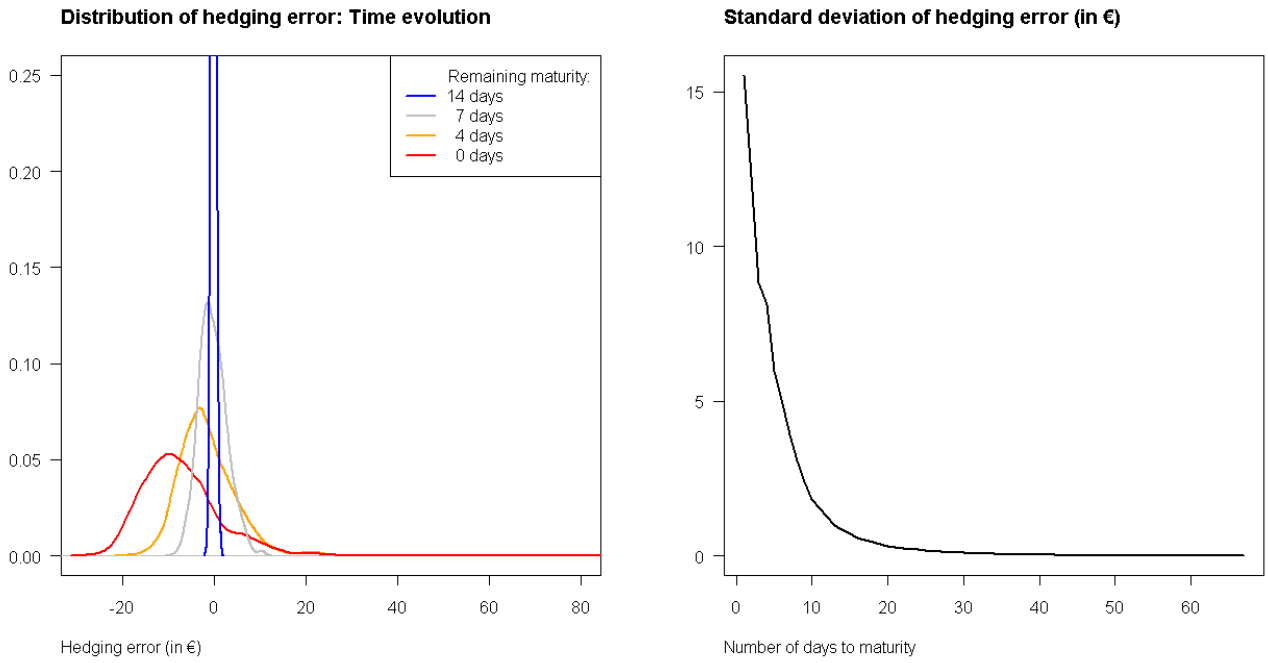

(a) Distribution

(b) Standard deviation

Figure 3.4.4: Hedging error

One can observe that up to two weeks before the expiration of the 3-month futures, the hedge generates hardly any error. However, during the last days of the product's life, the distribution of the error widens considerably. The standard deviation of the error, for instance, culminates at maturity, as shown by Figure 3.4.4b. Similarly, the asymmetry of the error culminates at maturity, where a positive skewness of 8.7 is measured (The maximum hedging error at maturity reaches 212€ on this sample. For the sake of readability, the hedging error on Figure 3.4.4a was limited to 80€). To sum up, depending on the time to maturity, two different behaviours can be observed:

- Far from maturity, the partial hedge is almost perfect. Indeed, recall equations (3.4.6) and (3.4.7). Because of the coefficients $e^{-\alpha_D(T-t)}$ and $e^{-\alpha_i(T-t)}$ (and the hypothesis of constant volatilities, see equation (3.4.3)), the weights $G_i^T(t, C_t, D_t)$ can be considered constant when $T - t \gg 0$. In this limit, the electricity futures are only driven by the fuels risks, ie. the electricity futures behaves like a basket of fuels futures. Consequently, the partial hedge turns out to be an almost perfect static hedge. With our estimates of α_D and α_i (around 70), this behaviour can be considered to hold up to two weeks before maturity.

- Close to maturity, the partial hedge is almost useless. Indeed, the coefficients $e^{-\alpha_D(T-t)}$ and $e^{-\alpha_i(T-t)}$ become no longer negligible, and consequently the unhedged risks stemming from demand and capacities start to drive the electricity futures. In fact, the partial derivatives of $G_i^T(t, C_t, D_t)$ can become huge, overwhelming the hedged fuel risks. It is such so that even if demand and capacities happened to be tradable assets, the necessary use of discrete hedging would lead to hedging errors similar to Figure 3.4.4. This behaviour is analogous to classical barrier options close to expiry while close to the barrier, in which case static hedging is to be preferred to dynamic hedging.

3.4.3.2 Spread options

Finally, we look at the prices of spread options obtained in our model. Using the results in Section 3.3.4.1, we compute the prices of options with payoff $\left(P_T - h_1 S_T^1 - K\right)^+$, $K > 0$, i.e. European calls on dark spread with instantaneous delivery period. We compute equation (3.3.19) using numerical integration, replacing the integration on $\mathbb{R}^3$ (the spread Y_T^1 being the exponential of a Gaussian random variable) by an integration on the hypercube $[0,1]^3$ using the bijective transformations $y = \frac{1}{1+\exp(-x)}$. Figure 3.4.5a illustrates the price surface for different strikes K and different instantaneous maturities T.

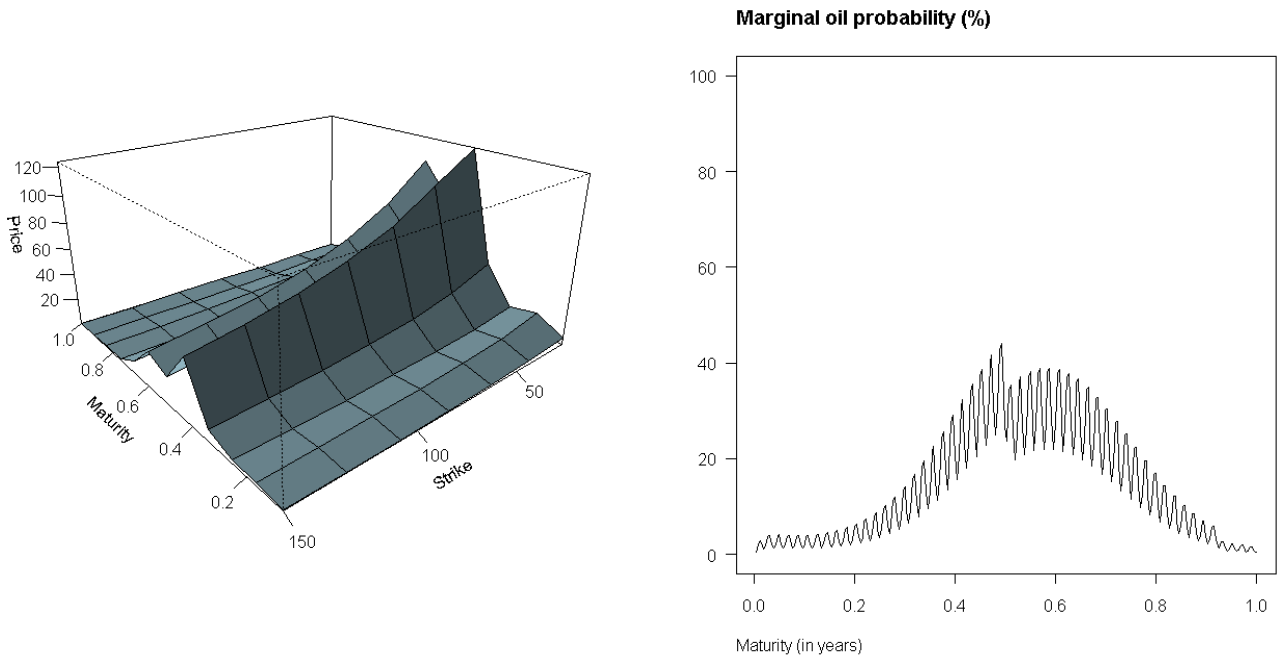

(a) Price surface

(b) Probability of marginal oil

Figure 3.4.5: Call on dark spread

The most visible effect is that of seasonality. The shape of the seasonality is driven by the quantity $\mathbb{P}\left(D_T > C_T^1\right)$, which represents the probability of using the most expensive fuel S_T^2 to produce electricity at time T. This quantity is depicted on Figure 3.4.5b for comparison. Recall the weekly seasonality on demand and the deterministic evolution of installed capacities (equation (3.4.4)). Here the middle of the time period covers the next winter, where the likelihood of using the expensive oil is at its highest.

3.5 Conclusion

This chapter is a contribution to the pricing of contingent claims on electricity markets. As opposed to the previous version of our model presented in [3], the introduction of a scarcity function allows to capture the electricity spot price spikes at a cheap cost (only two parameters more, that can be easily estimated). The electricity spot price model developed here is particulary well suited for spread options on the spot since it is based on the economic relation that holds between fuel prices and electricity spot prices. Precise evaluation of electricty spread options are of main importance since they are the building blocks for the pricing of generation assets. Our structural risk-neutral model with scarcity function should enable us to assess in a near future the problem of the optimal timing of investment in generation assets.

Moreover, for the sake of simplicity, we considered in this work only diffusion processes. While this framework seems reasonable for such processes as demand, capacities, and some fuels like oil and coal, it may be too restrictive, for example, for the gas price, which is noticeably prone to spikes. However, enriching the dynamics of some fuel prices, for instance by adding a jump component, would not alter the general methodology developed in this chapter. At the cost of more complex computations, it may refine the pricing of such electricity derivatives as spark spread options. This is a direction we plan to explore in our future research activities.

3.6 Appendices

3.6.1 Dataset

3.6.1.1 Retrieving the residual demand

RTE provides the following data, with their descriptions:

Demand:=Power Consumption+Network Losses −Pumping

Network Losses:=Power Production+Physical Import−Power Consumption−Physical Export

Interconnection:=Physical Export−Physical Import

(Incomplete) Production:=Hydro+Nuclear+Coal+Oil

In short, the following equations hold:

$$\begin{aligned} D &= C + L - Pump \\ L &= P - C - I \\ \widetilde{P} &= H + N + Coal + Oil \end{aligned}$$

The two first equations yield:

$$D + Pump = P - I$$

The available production data $\widetilde{P}$ covers, quoting RTE, "more than 90% of the generating units that are connected to the French transmission network". It covers therefore slightly less than the overall production P. Regarding the fact that RTE does not provide the production of the smallest units (<20MW), nor the production of non-hydroelectric renewables, we make the following reasonable assumption:

$$P \;=\; R + H + N + Coal + Oil \;=\; R + \widetilde{P}$$

where R denotes the production from non-hydroelectric renewable energies. Consequently, the residual demand can be retrieved as follows:

$$\begin{aligned}\text{Residual Demand} &:= (\text{Total Demand}) - (R + H + N)\\ &= (D + I + Pump) - (R + H + N)\\ &= P - R - H - N \;=\; \widetilde{P} - H - N\end{aligned}$$

Therefore computing the residual demand only requires $\widetilde{P}$, H and N, which are available from the same RTE record file.

3.6.1.2 Retrieving fuel capacity

RTE provides records of the actual production and the effective availability for each fuel. What corresponds to our capacity variables C_t^i, $1 \leq i \leq n$, is the effective availability. However, these data must be slightly adjusted first, as the comparison between actual production and effective availability reveals that even when electricity demand is at its highest, the actual production never reaches the effective availability. This phenomenon can be explained by the steady presence of primary and secondary reserves. For the accuracy of the model, in particular for the accurate detection of the marginal fuels, we correct this effect by the following adjustment:

$$C_t^1 := 0.94 \times (\text{coal effective availability}) \quad , \qquad C_t^2 := 0.80 \times (\text{oil effective availability})$$

The two coefficients were inferred from the data.

3.6.2 Proofs

3.6.2.1 Proposition 3.4.1

Proof. SDEs (3.4.2) have classical explicit solutions, which combined with equations (3.4.1) yield at time T, starting from time t:

$$\begin{aligned}D_T &= f_D(T) + e^{-\alpha_D(T-t)}(D_t - f_D(t)) + \int_t^T e^{-\alpha_D(T-s)} b dW_t^D\\ C_T^i &= f_i(T) + e^{-\alpha_i(T-t)}\left(C_t^i - f_i(t)\right) + \int_t^T e^{-\alpha_i(T-s)}\beta_i dW_t^i\end{aligned} \tag{3.6.1}$$

In particular, D_T and C_T^i are Gaussian random variables. Computing their expected values and their variances (using Itō's isometry) concludes the proof. □

3.6.2.2 Proposition 3.4.2

Proof. When $n = 1$, then $I_t =]-\infty, +\infty[$, and consequently equation (3.4.5) becomes $G^T(t, C_t, D_t) = \mathbb{E}\left[g(C_T - D_T)\,|\mathcal{F}_t^{D,C}\right]$. We use Proposition 3.4.1 and the independence between the processes C and D to deduce that, conditionally to $\mathcal{F}_t^{D,C}$, $C_T - D_T \sim \mathcal{N}(\bar{m}, \bar{\sigma})$ where $\bar{m} = m_{t,T} - m_{t,T}^D$ and $\bar{\sigma}^2 = (\sigma_{t,T})^2 + \left(\sigma_{t,T}^D\right)^2$. Therefore $G^T(t, C_t, D_t) = \mathbb{E}\left[g(C_T - D_T)\,|\mathcal{F}_t^{D,C}\right] = \mathcal{G}(\bar{m}, \bar{\sigma})$. □

3.6.2.3 Proposition 3.4.3

Proof. For every $1 \leq i \leq j \leq n$, define $\bar{C}_t^{i,j} := \sum_{k=i}^{j} C_t^k$. Thus, for $2 \leq i \leq n-1$:

$$\begin{aligned}
G_i^T(t, C_t, D_t) &= \mathbb{E}\left[g\left(\overline{C}_T^{1,n} - D_T\right)\mathbf{1}_{\left\{D_T \in I_T^i\right\}} | \mathcal{F}_t^{D,C}\right] = \mathbb{E}\left[g\left(\overline{C}_T^{1,n} - D_T\right)\mathbf{1}_{\left\{\overline{C}_T^{1,i-1} \leq D_T < \overline{C}_T^{1,i}\right\}} | \mathcal{F}_t^{D,C}\right] \\
&= \mathbb{E}\left[g\left(\overline{C}_T^{1,n} - D_T\right)\mathbf{1}_{\left\{D_T < \overline{C}_T^{1,i}\right\}} | \mathcal{F}_t^{D,C}\right] - \mathbb{E}\left[g\left(\overline{C}_T^{1,n} - D_T\right)\mathbf{1}_{\left\{D_T \leq \overline{C}_T^{1,i-1}\right\}} | \mathcal{F}_t^{D,C}\right] \\
&= \mathbb{E}\left[g\left(\left(\overline{C}_T^{1,i} - D_T\right) + \overline{C}_T^{i+1,n}\right)\mathbf{1}_{\left\{\overline{C}_T^{1,i} - D_T > 0\right\}} | \mathcal{F}_t^{D,C}\right] \qquad (3.6.2) \\
&- \mathbb{E}\left[g\left(\left(\overline{C}_T^{1,i-1} - D_T\right) + \overline{C}_T^{i,n}\right)\mathbf{1}_{\left\{\overline{C}_T^{1,i-1} - D_T \geq 0\right\}} | \mathcal{F}_t^{D,C}\right]
\end{aligned}$$

Using the same arguments as in Section 3.6.2.2, $\overline{C}_T^{1,i} - D_T \sim \mathcal{N}\left(\bar{m}_1^{i,D}, \bar{\sigma}_1^{i,D}\right)$ where $\bar{m}_1^{i,D} = \sum_{k=1}^{i} m_{t,T}^k - m_{t,T}^D$ and $\left(\bar{\sigma}_1^{i,D}\right)^2 = \sum_{k=1}^{i}\left(\sigma_{t,T}^k\right)^2 + \left(\sigma_{t,T}^D\right)^2$. Similarly, $\overline{C}_T^{i+1,n} \sim \mathcal{N}\left(\bar{m}_{i+1}^n, \bar{\sigma}_{i+1}^n\right)$ where $\bar{m}_{i+1}^n = \sum_{k=i+1}^{n} m_{t,T}^k$ and $\left(\bar{\sigma}_{i+1}^n\right)^2 = \sum_{k=i+1}^{n}\left(\sigma_{t,T}^k\right)^2$. Consequently:

$$\begin{aligned}
&\mathbb{E}\left[g\left(\left(\overline{C}_T^{1,i} - D_T\right) + \overline{C}_T^{i+1,n}\right)\mathbf{1}_{\left\{\overline{C}_T^{1,i} - D_T \geq 0\right\}} | \mathcal{F}_t^{D,C}\right] \\
&= \mathbb{E}\left[\mathbb{E}\left[g\left(\left(\overline{C}_T^{1,i} - D_T\right) + \overline{C}_T^{i+1,n}\right)\mathbf{1}_{\left\{\overline{C}_T^{1,i} - D_T \geq 0\right\}} | \overline{C}_T^{1,i} - D_T, \mathcal{F}_t^{D,C}\right] | \mathcal{F}_t^{D,C}\right] \\
&= \mathbb{E}\left[\mathcal{G}\left(\left(\overline{C}_T^{1,i} - D_T\right) + \bar{m}_{i+1}^n, \bar{\sigma}_{i+1}^n\right)\mathbf{1}_{\left\{\overline{C}_T^{1,i} - D_T \geq 0\right\}} | \mathcal{F}_t^{D,C}\right] \\
&= \int_0^{\infty} \mathcal{G}\left(x + \bar{m}_{i+1}^n, \bar{\sigma}_{i+1}^n\right) \Phi_N\left(x; \bar{m}_1^{i,D}, \bar{\sigma}_1^{i,D}\right) dx
\end{aligned}$$

□

where $\Phi_N(x; m, \sigma)$ is the density of a Gaussian random variable with mean m and variance $\sigma^2 > 0$. Performing the same calculation for the second expected value of equation (3.6.2), we get, for $2 \leq i \leq n-1$, the desired result. In the cases when $i = 1$ or $i = n$, few simplificating adjustments from the previous calculation give the desired results.

3.6.2.4 Proposition 3.4.4

Proof. The random variable $X \sim \mathcal{N}(m, \sigma)$ can be written as $X = m + \sigma N$ where $N \sim \mathcal{N}(0, 1)$ is a standard normal variable. Thus $\mathcal{G}(m, \sigma) = \mathbb{E}[g(X)] = \mathbb{E}[g(m + \sigma N)]$. Now, we remark that the function g can be written in the following way:

$$g(x) = \frac{\gamma}{x^{\nu}}\mathbf{1}_{\{x > \widetilde{x}\}} + M\mathbf{1}_{\{x \leq \widetilde{x}\}}$$

where $\widetilde{x} = \left(\frac{\gamma}{M}\right)^{\frac{1}{\nu}}$. Therefore:

$$\begin{aligned}
\mathcal{G}(m,\sigma) = \mathbb{E}[g(m+\sigma N)] &= \int_{\mathbb{R}} g(m+\sigma x)\frac{1}{\sqrt{2\pi}}e^{-\frac{x^2}{2}}dx \\
&= \int_{-\infty}^{\frac{\widetilde{x}-m}{\sigma}} M\frac{1}{\sqrt{2\pi}}e^{-\frac{x^2}{2}}dx + \int_{\frac{\widetilde{x}-m}{\sigma}}^{\infty}\frac{\gamma}{(m+\sigma x)^{\nu}}\frac{1}{\sqrt{2\pi}}e^{-\frac{x^2}{2}}dx \\
&= M\mathcal{N}\left(\frac{\widetilde{x}-m}{\sigma}\right) + \frac{\gamma}{\sigma^{\nu}\sqrt{2\pi}}\int_{\frac{\widetilde{x}-m}{\sigma}}^{\infty}\frac{1}{\left(\frac{m}{\sigma}+x\right)^{\nu}}e^{-\frac{x^2}{2}}dx \\
&= M\mathcal{N}\left(\frac{\widetilde{x}-m}{\sigma}\right) + \frac{\gamma}{\left(\sigma\sqrt{2}\right)^{\nu}\sqrt{\pi}}\widetilde{\mathcal{G}}\left(\frac{\widetilde{x}-m}{\sigma\sqrt{2}}, \frac{m}{\sigma\sqrt{2}}; \nu\right)
\end{aligned}$$

where $\widetilde{\mathcal{G}}$ is defined by equation (3.4.14). □

3.6.2.5 Interpretation of the EIGS integral

The EIGS integral has some probabilistic interpretations. The following proposition provides one of them, involving Pareto distributions.

Proposition 3.6.1. *If $\nu > 1$ and $x + y > 0$, then:*

$$\widetilde{\mathcal{G}}(x,y;\nu) = \frac{\sqrt{\pi}}{(\nu-1)(x+y)^{\nu-1}}\Phi_{N+P}\left(\sqrt{2}y\right) \tag{3.6.3}$$

where $P \sim Par\left(\nu-1, \sqrt{2}(x+y)\right)$ is a Pareto random variable, $N \sim \mathcal{N}(0,1)$ is a standard Gaussian random variable independent of P, and Φ_{N+P} is the density of the sum $N+P$.

Proof. We recall that, for $\alpha > 0$ and $x_m > 0$, the probability density function of the Pareto random variable $Par(\alpha, x_m)$ is given by $\Phi_P(x;\alpha,x_m) = \frac{\alpha x_m^{\alpha}}{x^{\alpha+1}}\mathbf{1}_{\{x>x_m\}}$. When $\nu > 1$ and $x+y > 0$, this Pareto density can be introduced as follows:

$$\begin{aligned}
\widetilde{\mathcal{G}}(x,y;\nu) &= \int_x^{\infty}\frac{1}{(y+z)^{\nu}}e^{-z^2}dz = \int_{\sqrt{2}(x+y)}^{\infty}\frac{1}{\left(\frac{u}{\sqrt{2}}\right)^{\nu}}e^{-\left(\frac{u}{\sqrt{2}}-y\right)^2}\frac{du}{\sqrt{2}} \\
&= \frac{1}{(\nu-1)\sqrt{2}(x+y)^{\nu-1}}\int_{\mathbb{R}}\Phi_P\left(u;\nu-1,\sqrt{2}(x+y)\right)e^{-\frac{(\sqrt{2}y-u)^2}{2}}du \\
&= \frac{\sqrt{\pi}}{(\nu-1)(x+y)^{\nu-1}}\int_{\mathbb{R}}\Phi_P\left(u;\nu-1,\sqrt{2}(x+y)\right)\Phi_N\left(\sqrt{2}y-u;0,1\right)du \\
&= \frac{\sqrt{\pi}}{(\nu-1)(x+y)^{\nu-1}}\Phi_{N+P}\left(\sqrt{2}y\right)
\end{aligned}$$

where $\Phi_N(x;m,\sigma)$ is the probability density function of a Gaussian random variable with mean m and variance σ^2, and Φ_{N+P} is the probability density function of the sum of a standard Gaussian random variable $N \sim N(0,1)$ and a Pareto random variable $P \sim Par\left(\nu-1,\sqrt{2}(x+y)\right)$, N and P being independent. We have used the change of variable $u = (y+z)\sqrt{2}$ for the second equality. □

3.6.2.6 Proposition 3.4.5

Proof. Suppose $\nu \neq 1$. Then:

$$\begin{aligned}
\widetilde{\mathcal{G}}(x,y;\nu) &= \int_x^\infty \left((y+z)^{-\nu}\right)\left(e^{-z^2}\right)dz \\
&= \left[\frac{(y+z)^{1-\nu}}{1-\nu}e^{-z^2}\right]_x^\infty - \int_x^\infty \left(\frac{(y+z)^{1-\nu}}{1-\nu}\right)\left(-2ze^{-z^2}\right)dz \\
&= -\frac{(x+y)^{1-\nu}}{1-\nu}e^{-x^2} + \frac{2}{1-\nu}\int_x^\infty (y+z)^{1-\nu} ze^{-z^2}dz \\
&= -\frac{(x+y)^{1-\nu}}{1-\nu}e^{-x^2} + \frac{2}{1-\nu}\int_{x+y}^\infty u^{1-\nu}(u-y)e^{-(u-y)^2}du \\
&= -\frac{(x+y)^{1-\nu}}{1-\nu}e^{-x^2} + \frac{2}{1-\nu}\left[\int_{x+y}^\infty u^{2-\nu}e^{-(u-y)^2}du - y\int_{x+y}^\infty u^{1-\nu}e^{-(u-y)^2}du\right] \\
&= -\frac{(x+y)^{1-\nu}}{1-\nu}e^{-x^2} + \frac{2}{1-\nu}\left[\int_x^\infty (y+z)^{2-\nu}e^{-z^2}dz - y\int_x^\infty (y+z)^{1-\nu}e^{-z^2}dz\right] \\
&= -\frac{(x+y)^{1-\nu}}{1-\nu}e^{-x^2} + \frac{2}{1-\nu}\left[\widetilde{\mathcal{G}}(x,y;\nu-2) - y\widetilde{\mathcal{G}}(x,y;\nu-1)\right]
\end{aligned}$$

□

3.6.2.7 Proposition 3.4.6

Proof. Using the definition of Γ (equation (3.4.17)), we obtain:

$$\Gamma\left(\frac{1-\nu}{2},(x+y)^2;-2y;-\frac{1}{2}\right) = \int_{(x+y)^2}^\infty \frac{1}{t^{\frac{1+\nu}{2}}}\exp\left(-t+2y\sqrt{t}\right)dt$$

$$= e^{y^2}\int_{(x+y)^2}^\infty \frac{1}{t^{\frac{1+\nu}{2}}}\exp\left(-\left[\sqrt{t}-y\right]^2\right)dt = 2e^{y^2}\int_{|x+y|-y}^\infty \frac{1}{(y+z)^\nu}e^{-z^2}dt = 2e^{y^2}\widetilde{\mathcal{G}}(|x+y|-y,y;\nu)$$

where we performed the change of variable $z = \sqrt{t} - y$. If $x + y \geq 0$, then $|x+y| - y = x$ and we get (3.4.16). If $x + y < 0$, then $|x+y| - y = -x - 2y$, and the change of variable $\{X = -x - 2y\,,\, Y = y\} \Leftrightarrow \{x = -X - 2Y\,,\, y = Y\}$ gives:

$$\begin{aligned}
\widetilde{\mathcal{G}}(X,Y;\nu) = \widetilde{\mathcal{G}}(-x-2y,y;\nu) &= \frac{1}{2}e^{-y^2}\Gamma\left(\frac{1-\nu}{2},(x+y)^2;-2y;-\frac{1}{2}\right) \\
&= \frac{1}{2}e^{-Y^2}\Gamma\left(\frac{1-\nu}{2},(-X-Y)^2;-2Y;-\frac{1}{2}\right) = \frac{1}{2}e^{-Y^2}\Gamma\left(\frac{1-\nu}{2},(X+Y)^2;-2Y;-\frac{1}{2}\right)
\end{aligned}$$

and we have recovered (3.4.16). Interestingly, the relation holds whatever the sign of $x + y$ is. However, we recall that, to ensure the definition of $\tilde{G}(x,y;\nu)$ for any ν, we imposed $x + y \geq 0$ (see equation (3.4.14)). The proposition provides an extension to $x + y < 0$, as long as the set of parameters ensure that the EIG function is well defined. □

3.6.2.8 Computing the EIG integral when $\beta < -1$

It can be shown that the condition $\beta > -1$ is required for the convergence of the series expansion (3.4.19). Therefore, before exploiting this expansion , the following ordered preliminary steps are to be performed:

Proposition 3.6.2. *The following relations hold:*

1. If $b = 0$, then $\Gamma(\alpha, x; 0; \beta) = \Gamma(\alpha, x)$
   a) If $\beta = 0$, then $\Gamma(\alpha, x; b; 0) = e^{-b}\Gamma(\alpha, x)$
   b) If $\beta = -1$ then $\Gamma(\alpha, x; b; -1) = \begin{cases} \frac{1}{(b+1)^\alpha}\Gamma(\alpha, (b+1)x) & \text{if } b+1 > 0 \\ \frac{1}{(-\alpha)x^{-\alpha}} & \text{if } b+1 = 0,\ \alpha < 0 \text{ and } x > 0 \\ +\infty & \text{else} \end{cases}$
   c) If $\beta < -1$ and $b > 0$, then $\Gamma(\alpha, x; b; \beta) = \frac{b^{\frac{\alpha}{\beta}}}{-\beta}\Gamma\left(-\frac{\alpha}{\beta}, bx^{-\beta}; b^{\frac{1}{\beta}}; \frac{1}{\beta}\right)$

Proof. The four relations are proved below: □

1. $\Gamma(\alpha, x; 0; \beta) = \int_x^\infty t^{\alpha-1}e^{-t}dt = \Gamma(\alpha, x)$
   a) $\Gamma(\alpha, x; b; 0) = \int_x^\infty t^{\alpha-1}e^{-t-b}dt = e^{-b}\Gamma(\alpha, x)$
   b) $\Gamma(\alpha, x; b; -1) = \int_x^\infty t^{\alpha-1}e^{-(b+1)t} = \begin{cases} \int_{(b+1)x}^\infty \left(\frac{u}{b+1}\right)^{\alpha-1} e^{-u}\frac{du}{b+1} = \frac{1}{(b+1)^\alpha}\Gamma(\alpha, (b+1)x) & \text{if } b+1>0 \\ \int_x^\infty t^{\alpha-1}dt = \frac{x^\alpha}{-\alpha} & \text{, if } b+1 = 0,\ \alpha < 0,\ x > 0 \\ \infty & \text{, else} \end{cases}$
   c) If $\beta < -1$ and $b > 0$, the change of variable $u = bt^{-\beta}$ provides the result $\Gamma(\alpha, x; b; \beta) = \frac{b^{\frac{\alpha}{\beta}}}{-\beta}\Gamma\left(-\frac{\alpha}{\beta}, bx^{-\beta}; b^{\frac{1}{\beta}}; \frac{1}{\beta}\right)$. More generally, the constraint $\beta < -1$ can be replaced by $\beta < 0$, but when $\beta \in [-1, 0[$, then $\frac{1}{\beta} \leq -1$ and thus the relation is of no help to compute $\Gamma(\alpha, x; b; \beta)$.

After these four preliminary checks, $\Gamma(\alpha, x; b; \beta)$ is either already computed, or is to be computed with $\beta > -1$ (as if $\beta < -1$ then $\frac{1}{\beta} > -1$), in which case the series expansion (3.4.19) can be exploited.

3.6.2.9 Large y asymptotics

Let y take large positive values:

$$\begin{aligned}
\widetilde{\mathcal{G}}(x, y; \nu) &= \int_x^\infty \frac{1}{(y+z)^\nu}e^{-z^2}dz = \frac{1}{y^\nu}\int_x^\infty \frac{1}{\left(1+\frac{z}{y}\right)^\nu}e^{-z^2}dz = \frac{1}{y^\nu}\int_x^\infty \exp\left(-\nu\ln\left(1+\frac{z}{y}\right) - z^2\right)dz \\
&\simeq \frac{1}{y^\nu}\int_x^\infty \exp\left(-\nu\frac{z}{y} - z^2\right)dz = \frac{1}{y^\nu}\exp\left(\frac{\nu^2}{4y^2}\right)\int_x^\infty \exp\left(-\left(z+\frac{\nu}{2y}\right)^2\right)dz \\
&= \frac{1}{y^\nu}\exp\left(\frac{\nu^2}{4y^2}\right)\int_{x+\frac{\nu}{2y}}^\infty e^{-z^2}dz = \frac{1}{y^\nu}\exp\left(\frac{\nu^2}{4y^2}\right)\sqrt{\pi}\left[1 - \mathcal{N}\left(\sqrt{2}\left(x+\frac{\nu}{2y}\right)\right)\right]
\end{aligned}$$

We have approximated $\ln\left(1+\frac{z}{y}\right)$ by its tangent in zero $\frac{z}{y}$ because of the Gaussian kernel that peaks at $z = 0$, which is where precision is needed. Ultimately, we have approximated $\frac{1}{\left(1+\frac{z}{y}\right)^\nu}$ by $e^{-\nu\frac{z}{y}}$. It seems preferable to the tangent in zero of $\frac{1}{\left(1+\frac{z}{y}\right)^\nu}$ directly (namely $1 - \nu\frac{z}{y}$) as $e^{-\nu\frac{z}{y}}$ behaves more like the initial function $\frac{1}{\left(1+\frac{z}{y}\right)^\nu}$. However, the difference, of course, vanishes for large y.

3.6.2.10 Proposition 3.4.8

Proof. Let $y < 0$ be given. We recall that, by definition of the function $\widetilde{\mathcal{G}}(x, y; \nu)$ (see equation (3.4.14)), $x + y \geq 0$. Thus, using equations (3.4.16) and (3.4.17):

$$\begin{aligned}
\widetilde{\mathcal{G}}(x,y;\nu) &= \frac{1}{2}e^{-y^2}\Gamma\left(\frac{1-\nu}{2}, (x+y)^2; -2y; -\frac{1}{2}\right) = \frac{1}{2}e^{-y^2}\int_{(x+y)^2}^{\infty} t^{-\frac{1+\nu}{2}}\exp\left(-t+2y\sqrt{t}\right)dt \\
&\leq \frac{1}{2}e^{-y^2}\exp\left(-(x+y)^2\right)\int_{(x+y)^2}^{\infty} t^{-\frac{1+\nu}{2}}\exp\left(2y\sqrt{t}\right)dt \\
&= e^{-y^2-(x+y)^2}(-2y)^{\nu-1}\int_{(-2y)(x+y)}^{\infty} u^{-\nu}e^{-u}du = e^{-y^2-(x+y)^2}(-2y)^{\nu-1}\Gamma\left(1-\nu, (-2y)(x+y)\right)
\end{aligned}$$

where we have performed the change of variable $u = -2y\sqrt{t}$ for the next-to-last equality. □

3.6.2.11 Proposition 3.4.9

Proof. Equation (3.4.13) yields:

$$\begin{aligned}
H(m_1,m_2,\sigma_1,\sigma_2) &= \int_0^{\infty}\mathcal{G}(x+m_1,\sigma_1)\,\Phi_N(x;m_2,\sigma_2)dx = \int_0^{\infty}\mathcal{G}(x+m_1,\sigma_1)\frac{1}{\sigma_2\sqrt{2\pi}}\exp\left(-\frac{(x-m_2)^2}{2\sigma_2^2}\right)dx \\
&= \int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^{\infty}\mathcal{G}\left(m_1+m_2+\sigma_2\sqrt{2}u,\sigma_1\right)\frac{1}{\sqrt{\pi}}\exp\left(-u^2\right)du = M\int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^{\infty}\mathcal{N}\left(\frac{\widetilde{x}-m_1-m_2-\sigma_2\sqrt{2}u}{\sigma_1}\right)\frac{1}{\sqrt{\pi}}e^{-u^2}du \\
&+ \frac{\gamma}{\pi\left(\sigma_1\sqrt{2}\right)^{\nu}}\int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^{\infty}\widetilde{\mathcal{G}}\left(\frac{\widetilde{x}-m_1-m_2-\sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}, \frac{m_1+m_2+\sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}; \nu\right)e^{-u^2}du \\
&= M\int_{-\infty}^{\frac{m_2}{\sigma_2}}\mathcal{N}\left(\frac{\widetilde{x}-m_1-m_2+\sigma_2 u}{\sigma_1}\right)\frac{1}{\sqrt{2\pi}}e^{-\frac{u^2}{2}}du + \frac{\gamma}{\pi\left(\sigma_1\sqrt{2}\right)^{\nu}}\widetilde{H}(m_1,m_2,\sigma_1,\sigma_2;\nu)
\end{aligned}$$

The integral above can be computed as follows:

$$\begin{aligned}
&\int_{-\infty}^{\frac{m_2}{\sigma_2}}\mathcal{N}\left(\frac{\widetilde{x}-m_1-m_2+\sigma_2 u}{\sigma_1}\right)\frac{1}{\sqrt{2\pi}}e^{-\frac{u^2}{2}}du = \int_{-\infty}^{\frac{m_2}{\sigma_2}}\int_{-\infty}^{\frac{\widetilde{x}-m_1-m_2+\sigma_2 u}{\sigma_1}}\frac{1}{\sqrt{2\pi}}e^{-\frac{z^2}{2}}dz\frac{1}{\sqrt{2\pi}}e^{-\frac{u^2}{2}}du \\
&= \frac{1}{2\pi}\int_{-\infty}^{\frac{m_2}{\sigma_2}}\int_{-\infty}^{\frac{\widetilde{x}-m_1-m_2}{\sigma_1}}\exp\left(-\frac{1}{2}\left[\left(v+\frac{\sigma_2}{\sigma_1}u\right)^2+u^2\right]\right)dvdu = \frac{1}{2\pi}\int_{-\infty}^{\frac{m_2}{\sigma_2}}\int_{-\infty}^{\frac{\widetilde{x}-m_1-m_2}{\sigma_1}}\exp\Bigg(\\
&-\frac{1}{2}\left[\left(1+\frac{\sigma_2^2}{\sigma_1^2}\right)u^2+2\frac{\sigma_2}{\sigma_1}uv+v^2\right]\Bigg)dvdu = \frac{1}{2\pi}\int_{-\infty}^{\frac{m_2}{\sigma_2}}\int_{-\infty}^{\frac{\widetilde{x}-m_1-m_2}{\sigma_1}}\exp\left(-\frac{1}{2}\begin{bmatrix} u & v\end{bmatrix}\Sigma^{-1}\begin{bmatrix} u \\ v\end{bmatrix}\right)dvdu
\end{aligned}$$

where $\Sigma^{-1} = \begin{bmatrix} 1+\frac{\sigma_2^2}{\sigma_1^2} & \frac{\sigma_2}{\sigma_1} \\ \frac{\sigma_2}{\sigma_1} & 1\end{bmatrix}$, and consequently $\Sigma = \begin{bmatrix} 1 & -\frac{\sigma_2}{\sigma_1} \\ -\frac{\sigma_2}{\sigma_1} & 1+\frac{\sigma_2^2}{\sigma_1^2}\end{bmatrix}$. Noting that $|\Sigma| = 1$, we have shown that $\int_{-\infty}^{\frac{m_2}{\sigma_2}}\mathcal{N}\left(\frac{\widetilde{x}-m_1-m_2+\sigma_2 u}{\sigma_1}\right)\frac{1}{\sqrt{2\pi}}e^{-\frac{u^2}{2}}du = \mathcal{N}\left(\frac{m_2}{\sigma_2}, \frac{\widetilde{x}-m_1-m_2}{\sigma_1}; 0, 0, \Sigma\right)$, which concludes the proof. □

3.6.2.12 Proposition 3.4.10

Proof. Using equations (3.4.24) and (3.4.21), and using the notation $\bar{m} = m_1 + m_2$ and $\bar{\sigma} = \sqrt{\sigma_1^2 + \sigma_2^2}$:

$$\begin{aligned}
\widetilde{H}\left(m_1, m_2, \sigma_1, \sigma_2; \nu\right) &= \int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^{\infty} \widetilde{\mathcal{G}}\left(\frac{\widetilde{x} - \bar{m} - \sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}, \frac{\bar{m} + \sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}; \nu\right) e^{-u^2} du \\
&= \int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^{\infty} \frac{1}{2} e^{-\left(\frac{\bar{m}+\sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}\right)^2} \sum_{n=0}^{\infty} \Gamma\left(\frac{1-\nu}{2} + \frac{n}{2}, \frac{\widetilde{x}^2}{2\sigma_1^2}\right) \frac{\left(2\frac{\bar{m}+\sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}\right)^n}{n!} e^{-u^2} du \\
&= \sum_{n=0}^{\infty} \frac{1}{2} \Gamma\left(\frac{1-\nu}{2} + \frac{n}{2}, \frac{\widetilde{x}^2}{2\sigma_1^2}\right) \frac{2^n}{n!} \int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^{\infty} \left(\frac{\bar{m} + \sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}\right)^n e^{-\left(\frac{\bar{m}+\sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}\right)^2 - u^2} du \qquad (3.6.4)
\end{aligned}$$

The integral can be computed as follows:

$$\begin{aligned}
&\int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^{\infty} \left(\frac{\bar{m} + \sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}\right)^n \exp\left(-\left(\frac{\bar{m} + \sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}\right)^2 - u^2\right) du \\
&= \int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^{\infty} \left(\frac{\bar{m} + \sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}\right)^n \exp\left(-\frac{\bar{\sigma}^2}{\sigma_1^2}\left(u + \frac{\bar{m}\sigma_2}{\bar{\sigma}^2\sqrt{2}}\right)^2 - \frac{\bar{m}^2}{2\bar{\sigma}^2}\right) du \\
&= e^{-\frac{\bar{m}^2}{2\bar{\sigma}^2}} \frac{\sigma_1}{\bar{\sigma}} \int_{\frac{m_1\sigma_2^2 - m_2\sigma_1^2}{\sigma_1\sigma_2\bar{\sigma}\sqrt{2}}}^{\infty} \left(\frac{\bar{m}\sigma_1}{\bar{\sigma}^2\sqrt{2}} + \frac{\sigma_2}{\bar{\sigma}} z\right)^n e^{-z^2} dz = e^{-\frac{\bar{m}^2}{2\bar{\sigma}^2}} \frac{\sigma_1}{\bar{\sigma}} \left(\frac{\sigma_2}{\bar{\sigma}}\right)^n \widetilde{\mathcal{G}}\left(\frac{m_1\sigma_2^2 - m_2\sigma_1^2}{\sigma_1\sigma_2\bar{\sigma}\sqrt{2}}, \frac{\bar{m}\sigma_1}{\sigma_2\bar{\sigma}\sqrt{2}}, -n\right)
\end{aligned}$$

where we have used the change of variable $z = \frac{\bar{\sigma}}{\sigma_1}\left(u + \frac{\bar{m}\sigma_2}{\bar{\sigma}^2\sqrt{2}}\right)$ for the next-to-last equality. Inserting this formula into equation (3.6.4) yields the desired result. □

3.6.2.13 Large $\bar{m}$ asymptotics

Let $\bar{m}$ take large positive values:

$$\begin{aligned}
\widetilde{H}\left(m_1, m_2, \sigma_1, \sigma_2; \nu\right) &= \int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^{\infty} \widetilde{\mathcal{G}}\left(\frac{\widetilde{x} - \bar{m} - \sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}, \frac{\bar{m} + \sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}; \nu\right) e^{-u^2} du \\
&= \int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^{\infty} \int_{\frac{\widetilde{x} - \bar{m} - \sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}}^{\infty} \frac{1}{\left(\frac{\bar{m}+\sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}} + z\right)^{\nu}} e^{-z^2} e^{-u^2} dz du \\
&= \frac{\left(\sigma_1\sqrt{2}\right)^{\nu}}{\bar{m}^{\nu}} \int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^{\infty} \int_{\frac{\widetilde{x} - \bar{m} - \sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}}^{\infty} \frac{1}{\left(1 + u\frac{\sigma_2\sqrt{2}}{\bar{m}} + z\frac{\sigma_1\sqrt{2}}{\bar{m}}\right)^{\nu}} e^{-z^2} e^{-u^2} dz du \\
&= \frac{1}{2} \frac{\left(\sigma_1\sqrt{2}\right)^{\nu}}{\bar{m}^{\nu}} \int_{-\infty}^{\frac{m_2}{\sigma_2}} \int_{-\infty}^{-\frac{\widetilde{x} - \bar{m} - \sigma_2 u}{\sigma_1}} \exp\left(-\nu \ln\left(1 - u\frac{\sigma_2}{\bar{m}} - z\frac{\sigma_1}{\bar{m}}\right) - \frac{z^2}{2} - \frac{u^2}{2}\right) dz du \\
&\simeq \frac{1}{2} \frac{\left(\sigma_1\sqrt{2}\right)^{\nu}}{\bar{m}^{\nu}} \int_{-\infty}^{\frac{m_2}{\sigma_2}} \int_{-\infty}^{-\frac{\widetilde{x} - \bar{m} - \sigma_2 u}{\sigma_1}} \exp\left(\nu\frac{\sigma_2}{\bar{m}}u + \nu\frac{\sigma_1}{\bar{m}}z - \frac{z^2}{2} - \frac{u^2}{2}\right) dz du \\
&= \frac{1}{2} \frac{\left(\sigma_1\sqrt{2}\right)^{\nu}}{\bar{m}^{\nu}} \int_{-\infty}^{\frac{m_2}{\sigma_2}} \int_{-\infty}^{\frac{\bar{m} - \widetilde{x}}{\sigma_1}} \exp\left(-\frac{1}{2}\left[u^2 + \left(v - \frac{\sigma_2}{\sigma_1}u\right)^2 - 2\nu\frac{\sigma_2}{\bar{m}}u - 2\nu\frac{\sigma_1}{\bar{m}}\left(v - \frac{\sigma_2}{\sigma_1}u\right)\right]\right) dv du \\
&= \frac{\pi\left(\sigma_1\sqrt{2}\right)^{\nu}}{\bar{m}^{\nu}} \exp\left(\frac{\nu^2\bar{\sigma}^2}{2\bar{m}^2}\right) \int_{-\infty}^{\frac{m_2}{\sigma_2}} \int_{-\infty}^{\frac{\bar{m} - \widetilde{x}}{\sigma_1}} \frac{1}{2\pi} \exp\left(-\frac{1}{2}\left[\begin{bmatrix} u - \bar{u} & v - \bar{v} \end{bmatrix} \Sigma^{-1} \begin{bmatrix} u - \bar{u} \\ v - \bar{v} \end{bmatrix}\right]\right) dv du
\end{aligned}$$

where $\bar{u} = \frac{\nu\sigma_2}{\bar{m}}$, $\bar{v} = \frac{\nu\bar{\sigma}^2}{\sigma_1\bar{m}}$ and $\Sigma^{-1} = \begin{bmatrix} 1+\frac{\sigma_2^2}{\sigma_1^2} & -\frac{\sigma_2}{\sigma_1} \\ -\frac{\sigma_2}{\sigma_1} & 1 \end{bmatrix}$, and consequently $\Sigma = \begin{bmatrix} 1 & \frac{\sigma_2}{\sigma_1} \\ \frac{\sigma_2}{\sigma_1} & 1+\frac{\sigma_2^2}{\sigma_1^2} \end{bmatrix}$.

In particular, $|\Sigma| = 1$. Thus, we have shown that, for large positive $\bar{m}$:

$$\widetilde{H}\left(m_1, m_2, \sigma_1, \sigma_2; \nu\right) \simeq \frac{\pi\left(\sigma_1\sqrt{2}\right)^{\nu}}{\bar{m}^{\nu}} \exp\left(\frac{\nu^2\bar{\sigma}^2}{2\bar{m}^2}\right) \mathcal{N}\left(\frac{m_2}{\sigma_2}, \frac{\bar{m}-\widetilde{x}}{\sigma_1}; \begin{bmatrix} \frac{\nu\sigma_2}{\bar{m}} & \frac{\nu\bar{\sigma}^2}{\sigma_1\bar{m}} \end{bmatrix}, \begin{bmatrix} 1 & \frac{\sigma_2}{\sigma_1} \\ \frac{\sigma_2}{\sigma_1} & 1+\frac{\sigma_2^2}{\sigma_1^2} \end{bmatrix}\right)$$

3.6.2.14 Proposition 3.4.11

Proof. Suppose $\nu \geq 0$. Then:

$$\begin{aligned}
&\widetilde{H}\left(m_1, m_2, \sigma_1, \sigma_2; \nu\right) = \int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^{\infty} \int_{\frac{\widetilde{x}-\bar{m}-\sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}}^{\infty} \frac{1}{\left(\frac{\bar{m}+\sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}+z\right)^{\nu}} e^{-z^2} e^{-u^2} dz du \\
&= \int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^{\infty} \int_{\frac{\widetilde{x}}{\sigma_1\sqrt{2}}}^{\infty} \frac{1}{v^{\nu}} \exp\left(-\left(v - \frac{\bar{m}+\sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}\right)^2 - u^2\right) dv du \\
&\leq \left(\frac{\sigma_1\sqrt{2}}{\widetilde{x}}\right)^{\nu} \int_{-\frac{m_2}{\sigma_2\sqrt{2}}}^{\infty} \int_{\frac{\widetilde{x}}{\sigma_1\sqrt{2}}}^{\infty} \exp\left(-\left(v - \frac{\bar{m}+\sigma_2\sqrt{2}u}{\sigma_1\sqrt{2}}\right)^2 - u^2\right) dv du \\
&= \left(\frac{\sigma_1\sqrt{2}}{\widetilde{x}}\right)^{\nu} \int_{-\infty}^{\frac{m_2}{\sigma_2}} \int_{-\infty}^{-\frac{\widetilde{x}}{\sigma_1}+\frac{\bar{m}}{\sigma_1}} \exp\left(-\left(-\frac{y}{\sqrt{2}} + \frac{\sigma_2 x}{\sigma_1\sqrt{2}}\right)^2 - \frac{x^2}{2}\right) dy dx \\
&= \left(\frac{\sigma_1\sqrt{2}}{\widetilde{x}}\right)^{\nu} \int_{-\infty}^{\frac{m_2}{\sigma_2}} \int_{-\infty}^{-\frac{\widetilde{x}}{\sigma_1}+\frac{\bar{m}}{\sigma_1}} \exp\left(-\frac{1}{2}\begin{bmatrix} x & y \end{bmatrix}\begin{bmatrix} 1+\frac{\sigma_2^2}{\sigma_1^2} & -\frac{\sigma_2}{\sigma_1} \\ -\frac{\sigma_2}{\sigma_1} & 1 \end{bmatrix}\begin{bmatrix} x \\ y \end{bmatrix}\right) dy dx \\
&= \pi\left(\frac{\sigma_1\sqrt{2}}{\widetilde{x}}\right)^{\nu} \mathcal{N}\left(\frac{m_2}{\sigma_2}; -\frac{\widetilde{x}}{\sigma_1}+\frac{\bar{m}}{\sigma_1}; \begin{bmatrix} 0 & 0 \end{bmatrix}; \begin{bmatrix} 1 & \frac{\sigma_2}{\sigma_1} \\ \frac{\sigma_2}{\sigma_1} & 1+\frac{\sigma_2^2}{\sigma_1^2} \end{bmatrix}\right)
\end{aligned}$$

□

3.6.2.15 Proposition 3.4.12

Proof. First, differentiating equation (3.4.14) yields $\frac{\partial\widetilde{\mathcal{G}}}{\partial x}\left(x, y; \nu\right) = -\frac{1}{(x+y)^{\nu}} e^{-x^2}$ and $\frac{\partial\widetilde{\mathcal{G}}}{\partial y}\left(x, y; \nu\right) = -\nu\widetilde{\mathcal{G}}\left(x, y; \nu+1\right)$. Using these expressions to differentiate equation (3.4.13):

$$\begin{aligned}
&\frac{\partial\mathcal{G}}{\partial m}\left(m, \sigma\right) = -\frac{M}{\sigma}\Phi_N\left(\frac{\widetilde{x}-m}{\sigma}; 0, 1\right) + \frac{\gamma}{\left(\sigma\sqrt{2}\right)^{\nu}\sqrt{\pi}}\left(-\frac{1}{\sigma\sqrt{2}}\frac{\partial\widetilde{\mathcal{G}}}{\partial x}\left(\frac{\widetilde{x}-m}{\sigma\sqrt{2}}, \frac{m}{\sigma\sqrt{2}}; \nu\right) + \right. \\
&\left.\frac{1}{\sigma\sqrt{2}}\frac{\partial\widetilde{\mathcal{G}}}{\partial y}\left(\frac{\widetilde{x}-m}{\sigma\sqrt{2}}, \frac{m}{\sigma\sqrt{2}}; \nu\right)\right) = -M\Phi_N\left(\widetilde{x}; m, \sigma\right) + \frac{\gamma}{\widetilde{x}^{\nu}\sigma\sqrt{2\pi}}\exp\left(-\frac{\left(\widetilde{x}-m\right)^2}{2\sigma^2}\right) - \\
&\frac{\nu\gamma}{\left(\sigma\sqrt{2}\right)^{\nu+1}\sqrt{\pi}}\widetilde{\mathcal{G}}\left(\frac{\widetilde{x}-m}{\sigma\sqrt{2}}, \frac{m}{\sigma\sqrt{2}}; \nu+1\right) = -\frac{\nu\gamma}{\left(\sigma\sqrt{2}\right)^{\nu+1}\sqrt{\pi}}\widetilde{\mathcal{G}}\left(\frac{\widetilde{x}-m}{\sigma\sqrt{2}}, \frac{m}{\sigma\sqrt{2}}; \nu+1\right)
\end{aligned}$$

because $\frac{\gamma}{\widetilde{x^\nu}} = M$. Thus, differentiating equation (3.4.10) and using equation (3.4.24):

$$\frac{\partial H}{\partial m_1}(m_1,m_2,\sigma_1,\sigma_2) = \int_0^\infty \frac{\partial \mathcal{G}}{\partial m}(x+m_1,\sigma_1)\,\Phi_N(x;m_2,\sigma_2)\,dx = -\frac{\nu\gamma}{\left(\sigma_1\sqrt{2}\right)^{\nu+1}\sqrt{\pi}}\times$$

$$\int_0^\infty \widetilde{\mathcal{G}}\left(\frac{\widetilde{x}-m_1-x}{\sigma_1\sqrt{2}}, \frac{m_1+x}{\sigma_1\sqrt{2}}; \nu+1\right)\Phi_N(x;m_2,\sigma_2)\,dx = -\frac{\nu\gamma}{\pi\left(\sigma_1\sqrt{2}\right)^{\nu+1}}\widetilde{H}(m_1,m_2,\sigma_1,\sigma_2;\nu+1)$$

$$\frac{\partial H}{\partial m_2}(m_1,m_2,\sigma_1,\sigma_2) = -\int_0^\infty \mathcal{G}(x+m_1,\sigma_1)\frac{\partial \Phi_N}{\partial x}(x;m_2,\sigma_2)\,dx$$
$$= \mathcal{G}(m_1,\sigma_1)\,\Phi_N(0;m_2,\sigma_2) + \int_0^\infty \frac{\partial \mathcal{G}}{\partial x}(x+m_1,\sigma_1)\,\Phi_N(x;m_2,\sigma_2)\,dx$$
$$= \mathcal{G}(m_1,\sigma_1)\,\Phi_N(0;m_2,\sigma_2) - \frac{\nu\gamma}{\pi\left(\sigma_1\sqrt{2}\right)^{\nu+1}}\widetilde{H}(m_1,m_2,\sigma_1,\sigma_2;\nu+1)$$

Now, for $1 \le i \le n$ and $1 \le k \le n$, we differentiate equations (3.4.11):

$$\frac{\partial \bar{m}_i^n}{\partial D_t} = 0 \quad , \quad \frac{\partial \bar{m}_i^n}{\partial C_t^k} = e^{-\alpha_k(T-t)}\mathbf{1}_{\{i\le k\}} \quad , \quad \frac{\partial \bar{m}_1^{i,D}}{\partial D_t} = -e^{-\alpha_D(T-t)} \quad , \quad \frac{\partial \bar{m}_1^{i,D}}{\partial C_t^k} = e^{-\alpha_k(T-t)}\mathbf{1}_{\{k\le i\}}$$

Using all these results, the differentiation of equations (3.4.9) and (3.4.12) is straightforward. □

3.6.3 Algorithms

This appendix provides the guidelines to compute efficiently the trickiest quantities from section 3.4.2. The other quantities are straightforward to compute from these building blocks.

3.6.3.1 Extended Incomplete Gamma Function

Following the algorithm described in Section 3.4.2.5, we give details a numerically stable computation of $\Gamma\left(\alpha,x;b;-\frac{1}{2}\right)$. We exploit the sum (3.4.19) and the recurrence relation (3.4.20). Denote as x_n the n-th term of the sum :

$$x_n := \Gamma\left(\alpha+\frac{n}{2},x\right)\frac{(-b)^n}{n!}$$

Then, for $n \ge 2$, using (3.4.20) :

$$x_n = \frac{b^2}{n(n-1)}a_n x_{n-2} + c_n$$

where $a_n = a + \frac{n-2}{2}$ and $c_n = e^{-x}x^{a+\frac{n-2}{2}}\frac{(-b)^n}{n!}$, computed using $a_n = a_{n-1} + \frac{1}{2}$ and $c_n = -\frac{b}{n}\sqrt{x}c_{n-2}$. Using these relations, the implementation of an algorithm to compute the EIG function is straightforward.

3.6.3.2 Extended Incomplete Goodwin-Staton Integral

Combining the different methods described in section 3.4.2.5, the implementation of an algorithm to compute $\widetilde{\mathcal{G}}(x,y;\nu)$ is also straightforward. The only remaining detail is to choose when to switch from one method to the other. We propose to use the large positive y approximation when $y > 10$, to test the large negative y upper bound whenever $y < 0$ (before using the series expansion if the upper bound is consider not small enough, say $> 10^{-30}$), and the series expansion in the remaining cases.

3.6.3.3 $\widetilde{H}(m_1, m_2, \sigma_1, \sigma_2; \nu)$

The first step is to build an efficient implementation of the sum (3.4.25). Let skip the multiplicative constant for the moment. Denote as x_n the n-th term of the sum :

$$x_n := \Gamma\left(a + \frac{n}{2}, x\right) \widetilde{\mathcal{G}}(X, Y; -n) \frac{b^n}{n!}$$

where x, b, X and Y are constants. In order to exploit the recurrence relations (3.4.15) and (3.4.20), for efficiency, while being numerically stable, we decompose x_n as:

$$x_n = p_n q_n \quad , \qquad p_n = \Gamma\left(a + \tfrac{n}{2}, x\right) \frac{\left(\sqrt{b}\right)^n}{\prod_{i=1}^n \sqrt{i}} \quad , \qquad q_n = \widetilde{\mathcal{G}}(X, Y; -n) \frac{\left(\sqrt{b}\right)^n}{\prod_{i=1}^n \sqrt{i}}$$

For $n \geq 2$, the recurrence relations result in:

$$p_n = b \frac{a-1+\frac{n}{2}}{\sqrt{n}\sqrt{n-1}} p_{n-2} + c_n \quad , \qquad q_n = \frac{b}{2} \frac{\sqrt{n-1}}{\sqrt{n}} q_{n-2} + Y \frac{\sqrt{b}}{\sqrt{n}} q_{n-1} + d_n$$

where $c_n := x^{a-1} e^{-x} \frac{\left(\sqrt{bx}\right)^n}{\prod_{i=1}^n \sqrt{i}}$ and $d_n := \frac{1}{2} \frac{e^{-X^2}}{X+Y} \frac{\left((X+Y)\sqrt{b}\right)^n}{\prod_{i=1}^n \sqrt{i}}$ are computed using $c_n = \frac{\sqrt{bx}}{\sqrt{n}} c_{n-1}$ and $d_n = \frac{(X+Y)\sqrt{b}}{\sqrt{n}} d_{n-1}$.

Using these relations, the implementation of an algorithm to compute the sum (3.4.25) is straightforward. Combining it with the different algorithms from Section 3.4.2.6, $\widetilde{H}(m_1, m_2, \sigma_1, \sigma_2; \nu)$ is easily computed. Again it remains to choose when to switch from one method to the other. We propose to use the large positive y approximation when $y > 10 \max(\sigma_1, \sigma_2)$, to test, if $\nu > 0$, the large negative y upper bound whenever $y < 0$ (before using the series expansion if the upper bound is consider not small enough, say $> 10^{-30}$), and the series expansion in the remaining cases.

4 A probabilistic numerical method for optimal multiple switching problem in high dimension

In this chapter, we present a probabilistic numerical algorithm combining dynamic programming, Monte Carlo simulations and local basis regressions to solve non-stationary optimal multiple switching problems in infinite horizon. We provide the rate of convergence of the method in terms of the time step used to discretize the problem, of the regression basis used to approximate conditional expectations, and of the truncating time horizon. To make the method viable for problems in high dimension and long time horizon, we extend a memory reduction method to the general Euler scheme, so that, when performing the numerical resolution, the storage of the Monte Carlo simulation paths is not needed. Then, we apply this algorithm to a model of optimal investment in power plants in dimension eight, i.e. with two different technologies and six random factors.

4.1 Introduction

This chapter presents a probabilistic numerical method for multiple switching problem. Our approach in this chapter takes advantage of the considerable progress made in the last ten years by numerical methods for high-dimensional American options valuation problems. For an up-to-date state of the art on this subject, the reader is referred to the recent book [31].

In this chapter, we first adapt the resolution of American options problems by Monte-Carlo methods and regression ([83, 102]), to the more general class of optimal switching problems. The crucial choice of regression basis is done here in the light of the work of [28], so as to obtain a stable algorithm suited to high-dimensional problems, aiming at the best possible numerical complexity. The memory complexity, often acknowledged as the major drawback of this Monte Carlo approach (see [32]), is drastically slashed by generalizing the memory reduction method from [36, 37, 38] to any stochastic differential equation. We provide a rigorous and comprehensive analysis of the rate of convergence of our algorithm, taking advantage of the works of, most notably, [27], [100] and [55]. Note that such unusual features as infinite horizon and non-stationarity are encompassed here.

Finally, we apply our algorithm to a long-term investment model for electricity generation based on a structural model for the spot price of electricity developed in [3] and [1]. This model has been shown to suitably reproduce the statistical and dynamical properties of the spot price of electricity. Nevertheless, to suit the purpose of long-term electricity price modeling, it has been adapted and extended in several directions (cointegrated fuels and CO_2 prices, stochastic availability rate of production capacities, new scarcity function). The resolution of this problem using our algorithm is illustrated on a simple numerical example with two different technologies, leading to an eight-dimensional problem (demand, CO_2 price, and, for each technology, fuel price, random outages and the controlled installed capacity). The time evolution of the distribution of power prices and of the generation mix is illustrated on a forty-year time horizon. To the knowledge of the authors, the highest dimension considered so far in the case of long-term investment models in electricity generation was three ([87, 23]).

The contribution of the chapter is twofold. Firstly, it provides a comprehensive analysis of convergence of a regression-based Monte-Carlo algorithm for a class of infinite horizon optimal multiple switching problems, large enough to handle realistic short term profit functions and investment cost structures with possible seasonality patterns. Secondly, we implement successfully our algorithm to a new stylized investment model for electricity generation, by adapting and generalizing a memory reduction method. A numerical resolution of this investment problem with our algorithm is illustrated on a specific example, providing, among many other outputs, an electricity spot price dynamics consistent with the investment decision process in power generation.

The outline of the chapter is the following. Section 4.2 presents the class of optimal switching problems studied here, including the detailed list of assumptions considered. Section 4.3 describes the resolution algorithm and analyzes its rate of convergence, in terms of the discretization step, of the choice of regression basis, and of the truncating time horizon. Section 4.4 details the computational complexity of the algorithm, as well as its memory complexity, along with the construction of the memory reduction method. In Section 4.5, we implement and illustrate numerically our algorithm on an investment model in electricity generation based on an extended structural model of power spot price. Finally, Section 4.6 concludes the chapter.

Notation

Here are some notation that will be used throughout the chapter:

- The notation $\mathbf{1}\{.\}$ stands for the indicator function.
- Throughout the chapter, $C > 0$ denotes a generic constant whose value may differ from line to line, but which does not depend on any parameter of our scheme.
- For any stochastic process $X = (X_s)_{s\geq 0}$ taking values in a given set $\mathcal{X}$, and any $(t,x) \in \mathbb{R}_+ \times \mathcal{X}$, we denote as $X^{t,x} = \left(X_s^{t,x}\right)_{s\geq t}$ the stochastic process with the same dynamics as X, but starting from x at time t: $X_t^{t,x} = x$.
- For any $(a,b) \in \mathbb{R}\times\mathbb{R}$, $a \wedge b := \min(a,b)$ and $a \vee b := \max(a,b)$.
- $\forall p \geq 1$, the norms $\|.\|_p$ and $\|.\|_{L_p}$ denote respectively the $p-$norm and the L_p- norm: $\forall x \in \mathbb{R}^n$ and any $\mathbb{R}$-valued random variable X such that $\mathbb{E}\left[|X|^p\right] < \infty$:

$$\|x\|_p = \left(\sum_{i=1}^n |x_i|^p\right)^{\frac{1}{p}} \quad , \quad \|X\|_{L_p} = \mathbb{E}\left[|X|^p\right]^{\frac{1}{p}}$$

  We recall that $\forall p \geq 1$, $\forall x \in \mathbb{R}^n$, $\|x\|_p \leq \|x\|_1 \leq n^{\frac{p-1}{p}} \|x\|_p$

4.2 Optimal switching problem

4.2.1 Formulation

Fix a filtered probability space $\left(\Omega, \mathcal{F}, \mathbb{F} = (\mathcal{F}_t)_{t\geq 0}, \mathbb{P}\right)$, where $\mathbb{F}$ satisfies the usual conditions of right-continuity and $\mathbb{P}$-completeness. We consider the following general class of (non-stationary) optimal switching problems:

$$v(t,x,i) = \sup_{\alpha \in \mathcal{A}_{t,i}} \mathbb{E}\left[\int_t^\infty f\left(s, X_s^{t,x}, I_s^\alpha\right) ds - \sum_{\tau_n \geq t} k\left(\tau_n, \zeta_n\right)\right] \tag{4.2.1}$$

where:

- $X^{t,x} = \left(X_s^{t,x}\right)_{s \geq t}$ is an $\mathbb{R}^d$-valued, $\mathbb{F}$-adapted Markovian diffusion starting from $X_t = x \in \mathbb{R}^d$, with generator $\mathcal{L}$.
- $I^\alpha = (I_s^\alpha)_{s \geq 0}$ is a càd-làg, $\mathbb{R}^{d'}$-valued, $\mathbb{F}$-adapted piecewise constant process. It is controlled by a strategy α, described below. We suppose it can only take values into a fixed finite set $\mathbb{I}_q = \{i_1, i_2, \ldots, i_q\}$, $q \in \mathbb{N}^*$ with $i_1 = 0 \left(\in \mathbb{R}^{d'}\right)$, which means that equation (4.2.1) corresponds to an optimal switching problem.
- An impulse control strategy α corresponds to a sequence $(\tau_n, \iota_n)_{n \in \mathbb{N}}$ of increasing stopping times $\tau_n \geq 0$, and $\mathcal{F}_{\tau_n}$-measurable random variables ι_n valued in $\mathbb{I}_q$. Using this sequence, $I^\alpha = (I_s^\alpha)_{s \geq 0}$ is defined as follows:

$$I_s^\alpha = \iota_0 \mathbf{1}\{0 \leq s < \tau_0\} + \sum_{n \in \mathbb{N}} \iota_n \mathbf{1}\{\tau_n \leq s < \tau_{n+1}\} \in \mathbb{I}_q$$

  Alternatively, α can be described by the sequence $(\tau_n, \zeta_n)_{n \in \mathbb{N}}$, where $\zeta_n := \iota_n - \iota_{n-1}$ (and $\zeta_0 := 0$). Using this alternative sequence, I^α can be written as follows:

$$I_s^\alpha = \iota_0 + \sum_{\tau_n \leq s} \zeta_n \in \mathbb{I}_q$$

- $\mathcal{A}$ is the set of admissible strategies: a strategy α belongs to $\mathcal{A}$ if $\tau_n \to +\infty$ a.s. as $n \to \infty$.
- For any $(t, i) \in \mathbb{R}_+ \times \mathbb{I}_q$, the set $\mathcal{A}_{t,i} \subset \mathcal{A}$ is defined as the subset of admissible strategies α such that $I_t^\alpha = i$.
- f and k are $\mathbb{R}$-valued measurable functions.

4.2.2 Assumptions

We complete the above formulation with the following relevant assumptions.

Assumption 2. *[Diffusion] The $\mathbb{R}^d$-valued uncontrolled process X is a diffusion process, governed by the dynamics*

$$\begin{aligned} dX_s &= b(s, X_s)\, ds + \sigma(s, X_s)\, dW_s \\ X_0 &= x_0 \in \mathbb{R}^d \end{aligned} \tag{4.2.2}$$

where W is a d-dimensional Brownian motion, and b and σ are respectively $\mathbb{R}^d$-valued and $\mathbb{R}^{d \times d}$-valued functions.

Assumption 3. *[Lipschitz] The functions $b : \mathbb{R}_+ \times \mathbb{R}^d \to \mathbb{R}^d$ and $\sigma : \mathbb{R}_+ \times \mathbb{R}^d \to \mathbb{R}^{d \times d}$ are Lipschitz-continuous (uniformly in t) with linear growth: $\exists C_b, C_\sigma > 0$ s.t. $\forall t \in \mathbb{R}_+$, $\forall (x, x') \in \left(\mathbb{R}^d\right)^2$:*

$$\begin{aligned} |b(t,x) - b(t,x')| &\leq C_b |x - x'| \\ |b(t,x)| &\leq C_b (1 + |x|) \\ |\sigma(t,x) - \sigma(t,x')| &\leq C_\sigma |x - x'| \\ |\sigma(t,x)| &\leq C_\sigma (1 + |x|) \end{aligned}$$

Remark 4.2.1. Assumption 3 is sufficient to prove the existence and uniqueness of a strong solution to the SDE (4.2.2) (see for instance Theorem 4.5.3 in [73]).

Remark 4.2.2. Under Assumption 3, there exist, for every $p \geq 1$, positive constants C_p and ρ_p such that $\forall s \geq t \geq 0$ and $\forall x \in \mathbb{R}^d$:

$$\mathbb{E}\left[\left|X_s^{t,x}\right|^p\right] \leq C_p\left(1+|x|^p\right)\exp\left(\rho_p\left(s-t\right)\right) \tag{4.2.3}$$

(use Burkholder-Davis-Gundy inequality and Gronwall's Lemma, see for instance [73] Theorem 4.5.4 for the even power case).

Assumption 4. *[Lipschitz&Discount] The functions f and k decrease exponentially in time: $\exists \rho > 0$ s.t. $\forall (t,x,i,j) \in \mathbb{R}_+ \times \mathbb{R}^d \times (\mathbb{I}_q)^2$:*

$$\begin{aligned} f\left(t,x,i\right) &= e^{-\rho t}\tilde{f}\left(t,x,i\right) \\ k\left(t,j-i\right) &= e^{-\rho t}\tilde{k}\left(t,j-i\right) \end{aligned}$$

where the functions $\tilde{f}$ and $\tilde{k}$ are Lipschitz continuous with linear growth:
$\exists C_f, C_k > 0$ s.t. $\forall \{(t,x,i,j),(t',x',i',j')\} \in \left\{\mathbb{R}_+ \times \mathbb{R}^d \times (\mathbb{I}_q)^2\right\}^2$:

$$\begin{aligned} \left|\tilde{f}\left(t,x,i\right)-\tilde{f}\left(t',x',i'\right)\right| &\leq C_f\left(|t-t'|+|x-x'|+|i-i'|\right) \\ \left|\tilde{f}\left(t,x,i\right)\right| &\leq C_f\left(1+|x|\right) \\ \left|\tilde{k}\left(t,j-i\right)-\tilde{k}\left(t',j'-i'\right)\right| &\leq C_k\left(|t-t'|+|(j-i)-(j'-i')|\right) \end{aligned}$$

Moreover, we assume in the following that $\rho > \rho_1$ where ρ_1 is defined in equation (4.2.3).

Assumption 5. *[Fixed costs] The cost function $k : \mathbb{R}_+ \times \mathbb{R}^{d'} \to \mathbb{R}_+$ is such that:*

- $\forall t \in \mathbb{R}_+$, $k\left(t,0\right) = 0$.
- $\exists \kappa > 0$ s.t. $\forall t \in \mathbb{R}_+$, $\forall (i,j) \in (\mathbb{I}_q)^2$, $\{i \neq j\} \Rightarrow \left\{\tilde{k}(t,j-i) \geq \kappa\right\}$.
- (triangular inequality) $\forall t \in \mathbb{R}_+$, $\forall (i,j,k) \in (\mathbb{I}_q)^3$ with $i \neq j$ and $j \neq k$:
$$k(t,k-i) < k(t,j-i) + k(t,k-j)\ .$$

Remark 4.2.3. The economic interpretations of Assumption 5 are the following:

1. There is no cost for not switching, but any switch incurs at least a positive fixed cost.
2. At any given date, it is always cheaper to switch directly from i to k than to switch first from i to j and then from j to k.

Remark 4.2.4. Under those standard assumptions, the value function v from equation (4.2.1) is well-defined and finite. Indeed, using equation (4.2.3), $\forall (t_0,t,x,i) \in \mathbb{R}_+ \times \mathbb{R}_+ \times \mathbb{R}^d \times \mathbb{R}^{d'}$ with $t_0 \leq t$ and $\forall \alpha \in \mathcal{A}_{t_0,i}$:

$$\begin{aligned} \mathbb{E}\left[\int_t^\infty \left|f\left(s,X_s^{t_0,x},I_s^\alpha\right)\right| ds\right] &\leq C_f \int_t^\infty e^{-\rho s}\left(1+\mathbb{E}\left[\left|X_s^{t_0,x}\right|\right]\right) ds \\ &\leq C_f\left(e^{-\rho t}+\left(1+|x|\right)\int_t^\infty e^{-\rho s}e^{\rho_1(s-t_0)}ds\right) \\ &\leq C_f\left(1+|x|\right)e^{-\bar{\rho}t-\rho_1 t_0} \end{aligned} \tag{4.2.4}$$

where $\bar{\rho} := \rho - \rho_1 > 0$ (Assumption 4). In particular, the costs being positive (Assumption 5), and recalling (4.2.1), it holds that:

$$v(t,x,i) \leq C_f (1 + |x|) e^{-\rho t} \tag{4.2.5}$$

4.2.3 Outline of the solution

From a theoretical point of view, the value functions $v_i := v(.,.,i)$, $i \in \mathbb{I}_q$ from equation (4.2.1) are known to satisfy (under suitable conditions on $f_i(.,.) := f(.,.,i)$ and k, see for instance [98] in a much more general setting) the following Hamilton-Jacobi-Bellman Quasi-Variational Inequalities (HJBQVI): $\forall (t,x,i) \in \mathbb{R}_+ \times \mathbb{R}^d \times \mathbb{I}_q$

$$\min \left\{ -\frac{\partial v_i}{\partial t}(t,x) - \mathcal{L} v_i(t,x) - f_i(t,x) \ , \ v_i(t,x) - \max_{j \in \mathbb{I}_P, j \neq i} \left(v_j(t,x) - k(t, j-i) \right) \right\} = 0 \tag{4.2.6}$$

together with suitable limit condition, which ensure existence and unicity of the solution to this system (cf. [63] for instance).

Alternatively, the process $v(t, X_t, i)$, $t \geq 0$ can be characterized as the solution of a particular Reflected Backward Stochastic Differential Equation ([64, 48]).

Moreover, the value function (4.2.1) satisfies the well-known dynamic programming principle, i.e., for any stopping time $\tau \geq t$:

$$v(t,x,i) = \sup_{\alpha \in \mathcal{A}_{t,i}} \mathbb{E} \left[\int_t^\tau f\left(s, X_s^{t,x}, I_s^\alpha\right) ds - \sum_{t \leq \tau_n \leq \tau} k(\tau_n, \zeta_n) + v\left(\tau, X_\tau^{t,x}, I_\tau^\alpha\right) \right] . \tag{4.2.7}$$

From a practical point of view, apart from a few simple examples in low-dimension, finding directly the solution of the HJBQVI (4.2.6) is usually infeasible, and the numerical PDE tools become cumbersome and inefficient in the multi-dimensional setting. Instead, probabilistic methods based on (4.2.7), in the spirit of [32], are usually more practical and versatile.

Indeed, as the diffusion X is not controlled, this optimal switching problem can be seen as an extended American option problem. This suggests that, up to some adjustments, the probabilistic numerical tools developed in this context (see [28] for instance) may be adapted to solve (4.2.1).

To be more specific, define a finite time grid $\Pi = \{t_0 = 0 < t_1 < \ldots < t_N = T\}$ for a fixed $T > 0$, and consider the function v^Π defined as v (equation (4.2.1)) but with the strategy set $\mathcal{A}$ replaced by $\mathcal{A}^\Pi \subset \mathcal{A}$, defined as the subset of strategies that can be modified only at the dates $t \in \Pi$. In other words, the switching decisions can now only take place on the time grid Π. Suppose, moreover, that the cost function k is such that at most one switch can occur on a given date t_k (triangular condition). Then $\forall i \in \mathbb{I}_q$, $\forall x \in \mathbb{R}^d$, and $\forall t_k \in \Pi$, the dynamic programming principle (4.2.7) becomes:

$$v^\Pi(t_k, x, i) = \max_{j \in \mathbb{I}_q} \left\{ E_j(t_k, x) - k(t_k, j-i) \mathbf{1}_{\{j \neq i\}} \right\} \tag{4.2.8}$$

where:

$$E_i(T,x) := \mathbb{E} \left[\int_T^\infty f_i\left(s, X_s^{T,x}\right) ds \right] \tag{4.2.9}$$

$$E_i(t_k,x) := \mathbb{E} \left[\int_{t_k}^{t_{k+1}} f_i\left(s, X_s^{t_k,x}\right) ds + v^\Pi\left(t_{k+1}, X_{t_{k+1}}^{t_k,x}, i\right) \right] \ , \ k = N-1, \ldots, 0 \tag{4.2.10}$$

which is explicit in the sense that $v^{\Pi}(t_k,.,.)$ directly depends on $v^{\Pi}(t_{k+1},.,.)$.

In practice, apart from the potential approximation of the stochastic process X and of the final values (4.2.9), the difficulty lies in the efficient computation of the conditional expectations (4.2.10).

In the American option literature, various approaches have been developed to solve (4.2.8) efficiently. Notable examples are the least-squares regression approach ([83, 102]), the quantization approach and the Malliavin calculus based formulation (see [28] for a thorough comparison and improvements of these techniques). In the spirit of [33], one may also consider non-parametric regression (see [74] and [101]) combined with speeding up techniques like Kd-trees or the Fast Gauss Transform in the case of kernel regression.

Here, we intend to solve (4.2.1) on numerical applications which bears the particularity of handling stochastic processes in high dimension ($\dim(X) = d \gg 3$, with however $\dim(I) = d' \approx 3$, see Section 4.5). For such problems, the most adequate technique so far seems to be the local regression method developed in [28]. We are thus going to make use of this specific method to solve (4.2.8) in practice.

In the following, we provide a detailed analysis of the above suggested computational method.

4.3 Numerical approximation and convergence analysis

This section is devoted to the precise description of the resolution of (4.2.1), along the lines of the discussions from Subsection 4.2.3. Moreover, the convergence rate of the proposed algorithm will be precisely assessed.

4.3.1 Approximations

Recall equation (4.2.1) defining the value function $v(t,x,i)$:

$$v(t,x,i) = \sup_{\alpha \in \mathcal{A}_{t,i}} \mathbb{E}\left[\int_t^{\infty} f\left(s, X_s^{t,x}, I_s^{\alpha}\right) ds - \sum_{\tau_n \geq t} k(\tau_n, \zeta_n)\right] \tag{4.3.1}$$

We are going to consider the following sequence of approximations:

- *[Finite time horizon]* The time horizon will be truncated to a finite horizon T.
- *[Time discretization]* The continuous state process X and investment process I will be discretized with a time step h.
- *[Space localization]* The $\mathbb{R}^d$- valued process X will be projected into a bounded domain $\mathcal{D}_\varepsilon$, parameterized by ε.
- *[Conditional expectation approximation]* The conditional expectation involved in the dynamic programming equation will be replaced by an empirical least-squares regression, computed on a bundle of M Monte Carlo trajectories, on a finite basis of local hypercubes with edges of size δ.

The rate of convergence of the algorithm will then be provided, as a function of these five numerical parameters: T, h, ε, M and δ.

4.3.1.1 Finite time horizon

The first step is to reduce the set of strategies to a finite horizon:

$$v_T(t,x,i) = \sup_{\alpha\in\mathcal{A}_{t,i}^T} \mathbb{E}\left[\int_t^T f\left(s, X_s^{t,x}, I_s^\alpha\right) ds - \sum_{t\le\tau_n\le T} k(\tau_n,\zeta_n) + g_f\left(T, X_T^{t,x}, I_T^\alpha\right)\right] \quad (4.3.2)$$

$$g_f(T,x,i) := \mathbb{E}\left[\int_T^\infty f\left(s, X_s^{T,x}, i\right) ds\right] \quad (4.3.3)$$

where $0 \le t \le T < +\infty$, and $\mathcal{A}_{t,i}^T \subset \mathcal{A}_{t,i}$ is the subset of strategies without switches strictly after time T. Hence the final value g_f corresponds to the remaining gain after T.

Alternatively, one may choose, for convenience, another final value g instead of g_f, as long as it is Lipschitz-continuous and satisfies a suitable condition (cf. equation (4.3.20)). The set of such functions will be denoted as Θ_{g_f}. The difference between the two value functions is quantified in Proposition 4.3.1.

This freedom on the final values will be used in practice to avoid a computation on an infinite interval $[T,\infty[$ as in the definition of g_f.

From now on, we choose and fix one such $g \in \Theta_{g_f}$.

4.3.1.2 Time discretization

Then, we discretize the time segment $[0,T]$. Introduce a time grid $\Pi = \{t_0 = 0 < t_1 < \ldots < t_N = T\}$ with constant mesh h. Consider the following approximation:

$$v_\Pi(t,x,i) = \sup_{\alpha\in\mathcal{A}_{t,i}^\Pi} \mathbb{E}\left[\int_t^T f\left(s, X_s^{t,x}, I_s^\alpha\right) ds - \sum_{t\le\tau_n\le T} k(\tau_n,\zeta_n) + g\left(T, X_T^{t,x}, I_T^\alpha\right)\right] \quad (4.3.4)$$

where $\mathcal{A}_{t,i}^\Pi \subset \mathcal{A}_{t,i}^T$ is the subset of strategies such that switches can only occur at dates $\tau_n \in \Pi \cap [t,T]$.

Now, with a slight abuse of notation, we can safely switch from the notation $\alpha = (\tau_n,\zeta_n)_{n\ge0}$ to the notation $\alpha = (\tau_n,\iota_n)_{n\ge0}$ (remember Subsection 4.2.1), replacing the quantity $\sum_{t\le\tau_n\le T} k(\tau_n,\zeta_n)$ by $\sum_{t\le\tau_n\le T} k\left(\tau_n, I_{\tau_n-h}^\alpha, I_{\tau_n}^\alpha\right)$ or by $\sum_{t\le\tau_n\le T} k(\tau_n,\iota_{n-1},\iota_n)$, where $k(t,i,j) = k(t,j-i)$. The error between v_T and v_Π is quantified in Proposition 4.3.2.

Next we also approximate the stochastic process X by its Euler scheme $\bar{X} = \left(\bar{X}_s\right)_{0\le s\le T}$, with dynamics:

$$\begin{aligned} d\bar{X}_s &= b\left(\pi(s), \bar{X}_{\pi(s)}\right) ds + \sigma\left(\pi(s), \bar{X}_{\pi(s)}\right) dW_s\,,\ 0 \le s \le T \\ \bar{X}_0 &= x_0 \in \mathbb{R}^d \end{aligned} \quad (4.3.5)$$

where $\forall s \in [0,T]$, $\pi(s) := \max\{t \in \Pi; t \le s\}$. More precisely, we substitute the piecewise constant $\bar{X}_{\pi(t)}$ for X_t. (Note that at this stage the process I^α is already piecewise constant). The new value function reads:

$$\bar{v}_\Pi(t,x,i) = \sup_{\alpha\in\mathcal{A}_{t,i}^\Pi} \mathbb{E}\left[\int_t^T f\left(\pi(s), \bar{X}_{\pi(s)}^{t,x}, I_s^\alpha\right) ds - \sum_{t\le\tau_n\le T} k(\tau_n,\iota_{n-1},\iota_n) + g\left(T, \bar{X}_T^{t,x}, I_T^\alpha\right)\right] \quad (4.3.6)$$

The error between v_Π and $\bar{v}_\Pi$ is computed in Proposition 4.3.3.

4.3.1.3 Space localization

In order to derive a rigorous convergence analysis, our subsequent choices in terms of conditional expectation approximation (Subsection 4.3.1.4 below) and specific choice of basis (Assumption 6) will require the underlying state process X to lie into a bounded set (cf. equation (4.3.16)). Thus, we explicitly build such an approximation and assess the associated error. Remark, though, that the usefulness of this step is more theoretical (for a proper convergence speed to hold) than practical (on a finite sample, this localization step would be somewhat redundant, and may safely be omitted).

Hence, let $\mathcal{D} = [-C, C]^d$, $C > 0$, be a bounded convex domain of $\mathbb{R}^d$ that contains x_0. For every $i = 1, \dots, d$, define the stopping time τ_i and the killed process $\bar{X}^{i,\mathcal{D}} = \left(\bar{X}_t^{i,\mathcal{D}}\right)_{0 \leq t \leq T}$ as follows:

$$\begin{aligned} \tau_i &= \inf\left\{t \in [0,T]\,; \bar{X}_t^i \notin [-C, C]\right\} \\ \bar{X}_t^{i,\mathcal{D}} &= \bar{X}_{t \wedge \tau_i}^i \,,\, t \in [0,T] \end{aligned}$$

In other words the $d-$dimensional process $\bar{X}^{\mathcal{D}} = \left(\bar{X}_t^{i,\mathcal{D}}\right)_{0 \leq t \leq T}^{1 \leq i \leq d}$ is equal to $\bar{X}$ most of the time (i.e. when $\bar{X}_t \in \mathcal{D}$), except when one component of $\bar{X}_t$ gets outside $\mathcal{D}$, in which case the corresponding component of $\bar{X}_t^{\mathcal{D}}$ is killed and remains on the border of the domain $\mathcal{D}$ (the other components being unaffected). In particular, the killed process $\bar{X}^{\mathcal{D}}$ is bounded and Markovian.

Finally, one can choose C sufficiently large such that

$$\sup_{t \in [0,T]} \mathbb{E}\left[\left|\bar{X}_t - \bar{X}_t^{\mathcal{D}}\right|\right] \leq \varepsilon \tag{4.3.7}$$

for some $\varepsilon > 0$ (in which case $C = C(T, \varepsilon)$). This is the parameterization of the domain $\mathcal{D} = \mathcal{D}_\varepsilon$ that we adopt in the following.

Define $\bar{v}_\Pi^\varepsilon$ as the value function $\bar{v}_\Pi$ from equation (4.3.6) with $\left(\bar{X}_{\pi(t)}\right)_{0 \leq t \leq T}$ replaced by $\left(\bar{X}_{\pi(t)}^{\mathcal{D}_\varepsilon}\right)_{0 \leq t \leq T}$. The error between those two value functions is computed in Proposition 4.3.4.

Example 4.3.1. To clarify this construction of space localization, we explicit it on the very simple example of a d-dimensional standard brownian motion $(W_t)_{t \in [0,T]}$. In this case, $\bar{X}_t = X_t = W_t$. In this example, equation (4.3.7) can be shown to hold by choosing for instance $C(T, \varepsilon) = \sqrt{2T \log\left(\frac{2dT}{\pi\varepsilon}\right)}$.

4.3.1.4 Conditional expectation approximation

From now on, in order to prevent the notation from becoming too cumbersome and clumsy, we are going to drop the ε index in the following, i.e. $\bar{X}_t$ will stand for $\bar{X}_t^\varepsilon$, and $\bar{v}_\Pi$ for $\bar{v}_\Pi^\varepsilon$.

For the fully discretized problem (4.3.6), the dynamic programming principle (4.2.8) becomes:

$$\begin{aligned} &\bar{v}_\Pi(T, x, i) = g(T, x, i) \\ &\bar{v}_\Pi(t_n, x, i) = \max_{j \in \mathbb{I}_q}\left\{hf(t_n, x, j) - k(t_n, i, j) + \mathbb{E}\left[\bar{v}_\Pi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n, x}, j\right)\right]\right\} \,, n = N-1, \dots, 0 \end{aligned} \tag{4.3.8}$$

The last step is to approximate the conditional expectation appearing in equation (4.3.8). As discussed in Subsection 4.2.3, we choose to approximate it by least-squares regression. Consider

basis functions $(e_k(x))_{1\leq k\leq K}$, $K\in\mathbb{N}\cup\{+\infty\}$, $x\in\mathbb{R}^d$. For suitable functions $\varphi:\Pi\times\mathbb{R}^d\times\mathbb{I}_q\to\mathbb{R}$, define:

$$\tilde{\lambda}=\tilde{\lambda}_i^{t_n}(\varphi):=\arg\min_{\lambda\in\mathbb{R}^K}\mathbb{E}\left[\left(\varphi\left(t_{n+1},\bar{X}_{t_{n+1}},i\right)-\sum_{k=1}^{K}\lambda_k e_k\left(\bar{X}_{t_n}\right)\right)^2\right] \tag{4.3.9}$$

Now, before using this projection, it is more cautious to truncate it within known bounds (see [27, 57, 100]). Hence, suppose that there exist known bounds $\underline{\Gamma}^{t_n,x}(\varphi)$ and $\overline{\Gamma}^{t_n,x}(\varphi)$ around $\mathbb{E}\left[\varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x},i\right)\right]$:

$$\underline{\Gamma}^{t_n,x}(\varphi)\leq\mathbb{E}\left[\varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x},i\right)\right]\leq\overline{\Gamma}^{t_n,x}(\varphi) \tag{4.3.10}$$

Then, $\forall i\in\mathbb{I}_q$ the quantity $\mathbb{E}\left[\varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x},i\right)\right]$ is approximated by:

$$\tilde{\mathbb{E}}\left[\varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x},i\right)\right]:=\underline{\Gamma}^{t_n,x}(\varphi)\vee\sum_{k=1}^{K}\tilde{\lambda}_k e_k(x)\wedge\overline{\Gamma}^{t_n,x}(\varphi) \tag{4.3.11}$$

which is used to define the next approximation $\tilde{v}_\Pi$ of the value function:

$$\begin{aligned}\tilde{v}_\Pi(T,x,i)&=g(T,x,i)\\ \tilde{v}_\Pi(t_n,x,i)&=\max_{j\in\mathbb{I}_q}\left\{hf(t_n,x,j)-k(t_n,i,j)+\tilde{\mathbb{E}}\left[\tilde{v}_\Pi\left(t_{n+1},\bar{X}_{t_{n+1}}^{t_n,x},j\right)\right]\right\},\ n=N-1,\dots,0\end{aligned} \tag{4.3.12}$$

Interesting discussions on the choice of function basis can be found in [28]. In particular they advocate bases of local polynomials, which is numerically efficient and well-suited to tackle large-dimensional problems (see Subsection 4.4.1). However, for the sake of simplicity, we will restrict our study in this section to a basis of indicator functions on local hypercubes (cf. [100] and the numerical experiments of [57]) (which is the simplest example of local polynomials). Assumption 6 below states this specific choice.

Assumption 6. *The regression basis is set to a basis of indicator function on disjoint local hypercubes, as described in Definition 4.3.1 below.*

Definition 4.3.1. For every $t_n\in\Pi$, consider a partition of the domain $\mathcal{D}_\varepsilon$ into hypercubes $\left(B_{t_n}^k\right)_{k=1,\dots,K_\varepsilon}$, i.e., $\cup_{k=1,\dots,K_\varepsilon}B_{t_n}^k=\mathcal{D}_\varepsilon$ and $B_{t_n}^i\cap B_{t_n}^j=\emptyset\ \forall i\neq j$. It may be deterministic, or computed from a sample of $\bar{X}$. We only assume that there exists $(\underline{\delta},\delta)\in\mathbb{R}_+^2$ with $\underline{\delta}\leq\delta$ such that the lengths of the edges of the hypercubes, in each dimension, belong to $[\underline{\delta},\delta]$ (in particular, the volume of each hypercube $B_{t_n}^k$ belongs to $\left[\underline{\delta}^d,\delta^d\right]$). This freedom over the definition of the partition enables to encompass to some extend the kind of adaptative partition described in [28]. Then, the basis functions considered here are defined by $e_{t_n}^k(x):=\mathbf{1}\left\{x\in B_{t_n}^k\right\}$, $x\in\mathbb{R}^d$, $1\leq k\leq K_\varepsilon$.

Under Assumption 6, the error between $\bar{v}_\Pi$ and $\tilde{v}_\Pi$ is computed in Proposition 4.3.5.

Finally, let $\left(\bar{X}_{t_n}^m\right)_{1\leq n\leq N}^{1\leq m\leq M}$ be a finite sample of size M of paths of the process $\bar{X}$. The final step is to replace the regression (4.3.9) by a regression on this sample:

$$\hat{\lambda}=\hat{\lambda}_i^{t_n}(\varphi):=\arg\min_{\lambda\in\mathbb{R}^K}\frac{1}{M}\sum_{m=1}^{M}\left[\left(\varphi\left(t_{n+1},\bar{X}_{t_{n+1}}^m,i\right)-\sum_{k=1}^{K}\lambda_k e_k\left(\bar{X}_{t_n}^m\right)\right)^2\right]. \tag{4.3.13}$$

Then $\forall i \in \mathbb{I}_q$ the quantity $\mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x}, i\right)\right]$ is approximated by:

$$\hat{\mathbb{E}}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x}, i\right)\right] := \underline{\Gamma}^{t_n,x}(\varphi) \vee \sum_{k=1}^{K} \hat{\lambda}_k e_k(x) \wedge \overline{\Gamma}^{t_n,x}(\varphi) \tag{4.3.14}$$

leading to the final, computable approximation $\hat{v}_\Pi$ of the value function:

$$\begin{aligned} \hat{v}_\Pi(T,x,i) &= g(T,x,i) \\ \hat{v}_\Pi(t_n,x,i) &= \max_{j\in\mathbb{I}_q}\left\{hf(t_n,x,j) - k(t_n,i,j) + \hat{\mathbb{E}}\left[\hat{v}_\Pi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x}, j\right)\right]\right\}, \ n = N-1,\dots,0 \end{aligned} \tag{4.3.15}$$

Under Assumption 6, the error between $\tilde{v}_\Pi$ and $\hat{v}_\Pi$ is given in Proposition 4.3.6. This proposition will make use of the following quantity:

$$p(T,\delta,\varepsilon) := \min_{t\in\Pi} \min_{B_t^k \subset \mathcal{D}^\varepsilon} \mathbb{P}\left(\bar{X}_t \in B_t^k\right) \tag{4.3.16}$$

which is strictly positive, as the domain $\mathcal{D}_\varepsilon$ is (purposely) bounded.

Example 4.3.2. Carrying on with Example 4.3.1 of a d-dimensional Brownian motion, we explicit a lower bound for $p(T,\delta,\varepsilon)$ in this simple case. First, $\mathbb{P}\left(W_T \in B_T^k\right) = \int_{B_T^k} f_{W_T}(x)\,dx$ where f_{W_T} is the density of W_T. As $\forall k$, $B_T^k \subset \mathcal{D}^\varepsilon$, with $C(T,\varepsilon) = \sqrt{2T\log\left(\frac{2dT}{\pi\varepsilon}\right)}$, it holds that $\forall x \in \mathcal{D}^\varepsilon$, $f_{W_T}(x) \geq \left(f_{W_T^1}(C(T,\varepsilon))\right)^d = \frac{\varepsilon^d}{(2d)^d T^{\frac{3d}{2}}}$. Hence $\mathbb{P}\left(W_t \in B_t^k\right) \geq \frac{\varepsilon^d}{(2d)^d T^{\frac{3d}{2}}}\mathrm{Vol}\left(B_t^k\right) \geq \frac{\varepsilon^d}{(2d)^d T^{\frac{3d}{2}}}\underline{\delta}^d$. As a conclusion, $p(T,\delta,\varepsilon) \geq \frac{\varepsilon^d}{(2d)^d T^{\frac{3d}{2}}}\underline{\delta}^d$. Remark however that this lower bound is very crude, and that it can be very far below $p(T,\delta,\varepsilon)$ for large δ.

Combining all these results, we obtain a rate of convergence of $\hat{v}_\Pi$ towards v:

Theorem 4.3.1. *$\forall p \geq 1$, $\exists C_p > 0$ such that:*

$$\begin{aligned} &\left\|\max_{i\in\mathbb{I}_q} |v(t_0,x_0,i) - \hat{v}_\Pi(t_0,x_0,i)|\right\|_{L_p} \\ \leq\ & C_p\left\{(1+|x_0|)e^{-\bar{\rho}T} + (1+|x_0|)^{\frac{3}{2}}\sqrt{h} + \varepsilon + \frac{\delta}{h} + \frac{1+C(T,\varepsilon)}{h\sqrt{M}p(T,\delta,\varepsilon)^{1-\frac{1}{p\vee 2}}} + \frac{1+C(T,\varepsilon)}{hMp(T,\delta,\varepsilon)}\right\} \end{aligned}$$

In particular, $\hat{v}_\Pi(0,x_0,i) \to_{L_p} v(0,x_0,i)$ uniformly in $i \in \mathbb{I}_q$ when $T \to \infty$, $h \to 0$, $\varepsilon \to 0$, $\delta \to 0$ and $M \to \infty$ with $\frac{\delta}{h} \to 0$, $\frac{1+C(T,\varepsilon)}{h\sqrt{M}p(T,\delta,\varepsilon)^{1-\frac{1}{p\vee 2}}} \to 0$ and $\frac{1+C(T,\varepsilon)}{hMp(T,\delta,\varepsilon)} \to 0$.

The proof of Theorem 4.3.1 will be given at the end of the next Subsection 4.3.2.

Remark 4.3.1. If the cost function k (recall Assumption 4) were to depend on x, then, under a usual Lipschitz condition on k (similar to that of f), Theorem 4.3.1 would still hold, replacing only the term $(1+|x_0|)^{\frac{3}{2}}\sqrt{h}$ by $\left(1+|x_0|^{\frac{5}{2}}\right)\sqrt{h\log\left(\frac{2T}{h}\right)}$ (recalling Remark 4.3.4).

Remark 4.3.2. The adaptative local basis can be such that each hypercube contains approximately the same number of Monte Carlo trajectories (see [28]). This means that $\frac{1}{p(T,\delta,\varepsilon)} \sim b$ where b is the number of functions in the regression basis. With this remark in mind, the leading error term in Theorem 4.3.1 behaves like $\frac{\sqrt{b}}{h\sqrt{M}}$ for $p=2$. This is close to the corresponding statistical error term in [82] ($\sqrt{\frac{b\log(M)}{hM}}$) in the context of BSDEs. The advantage of their approach is that it can handle any (orthonormal) regression basis, while our approach (in the context of optimal switching) provides a bound on the L_p error for every $p \geq 1$.

Example 4.3.3. In the case of a d-dimensional Brownian motion, the rate of convergence of Theorem 4.3.1 can be explicited further, using the upper bound on $C(T,\varepsilon)$ from Example 4.3.1 and the lower bound on $p(T,\delta,\varepsilon)$ from Example 4.3.2. Moreover, one can express the rate of convergence as a function of only one parameter, choosing the five numerical parameters T, h, ε, δ and M accordingly. For instance, assuming $\underline{\delta} = \delta$, and minimizing over δ, h, ε and T, one can get a convergence rate upper bounded by $C_p(1+|x|)^{\frac{3}{2}} z$ by choosing $M \sim z^{-\frac{1}{2}[6(d+1)]^2}$. This is admittedly highly demanding in terms of sample size M, but remember that this expression suffers from the crude lower bound on $p(T,\delta,\varepsilon)$ we chose previously.

4.3.2 Convergence analysis

From now on, we suppose that all the assumptions from Subsection 4.2.2 are in force.

4.3.2.1 Finite time horizon

Lemma 4.3.1. *There exists $C > 0$ such that $\forall (t,x,i) \in [0,T] \times \mathbb{R}^d \times \mathbb{R}^{d'}$:*

$$0 \leq v(t,x,i) - v_T(t,x,i) \leq C(1+|x|)\, e^{-\bar{\rho} t \vee T - \rho_1 t}.$$

Proof. First, we introduce the following notations:

$$H(t,T,x,\alpha) := \int_t^T f\left(s, X_s^{t,x}, I_s^\alpha\right) ds - \sum_{t<\tau_n \leq T} k(\tau_n, \zeta_n) \tag{4.3.17}$$

$$J(t,T,x,\alpha) := \mathbb{E}\left[H(t,T,x,\alpha)\right] \tag{4.3.18}$$

for any admissible strategy $\alpha \in \mathcal{A}_{t,i}$. In particular:

$$v(t,x,i) = \sup_{\alpha \in \mathcal{A}_{t,i}} J(t,+\infty,x,\alpha) \quad , \quad v_T(t,x,i) = \sup_{\alpha \in \mathcal{A}_{t,i}^T} J(t,+\infty,x,\alpha)\,. \tag{4.3.19}$$

Fix $(t,x,i) \in \mathbb{R}_+ \times \mathbb{R}^d \times \mathbb{R}^{d'}$. Using equation (4.3.19):

$$v_T(t,x,i) = \sup_{\alpha \in \mathcal{A}_{t,i}^T} J(t,\infty,x,\alpha) \leq \sup_{\alpha \in \mathcal{A}_{t,i}} J(t,\infty,x,\alpha) = v(t,x,i)$$

which provides the first inequality. Consider now the second inequality. Choose $\varepsilon > 0$. From the definition of v (equation (4.3.1)) there exists a strategy $\alpha^\varepsilon \in \mathcal{A}_{t,i}$ such that:

$$v(t,x,i) - \varepsilon \leq J(t,\infty,x,\alpha^\varepsilon) \leq v(t,x,i)$$

Define the truncated strategy $\alpha_T^\varepsilon \in \mathcal{A}_{t,i}^T$ such that $\forall s \in [t,T]$, $I_s^{\alpha_T^\varepsilon} = I_s^{\alpha^\varepsilon}$ and $\forall s > T$, $I_s^{\alpha_T^\varepsilon} = I_T^{\alpha^\varepsilon}$. In order not to mix up the variables τ_n and ζ_n from different strategies, we add the name of the

strategy in index when needed. Then:

$$
\begin{aligned}
&H\left(t,\infty,x,\alpha^{\varepsilon}\right)-H\left(t,\infty,x,\alpha_{T}^{\varepsilon}\right)\\
&=\left\{\int_{t}^{\infty}f\left(s,X_{s}^{t,x},I_{s}^{\alpha^{\varepsilon}}\right)ds-\sum_{\tau_{n}^{\alpha^{\epsilon}}\geq t}k\left(\tau_{n}^{\alpha^{\epsilon}},\zeta_{n}^{\alpha^{\epsilon}}\right)\right\}-\left\{\int_{t}^{\infty}f\left(s,X_{s}^{t,x},I_{s}^{\alpha_{T}^{\varepsilon}}\right)ds-\sum_{\tau_{n}^{\alpha_{T}^{\epsilon}}\geq t}k\left(\tau_{n}^{\alpha_{T}^{\epsilon}},\zeta_{n}^{\alpha_{T}^{\epsilon}}\right)\right\}\\
&=\left\{\int_{t}^{\infty}f\left(s,X_{s}^{t,x},I_{s}^{\alpha^{\varepsilon}}\right)ds-\sum_{\tau_{n}^{\alpha^{\epsilon}}\geq t}k\left(\tau_{n}^{\alpha^{\epsilon}},\zeta_{n}^{\alpha^{\epsilon}}\right)\right\}\\
&-\left\{\int_{t}^{t\vee T}f\left(s,X_{s}^{t,x},I_{s}^{\alpha^{\varepsilon}}\right)ds+\int_{t\vee T}^{\infty}f\left(s,X_{s}^{t,x},I_{t\vee T}^{\alpha^{\varepsilon}}\right)ds-\sum_{t\vee T\geq\tau_{n}^{\alpha^{\epsilon}}\geq t}k\left(\tau_{n}^{\alpha^{\epsilon}},\zeta_{n}^{\alpha^{\epsilon}}\right)\right\}\\
&=\int_{t\vee T}^{\infty}f\left(s,X_{s}^{t,x},I_{s}^{\alpha^{\varepsilon}}\right)ds-\int_{t\vee T}^{\infty}f\left(s,X_{s}^{t,x},I_{t\vee T}^{\alpha^{\varepsilon}}\right)ds-\sum_{\tau_{n}^{\alpha^{\epsilon}}\geq t\vee T}k\left(\tau_{n}^{\alpha^{\epsilon}},\zeta_{n}^{\alpha^{\epsilon}}\right)\\
&\leq\int_{t\vee T}^{\infty}f\left(s,X_{s}^{t,x},I_{s}^{\alpha^{\varepsilon}}\right)ds-\int_{t\vee T}^{\infty}f\left(s,X_{s}^{t,x},I_{t\vee T}^{\alpha^{\varepsilon}}\right)ds
\end{aligned}
$$

as $k\left(s,0\right)=0$ and $k\geq0$ (Assumption 5). Hence, using Jensen's inequality and equation (4.2.4), $\exists C>0$ such that

$$
\begin{aligned}
\left|J\left(t,\infty,x,\alpha^{\varepsilon}\right)-J\left(t,\infty,x,\alpha_{T}^{\varepsilon}\right)\right| &\leq \mathbb{E}\left[\left|H\left(t,\infty,x,\alpha^{\varepsilon}\right)-H\left(t,\infty,x,\alpha_{T}^{\varepsilon}\right)\right|\right]\\
&\leq \mathbb{E}\left[\int_{t\vee T}^{\infty}\left|f\left(s,X_{s}^{t,x},I_{s}^{\alpha^{\varepsilon}}\right)\right|ds\right]+\mathbb{E}\left[\int_{t\vee T}^{\infty}\left|f\left(s,X_{s}^{t,x},I_{t\vee T}^{\alpha^{\varepsilon}}\right)\right|ds\right]\\
&\leq C\left(1+\left|x\right|\right)e^{-\bar{\rho}t\vee T-\rho_{1}t}
\end{aligned}
$$

Finally, given that $v\left(t,x,i\right)\leq\varepsilon+J\left(t,\infty,x,\alpha^{\varepsilon}\right)$ and $v_{T}\left(t,x,i\right)\geq J\left(t,\infty,x,\alpha_{T}^{\varepsilon}\right)$, the following holds:

$$
\begin{aligned}
v\left(t,x,i\right)-v_{T}\left(t,x,i\right) &\leq \varepsilon+J\left(t,\infty,x,\alpha^{\varepsilon}\right)-J\left(t,\infty,x,\alpha_{T}^{\varepsilon}\right)\\
&\leq \varepsilon+C\left(1+\left|x\right|\right)e^{-\bar{\rho}t\vee T-\rho_{1}t}.
\end{aligned}
$$

Since this is true for any $\varepsilon>0$, and that C, ρ and ρ_1 do not depend on ε, the proposition is proved. □

Now, we focus on the final boundary g_f. For the time being, denote the value function (4.3.2) as $v_{T}^{g_f}$ to emphasize the dependence of v on the terminal condition. As a consequence of equation (4.2.4), $\forall\left(x,i\right)\in\mathbb{R}^{d}\times\mathbb{I}_{q}$:

$$
\left|g_{f}(T,x,i)\right|\leq C\left(1+\left|x\right|\right)e^{-\rho T} \tag{4.3.20}
$$

Hence, define the class Θ_{g_f} of Lipschitz functions from $\mathbb{R}_{+}\times\mathbb{R}^{d}\times\mathbb{I}_{q}$ into $\mathbb{R}$ such that $\forall g\in\Theta_{g_f}$, $\forall\left(T,x,x',i\right)\in\mathbb{R}_{+}\times\mathbb{R}^{d}\times\mathbb{R}^{d}\times\mathbb{I}_{q}$:

$$
\begin{aligned}
\left|g(T,x,i)-g(T,x',i)\right| &\leq C_{g}e^{-\rho T}\left|x-x'\right| &(4.3.21)\\
\left|g(T,0,i)\right| &\leq C_{g}e^{-\rho T} &(4.3.22)
\end{aligned}
$$

for some $C_g>0$. In particular, the growth rate of such functions is at most linear in x:

$$
\left|g(T,x,i)\right|\leq C_{g}e^{-\rho T}\left(1+\left|x\right|\right). \tag{4.3.23}
$$

Obviously $g_f\in\Theta_{g_f}$. Now, for any $g\in\Theta_{g_f}$, denote as v_{T}^{g} the value function defined as in equation (4.3.2) with g instead of g_f. We are going to show that the precise approximation error due to the choice of final value g does not matter much as long as g is chosen in this class Θ_{g_f}.

Lemma 4.3.2. *There exists $C > 0$ such that $\forall\,(t,x,i) \in \mathbb{R}_+ \times \mathbb{R}^d \times \mathbb{I}_q$:*

$$\left|v_T^{g_f}(t,x,i) - v_T^{g}(t,x,i)\right| \leq C\,(1+|x|)\,e^{-\bar{\rho}t\vee T-\rho_1 t}$$

Proof. Fix $(t,x,i) \in \mathbb{R}_+ \times \mathbb{R}^d \times \mathbb{I}_q$. To shorten the proof, we assume that $v_T^{g_f}$ (resp. v_T^{g}) admits an optimal strategy $\alpha_f^* \in \mathcal{A}_{t,i}^T$ (resp. $\alpha^* \in \mathcal{A}_{t,i}^T$) (this assumption can then be relaxed using ε-optimal strategies as in the proof of Proposition 4.3.1)[1]. Therefore, recalling the notations H (equation (4.3.17)) and J (equation (4.3.18)) introduced in the proof of Lemma 4.3.1:

$$\begin{aligned}
v_T^{g_f}(t,x,i) - v_T^{g}(t,x,i) &= J\left(t,T,x,\alpha_f^*\right) + \mathbb{E}\left[g_f\left(T,X_T^{t,x},I_T^{\alpha_f^*}\right)\right] - J(t,T,x,\alpha^*) - \mathbb{E}\left[g\left(T,X_T^{t,x},I_T^{\alpha^*}\right)\right] \\
&= J\left(t,T,x,\alpha_f^*\right) + \mathbb{E}\left[g\left(T,X_T^{t,x},I_T^{\alpha_f^*}\right)\right] - J\left(t,T,x,\alpha^*\right) - \mathbb{E}\left[g\left(T,X_T^{t,x},I_T^{\alpha^*}\right)\right] \\
&+ \mathbb{E}\left[g_f\left(T,X_T^{t,x},I_T^{\alpha_f^*}\right) - g\left(T,X_T^{t,x},I_T^{\alpha^*}\right)\right] \\
&\leq \mathbb{E}\left[g_f\left(T,X_T^{t,x},I_T^{\alpha_f^*}\right) - g\left(T,X_T^{t,x},I_T^{\alpha^*}\right)\right] \\
&\leq C\left(1+\mathbb{E}\left[\left|X_T^{t,x}\right|\right)\right)e^{-\rho T} \;\leq\; C\,(1+|x|)\,e^{-\bar{\rho}t\vee T-\rho_1 t}
\end{aligned}$$

Symmetrically, the same inequality holds for $v_T^{g}(t,x,i) - v_T^{g_f}(t,x,i)$, ending the proof. □

Proposition 4.3.1. *There exists $C > 0$ independent of T such that $\forall\,(t,x,i) \in \mathbb{R}_+ \times \mathbb{R}^d \times \mathbb{I}_q$ and $\forall g \in \Theta_{g_f}$:*

$$\left|v(t,x,i) - v_T^{g}(t,x,i)\right| \leq C\,(1+|x|)\,e^{-\bar{\rho}t\vee T-\rho_1 t}$$

Proof. Combine Lemmas 4.3.1 and 4.3.2. □

From now on, we choose and keep one final value function $g \in \Theta_{g_f}$, and remove the index g from the notation of v and its subsequent approximations.

4.3.2.2 Time Discretization

Proposition 4.3.2. *There exists a positive constant C such that for any $(t,x,i) \in \Pi \times \mathbb{R}^d \times \mathbb{I}_q$:*

$$\left|v_T(t,x,i) - v_\Pi(t,x,i)\right| \leq Ce^{-\rho t}\left(1+|x|^{\frac{3}{2}}\right)h^{\frac{1}{2}} \tag{4.3.24}$$

Proof. Under the assumptions from Subsection 4.2.2, one can apply Theorem 3.1 in [55] to prove (4.3.24), noticing that the cost function k does not depend on the state variable x.

Use the discounting factor in the definition of f to factor the $e^{-\rho t}$ term and to get that C does not depend on T. □

Remark 4.3.3. Another alternative to get this rate of $h^{\frac{1}{2}}$ is to work with the reflected BSDE representation of v_Π, as in [32] (adapting [27]) or [39].

[1] Note that under the assumptions from Subsection 4.2.2, one may use Theorem 3.1 from [67] to get the existence of a unique optimal strategy α^* for the value function (4.3.2), satisfying $\mathbb{E}\left[\left|\sum_{0\leq\tau_n^{\alpha^*}\leq T} k\left(\tau_n^{\alpha^*},\zeta_n^{\alpha^*}\right)\right|^2\right] < \infty$

Remark 4.3.4. Were the cost function k to depend on the state variable, the upper bound in Proposition 4.3.2 would only be $Ce^{-\rho t}\left(1+|x|^{\frac{5}{2}}\right)\left(h\log\left(\frac{2T}{h}\right)\right)^{\frac{1}{2}}$, as stated in [55] (making use of results from [50]).

Proposition 4.3.3. *There exists $C>0$ such that for any $(t,x,i)\in\Pi\times\mathbb{R}^d\times\mathbb{I}_q$:*

$$\left|v_\Pi(t,x,i)-\bar{v}_\Pi(t,x,i)\right|\leq Ce^{-\rho t}h^{\frac{1}{2}}$$

Proof. T and g being fixed, we can define, in the spirit of equations (4.3.17) and (4.3.18), the following quantities:

$$H(t,x,\alpha):=\int_t^T f\left(s,X_s^{t,x},I_s^\alpha\right)ds-\sum_{t\leq\tau_n\leq T}k\left(\tau_n,\iota_{n-1},\iota_n\right)+g\left(T,X_T^{t,x},I_T^\alpha\right) \tag{4.3.25}$$

$$J(t,x,\alpha):=\mathbb{E}\left[H(t,x,\alpha)\right] \tag{4.3.26}$$

$$\bar{H}(t,x,\alpha):=\int_t^T f\left(\pi(s),\bar{X}_{\pi(s)}^{t,x},I_s^\alpha\right)ds-\sum_{t\leq\tau_n\leq T}k\left(\tau_n,\iota_{n-1},\iota_n\right)+g\left(T,\bar{X}_T^{t,x},I_T^\alpha\right) \tag{4.3.27}$$

$$\bar{J}(t,x,\alpha):=\mathbb{E}\left[\bar{H}(t,x,\alpha)\right] \tag{4.3.28}$$

for any admissible strategy $\alpha\in\mathcal{A}_{t,i}^\Pi$. For these discretized problems, the existence of optimal controls α^* and $\bar{\alpha}^*$ is granted. Hence:

$$\begin{aligned}
v_\Pi(t,x,i)-\bar{v}_\Pi(t,x,i)&=J(t,x,\alpha^*)-\bar{J}(t,x,\bar{\alpha}^*)\\
&=J(t,x,\alpha^*)-\bar{J}(t,x,\alpha^*)+\left\{\bar{J}(t,x,\alpha^*)-\bar{J}(t,x,\bar{\alpha}^*)\right\}\\
&\leq J(t,x,\alpha^*)-\bar{J}(t,x,\alpha^*)\\
&=\int_t^T e^{-\rho s}\mathbb{E}\left[\tilde{f}\left(s,X_s^{t,x},I_s^{\alpha^*}\right)-\tilde{f}\left(\pi(s),\bar{X}_{\pi(s)}^{t,x},I_s^{\alpha^*}\right)\right]ds\\
&+\mathbb{E}\left[g\left(T,X_T^{t,x},I_T^{\alpha^*}\right)-g\left(T,\bar{X}_T^{t,x},I_T^{\alpha^*}\right)\right]\\
&\leq C_f\int_t^T e^{-\rho s}\mathbb{E}\left[\left|X_s^{t,x}-\bar{X}_{\pi(s)}^{t,x}\right|\right]ds+C_g e^{-\rho T}\mathbb{E}\left[\left|X_T^{t,x}-\bar{X}_T^{t,x}\right|\right]\\
&\leq Ce^{-\rho t}\mathbb{E}\left[\sup_{t\leq s\leq T}\left|X_s^{t,x}-\bar{X}_{\pi(s)}^{t,x}\right|\right]\leq Ce^{-\rho t}h^{\frac{1}{2}}
\end{aligned}$$

using the strong convergence speed of the Euler scheme on $[t,T]$. Symmetrically, the same inequality holds for $\bar{v}_\Pi(t,x,i)-v_\Pi(t,x,i)$, ending the proof. □

4.3.2.3 Space localization

Recall from Subsection 4.3.1.3 the definition of the bounded domain $\mathcal{D}^\varepsilon$, $t\in[0,T]$.

Proposition 4.3.4. *$\exists C>0$ such that $\forall\varepsilon>0$ and $\forall\,(x,i)\in\mathbb{R}^d\times\mathbb{I}_q$:*

$$\left|\bar{v}_\Pi(0,x,i)-\bar{v}_\Pi^\varepsilon(0,x,i)\right|\leq C\varepsilon$$

Proof. Recall the definitions of $\bar{H}(t,x,\alpha)$ (equation (4.3.27)) and $\bar{J}(t,x,\alpha)$ (equation (4.3.28)), and define the quantities $\bar{H}^\varepsilon(t,x,\alpha)$ and $\bar{J}^\varepsilon(t,x,\alpha)$ like $\bar{H}(t,x,\alpha)$ and $\bar{J}(t,x,\alpha)$ but with $\bar{X}_{\pi(.)}$

replaced by $\bar{X}^{\varepsilon}_{\pi(.)}$. Then, for every $(t,x,i) \in \Pi \times \mathbb{R}^d \times \mathbb{I}_q$ and $\alpha \in \mathcal{A}^{\Pi}_{t,i}$:

$$\begin{aligned}\bar{J}(t,x,\alpha) = \bar{J}^{\varepsilon}(t,x,\alpha) + \int_t^T \mathbb{E}\left[f\left(\pi(s), \bar{X}^{t,x}_{\pi(s)}, I^{\alpha}_s\right) - f\left(\pi(s), \bar{X}^{\varepsilon,t,x}_{\pi(s)}, I^{\alpha}_s\right)\right] ds \\ + \mathbb{E}\left[g\left(T, \bar{X}^{t,x}_T, I^{\alpha}_T\right) - g\left(T, \bar{X}^{\varepsilon,t,x}_T, I^{\alpha}_T\right)\right]\end{aligned}$$

But:

$$\begin{aligned}&\left|\int_t^T \mathbb{E}\left[f\left(\pi(s), \bar{X}^{t,x}_{\pi(s)}, I^{\alpha}_s\right) - f\left(\pi(s), \bar{X}^{\varepsilon,t,x}_{\pi(s)}, I^{\alpha}_s\right)\right] ds + \mathbb{E}\left[g\left(T, \bar{X}^{t,x}_T, I^{\alpha}_T\right) - g\left(T, \bar{X}^{\varepsilon,t,x}_T, I^{\alpha}_T\right)\right]\right| \\ &\leq C_f \int_t^T e^{-\rho s}\mathbb{E}\left[\left|\bar{X}^{t,x}_{\pi(s)} - \bar{X}^{\varepsilon,t,x}_{\pi(s)}\right|\right] ds + C_g e^{-\rho T}\mathbb{E}\left[\left|\bar{X}^{t,x}_T - \bar{X}^{\varepsilon,t,x}_T\right|\right]\end{aligned}$$

It follows that:

$$|\bar{v}_{\Pi}(t,x,i) - \bar{v}^{\varepsilon}_{\Pi}(t,x,i)| \leq C_f \int_t^T e^{-\rho s}\mathbb{E}\left[\left|\bar{X}^{t,x}_{\pi(s)} - \bar{X}^{\varepsilon,t,x}_{\pi(s)}\right|\right] ds + C_g e^{-\rho T}\mathbb{E}\left[\left|\bar{X}^{t,x}_T - \bar{X}^{\varepsilon,t,x}_T\right|\right]$$

In particular, at $t = 0$, using equation (4.3.7), $\exists C > 0$ such that:

$$|\bar{v}_{\Pi}(0,x,i) - \bar{v}^{\varepsilon}_{\Pi}(0,x,i)| \leq C\varepsilon$$

□

4.3.2.4 Conditional expectation approximation

From now on the domain $\mathcal{D}_{\varepsilon}$ is fixed once and for all, and, with a slight abuse of notation, we will drop ε from the subsequent notations.

We start with preliminary remarks. First, regarding the choice of regression basis, Assumption 6 is now supposed to hold. Then, recalling Subsection 4.3.1.4, and taking advantage of the orthogonality of the basis, one can easily compute the explicit solution of the minimisation equations that define the regression coefficients $\tilde{\lambda}^{t_n}_i(\varphi) = \left(\tilde{\lambda}^{t_n}_{i,k}(\varphi)\right)_{1\leq k\leq K}$ (equation (4.3.9)) and $\hat{\lambda}^{t_n}_i(\varphi) = \left(\hat{\lambda}^{t_n}_{i,k}(\varphi)\right)_{1\leq k\leq K_{\varepsilon}}$ (equation (4.3.13)). Namely:

$$\tilde{\lambda}^{t_n}_{i,k}(\varphi) = \frac{\mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, i\right)\mathbf{1}\left\{\bar{X}_{t_n} \in B^k_{t_n}\right\}\right]}{\mathbb{P}\left(\bar{X}_{t_n} \in B^k_{t_n}\right)} = \mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, i\right)\middle|\, \bar{X}_{t_n} \in B^k_{t_n}\right],\ 1 \leq k \leq K$$

$$\hat{\lambda}^{t_n}_{i,k}(\varphi) = \frac{\frac{1}{M}\sum_{m=1}^M \varphi\left(t_{n+1}, \bar{X}^m_{t_{n+1}}, i\right)\mathbf{1}\left\{\bar{X}^m_{t_n} \in B^k_{t_n}\right\}}{\frac{1}{M}\sum_{m=1}^M \mathbf{1}\left\{\bar{X}^m_{t_n} \in B^k_{t_n}\right\}},\ 1 \leq k \leq K$$

Extending these equations, define

$$\tilde{\lambda}^{t_n,x}_i(\varphi) := \frac{\mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, i\right)\mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}(x)\right\}\right]}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}(x)\right)} = \mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, i\right)\middle|\, \bar{X}_{t_n} \in B_{t_n}(x)\right] \tag{4.3.29}$$

$$\hat{\lambda}^{t_n,x}_i(\varphi) := \frac{\frac{1}{M}\sum_{m=1}^M \varphi\left(t_{n+1}, \bar{X}^m_{t_{n+1}}, i\right)\mathbf{1}\left\{\bar{X}^m_{t_n} \in B_{t_n}(x)\right\}}{\frac{1}{M}\sum_{m=1}^M \mathbf{1}\left\{\bar{X}^m_{t_n} \in B_{t_n}(x)\right\}} = \frac{1}{M^x_{t_n}}\sum_{m\in\mathcal{M}^x_{t_n}} \varphi\left(t_{n+1}, \bar{X}^m_{t_{n+1}}, i\right) \tag{4.3.30}$$

for every $(t_n, x, i) \in \Pi \times \mathcal{D} \times \mathbb{I}_q$, where $\forall x \in \mathcal{D}$, $B_{t_n}(x)$ is the unique hypercube in the partition which contains x at time t_n, $\mathcal{M}_{t_n}^x := \left\{ m \in [1, M],\ \bar{X}_{t_n}^m \in B_{t_n}(x) \right\}$ and $M_{t_n}^x := \#\mathcal{M}_{t_n}^x$.

Finally, recalling the approximated conditional expectations (4.3.11) and (4.3.14),

define for any $(t_n, x, j) \in \Pi \times \mathcal{D} \times \mathbb{I}_q$ and any measurable function $\varphi : \Pi \times \mathbb{R}^d \times \mathbb{I}_q \to \mathbb{R}$, the following quantities:

$$\Phi_j^{t_n,x}(\varphi) := \mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x}, j\right)\right] \tag{4.3.31}$$

$$\tilde{\Phi}_j^{t_n,x}(\varphi) := \tilde{\mathbb{E}}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x}, j\right)\right] = \underline{\Gamma}^{t_n,x}(\varphi) \vee \tilde{\lambda}_j^{t_n,x}(\varphi) \wedge \overline{\Gamma}^{t_n,x}(\varphi) \tag{4.3.32}$$

$$\hat{\Phi}_j^{t_n,x}(\varphi) := \hat{\mathbb{E}}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x}, j\right)\right] = \underline{\Gamma}^{t_n,x}(\varphi) \vee \hat{\lambda}_j^{t_n,x}(\varphi) \wedge \overline{\Gamma}^{t_n,x}(\varphi) \tag{4.3.33}$$

where (recalling equation 4.3.10) $\underline{\Gamma}^{t_n,x}(\varphi)$ and $\overline{\Gamma}^{t_n,x}(\varphi)$ are lower and upper bounds on $\Phi_j^{t_n,x}(\varphi)$:

$$\underline{\Gamma}^{t_n,x}(\varphi) \le \Phi_j^{t_n,x}(\varphi) \le \overline{\Gamma}^{t_n,x}(\varphi)$$

Remark 4.3.5. These definitions are useful to express the dynamic programming equations (4.3.8), (4.3.12) and (4.3.15). Indeed, these equations become:

$$\begin{aligned}
\bar{v}_\Pi(T, x, i) &= g(T, x, i) \\
\bar{v}_\Pi(t_n, x, i) &= \max_{j \in \mathbb{I}_q} \left\{ hf(t_n, x, j) - k(t_n, i, j) + \Phi_j^{t_n,x}(\bar{v}_\Pi) \right\},\ n = N-1, \dots, 0
\end{aligned}$$

$$\begin{aligned}
\tilde{v}_\Pi(T, x, i) &= g(T, x, i) \\
\tilde{v}_\Pi(t_n, x, i) &= \max_{j \in \mathbb{I}_q} \left\{ hf(t_n, x, j) - k(t_n, i, j) + \tilde{\Phi}_j^{t_n,x}(\tilde{v}_\Pi) \right\},\ n = N-1, \dots, 0
\end{aligned}$$

$$\begin{aligned}
\hat{v}_\Pi(T, x, i) &= g(T, x, i) \\
\hat{v}_\Pi(t_n, x, i) &= \max_{j \in \mathbb{I}_q} \left\{ hf(t_n, x, j) - k(t_n, i, j) + \hat{\Phi}_j^{t_n,x}(\hat{v}_\Pi) \right\},\ n = N-1, \dots, 0
\end{aligned}$$

Remark 4.3.6. For $\varphi = \bar{v}_\Pi$, we can easily explicit bounding functions $\underline{\Gamma}^{t_n,x}(\bar{v}_\Pi)$ and $\overline{\Gamma}^{t_n,x}(\bar{v}_\Pi)$ of $\Phi_j^{t_n,x}(\bar{v}_\Pi)$. Indeed, using the growth conditions on f and g, the nonnegativity of k and the definition of $C(T, \varepsilon)$ (see Paragraph 4.3.1.3), there exists $C > 0$ such that $\forall (t_n, x, j) \in \Pi \times \mathcal{D} \times \mathbb{I}_q$:

$$|\bar{v}_\Pi(t_n, x, j)| \le Ce^{-\rho t_n}(1 + C(T, \varepsilon)) \tag{4.3.34}$$

$$\left|\Phi_j^{t_n,x}(\bar{v}_\Pi)\right| \le \Gamma^{t_n}(\bar{v}_\Pi) := Ce^{-\rho t_n}(1 + C(T, \varepsilon)) \tag{4.3.35}$$

Moreover, the same is true for $\varphi = \tilde{v}_\Pi$: there exists $C > 0$ such that $\forall (t_n, x, j) \in \Pi \times \mathcal{D} \times \mathbb{I}_q$:

$$|\tilde{v}_\Pi(t_n, x, j)| \le Ce^{-\rho t_n}(1 + C(T, \varepsilon)) \tag{4.3.36}$$

$$\left|\tilde{\Phi}_j^{t_n,x}(\tilde{v}_\Pi)\right| \le \Gamma^{t_n}(\tilde{v}_\Pi) := Ce^{-\rho t_n}(1 + C(T, \varepsilon)) \tag{4.3.37}$$

Finally, we impose the same bound for the definition of $\hat{v}_\Pi$, i.e. $\Gamma^{t_n}(\hat{v}_\Pi) := \Gamma^{t_n}(\bar{v}_\Pi)$.

Now we can start the assessment of the regression error.

Lemma 4.3.3. *Consider a measurable function $\varphi : \Pi \times \mathbb{R}^d \times \mathbb{I}_q \to \mathbb{R}$. Suppose that, for a fixed $t_{n+1} \in \Pi$, it is Lipschitz with constant C_{n+1}, uniformly in j: $\forall (x_1, x_2, j) \in \mathbb{R}^d \times \mathbb{R}^d \times \mathbb{I}_q$*

$$|\varphi(t_{n+1}, x_1, j) - \varphi(t_{n+1}, x_2, j)| \le C_{n+1}|x_1 - x_2| \tag{4.3.38}$$

Then $\Phi_j^{t_n,x}(\varphi)$ is Lipschitz with constant $C_{n+1}(1 + Lh)$, uniformly in j, where $L := C_b + \frac{C_\sigma^2}{2} > 0$.

Proof. Choose $(t_n, j, x_1, x_2) \in \Pi \times \mathbb{I}_q \times \mathbb{R}^d \times \mathbb{R}^d$. Then:

$$\begin{aligned}\left|\Phi_j^{t_n,x_1}(\varphi) - \Phi_j^{t_n,x_2}(\varphi)\right| &= \left|\mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x_1}, j\right) - \varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x_2}, j\right)\right]\right| \\ &\leq \left\|\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x_1}, j\right) - \varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x_2}, j\right)\right\|_{L_1} \\ &\leq \left\|\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x_1}, j\right) - \varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x_2}, j\right)\right\|_{L_2}\end{aligned}$$

Now, using equations (4.3.38) and (4.3.5), and G denoting a d-dimensional standard Gaussian random variable, we have

$$\begin{aligned}&\mathbb{E}\left[\left(\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x_1}, j\right) - \varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x_2}, j\right)\right)^2\right] \\ \leq\ & C_{n+1}^2 \mathbb{E}\left[\left(\bar{X}_{t_{n+1}}^{t_n,x_1} - \bar{X}_{t_{n+1}}^{t_n,x_2}\right)^2\right] \\ \leq\ & C_{n+1}^2 \mathbb{E}\left[\left(x_1 - x_2 + h\left(b\left(t_n, x_1\right) - b\left(t_n, x_2\right)\right) + \sqrt{h}\left(\sigma\left(t_n, x_1\right) - \sigma\left(t_n, x_2\right)\right) G\right)^2\right] \\ =\ & C_{n+1}^2 \left\{\left(x_1 - x_2 + h\left(b\left(t_n, x_1\right) - b\left(t_n, x_2\right)\right)\right)^2 + h\mathbb{E}\left[\left(\left(\sigma\left(t_n, x_1\right) - \sigma\left(t_n, x_2\right)\right) G\right)^2\right]\right\} \\ \leq\ & C_{n+1}^2 \left(x_1 - x_2\right)^2 \left\{1 + \left(2C_b + C_\sigma^2\right) h + C_b^2 h^2\right\} \\ \leq\ & C_{n+1}^2 \left(x_1 - x_2\right)^2 \left(C_b + \frac{C_\sigma^2}{2}\right)^2 \left(\frac{1}{C_b + \frac{C_\sigma^2}{2}} + h\right)^2.\end{aligned}$$

Thus:

$$\left|\Phi_j^{t_n,x_1}(\varphi) - \Phi_j^{t_n,x_2}(\varphi)\right| \leq C_{n+1}\left(1 + \left(C_b + \frac{C_\sigma^2}{2}\right) h\right) |x_1 - x_2|$$

□

Lemma 4.3.4. *Consider again a function $\varphi : \Pi \times \mathbb{R}^d \times \mathbb{I}_q \to \mathbb{R}$ such that (4.3.38) holds for a given $t_{n+1} \in \Pi$. Then, $\forall (x, j) \in \mathcal{D} \times \mathbb{I}_q$:*

$$\left|\tilde{\lambda}_j^{t_n,x}(\varphi) - \Phi_j^{t_n,x}(\varphi)\right| \leq C_{n+1}\delta\left(1 + Lh\right).$$

In particular:

$$\left|\tilde{\Phi}_j^{t_n,x}(\varphi) - \Phi_j^{t_n,x}(\varphi)\right| \leq C_{n+1}\delta\left(1 + Lh\right) \tag{4.3.39}$$

Proof. Recalling the definitions of $B_{t_n}(x)$, of $\tilde{\lambda}_j^{t_n,x}(\varphi)$ (equation (4.3.29)) and of $\Phi_j^{t_n,x}(\varphi)$ (equation (4.3.31)), simply remark that:

$$\begin{aligned}\min_{\tilde{x} \in B_{t_n}(x)} \Phi_j^{t_n,\tilde{x}}(\varphi) &\leq \Phi_j^{t_n,x}(\varphi) \leq \max_{\tilde{x} \in B_{t_n}(x)} \Phi_j^{t_n,\tilde{x}}(\varphi) \\ \min_{\tilde{x} \in B_{t_n}(x)} \Phi_j^{t_n,\tilde{x}}(\varphi) &\leq \tilde{\lambda}_j^{t_n,x}(\varphi) \leq \max_{\tilde{x} \in B_{t_n}(x)} \Phi_j^{t_n,\tilde{x}}(\varphi).\end{aligned}$$

Now, using Lemma 4.3.3:

$$\begin{aligned}\left|\tilde{\lambda}_j^{t_n,x}(\varphi) - \Phi_j^{t_n,x}(\varphi)\right| &\leq \max_{\tilde{x} \in B_{t_n}(x)} \Phi_j^{t_n,\tilde{x}}(\varphi) - \min_{\tilde{x} \in B_{t_n}(x)} \Phi_j^{t_n,\tilde{x}}(\varphi) \\ &\leq C_{n+1}\left(1 + Lh\right) \max_{(x_1,x_2) \in B_{t_n}(x)^2} |x_1 - x_2| \\ &\leq C_{n+1}\left(1 + Lh\right)\delta\end{aligned}$$

□

Lemma 4.3.5. $\forall (t_n, x_1, x_2, i) \in \Pi \times \left(\mathbb{R}^d\right)^2 \times \mathbb{I}_q$*:*

$$|\bar{v}_\Pi(t_n, x_1, i) - \bar{v}_\Pi(t_n, x_2, i)| \leq C_n |x_1 - x_2| \tag{4.3.40}$$

where:

$$\begin{aligned} C_N &= e^{-\rho t_N} C_g \\ C_n &= hC_f e^{-\rho t_n} + C_{n+1}(1 + Lh)\ ,\ n = N-1, \ldots, 0 \end{aligned} \tag{4.3.41}$$

In particular, $\exists C > 0$ *such that* $\forall n = 0, 1, \ldots, N$*:*

$$C_n \leq Ce^{-\rho t_n} e^{L(T-t_n)} \tag{4.3.42}$$

Proof. Recall Remark 4.3.5. We prove the lemma by induction. First, remark that, using hypothesis (4.3.21), it holds for $n = N$. Now, suppose that it holds for some $(n+1) \in [1, \ldots, N]$. Then, using Lemma 4.3.3:

$$\begin{aligned} &\bar{v}_\Pi(t_n, x_1, i) \\ &= \max_{j \in \mathbb{I}_q} \left\{ hf(t_n, x_1, j) - k(t_n, i, j) + \Phi_j^{t_n, x_1}(\bar{v}_\Pi) \right\} \\ &= \max_{j \in \mathbb{I}_q} \left\{ hf(t_n, x_2, j) - k(t_n, i, j) + \Phi_j^{t_n, x_2}(\bar{v}_\Pi) + h(f(t_n, x_1, j) - f(t_n, x_2, j)) + \left(\Phi_j^{t_n, x_1}(\bar{v}_\Pi) - \Phi_j^{t_n, x_2}(\bar{v}_\Pi) \right) \right\} \\ &\leq \max_{j \in \mathbb{I}_q} \left\{ hf(t_n, x_2, j) - k(t_n, i, j) + \Phi_j^{t_n, x_2}(\bar{v}_\Pi) + he^{-\rho t_n} C_f |x_1 - x_2| + C_{n+1}(1 + Lh) |x_1 - x_2| \right\} \\ &= \bar{v}_\Pi(t_n, x_2, i) + \left(he^{-\rho t_n} C_f + C_{n+1}(1 + Lh) \right) |x_1 - x_2| \end{aligned}$$

Symmetrically, the same inequality holds for $\bar{v}_\Pi(t_n, x_2, i) - \bar{v}_\Pi(t_n, x_1, i)$, yielding equations (4.3.40) and (4.3.41). Finally, use the discrete version of Gronwall's inequality to obtain equation (4.3.42) □

Proposition 4.3.5. $\exists C > 0$ *s.t.* $\forall (t, x, i) \in \Pi \times \mathbb{R}^d \times \mathbb{I}_q$ *:*

$$|\bar{v}_\Pi(t, x, i) - \tilde{v}_\Pi(t, x, i)| \leq C \frac{\delta}{h} e^{-\rho t}\,.$$

Proof. For each $t_n \in \Pi$, we look for an upper bound E_n, independent of x and i, of the quantity $|\bar{v}_\Pi(t_n, x, i) - \tilde{v}_\Pi(t_n, x, i)|$. First:

$$|\bar{v}_\Pi(T, x, i) - \tilde{v}_\Pi(T, x, i)| = |g(T, x, i) - g(T, x, i)| = 0$$

Hence $E_N = 0$. Fix now $n \in [0, N-1]$. Using Remark 4.3.5:

$$\begin{aligned} \tilde{v}_\Pi(t_n, x, i) &= \max_{j \in \mathbb{I}_q} \left\{ hf(t_n, x, j) - k(t_n, i, j) + \tilde{\Phi}_j^{t_n, x}(\tilde{v}_\Pi) \right\} \\ &= \max_{j \in \mathbb{I}_q} \Big\{ hf(t_n, x, j) - k(t_n, i, j) + \Phi_j^{t_n, x}(\bar{v}_\Pi) \\ &\qquad + \tilde{\Phi}_j^{t_n, x}(\bar{v}_\Pi) - \Phi_j^{t_n, x}(\bar{v}_\Pi) \\ &\qquad + \tilde{\Phi}_j^{t_n, x}(\tilde{v}_\Pi) - \tilde{\Phi}_j^{t_n, x}(\bar{v}_\Pi) \Big\} \end{aligned}$$

Using Lemmas 4.3.4 and 4.3.5, $\tilde{\Phi}_j^{t_n, x}(\bar{v}_\Pi) - \Phi_j^{t_n, x}(\bar{v}_\Pi) \leq C_{n+1} \delta (1 + Lh)$ where C_{n+1} is the Lipschitz constant of $\bar{v}_\Pi$ at time t_{n+1} (see Lemma 4.3.5). Moreover,

$$\begin{aligned} \tilde{\Phi}_j^{t_n, x}(\tilde{v}_\Pi) - \tilde{\Phi}_j^{t_n, x}(\bar{v}_\Pi) &\leq \mathbb{E}\left[\tilde{v}_\Pi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right) - \bar{v}_\Pi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right) \middle| X_{t_n} \in B_{t_n}(x) \right] \\ &\leq E_{n+1}\,. \end{aligned}$$

Hence:

$$\tilde{v}_{\Pi}(t_n, x, i) \leq \bar{v}_{\Pi}(t_n, x, i) + C_{n+1}\delta(1 + Lh) + E_{n+1}$$

Symmetrically, the same inequality holds for $\bar{v}_{\Pi}(T, x, i) - \tilde{v}_{\Pi}(t_n, x, i)$, leading to:

$$|\bar{v}_{\Pi}(t_n, x, i) - \tilde{v}_{\Pi}(t_n, x, i)| \leq E_n$$

where:

$$\begin{aligned} E_N &= 0 \\ E_n &= C_{n+1}\delta(1 + Lh) + E_{n+1}. \end{aligned}$$

Consequently, using equation (4.3.42):

$$E_n = \delta(1 + Lh) \sum_{k=n+1}^{N} C_k \leq C\frac{\delta}{h}e^{-\rho t_n}$$

where $C > 0$ does not depend on t_n nor T. □

The following lemma measures the regression error. It is an extension of Lemma 3.8 in [100] (itself inspired by Theorem 5.1 in [27]).

Lemma 4.3.6. *Consider a measurable function $\varphi : \Pi \times \mathbb{R}^d \times \mathbb{I}_q \to \mathbb{R}$. For any $p \geq 1$, there exists $C_p \geq 0$ such that $\forall (t_n, l, j) \in \Pi \times [1, M] \times \mathbb{I}_q$:*

$$\left\| \hat{\Phi}_j^{t_n, \bar{X}_{t_n}^l}(\varphi) - \tilde{\Phi}_j^{t_n, \bar{X}_{t_n}^l}(\varphi) \right\|_{L_p} \leq \frac{C_p}{\sqrt{M}} \frac{\Gamma^{t_n}(\varphi) + \bar{\varphi}^{t_n}}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)^{1 - \frac{1}{p \vee 2}}} + \frac{C_p}{M} \frac{\bar{\varphi}^{t_n}}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)} \tag{4.3.43}$$

where $\bar{\varphi}^{t_n} \in \mathbb{R}_+$ is such that $\left|\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right)\right| \leq \bar{\varphi}^{t_n}$ a.s. .

Proof. Define the following centered random variables:

$$\varepsilon_j^{t_n, \bar{X}_{t_n}^l}(\varphi) := \frac{1}{M} \sum_{m=1}^{M} \varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^m, j\right) \mathbf{1}\left\{\bar{X}_{t_n}^m \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right\} - \mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}^m, j\right) \mathbf{1}\left\{\bar{X}_{t_n}^m \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right\}\right]$$

$$\varepsilon^{t_n, \bar{X}_{t_n}^l}(1) := \frac{1}{M} \sum_{m=1}^{M} \mathbf{1}\left\{\bar{X}_{t_n}^m \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right\} - \mathbb{P}\left(\bar{X}_{t_n}^m \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)$$

Then:

$$\begin{aligned} &\left|\hat{\Phi}_j^{t_n, \bar{X}_{t_n}^l}(\varphi) - \tilde{\Phi}_j^{t_n, \bar{X}_{t_n}^l}(\varphi)\right| = \left|\hat{\Phi}_j^{t_n, \bar{X}_{t_n}^l}(\varphi) - \tilde{\Phi}_j^{t_n, \bar{X}_{t_n}^l}(\varphi)\right| \wedge 2\Gamma^{t_n}(\varphi) \\ &\leq \left|\hat{\Phi}_j^{t_n, \bar{X}_{t_n}^l}(\varphi) - \tilde{\Phi}(\varphi)\right| \mathbf{1}\left\{\frac{\left|\varepsilon^{t_n, \bar{X}_{t_n}^l}(1)\right|}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)} \leq \frac{1}{2}\right\} + 2\Gamma^{t_n}(\varphi) \mathbf{1}\left\{\frac{\left|\varepsilon^{t_n, \bar{X}_{t_n}^l}(1)\right|}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)} > \frac{1}{2}\right\} \end{aligned}$$

and:

$$\left|\hat{\Phi}_j^{t_n, \bar{X}_{t_n}^l}(\varphi) - \tilde{\Phi}_j^{t_n, \bar{X}_{t_n}^l}(\varphi)\right| \mathbf{1}\left\{\frac{\left|\varepsilon^{t_n, \bar{X}_{t_n}^l}(1)\right|}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}_{t_n}^l\right)\right)} \leq \frac{1}{2}\right\}$$

$$
\begin{aligned}
&= \left| \hat{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi) - \tilde{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi) \frac{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)}{\frac{1}{M}\sum_{m=1}^M \mathbf{1}\left\{\bar{X}^m_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right\}} - \right. \\
&\left. \tilde{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi) \frac{\varepsilon^{t_n,\bar{X}^l_{t_n}}(1)}{\frac{1}{M}\sum_{m=1}^M \mathbf{1}\left\{\bar{X}^m_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right\}} \right| \mathbf{1}\left\{ \frac{\left|\varepsilon^{t_n,\bar{X}^l_{t_n}}(1)\right|}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)} \leq \frac{1}{2} \right\} \\
&\leq \left\{ \frac{\left|\varepsilon_j^{t_n,\bar{X}^l_{t_n}}(\varphi)\right|}{\frac{1}{M}\sum_{m=1}^M \mathbf{1}\left\{\bar{X}^m_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right\}} \wedge 3\Gamma^{t_n}(\varphi) + \right. \\
&\left. \left|\tilde{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi)\right| \frac{\left|\varepsilon^{t_n,\bar{X}^l_{t_n}}(1)\right|}{\frac{1}{M}\sum_{m=1}^M \mathbf{1}\left\{\bar{X}^m_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right\}} \right\} \mathbf{1}\left\{ \frac{\left|\varepsilon^{t_n,\bar{X}^l_{t_n}}(1)\right|}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)} \leq \frac{1}{2} \right\} \\
&\leq \frac{2}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)} \left\{ \left|\varepsilon_j^{t_n,\bar{X}^l_{t_n}}(\varphi)\right| \wedge 5\Gamma^{t_n}(\varphi) + \left|\varepsilon^{t_n,\bar{X}^l_{t_n}}(1)\right| \Gamma^{t_n}(\varphi) \right\} \mathbf{1}\left\{ \frac{\left|\varepsilon^{t_n,\bar{X}^l_{t_n}}(1)\right|}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)} \leq \frac{1}{2} \right\}
\end{aligned}
$$

Now, for any $p \geq 1$:

$$
\begin{aligned}
&\left| \hat{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi) - \tilde{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi) \right|^p \\
&\leq \frac{2^{3p-2}}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)^p} \left\{ \left\{ \left|\varepsilon_j^{t_n,\bar{X}^l_{t_n}}(\varphi)\right| \wedge 5\Gamma^{t_n}(\varphi) \right\}^p + \left\{ \left|\varepsilon^{t_n,\bar{X}^l_{t_n}}(1)\right| \Gamma^{t_n}(\varphi) \right\}^p \right\} \times \\
&\mathbf{1}\left\{ \frac{\left|\varepsilon^{t_n,\bar{X}^l_{t_n}}(1)\right|}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)} \leq \frac{1}{2} \right\} + 2^{2p-1} \left(\Gamma^{t_n}(\varphi)\right)^p \mathbf{1}\left\{ \frac{\left|\varepsilon^{t_n,\bar{X}^l_{t_n}}(1)\right|}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)} > \frac{1}{2} \right\}
\end{aligned}
$$

and:

$$
\begin{aligned}
&\mathbb{E}\left[\left| \hat{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi) - \tilde{\Phi}_j^{t_n,\bar{X}^l_{t_n}}(\varphi) \right|^p \right] \\
&\leq \frac{2^{3p-2}}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)^p} \left\{ \mathbb{E}\left[\left\{ \left|\varepsilon_j^{t_n,\bar{X}^l_{t_n}}(\varphi)\right| \wedge 5\Gamma^{t_n}(\varphi) \right\}^p \right] + \left(\Gamma^{t_n}(\varphi)\right)^p \mathbb{E}\left[\left|\varepsilon^{t_n,\bar{X}^l_{t_n}}(1)\right|^p \right] \right\} \\
&+ 2^{2p-1} \left(\Gamma^{t_n}(\varphi)\right)^p \mathbb{P}\left(\left|\varepsilon^{t_n,\bar{X}^l_{t_n}}(1)\right|^p > \frac{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)^p}{2^p} \right) \\
&\leq \frac{8^p}{\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)^p} \left\{ \mathbb{E}\left[\left\{ \left|\varepsilon_j^{t_n,\bar{X}^l_{t_n}}(\varphi)\right| \wedge 5\Gamma^{t_n}(\varphi) \right\}^p \right] + \left\{\Gamma^{t_n}(\varphi)\right\}^p \mathbb{E}\left[\left|\varepsilon_j^{t_n,\bar{X}^l_{t_n}}(1)\right|^p \right] \right\}
\end{aligned} \tag{4.3.44}
$$

using Markov's inequality. We then obtain upper bounds for $\mathbb{E}\left[\left|\varepsilon^{t_n,\bar{X}^l_{t_n}}(1)\right|^p\right]$ and $\mathbb{E}\left[\left|\varepsilon_j^{t_n,\bar{X}^l_{t_n}}(\varphi)\right|^p\right]$ using Lemma 4.7.1 in Appendix 4.7.1. Suppose that $\exists \bar{\varphi}^{t_n} \in \mathbb{R}_+$ s.t. $\left|\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right)\right| \leq \bar{\varphi}^{t_n}$ a.s. . Then, using Lemma 4.7.1, $\exists C_p > 0$ such that:

$$
\mathbb{E}\left[\left|\varepsilon^{t_n,\bar{X}^l_{t_n}}(1)\right|^p\right] \leq \frac{C_p}{M^{\frac{p}{2}}} \mathbb{E}\left[\left| \mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right\} - \mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right) \right|^{p \vee 2} \right]^{\frac{p}{p \vee 2}} \tag{4.3.45}
$$

$$
\begin{aligned}
\mathbb{E}\left[\left|\varepsilon_j^{t_n,\bar{X}^l_{t_n}}(\varphi)\right|^p\right] \leq\ & C_p \left\{ \frac{(\bar{\varphi}^{t_n})^p}{M^p} + \frac{1}{M^{\frac{p}{2}}} \mathbb{E}\left[\left| \varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right) \mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right\} \right.\right. \\
& \left.\left. - \mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right) \mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right\} \right] \right|^{p \vee 2} \right]^{\frac{p}{p \vee 2}} \right\}
\end{aligned} \tag{4.3.46}
$$

where, for the second inequality, the term $m = l$ in the sum was treated separately. Then:

$$\begin{aligned}
&\mathbb{E}\left[\left|\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right)\mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right\} - \mathbb{E}\left[\varphi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right)\mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right\}\right]\right|^{p\vee 2}\right]^{\frac{p}{p\vee 2}} \\
&\leq \left(2^{p\vee 2-1}\mathbb{E}\left[\left(\bar{\varphi}^{t_n}\right)^{p\vee 2}\mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right\} + \mathbb{E}\left[\left(\bar{\varphi}^{t_n}\right)^{p\vee 2}\mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right\}\right]\right]\right)^{\frac{p}{p\vee 2}} \\
&\leq 2^p\left(\bar{\varphi}^{t_n}\right)^p \mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)^{\frac{p}{p\vee 2}} \qquad (4.3.47)
\end{aligned}$$

In a similar manner:

$$\mathbb{E}\left[\left|\mathbf{1}\left\{\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right\} - \mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)\right|^{p\vee 2}\right]^{\frac{p}{p\vee 2}} \leq 2^p\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)^{\frac{p}{p\vee 2}} \qquad (4.3.48)$$

Finally, the combination of inequalities (4.3.44), (4.3.45), (4.3.46), (4.3.47) and (4.3.48) proves equation (4.3.43). □

We now apply Lemma 4.3.6 to $\bar{v}_\Pi$ in the following Corollary:

Corollary 4.3.1. *For every $p \geq 1$, there exists $C_p \geq 0$ s.t. $\forall\,(t_n, l, j) \in \Pi \times [1, M] \times \mathbb{I}_q$:*

$$\left\|\hat{\Phi}_j^{t_n, \bar{X}^l_{t_n}}\left(\tilde{v}_\Pi\right) - \tilde{\Phi}_j^{t_n, \bar{X}^l_{t_n}}\left(\tilde{v}_\Pi\right)\right\|_{L_p} \leq C_p e^{-\rho t_n}\frac{1 + C\left(T, \varepsilon\right)}{\sqrt{M}p\left(T, \delta, \varepsilon\right)^{1-\frac{1}{p\vee 2}}}\left(1 + \frac{1}{\sqrt{M}p\left(T, \delta, \varepsilon\right)^{\frac{1}{p\vee 2}}}\right)$$

Proof. First, recall from equation (4.3.36) and (4.3.37) that there exists $C > 0$ such that for every $(t_n, j) \in \Pi \times \mathbb{I}_q$:

$$\begin{aligned}
\Gamma_j^{t_n}\left(\tilde{v}_\Pi\right) &= Ce^{-\rho t_n}\left(1 + C\left(T, \varepsilon\right)\right) \\
\left|\tilde{v}_\Pi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right)\right| &\leq Ce^{-\rho t_n}\left(1 + C\left(T, \varepsilon\right)\right)
\end{aligned}$$

Hence one can apply Lemma 4.3.6 to $\tilde{v}_\Pi$ with these upper bounds. The final step is to recall that the minimum probability $p\left(T, \delta, \varepsilon\right)$ defined in equation (4.3.16) is a lower bound on $\mathbb{P}\left(\bar{X}_{t_n} \in B_{t_n}\left(\bar{X}^l_{t_n}\right)\right)$ for any $(t_n, l) \in \Pi \times [1, M]$. □

Using this result, we can now assess the error between $\tilde{v}_\Pi$ and $\hat{v}_\Pi$.

Proposition 4.3.6. *$\forall p \geq 1$, $\exists C_p > 0$ s.t. $\forall\,(t_n, l) \in \Pi \times [1, M]$:*

$$\left\|\sup_{i\in\mathbb{I}_q^N}\left|\tilde{v}_\Pi\left(t, \bar{X}^l_{t_n}, i\right) - \hat{v}_\Pi\left(t, \bar{X}^l_{t_n}, i\right)\right|\right\|_{L_p} \leq C_p e^{-\rho t_n}\frac{1 + C\left(T, \varepsilon\right)}{h\sqrt{M}p\left(T, \delta, \varepsilon\right)^{1-\frac{1}{p\vee 2}}}\left(1 + \frac{1}{\sqrt{M}p\left(T, \delta, \varepsilon\right)^{\frac{1}{p\vee 2}}}\right)$$

where $\mathbb{I}_q^N$ is the set of $\mathcal{F}_{t_N}$-measurable random variables taking values in $\mathbb{I}_q$.

Proof. For each $t_n \in \Pi$, we look for an upper bound E_n, independent of l, such that:

$$\left\|\sup_{i\in\mathbb{I}_q^N}\left|\tilde{v}_\Pi\left(t, \bar{X}^l_{t_n}, i\right) - \hat{v}_\Pi\left(t, \bar{X}^l_{t_n}, i\right)\right|\right\|_{L_p} \leq E_n\ .$$

First:

$$\left\|\sup_{i\in\mathbb{I}_q^N}\left|\tilde{v}_\Pi\left(T, \bar{X}^l_T, i\right) - \hat{v}_\Pi\left(T, \bar{X}^l_T, i\right)\right|\right\|_{L_p} = \left\|\sup_{i\in\mathbb{I}_q^N}\left|g\left(T, \bar{X}^l_T, i\right) - g\left(T, \bar{X}^l_T, i\right)\right|\right\|_{L_p} = 0$$

Hence $E_N = 0$. Fix now $n \in [0, N-1]$. Recall the dynamic programming equations from Remark 4.3.5, and, for every $(i,l) \in \mathbb{I}_q^N \times [1,M]$, introduce $\tilde{j}^*$ (resp. $\hat{j}^*$) the arg max for $\tilde{v}_\Pi$ (resp. $\hat{v}_\Pi$) at point $\bar{X}_{t_n}^l$, i.e.:

$$\begin{aligned}
\tilde{v}_\Pi\left(t_n, \bar{X}_{t_n}^l, i\right) &= hf\left(t_n, \bar{X}_{t_n}^l, \tilde{j}^*\right) - k\left(t_n, i, \tilde{j}^*\right) + \tilde{\Phi}_{\tilde{j}^*}^{t_n, \bar{X}_{t_n}^l}\left(\tilde{v}_\Pi\right) \\
\hat{v}_\Pi\left(t_n, \bar{X}_{t_n}^l, i\right) &= hf\left(t_n, \bar{X}_{t_n}^l, \hat{j}^*\right) - k\left(t_n, i, \hat{j}^*\right) + \hat{\Phi}_{\hat{j}^*}^{t_n, \bar{X}_{t_n}^l}\left(\hat{v}_\Pi\right)
\end{aligned}$$

Now:

$$\begin{aligned}
\hat{v}_\Pi\left(t_n, \bar{X}_{t_n}^l, i\right) &= hf\left(t_n, \bar{X}_{t_n}^l, \hat{j}^*\right) - k\left(t_n, i, \hat{j}^*\right) + \hat{\Phi}_{\hat{j}^*}^{t_n, \bar{X}_{t_n}^l}\left(\hat{v}_\Pi\right) \\
&= \left\{hf\left(t_n, \bar{X}_{t_n}^l, \hat{j}^*\right) - k\left(t_n, i, \hat{j}^*\right) + \tilde{\Phi}_{\hat{j}^*}^{t_n, \bar{X}_{t_n}^l}\left(\tilde{v}_\Pi\right)\right\} \\
&\quad + \left\{\hat{\Phi}_{\hat{j}^*}^{t_n, \bar{X}_{t_n}^l}\left(\tilde{v}_\Pi\right) - \tilde{\Phi}_{\hat{j}^*}^{t_n, \bar{X}_{t_n}^l}\left(\tilde{v}_\Pi\right)\right\} + \left\{\hat{\Phi}_{\hat{j}^*}^{t_n, \bar{X}_{t_n}^l}\left(\hat{v}_\Pi\right) - \hat{\Phi}_{\hat{j}^*}^{t_n, \bar{X}_{t_n}^l}\left(\tilde{v}_\Pi\right)\right\} \\
&\leq \tilde{v}_\Pi\left(t_n, \bar{X}_{t_n}^l, i\right) + \sum_{j \in \mathbb{I}_q}\left|\hat{\Phi}_j^{t_n, \bar{X}_{t_n}^l}\left(\tilde{v}_\Pi\right) - \tilde{\Phi}_j^{t_n, \bar{X}_{t_n}^l}\left(\tilde{v}_\Pi\right)\right| \\
&\quad + \sup_{j \in \mathbb{I}_q^N}\left|\hat{\Phi}_j^{t_n, \bar{X}_{t_n}^l}\left(\hat{v}_\Pi\right) - \hat{\Phi}_j^{t_n, \bar{X}_{t_n}^l}\left(\tilde{v}_\Pi\right)\right|
\end{aligned}$$

Symmetrically:

$$\begin{aligned}
\tilde{v}_\Pi\left(t_n, \bar{X}_{t_n}^l, i\right) &\leq \hat{v}_\Pi\left(t_n, \bar{X}_{t_n}^l, i\right) + \sum_{j \in \mathbb{I}_q}\left|\tilde{\Phi}_j^{t_n, \bar{X}_{t_n}^l}\left(\tilde{v}_\Pi\right) - \hat{\Phi}_j^{t_n, \bar{X}_{t_n}^l}\left(\tilde{v}_\Pi\right)\right| \\
&\quad + \sup_{j \in \mathbb{I}_q^N}\left|\hat{\Phi}_j^{t_n, \bar{X}_{t_n}^l}\left(\tilde{v}_\Pi\right) - \hat{\Phi}_j^{t_n, \bar{X}_{t_n}^l}\left(\hat{v}_\Pi\right)\right|
\end{aligned}$$

Combining these two inequalities:

$$\begin{aligned}
\sup_{i \in \mathbb{I}_q^N}\left|\tilde{v}_\Pi\left(t_n, \bar{X}_{t_n}^l, i\right) - \hat{v}_\Pi\left(t_n, \bar{X}_{t_n}^l, i\right)\right| &\leq \sum_{j \in \mathbb{I}_q}\left|\hat{\Phi}_j^{t_n, \bar{X}_{t_n}^l}\left(\tilde{v}_\Pi\right) - \tilde{\Phi}_j^{t_n, \bar{X}_{t_n}^l}\left(\tilde{v}_\Pi\right)\right| \\
&\quad + \sup_{j \in \mathbb{I}_q^N}\left|\hat{\Phi}_j^{t_n, \bar{X}_{t_n}^l}\left(\hat{v}_\Pi\right) - \hat{\Phi}_j^{t_n, \bar{X}_{t_n}^l}\left(\tilde{v}_\Pi\right)\right|
\end{aligned}$$

Hence, using the triangular inequality, Corollary 4.3.1, equation (4.3.30), and the induction hypothesis:

$$\begin{aligned}
\left\|\sup_{i \in \mathbb{I}_q^N}\left|\tilde{v}_\Pi\left(t_n, \bar{X}_{t_n}^l, i\right) - \hat{v}_\Pi\left(t_n, \bar{X}_{t_n}^l, i\right)\right|\right\|_{L_p} \leq E_n := \; & C_p e^{-\rho t_n} \frac{1 + C\left(T, \varepsilon\right)}{\sqrt{M} p\left(T, \delta, \varepsilon\right)^{1 - \frac{1}{p \vee 2}}} \\
& + C_p e^{-\rho t_n} \frac{1 + C\left(T, \varepsilon\right)}{M p\left(T, \delta, \varepsilon\right)} + E_{n+1}
\end{aligned}$$

for some constant $C_p > 0$ which depends only on p. Consequently:

$$E_n \leq C_p e^{-\rho t_n} \frac{1 + C\left(T, \varepsilon\right)}{h\sqrt{M} p\left(T, \delta, \varepsilon\right)^{1 - \frac{1}{p \vee 2}}} \left(1 + \frac{1}{\sqrt{M} p\left(T, \delta, \varepsilon\right)^{\frac{1}{p \vee 2}}}\right)$$

where $C_p > 0$ depends only on p. □

Finally, the combination of Propositions 4.3.1 4.3.2, 4.3.3, 4.3.5 and 4.3.6 at time $t = t_0$ proves Theorem 4.3.1.

4.4 Complexity analysis and memory reduction

4.4.1 Complexity

4.4.1.1 Computational complexity

The number of operations required by the algorithm described below is in $\mathcal{O}(q^2 \cdot N \cdot M)$, where we recall that q is the number of possible switches, N is the number of time steps and M is the number of Monte Carlo trajectories.

- The q^2 term stems from the fact that for every $i \in \mathbb{I}_q$, one has to compute a maximum on $j \in \mathbb{I}_q$ (see equation (4.3.15)). However, this q^2 can be reduced to q as soon as the two following conditions are satisfied:

1. (Irreversibility) The controlled variable can only be increased (or, symmetrically, can only be decreased)

2. (Cost Separability) There exists two functions k_1 and k_2 such that $\forall\,(t,i,j) \in \mathbb{R}_+ \times \mathbb{I}_q \times \mathbb{I}_q$, $k\,(t,i,j) = k_1\,(t,i) + k_2\,(t,j)$. For instance, this is true of affine costs.

Indeed, under those two conditions, equation (4.3.15) becomes:

$$\hat{v}_\Pi(t_n,x,i)+k_1(t_n,i) = \max_{j\in\mathbb{I}_q,\, j\geq i} \left\{ hf(t_n,x,j) - k_2(t_n,j) + \hat{\mathbb{E}}\left[\hat{v}_\Pi\left(t_{n+1}, \bar{X}_{t_{n+1}}^{t_n,x}, j\right)\right]\right\}\ ,\ n = N-1,\dots,0$$

These maxima can be computed in $\mathcal{O}(q)$ instead of $\mathcal{O}(q^2)$ by starting from the biggest element $i = i_q$ down to the smallest element $i = i_1$ (in lexicographical order) and keeping track of the partial maxima.

Note that these two conditions hold for the numerical application from Section 4.5, providing the improved complexity $\mathcal{O}(q \cdot N \cdot M)$.

- The N term comes from the backward time induction.

- The M term corresponds to the cost of a regression, which is in $\mathcal{O}\,(M)$ (by using either the Cholesky decomposition or the more stable Thin SVD decomposition) .

4.4.1.2 Memory complexity

The memory size required for solving optimal switching problems (as well as the simpler American option problems and the more complex BSDE problems) by Monte Carlo methods is often said to be in $\mathcal{O}(N \cdot M)$, because, as the Euler scheme is a forward scheme and the dynamic programming principle is a backward scheme, the storage of the Monte Carlo trajectories seems inescapable. This fact is the major limitation of such methods, as acknowledged in [32] for instance.

Since such a complexity would be unbearable in high dimension, we describe below a general memory reduction method to obtain a much more amenable $\mathcal{O}(N + M)$ complexity (or, more precisely, of $\mathcal{O}(m \cdot N + q \cdot M)$ with $m \ll M$). This improvement really opens the door to the use of Monte Carlo methods for American options, optimal switching and BSDEs on high-dimensional practical applications. Note that this tool can be combined with all the existing Monte Carlo backward methods which (seem to) require the storage of all the trajectories.

A drawback of this tool is that it is limited to Markovian processes. However, one can usually circumvent this restriction by increasing the dimension of the state variable.

4.4.2 General memory reduction method

4.4.2.1 Description

The memory reduction method for Monte Carlo pricing of American options was pioneered by [36] for the geometric Brownian motion, and was subsequently extended to multi-dimensional geometric Brownian motions ([37]) as well as exponential Lévy processes ([38]). These papers take advantage of the additivity property of the processes considered. However, as briefly hinted in [104], the memory reduction trick can be extended to more general processes. In particular, it can be combined with any discretization scheme, for instance the Euler scheme or Milstein scheme, as long as the value of the stochastic process at one time step can be expressed via its value at the subsequent time step.

From a practical point of view, the production of "random" sequences usually involves wisely chosen deterministic sequences, with statistical properties as close as possible to true randomness (cf. [76] for instance for an overview). These sequences are usually set using a *seed*, i.e. a (possibly multidimensional) fixed value aimed at initializing the algorithm which produces the sequence:

$$\begin{array}{cccccccccc} \{\text{set seed } s\} & \to & s_1 & \to & s_2 & \to & \dots & \to & s_n \\ & \text{rand}() & & \text{rand}() & & \text{rand}() & & \text{rand}() & \\ & \downarrow & & \downarrow & & \downarrow & & \downarrow & \\ & \varepsilon_1 & & \varepsilon_2 & & \varepsilon_3 & & \varepsilon_n & \end{array} \tag{4.4.1}$$

The rand() produces a new random value ε and changes the internal seed value s. The internal value of the seed can be read (getseed()) and changed (setseed()). Now two useful aspects can be stressed. The first is that one can usually recover the current seed at any stage of the sequence. The second is that, if the seed is set later to, say, once again the seed s from equation (4.4.1), then the following elements of the sequence will be once again ε_1, ε_2, ... In other words, one can recover any previously produced subsequence of the sequence $(\varepsilon_n)_{n\geq 1}$, provided one stored beforehand the seed at the beginning of the subsequence. This feature is at the core of the memory reduction method, which we are going to discuss below in a general setting.

Consider a Markovian stochastic process $(X_t)_{t\geq 0}$, for instance the solution of the stochastic differential equation (4.2.2), recalled below:

$$\begin{aligned} X_0 &= x_0 \in \mathbb{R}^d \\ dX_s &= b(s, X_s)\, ds + \sigma(s, X_s)\, dW_s \end{aligned}$$

The application of the Euler scheme to this equation can be denoted as follows:

$$x^j_{t_{i+1}} = f\left(x^j_{t_i}, \varepsilon^j_i\right) \tag{4.4.2}$$

$$f(x, \varepsilon) := x + b(t_i, x)\, h + \sigma(t_i, x)\, \varepsilon\sqrt{h} \tag{4.4.3}$$

where $\forall i \in [0, N-1]$ and $\forall j \in [1, M]$, $\varepsilon^j_i \in \mathbb{R}^d$ is drawn from a d-dimensional Gaussian random variable. Suppose that for any $\varepsilon \in \mathbb{R}^d$, the function $x \mapsto f(x, \varepsilon)$ is invertible (call f_{inv} its inverse). Then, starting from the final value $x^j_{t_N}$ of the sequence (4.4.2), one can recover the whole trajectory of X:

$$x^j_{t_i} = f_{\text{inv}}\left(x^j_{t_{i+1}}, \varepsilon^j_i\right) \tag{4.4.4}$$

as long as one can recover the previous draws ε^j_{N-1}, ..., ε^j_0. The following pseudo-code describes an easy way to do it.

Algorithm 4.1 Euler Scheme

```
% Initialization
for j from 1 to M
        X[j] <- xj
end for

% LOOP 1: Euler scheme
for i from 0 to N-1
   S[i] <- getseed()
   for j from 1 to M
      E <- rand(d)
      X[j] <- f(X[j],E)
   end for
end for
S[N] <- getseed()
```

Inverse Euler Scheme

```
% LOOP 2: Inverse Euler scheme
for i from N-1 down to 0
   setseed(S[i])
   for j from 1 to M
      E <- rand(d)
      X[j] <- finv(X[j],E)
   end for
end for
setseed(S[N])
```

The first column of Algorithm 4.1 corresponds to the Euler scheme, with the addition of the storage of the seeds. At the end of the first column, the vector X contains the last values X_T^j, $j = 1, \dots, M$. From this point, one can recover the previous values $X_{t_i}^j$, $i = N-1, \dots, 0$, $j = 1, \dots, M$ as done in the second column.

Inside this last loop, one can perform the estimation of the conditional expectations required by the resolution algorithm of our stochastic control problem (equation (4.2.10)). Compared to the standard storage of the full trajectories $X_{t_i}^j$, $i = 0, \dots, N$, $j = 1, \dots, M$, the pros and cons are the following:

- The number of calls to the rand() function is doubled.
- The memory needed is brought down from $\mathcal{O}(M \times N)$ to $\mathcal{O}(M + N)$ (storage of the vector space and the seeds).

In other words, at the price of doubling the computation time, one can bring down the required memory storage by the factor $\min(M, N)$, which is a very significant saving. Moreover, the theoretical additional computation time can be insignificant in practice, as the availability of much more physical memory makes the resort to the slower virtual memory much less likely.

Remark 4.4.1. Even though the storage of the seeds does take $\mathcal{O}(N)$ in memory size, the constant may be much greater than 1. For instance, on Matlab®, a seed from the Combined Multiple Recursive algorithm (refer for instance to [76] for a description of several random generators) is made of 12 uint32 (32-bit unsigned integer), a seed from the Multiplicative Lagged Fibonacci algorithm is made of 130 uint64, and a seed from the popular Mersenne Twister algorithm is made of 625 uint32.

In order to relieve the storage of the seeds, we now provide a finer memory reduction algorithm (Algorithm 4.2). Although Algorithm 4.2 requires three main loops, it enables to perform the last loop without fiddling the seed of the random generator, and without any vector of seeds locked in memory, which will thus be fully dedicated to the regressions and other resolution operations. Moreover, the first two main loops can be performed beforehand once and for all, storing only the last values of the vector X as well as the first seed S[0]. Finally, if the random generator is able to leapfrop a given number of steps, the first loop can be drastically reduced.

Algorithm 4.2 General Memory Reduction Method

```
% LOOP 1: Seeds storage
for i from 0 to N-1
   S[i] <- getseed()
   for j from 1 to M
      E <- rand(d)
   end for
end for

% Initialization
for j from 1 to M
        X[j] <- xj
end for
%
%
%
%
```

```
% LOOP 2: Euler scheme
for i from 0 to N-1
   setseed(S[N-i-1])
   for j from 1 to M
      E <- rand(d)
      X[j] <- f(X[j],E)
   end for
end for
setseed(S[0]) ; free(S)

% LOOP 3: Inverse Euler scheme
for i from N-1 down to 0
   for j from 1 to M
      E <- rand(d)
      X[j] <- finv(X[j],E)
   end for
end for
```

4.4.2.2 Numerical stability

Theoretically, the trajectories produced by the Euler scheme (4.4.2) and the inverse Euler scheme (4.4.4) are exactly the same. In practice however, a discrepancy may appear, the cause of which is discussed below.

On a computer, not all real numbers can be reproduced. Indeed, they must be stored on a finite number of bits, using a predefined format (usually the IEEE Standard for Floating-Point Arithmetic (IEEE 754)). In particular, there exists an incompressible distance $\varepsilon > 0$ between two different numbers stored. This causes rounding errors when performing operations on real numbers.

For instance, consider $x \in \mathbb{R}$ and an invertible function $f : \mathbb{R} \mapsto \mathbb{R}$. Compute $y = f(x)$ and then compute $\hat{x} = f_{\text{inv}}(y)$. One would expect that $\hat{x} = x$, but in practice, because of rounding effects, one may get $\hat{x} = x + \epsilon z$ for a small $\epsilon > 0$, where z is a discrete variable, which can be deemed random, taking values around zero. This phenomenon is illustrated on Figure 4.4.1, which displays a histogram of $\hat{x} - x$ for $n = 10^7$ different values of $x \in [0, 1]$ and for the simple linear function $f(x) = 2x + 3$.

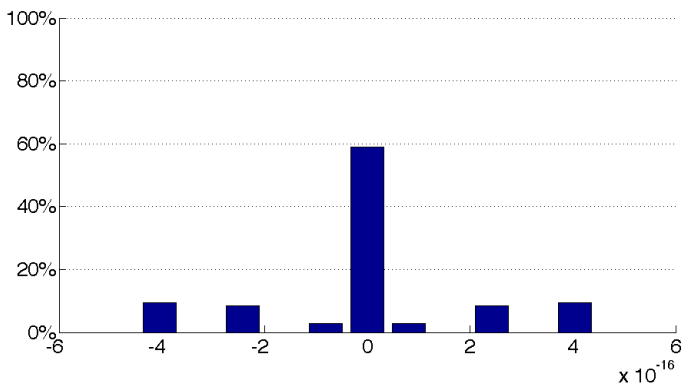

Figure 4.4.1: Histogram of rounding errors

We now describe how this affects our memory reduction method. Recall equation 4.4.2:

$$x^j_{t_{i+1}} = f\left(x^j_{t_i}, \varepsilon^j_i\right)$$

Now, instead of equation (4.4.4), the inverse Euler scheme will provide something like:

$$\begin{aligned} y^j_{t_N} &= x^j_{t_N} \\ y^j_{t_i} &= f_{\text{inv}}\left(y^j_{t_{i+1}}, \varepsilon^j_i\right) + \epsilon z^j_i \end{aligned} \tag{4.4.5}$$

for a small $\epsilon > 0$, where z^j_i, $i = 0, \dots, N$, $j = 1, \dots, M$, can be deemed realizations of a discrete random variable Z, independent of W. The distribution of Z is unknown, but data suggests it may be innocuously assumed centered, symmetric, and with finite moments.

We are now interested in studying the compound rounding error $y_{t_i} - x_{t_i}$ as a function of ϵ. Of course, its behaviour depends on the choice of f (equation (4.4.3)). Below, we explicit this error on two simple examples: an arithmetic Brownian motion and an Ornstein-Uhlenbeck process. These two examples illustrate how the compound rounding error can vary dramatically w.r.t. f.

First example: arithmetic Brownian motion Consider first the case of an arithmetic Brownian motion with drift parameter μ and volatility parameter σ. Here f and its inverse are given by:

$$\begin{aligned} f(x, \varepsilon) &= x + \mu h + \sigma\sqrt{h}\varepsilon \\ f_{\text{inv}}(x, \varepsilon) &= x - \mu h - \sigma\sqrt{h}\varepsilon \end{aligned}$$

Hence, using equation (4.4.5), for every $j = 1, \dots, M$:

$$y^j_{t_i} - x^j_{t_i} = \epsilon \sum_{k=i}^{N-1} z^j_k$$

In other words, the compound rounding error behaves as a random walk, multiplied by the small parameter ϵ. Hence, as long as $\epsilon \ll h$ (which is always the case as real numbers smaller than ϵ cannot be handled properly on a computer), this numerical error is harmless.

Remark that a similar numerical error arises from the algorithms proposed in [36] , [37] and [38], but, fortunately, as discussed above, this error is utterly negligible.

Second example: Ornstein-Uhlenbeck process Now, consider the case of an Ornstein-Uhlenbeck process with mean reversion $\alpha > 0$, long-term mean μ and volatility σ. Here:

$$\begin{aligned} f(x, \varepsilon) &= x + \alpha(\mu - x)h + \sigma\sqrt{h}\varepsilon \\ f_{\text{inv}}(x, \varepsilon) &= \frac{1}{1 - \alpha h}\left(x - \alpha\mu h - \sigma\sqrt{h}\varepsilon\right) \end{aligned}$$

Using equation (4.4.5), for every $j = 1, \dots, M$ the compound error is given by:

$$y^j_{t_i} - x^j_{t_i} = \epsilon \sum_{k=i}^{N-1} \frac{1}{(1 - \alpha h)^{k-i}} z^j_k$$

As $(1 - \alpha h)^{-N} \sim \exp(\alpha T)$ when $h \to 0$, one can see that, as soon as $T > -\frac{\ln(\epsilon)}{\alpha}$, this error may become overwhelming. This phenomenon is illustrated on Figure 4.4.2a on a sample of 100 trajectories.

In order to mitigate this effect, we propose to modify the Algorithm 4.2 as follows: in its second loop (usual Euler scheme), instead of saving only the last values x_T^j, one may define a small subset $\tilde{\Pi} \subset \Pi$ and save the intermediate values $x_{t_i}^j$, $t_i \in \tilde{\Pi}$. Then, in the last loop (inverse Euler scheme), every time that $t_i \in \tilde{\Pi}$, the current value of the set $x_{t_i}^j$ may be recovered from this previous storage.

Figure 4.4.2b illustrates the new behaviour of the compound rounding error with this mended algorithm, on an example with $T = 10$ years and 4 intermediate saves (in addition to the final values).

The drawback of this modification, of course, is that it multiplies the required storage space by the factor $\#\tilde{\Pi}$. However, this remains much smaller than the $\mathcal{O}(M \times N)$ required by the naive full storage algorithm.

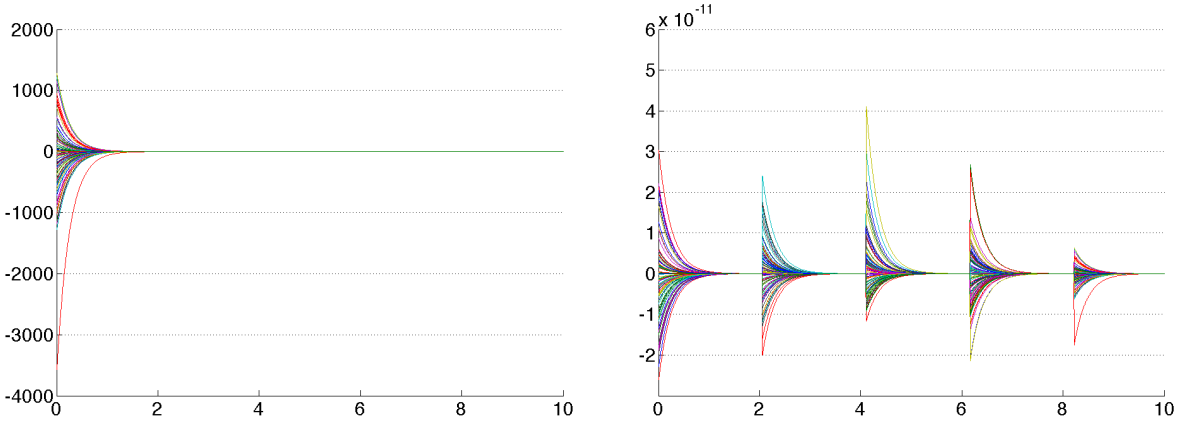

(a) Without intermediate saves

(b) With intermediate saves

Figure 4.4.2: Compound rounding error for the Ornstein-Uhlenbeck process

4.5 Application to investment in electricity generation

This section is devoted to an application of the resolution method studied in Section 4.3 to an investment problem in electricity generation.

Since our intention here is to show that the algorithm described in Section 4.3 can handle high-dimensional problems, our modeling of the electric system focuses on the various fundamental drivers of the electricity spot price formation mechanism that are electricity demand, available capacities and above all fuel prices.

Thus, were neglected some strategic aspects of investment, like construction delays and network constraints. We did not consider dynamic constraints of production either, which are known to increase spot price during peak hours and to decrease them during off peak hours (see [80]), as we consider these effects to be negligible compared to the effect induced by a lack or an excess of capacity.

We based our model on [3, 1] where the electricity spot price is defined as a combination of fuel prices adjusted by a scarcity factor. This model exhibits the main feature wanted here, which is that the spot price, being determined both by the fuel prices and the residual capacity, is directly affected by the evolution of the installed capacity. When the residual capacity tends to decrease, spot prices will tend to increase, making investment valuable. Thus, in this model,

investments are undertaken not on the specific purpose of satisfying the demand but as soon as they are profitable. Energy non-served and loss of load probability may still be adjusted through the price cap on the spot market.

In this section, we first detail the chosen modeling and objective function (which will be shown to be encompassed in the general optimal multiple switching problem (4.2.1)), and then solve it numerically using the general algorithm developed in the previous sections.

4.5.1 Modeling

The key variable in order to describe our electricity generation investment problem is the price of electricity. More precisely, the key quantities are the spreads between the prices of electricity and other energies. To model these spreads accurately, it may be worth considering a structural model for electricity (cf. the survey [30]). Here we choose such a model, mainly inspired by those introduced in [3] and [1], albeit amended and customized for a long-term time horizon. All the variables involved are detailed below.

4.5.1.1 Electricity demand

The electricity demand, or electricity load, at time t on the given geographical zone considered is modelled by an exogenous stochastic process $(D_t)_{t\geq 0}$:

$$D_t = f_0(t) + Z_t^0 \tag{4.5.1}$$

where Z^0 is an Ornstein-Uhlenbeck (henceforth O.U.) process:

$$dZ_t^0 = -\alpha_0 Z_t^0 dt + \beta_0 dW_t^D$$

where α_0 and β_0 are constants, and f_0 is a deterministic function that takes into account demand seasonalities.

4.5.1.2 Production capacities

Let d' be the number of different production technologies. Denote as $I_t = \left(I_t^1, \dots, I_t^{d'}\right)$ the installed production capacities at time t. They represent the maximum amount of electricity that is physically possible to produce. These fleets can be modified: at a given time τ_n, one can decide to build (or dismantle) an amount ζ_n of capacities:

$$I_{\tau_n} = I_{\tau_n^-} + \zeta_n\,,\ n \geq 0 \tag{4.5.2}$$

Denote as $\alpha = (\tau_n, \zeta_n)_{n\geq 1}$ the corresponding impulse control strategy, where $(\tau_n)_{n\geq 0}$ is an increasing sequence of stopping times with $\tau_n \nearrow \infty$ when $n \to \infty$, and $(\zeta_n)_{n\geq 0}$ is a sequence of vectors corresponding to the increases (or decreases) in capacities. Apart from these variations, I_t will be deemed constant, i.e.:

$$I_t = I_{0-} + \sum_{n,\, \tau_n \leq t} \zeta_n\,. \tag{4.5.3}$$

Now, denote as $C_t = \left(C_t^1, \dots, C_t^{d'}\right)$ the available production capacities. Because of spinning reserves, maintenance and random outages, these quantities are lower than the installed capacities I_t, which represent their physical maximum. In other terms, C_t is a fraction of I_t:

$$C_t^i = I_t^i \times A_t^i \tag{4.5.4}$$

for every $1 \leq i \leq d'$, where A_t^i corresponds to the rate of availability of the i$^{\text{th}}$ production technology. Therefore one must choose a model for the process A_t that ensures that it stays within the interval $[0, 1]$.

One possibility is to make use of the bounded Jacobi process (see for instance [103],where it is used to model stochastic correlations, and the references therein for more information on this process). This process is however tricky to estimate and simulate (see [59] for the description of some possible methods). Moreover, its main simulation method, the truncated Euler scheme, prevents the use of the memory reduction method described in Subsection 4.4.2. Hence we look for a simpler model.

Adapting the (bounded) wind power infeed efficiency model from [105], we model $(A_t^i)_{t\geq 0}^{1\leq i\leq d'}$ as follows:

$$A_t^i := \mathcal{T}\left(f_i(t) + Z_t^i\right) \tag{4.5.5}$$

where Z, f and $\mathcal{T}$ are chosen as follows:

- Z^i is an O.U. process :
$$dZ_t^i = -\alpha_i Z_t^i dt + \beta_i dW_t^{Z^i}$$
where $\alpha_i > 0$, $\beta_i > 0$ and $\left(W_t^{Z^i}\right)_{t\geq 0}$ is a Brownian motion.

- The deterministic function f_i accounts for the seasonality in the availability of production capacities, which stems from the maintenance plannings, which usually mimic the long term seasonality of demand (which in turn originates in the seasonality of temperature).

- The mapping $\mathcal{T} : \mathbb{R} \to [0, 1]$ is here to ensure that $\forall t \geq 0$, $A_t \in [0, 1]^{d'}$. One can choose the versatile logit function as in [105], or any other mapping of $\mathbb{R}$ into $[0, 1]$. For instance, any cumulative distribution function would be suitable. As the process Z is Gaussian and asymptotically stationary, we choose for $\mathcal{T}$ the (standard) normal cumulative distribution function. In particular, this choice makes the calibration process instantaneous.

4.5.1.3 Fuels and CO_2 prices

For each technology i, denote as S_t^i the price of the fuel i to produce electricity at time t. In the particular case of renewable energies, which, *per se,* do not involve traded fuels, the corresponding S_t^i can be chosen to be zero. Moreover, define S_t^0 as the price of CO_2. Denote as S_t the full vector $\left(S_t^0, S_t^1, \ldots, S_t^{d'}\right)$.

Now, we introduce the multiplicative constants needed to convert theses quantities into €/MWh. For each technology $i = 1, \ldots, d'$, let h_i denote its heat rate, and h_i^0 denote its CO_2 emission rate. Hence, the quantity

$$\tilde{S}_t^i := h_i^0 S_t^0 + h_i S_t^i \tag{4.5.6}$$

expressed in €/MWh, corresponds to the price in € to pay in order to produce 1MWh of electricity using the ith technology. We note $h^0 = (h_1^0, \ldots, h_{d'}^0) \in \mathbb{R}^{d'}$ and $h = (h_1, \ldots, h_{d'}) \in \mathbb{R}^{d'}$.

Remark 4.5.1. One can choose to add a fixed cost into the definition of $\tilde{S}_t^i$. This is all the more so relevant for technologies whose fixed costs outweigh the cost of fuel (e.g. nuclear).

Over long time horizons, it is crucial to take into account the existence of long-term relationships between energy prices (c.f.[90] for instance). Thus, extending the model of cointegrated Brownian motions from [14], we model S_t as cointegrated geometric Brownian motions:

$$dS_t = \Xi S_t dt + \operatorname{diag}(S_t)\, \Sigma dW_t^S$$

where Ξ is the $(d'+1)\times(d'+1)$ cointegration matrix (which models the long term relations), Σ is the $(d'+1)\times(d'+1)$ covariance matrix (which models the short term behaviour), and $\left(W_t^S\right)_{t\geq 0}$ is a $(d'+1)$-dimensional Brownian motion. We assume that $1\leq \operatorname{rank}(\Xi) < d'$ (so as to produce "true" cointegration, see [14]), and that for every $i\neq j$, $\Xi_{i,j}\geq 0$ (so as to ensure that the process S stays positive, see Appendix 4.7.2).

4.5.1.4 Electricity price

We model the price of electricity using a long-term structural model. We model it as the sum of two building blocks: the marginal cost of producing electricity (cf. [3] for more details) plus a power law *scarcity premium* (along the lines of [1]), this sum being capped at a fixed upper bound [1].

For any time $t\geq 0$, define the permutation $(1),\dots(d')$ of the numbers $1,\dots,d'$, such that $\widetilde{S}_t^{(1)}\leq\dots\leq\widetilde{S}_t^{(d')}$. Then, define $\overline{C}_t^{(i)}$ as the total capacity available at time t from the i first technologies, i.e. $\overline{C}_t^{(i)} := \sum_{j\leq i} C_t^{(j)}$.

Now, from two points (x_1,y_1) and (x_2,y_2) in $\mathbb{R}^2$, one can always find three positive constants $a := a(x_1,x_2,y_1,y_2)$, $b := b(x_1,x_2,y_1,y_2)$ and $c := c(x_1,x_2,y_1,y_2)$ such that the function:

$$p(x) := p(x;x_1,x_2,y_1,y_2) = \frac{x}{a-bx} + c \tag{4.5.7}$$

satisfies $p(x_1)=y_1$ and $p(x_2)=y_2$ [2].

Using this notation, we model the price P_t of electricity as follows:

$$\begin{aligned} P_t \;:=\; & \widetilde{S}_t^{(1)}\mathbf{1}\left\{D_t<0\right\} + \left\{\widetilde{S}_t^{(1)} + p\left(D_t;0,\overline{C}_t^{(1)},\widetilde{S}_t^{(1)},\widetilde{S}_t^{(2)}\right)\right\}\mathbf{1}\left\{0\leq D_t<\overline{C}_t^{(1)}\right\} \\ & \sum_{i=2}^{d'-1}\left\{\widetilde{S}_t^{(i)} + p\left(D_t;\overline{C}_t^{(i-1)},\overline{C}_t^{(i)},\widetilde{S}_t^{(i)},\widetilde{S}_t^{(i+1)}\right)\right\}\mathbf{1}\left\{\overline{C}_t^{(i-1)}\leq D_t<\overline{C}_t^{(i)}\right\} \\ & + \left\{\widetilde{S}_t^{(d')} + p\left(D_t;\overline{C}_t^{(d'-1)},\overline{C}_t^{(d')},\widetilde{S}_t^{(d')},P_{\max}\right)\right\}\mathbf{1}\left\{\overline{C}_t^{(d'-1)}\leq D_t\right\} \end{aligned} \tag{4.5.8}$$

where $P_{\max}>0$ is a fixed upper bound on the price of electricity. In particular, the last term, the one involving $P_{\max}$, enables price spikes to occur (when the residual capacity is small). Remark that the price of CO_2 emissions is explicitly included in the marginal cost (through equation (4.5.6)). Finally, remark that thanks to the knitting function (4.5.7), the electricity price P is a Lipschitz continuous function of the structural variables D, C and S[3].

Remark 4.5.2. Remark that the price model (4.5.8) is different from the one proposed in [1]. If both model are indeed based on a penalization of the marginal cost by a scarcity factor, this factor is multiplied to the marginal cost in [1], while here it is added to the marginal cost, in a manner that ensures the continuity of the structural function that generates the power price. Without this modification, one could try to preserve the continuity of the structural function by including the electricity price P, regarded as a jump diffusion process, among the state variables. However, the type of discontinuity (composition of an indicator function with a diffusion) is such that P would not be a jump diffusion process (it would not admit any jump

[1] Indeed, in the French, German and Austrian markets for instance, power prices cannot be set outside the $[-3000,3000]$€/MWh range, see http://www.epexspot.com/en/product-info/auction.

[2] For instance, fix $a>0$, then define $b=\frac{1}{2}\left(x_1+x_2+\sqrt{(x_2-x_1)^2+4a\frac{x_2-x_1}{y_2-y_1}}\right)$ and finally $c=y_1-\frac{a}{b-x_1}$.

[3] Rigorously, this property requires that C does not reach zero. One can, for instance, add a fixed minimum availability rate $1\gg a_{\min}>0$ to the definition (4.5.5), replacing $\mathcal{T}$ by $a_{\min}+(1-a_{\min})\mathcal{T}$

measure, cf. Appendix 4.7.3). Therefore, we opted for the simple additive model (4.5.8). This modification does not alter the calibration process very much, and seems to be better suited to long term modeling.

4.5.1.5 Objective function

We now explicit the objective function of the investor in electricity generation. Suppose that, at time t, the level of installed capacity of type $j \in [1, d']$ is changed from I^j_{t-} to $I^j_s = I^j_{t-} + \zeta^j$, $s \geq t$. It generates the cost:

$$k\left(\zeta^j\right) := \begin{cases} \kappa_j^{f+} + \zeta^j \kappa_j^{p+} & , \zeta^j > 0 \\ 0 & , \zeta^j = 0 \\ \kappa_j^{f-} - \zeta^j \kappa_j^{p-} & , \zeta^j < 0 \end{cases}$$

where κ_j^{f+} and κ_j^{p+} are the fixed and proportional costs of building new plants of type j, and κ_j^{f-} and κ_j^{p-} are the fixed and proportional costs of dismantling old plants of type j.

Consider the case of new plants ($\zeta^j > 0$). Assuming that the global availability rate(4.5.5) of technology j applies to the new plants, they can then produce up to $\zeta^j A^j_s$, $s \geq t$, or, more precisely, according to the stack order principle, up to

$$\min\left\{\zeta^j A^j_s,\ \left(D_s - \overline{C}_s^{(j-1)}\right)^+\right\}$$

assuming that, in the stack order, the new plants are called before the older plants I_{t-} of the same technology (as their efficiency rate can be expected to be at least slightly better than the older plants of the same technology, a phenomenon that can be seen as a partial interpretation of the knitting function (4.5.7)).

At time $s \geq t$, this production is sold at price P_s, but costs $\tilde{S}_s$ to produce (if $P_s < \tilde{S}_s$, then of course the producer chooses not to produce). In addition, regardless of the output level, there may exist a fixed maintenance cost κ_j. Summing up all these gains, discounted to time t using a constant interest rate $\rho > 0$, the new plants yield a (random) revenue of:

$$\int_t^\infty e^{-\rho s}\left(\min\left\{\zeta^j A^j_s,\ D_s - \overline{C}_s^{(j-1)}\right\} \times \left(P_s - \widetilde{S}^j_s\right)^+ - \kappa_i\right) ds$$

(noticing that our power price model is such that $\left\{D_s - \overline{C}_s^{(j-1)} \leq 0\right\} \Leftrightarrow \left\{P_s - \widetilde{S}^j_s \leq 0\right\}$).

Now, after this cost-benefit analysis for one quantity ζ^j of new plants, consider the gains of the whole fleet of power plants on a given geographical zone. The maximization of the expected discounted gains along the potential new plants yields the following stochastic control problem:

$$v\left(t,x,i\right) = \sup_{\alpha \in \mathcal{A}_{t,i}} \mathbb{E}\left[\sum_{j=1}^{d'} \int_t^\infty e^{-\rho s}\left(\min\left\{C^j_s, D_s - \overline{C}_s^{(j-1)}\right\} \times \left(P_s - \widetilde{S}^j_s\right)^+ - \kappa_i\right) ds - \sum_{\tau_n \geq t} e^{-\rho \tau_n} k\left(\zeta^j\right)\right] \tag{4.5.9}$$

where the strategies α affect the installed capacities (equations (4.5.3)), hence also the available capacities (equation (4.5.4)) as well as the power price (equation (4.5.8)), and where the cash flows are purposely discounted up to time 0, which is the time of interest.

Remark 4.5.3. Replacing P in (4.5.9) by its definition (4.5.8), it is patent that this objective function fits into the mould studied thoroughly in Section 4.3. In Subsection 4.5.2 below, we apply our algorithm to this specific objective function.

Remark 4.5.4. Remark that under this modeling, the demand is satisfied as long as it does not exceed the total available capacity. Indeed, the effective output of the plant ζ^j is equal to $\mathbf{1}\left\{P_s - \widetilde{S}_s^j > 0\right\} \times \min\left\{\zeta^j A_s^j, \left(D_s - \overline{C}_s^{(j-1)}\right)^+\right\}$. It is indeed governed by the electricity spot price level, but, as under our modelling $\left\{P_s - \widetilde{S}_s^j > 0\right\} = \left\{D_s - \overline{C}_s^{(j-1)} > 0\right\}$ a.s., summing up the effective outputs of all the power plants yields $\sum_{j=1}^{d'} \min\left\{C_s^j, D_s - \overline{C}_s^{(j-1)}\right\} \times \mathbf{1}\left\{D_s - \overline{C}_s^{(j-1)} > 0\right\} = \min\left\{D_s, \overline{C}_s^{(d')}\right\}$.

4.5.2 Numerical results

Finally, we solve the control problem described in Subsection 4.5.1 on a numerical example, using the algorithm detailed in Subsection 4.3 combined with the general memory reduction method described in Subsection 4.4.2.

Our purpose here is not to perform a full study of investments in electricity markets, but a more modest attempt at illustrating the practical feasibility of our approach, with some possible outputs that the algorithm can provide.

We consider a numerical example including two cointegrated fuels (in addition to the price of CO_2): one "base fuel" and one "peak fuel", starting respectively from 40€/MWh and 80€/MWh. Hence, using the notations from Subsection 4.5.1, $d' = 2$ (two technologies) and $d = 6$ (electricity demand, CO_2 price, two fuel prices and two availability rates). The main choices of parameters for this application (initial fuel prices and volatilities, initial fleet and proportional costs of new power plants) are summed up in Table 4.5.1. Moreover, the demand process starts from $D_0 = 70$GW and does not integrate any linear trend.

i	S_0^i	σ^i	I_0^i	κ_i^{p+}
1	40€/MWh	5%	67GW	2.00 10^9€/GW
2	80€/MWh	15%	33GW	0.24 10^9€/GW

Table 4.5.1: Model parameters

In order to take into account the minimum size of one power plant we restrict the values of the installed capacity process(4.5.3) to a (bi-dimensional) fixed grid $\Lambda^{d'}$, with a mesh of 1GW. We make the simplifying assumptions that investments are irreversible, and that no dismantling can occur (recall from Subsection 4.4.1 the computational gain provided by this assumption).

Remark 4.5.5. If such a grid is indeed manageable in dimension $d' = 2$, it may less be the case if additional technologies were considered. However, as discussed in [100] equation (3.2), instead of performing one regression for each $i \in \Lambda^{d'}$, one can solve equation (4.3.15) at time t_i by only one $(d + d')$-dimensional regression, by choosing an a priori law for the randomized control ζ_{t_i}. The error analysis from Section 4.2 can be generalized to such regressions in higher dimension.

Finally, we consider the following numerical parameters. We choose a time horizon $T = 40$ years and a time step $h = \frac{1}{730}$ (i.e. two time steps per day, allowing for some intraday pattern in the demand process) but allow for only one investment decision per year. For the regression, we consider a basis of $b = 2^d = 64$ adaptative local functions, chosen piecewise linear on each hypercube (which is a bit more refined than the piecewise constant basis studied in Section 4.3) on a sample of $M = 5000$ trajectories.

With these numerical parameters, we obtain a non-parametric confidence interval of $[3.731, 3.752] \times 10^8$ for the value function $v(0, x_0, i_0)$ at time 0 (cf. Appendix 4.7.4 on how these bounds are

computed), i.e. a relative error smaller than 1%, which is sufficiently small for the numerical results obtained, displayed on Figures 4.5.1 and 4.5.2, to be considered relevant.

First, Figure 4.5.1 deals with the optimal strategies. Figure 4.5.1a displays the time evolution of the average as well as the variability of the optimal fleet (only the new plants are shown). One can distinguish a first short phase characterised by the construction of several GW of peak load assets, followed by a much slower second phase involving the construction of both base load and peak load assets. Moreover, the variability of the optimal fleet increases over time. The detailed histogram of the optimal strategy at time $T = 40$ years is displayed on Figure 4.5.1b, where it is combined with the price of fuel. One can see that the more the peak fuel is expensive (and hence both fuels are expensive on average, as they are cointegrated), the more constructions of base load plants occur.

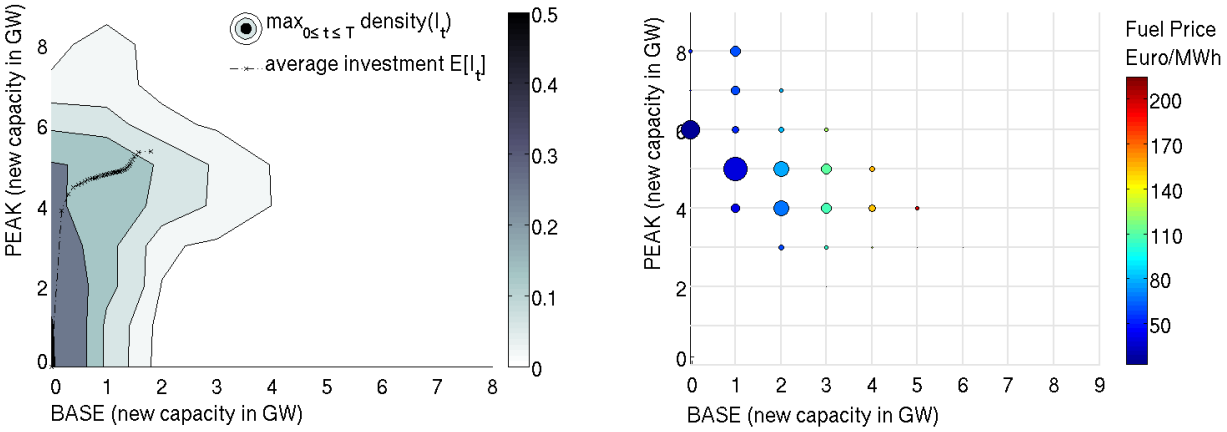

(a) Time evolution of new capacities

(b) Final fleet distribution

Figure 4.5.1: Optimal strategies

The fact that the average fleet seem to converge is related to the fact that this numerical example does not consider any growth trend in the electricity demand. Otherwise, more investments would occur, indeed, to keep pace with consumption.

Then Figure 4.5.2 provides information on the price of electricity. Figure 4.5.2a displays the time evolution of the electricity spot price density. For better readability, each density covers one whole year. One can see how the density moves away from the initial bimodal density (with prices clustering around the initial prices of the two fuels) towards a more diffuse density. Moreover, the downward effect of investments on prices can be noticed.

This downward effect is even more visible on Figure 4.5.2b. It compares the effect on electricity prices of three different strategies: the optimal strategy, the optimal deterministic strategy (computed as the average of the optimal strategy), and the do-nothing strategy. For each strategy, the joint time-evolution of the yearly median price and the yearly interquartile range are drawn. As expected, prices tend to be higher and more scattered without any new plant. Nevertheless, on this specific example, the price distribution under the optimal deterministic strategy is close to that under the optimal strategy (only slightly more scattered).

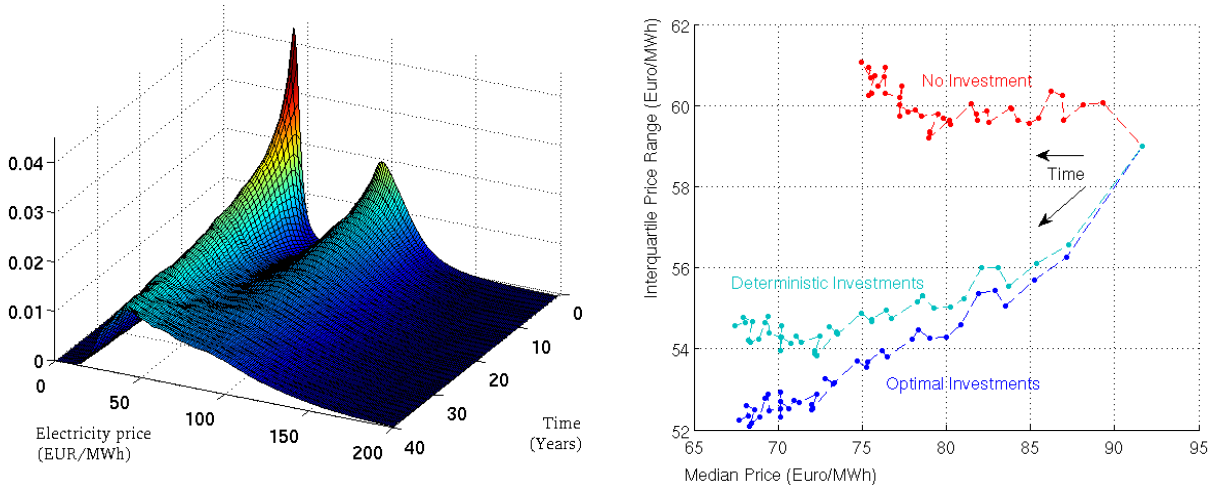

(a) Time evolution of electricity spot price density

(b) Comparison between investment strategies

Figure 4.5.2: Electricity spot price

Figure 4.5.3 provides an alternative graphical representation of the time evolution of the distribution of the electricity spot price. The construction of these graphics is detailed in Appendix 4.7.5. Two strategies are compared: the do-nothing strategy (Figure 4.5.3a) and the optimal strategy (Figure 4.5.3b). One can clearly see that the new constructions of plants decrease the probability of extremely high prices (as shown by the extent of the 99% areas), and decrease the use of the peak plants (to such an extent that the distribution turns unimodal before the end of the time period).

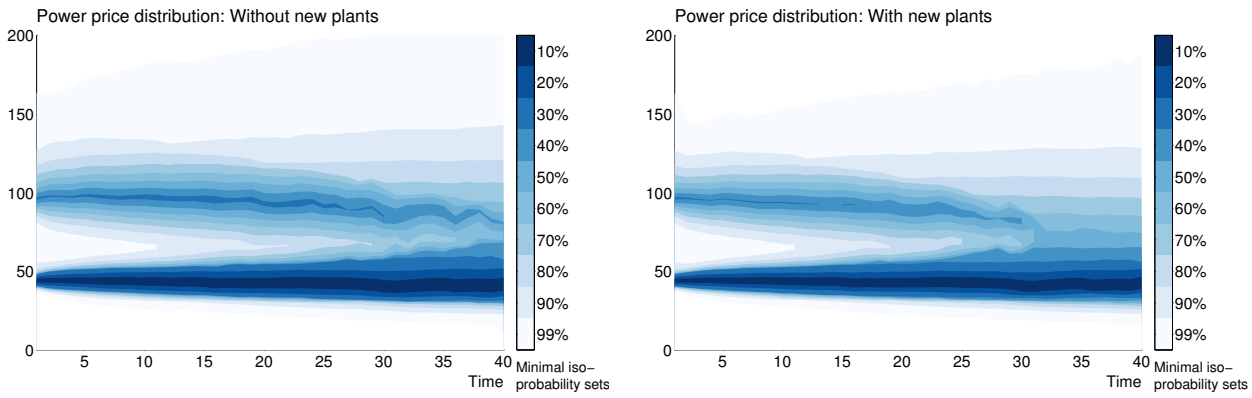

(a) Without new plants (no investments)

(b) With new plants (optimal investments)

Figure 4.5.3: Time-evolution of electricity spot price

These few pictures illustrate the kind on information that can be be extracted from the resolution of this control problem. Of course, as a by-product of the resolution, much more can be extracted and analyzed (distribution of income, CO_2 emissions, optimal exercise frontiers, etc) if needed.

4.6 Conclusion

In this chapter, we presented a probabilistic method to solve optimal multiple switching problems. We showed on a realistic investment model for electricity generation that it can efficiently provide insight into the distribution of future generation mixes and electricity spot prices. We intend to develop this work in several directions in the future. First, we wish to take into account more generation technologies, most notably wind farms, nuclear production, as well as solar distributed production. These additions would raise the dimension of the problem from eight to fifteen. Yet another range of innovations in numerical methods will be necessary to overcome this increase in dimension. Second, we wish to take time-to-build into account. And last but not least, we wish to adapt the problem to a continuous-time multiplayer game and contribute to the quest for an efficient algorithm to solve it.

4.7 Appendices

4.7.1 L_p convergence speed of empirical mean

Lemma 4.7.1. *For every $p \geq 1$, there exists $C_p > 0$ such that for any i.i.d. sample $X_1, \dots, X_M$ of $\mathbb{R}$-valued random variables such that $\mathbb{E}[X_1] = 0$ and $\mathbb{E}\left[|X_1|^{p\vee 2}\right] < \infty$, the following holds:*

$$\left\| \frac{1}{M} \sum_{m=1}^{M} X_m \right\|_{L_p} \leq \frac{C_p}{\sqrt{M}} \|X_1\|_{L_{p\vee 2}} \tag{4.7.1}$$

Proof. Using Marcinkiewicz-Zygmund's inequality, there exists $C_p > 0$ such that:

$$\mathbb{E}\left[\left| \sum_{m=1}^{M} X_m \right|^p \right] \leq C_p \mathbb{E}\left[\left(\sum_{m=1}^{M} |X_m|^2 \right)^{\frac{p}{2}} \right]$$

Multiplying both sides by $\frac{1}{M^p}$:

$$\mathbb{E}\left[\left| \frac{1}{M} \sum_{m=1}^{M} X_m \right|^p \right] \leq \frac{C_p}{M^{\frac{p}{2}}} \mathbb{E}\left[\left(\frac{1}{M} \sum_{m=1}^{M} |X_m|^2 \right)^{\frac{p}{2}} \right] \tag{4.7.2}$$

If $p \geq 2$, then $\frac{p}{2} \geq 1$ and, using Jensen's inequality:

$$\left(\frac{1}{M} \sum_{m=1}^{M} |X_m|^2 \right)^{\frac{p}{2}} \leq \frac{1}{M} \sum_{m=1}^{M} \left(|X_m|^2 \right)^{\frac{p}{2}} = \frac{1}{M} \sum_{m=1}^{M} |X_m|^p$$

Taking expectations on both sides:

$$\mathbb{E}\left[\left(\frac{1}{M} \sum_{m=1}^{M} |X_m|^2 \right)^{\frac{p}{2}} \right] \leq \mathbb{E}\left[|X_1|^p \right] \tag{4.7.3}$$

Now, if $p < 2$, then $\frac{p}{2} < 1$ and, using Jensen's inequality:

$$\mathbb{E}\left[\left(\frac{1}{M} \sum_{m=1}^{M} |X_m|^2 \right)^{\frac{p}{2}} \right] \leq \mathbb{E}\left[\left(\frac{1}{M} \sum_{m=1}^{M} |X_m|^2 \right) \right]^{\frac{p}{2}} = \mathbb{E}\left[|X_1|^2 \right]^{\frac{p}{2}} \tag{4.7.4}$$

Then combine inequalities (4.7.2), (4.7.3) and (4.7.4) and take the power $\frac{1}{p}$ to obtain inequality (4.7.1). □

4.7.2 Positivity of cointegrated geometric Brownian motions

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space. Consider the following d-dimensional process:

$$\begin{aligned} dS_t &= \Xi S_t dt + \text{diag}\left(S_t\right) \Sigma dW_t \\ S_0 &> 0 \end{aligned}$$

where W is a $\mathcal{F}$-adapted, d-dimensional Brownian motion, Ξ is a $d \times d$ cointegration matrix, and Σ is a $d \times d$ covariance matrix.

Proposition 4.7.1. *$S > 0$ a.s. if and only if $\forall i \neq j$, $\Xi_{i,j} \geq 0$*

Proof. First, suppose that $\forall i, j = 1, \ldots, d$, $i \neq j$, $\Xi_{i,j} \geq 0$. Consider the following stopping time:

$$\tau = \inf\left\{t \geq 0\,;\, \exists j \in [1,d] \text{ } s.t. \, S_\tau^j = 0\right\}$$

i.e. τ is the first time when one component of S reaches 0. In particular, $S_t \geq 0$ a.s. $\forall t \in [0, \tau]$.

Now, suppose that $\tau < \infty$. There exists at least one component i such that $S_\tau^i = 0$. Recall the dynamics of S^i:

$$dS_t^i = \left(\sum_{j=1}^d \Xi_{i,j} S_t^j\right) dt + S_t^i \left(\sum_{j=1}^d \Sigma_{i,j} dW_t^j\right)$$

By Girsanov's theorem, there exists a probability measure $\mathbb{Q}^i$, equivalent to $\mathbb{P}$, such that

$$dS_t^i = \left(\sum_{1 \leq j \leq d; j \neq i} \Xi_{i,j} S_t^j\right) dt + S_t^i \left(\sum_{j=1}^d \Sigma_{i,j} d\tilde{W}_t^j\right)$$

where $\tilde{W}$ is a d-dimensional $\mathbb{Q}^i$-Brownian motion. Then, using Proposition (2.3) from [95] (Chapter IX):

$$S_t^i = \mathcal{E}\left(\tilde{X}^i\right)_t \left\{S_0^i + \int_0^t \mathcal{E}\left(\tilde{X}^i\right)_s^{-1} \left(\sum_{1 \leq j \leq d; j \neq i} \Xi_{i,j} S_s^j\right) ds\right\} \tag{4.7.5}$$

where $\tilde{X}_t^i := \sum_{j=1}^d \Sigma_{i,j} d\tilde{W}_t^j$, and $\mathcal{E}\left(\tilde{X}^i\right)$ denotes the exponential martingale of $\tilde{X}^i$. At time τ, it yields:

$$0 = S_\tau^i = \mathcal{E}\left(\tilde{X}^i\right)_\tau \left\{S_0^i + \int_0^\tau \mathcal{E}\left(\tilde{X}^i\right)_s^{-1} \left(\sum_{1 \leq j \leq d; j \neq i} \Xi_{i,j} S_s^j\right) ds\right\} > 0$$

using the positivity of S_0^i and of the exponential martingale, as well as the non-negativity of $\Xi_{i,j}$, $i \neq j$, and of S before τ. This contradiction means that $S_\tau^i > 0$. As the same reasoning can be applied for every $i \in [1, d]$, this means that $\tau = \infty$, i.e. that $S > 0$ a.s. .

Next, suppose that $S > 0$ a.s. . Choose $i \in [1, d]$. Using equation (4.7.5) and the positivity of S, we obtain:

$$S_0^i + \sum_{1 \leq j \leq d; j \neq i} \Xi_{i,j} \int_0^t \mathcal{E}\left(\tilde{X}^i\right)_s^{-1} S_s^j ds > 0 \text{ a.s.}$$

As $S_0^i > 0$ and the coefficients $\int_0^t \mathcal{E}\left(\tilde{X}^i\right)_s^{-1} S_s^j ds$ are a.s. positive with support $\mathbb{R}_+$, the only possibility for the above inequality to hold a.s. is that $\Xi_{i,j} \geq 0$ for all $i \neq j$. □

4.7.3 No jump measure for diffusion-based discontinuities

This Appendix deals with discontinuous jump processes of the type $\mathbf{1}_{\{W_t>0\}}$, ie. the composition of an indicator function of a subset of $\mathbb{R}$ with a $\mathbb{R}-$valued diffusion. Such processes arise for instance in the electricity spot model studied in [1]. We show below (on the representative example $\mathbf{1}_{\{W_t>0\}}$) that such processes cannot admit any jump measure.

Theorem 4.7.1. *Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space, $T > 0$ be a fixed time horizon, and $W = (W_t)_{t\in[0,T]}$ be a standard Brownian motion for the measure $\mathbb{P}$. The filtration $\mathcal{F} = (\mathcal{F}_t)_{t\in[0,T]}$ is the natural filtration, satisfying the usual conditions, generated by W. Then, the stochastic process $J^W := \left(J_t^W\right)_{t\in[0,T]}$ defined by $J_t^W := \mathbf{1}_{\{W_t>0\}}$, $t \in [0,T]$, does not admit any jump measure.*

Proof. Suppose that J^W admits a jump measure, i.e. there exists a random measure M such that:

$$\forall\omega \in \Omega,\ \forall t \in [0,T],\ J_t^W(\omega) = M(\omega, [0,t]) = \int_0^t M(\omega, ds)$$

Given that the stochastic process J^W is a.s. non-monotone, the random measure M must be a signed random measure, i.e. for a.e. $\omega \in \Omega$, $M(\omega, .)$ is a countably additive real function defined on the σ-algebra $\mathcal{B}([0,T])$. By simple properties of signed measures, M need to be a.s. finite, and even bounded (cf. for instance [20], Corollary 3.1.3 p.176). Consequently, the Jordan decomposition holds, i.e. there exists two positive (random) measures M^+ and M^- such that:

$$M = M^+ - M^-$$

Moreover, both measures M^+ and M^- are a.s. finite.

Now, we are going to show that the finite measure M^+ is not countably additive. The same reasoning will hold for M^-.

For every $n \geq 1$, define $A_n := \left]0, \frac{1}{n}\right]$. Then, for every $n \geq 1$ $A_n \subset A_{n+1}$, and $\cap_{n=1}^{\infty} A_n = \emptyset$. Now, the law of the iterated logarithm (cf. [95]) implies that a.s. $\limsup_{t\searrow 0} W_t > 0$. Thus, for every $n \geq 1$, $M^+(A_n) \geq 1$ a.s., and consequently $\lim_{n\to\infty} \mathbb{P}\left(M^+(A_n) \geq 1\right) = 1$. In particular the relation $\lim_{n\to\infty} M^+(A_n) = 0$ a.s. does not hold. This shows that for a.e. $\omega \in \Omega$, the function $M^+(\omega, .)$ is not continuous at zero, and thus is not countably additive (cf. [20], Proposition 1.3.3 p.9), leading to a contradiction.

As a consequence, J^W does not admit any jump measure. □

4.7.4 Empirical confidence intervals

This Appendix describes how to obtain an empirical confidence interval for $v(0, x_0, i_0)$. Here we adapt arguments from [28] to our optimal switching problem.

We assume that the parameters T (time localisation) and h (discretization) are chosen such that the error between v and $\bar{v}_\Pi$ is negligible (the space localization being redundant in practice), and focus on the error between $\bar{v}_\Pi$ and $\hat{v}_\Pi$.

First, from equation (4.3.8), the dynamic programming principle for the process $\bar{v}_\Pi\left(t_n, \bar{X}_{t_n}, i\right)$ reads:

$$\begin{aligned}
&\bar{v}_\Pi\left(T, \bar{X}_T, i\right) = g\left(T, \bar{X}_T, i\right)\\
&\bar{v}_\Pi\left(t_n, \bar{X}_{t_n}, i\right) = \sup_{j\in\mathbb{I}_q^n}\left\{hf\left(t_n, \bar{X}_{t_n}, j\right) - k(t_n, i, j) + \mathbb{E}\left[\bar{v}_\Pi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right)|\mathcal{F}_{t_n}\right]\right\}, n = N-1, \ldots, 0
\end{aligned} \tag{4.7.6}$$

where $\mathbb{I}_q^n$ is the set of $\mathcal{F}_{t_n}$-measurable random variables taking values in $\mathbb{I}_q$. Suppose that the approximated conditional expectation $\hat{\mathbb{E}}[.|\mathcal{F}_{t_n}]$ is unbiased, i.e. that

$$\mathbb{E}\left[\hat{\mathbb{E}}[.|\mathcal{F}_{t_n}]\right] = \mathbb{E}[.|\mathcal{F}_{t_n}]$$

Then, using equation (4.3.15) and Jensen's inequality, the following holds:

$$\mathbb{E}\left[\hat{v}_\Pi\left(t_n, \bar{X}_{t_n}, i\right)\right] \geq \sup_{j\in\mathbb{I}_q^n}\left\{hf\left(t_n, \bar{X}_{t_n}, j\right) - k(t_n, i, j) + \mathbb{E}\left[\hat{v}_\Pi\left(t_{n+1}, \bar{X}_{t_{n+1}}, j\right)|\mathcal{F}_{t_n}\right]\right\}, n = N-1, \ldots, 0 \tag{4.7.7}$$

Combining equations (4.7.6) and (4.7.7), an induction argument yields:

$$\mathbb{E}\left[\hat{v}_\Pi\left(t_n, \bar{X}_{t_n}, i\right)\right] \geq \mathbb{E}\left[\bar{v}_\Pi\left(t_n, \bar{X}_{t_n}, i\right)\right]$$

In particular, $\mathbb{E}[\hat{v}_\Pi(0, x_0, i)] \geq \bar{v}_\Pi(0, x_0, i)$. This reasoning means that $\hat{v}_\Pi(0, x_0, i)$ can be used approximatively as an asymptotic upper bound for $\bar{v}_\Pi(0, x_0, i)$.

For the lower bound, simply use the estimated optimal control $\hat{\alpha}$, which is a side-product of the computation of $\hat{v}_\Pi$, and compute equation (4.3.6) by replacing the supremum over every control α by this specific $\hat{\alpha}$. By definition of the supremum, this yields a lower bound for $\bar{v}_\Pi(0, x_0, i)$.

4.7.5 Graphical representation of random processes

The purpose of this Appendix is to provide a simple one-dimensional graphical representation of the distribution of a real-valued random variable. Such a representation can then be used iteratively in order to display conveniently the time evolution of a random process.

Consider a real-valued random variable X. Suppose that its distribution is known, or that an estimate is available. For simplicity, we suppose that X is a continuous random variable. Let f_X denote its probability distribution function. Then, proceed as follows:

- Choose a probability $p \in]0, 1[$, and compute the smallest set A_p such that $\mathbb{P}(A_p) = p$. To do so, it suffices to find the highest $y > 0$ such that the set $A := \{x \in \mathbb{R}; f_X(x) \geq y\}$ suits the condition. Call y_p the level such that $A_p = \{x \in \mathbb{R}; f_X(x) \geq y_p\}$[1]. We call A_p the *$p-$minimal isoprobability (Borel) set*. Figure 4.7.1a illustrates the area (A_p, y_p) for $p = 99\%$ when X is a mixture of two gaussian random variables ($(m_1, \sigma_1, p_1) = (-5, 2, 50\%)$ and $(m_2, \sigma_2, p_2) = (5, 3, 50\%)$).

[1] To be more precise, this construction works if the function $y \mapsto \mathbb{P}(A)$ is continuous. If the density has flat parts (eg. a uniform distribution), then this is not true anymore, and many sets A may satisfy $\mathbb{P}(A_p) = p$. To keep the uniqueness of A_p, one can simply choose to "consume" those flat parts by starting from their middle.

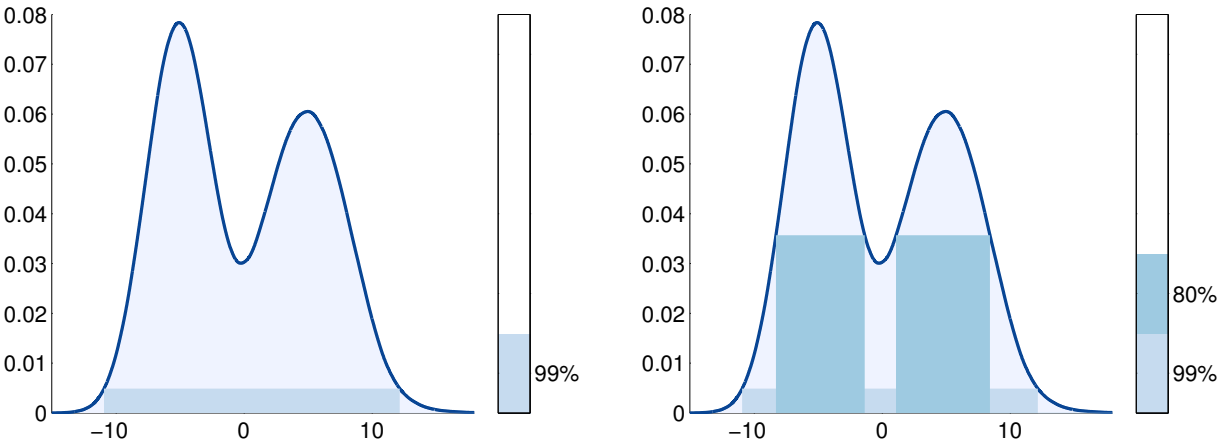

(a) Step 1/5: $p = 99\%$,$y_p = 0.005$

(b) Step 2/5: $p = 80\%$, $y_p = 0.036$

Figure 4.7.1: Density slicing

- Compute the area (A_p, y_p) for several other values of p. On the same example, we illustrate additional areas computed for $p = 80\%$ (Figure 4.7.1b), $p = 60\%$ (Figure 4.7.2a), $p = 40\%$ (Figure 4.7.2b) and finally $p = 20\%$ (Figure 4.7.3a). At each step, we change the color of the current bar according to the value of p.[2] Remark that, unlike a simple interquantile range, a minimal isoprobability set can be a union of disjoints intervals.

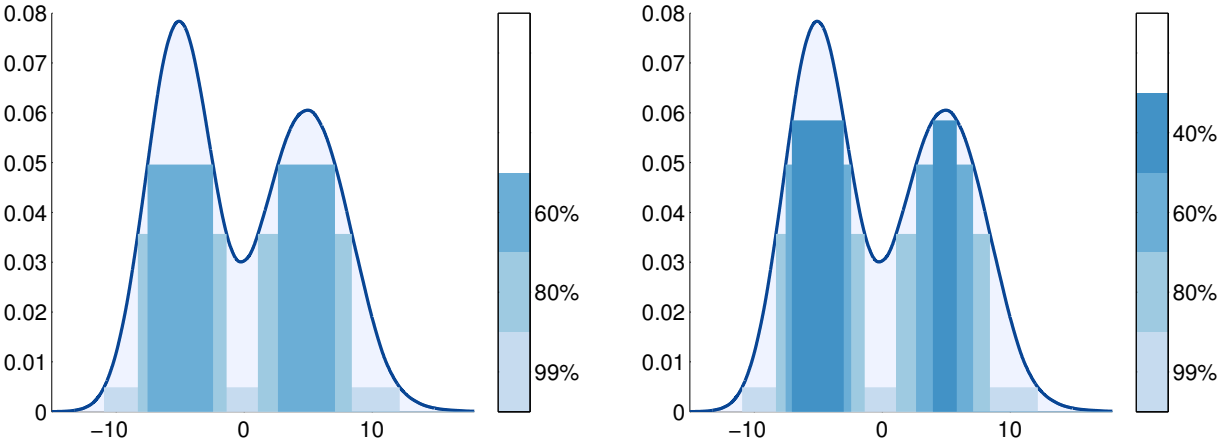

(a) Step 3/5: $p = 60\%$, $y_p = 0.050$

(b) Step 4/5: $p = 40\%$, $y_p = 0.059$

Figure 4.7.2: Density slicing

[2] On these examples, we used the sequential colormaps from the MATLAB package cbrewer (http://www.mathworks.com/matlabcentral/fileexchange/34087-cbrewer-colorbrewer-schemes-for-matlab)

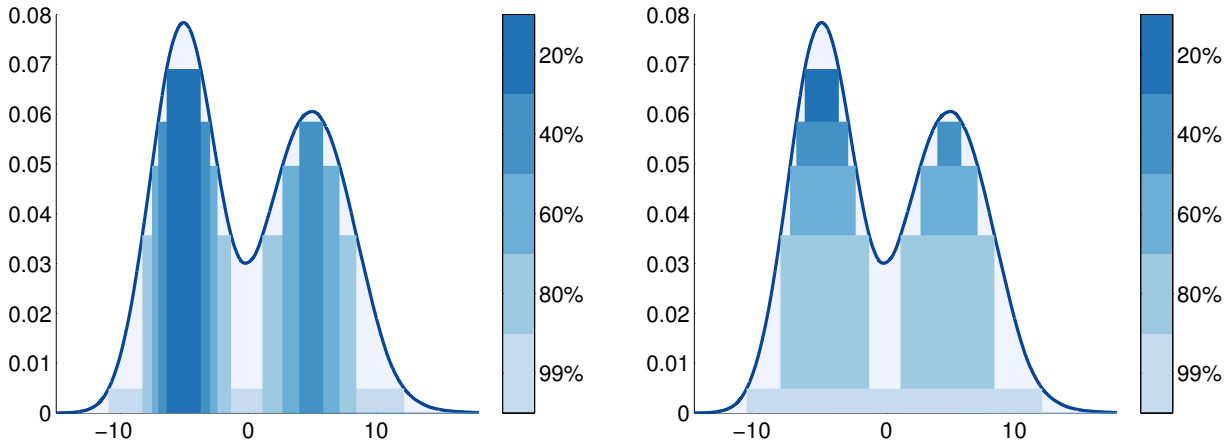

(a) Skyline

(b) Wedding cake

Figure 4.7.3: Inclusive histograms

Figure 4.7.3a does not constitute a histogram per se, as the classes A_p are imbricated (for any $0 \leq p_1 \leq p_2 \leq 1$, $A_{p_2} \subset A_{p_1}$). We call this construction an *inclusive histogram*. There are two variants of it. Proceeding from the lowest p to the highest p (ie. adding taller but thinner bars in front of the previous ones) generates the *"skyline"* version (as in Figure 4.7.3a), while the other way around generates the *"wedding cake"* version (as in Figure 4.7.3b).

However, the real interest of these inclusive histograms is that they allow for an intuitive unidimensional summary of the corresponding distribution. Indeed a look from above at the slices from Figure 4.7.3 leads to the barcode-like summary of the distribution of X displayed on Figure 4.7.4a:

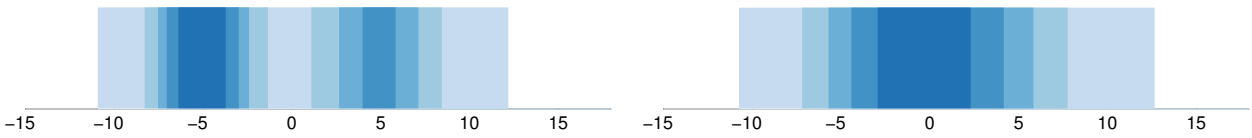

(a) Minimal isoprobability sets

(b) Interquantile ranges

Figure 4.7.4: Barcodes

For comparison, Figure 4.7.4b displays the simple interquantile ranges of the same random variables (for the same levels of probability, ie. the (centered) intervals $[0.5\%, 99.5\%]$, $[10\%, 90\%]$, $[20\%, 80\%]$,$[30\%, 70\%]$ and $[40\%, 60\%]$).

Figure 4.7.4a provides a summary of the distribution of X in one glance: two modes centered around -5 and 5, the second one being more scattered. On the contrary, these informations are not readily available from the interquantile ranges displayed on Figure 4.7.4b.

Basically, the reason for this is that the minimal isoprobability sets are based on the level curves of the probability distribution function, while the interquantile ranges are based on the level curves of the cumulative distribution function. This difference is minor for unimodal distributions (for unimodal and symmetric distributions, like the gaussian distribution, both constructions are identical), but is very significant for multimodal distributions, are shown by Figure 4.7.4.

To sum up, our construction provides, like interquantile ranges, a function $p \in [0,1] \to A_p \in \mathcal{B}(\mathbb{R})$, but this function, unlike interquantile ranges, can systematically account for the shape of the distribution at hand.

Finally, one can use this construction to produce a graphical representation of a given random process. Indeed, computing such barcodes, based on the minimal isoprobability sets, on a given time grid provides an easy-to-grasp graphical representation of the time evolution of the distribution of the random process. Figure 4.5.3 provides two examples of such graphics.

5 A numerical algorithm for fully nonlinear HJB equations via BSDEs with nonpositive jumps

We propose a new probabilistic numerical scheme for fully nonlinear equation of Hamilton-Jacobi-Bellman (HJB) type associated to stochastic control problem, which is based on the Feynman-Kac representation in [72] by means of control randomization and backward stochastic differential equation with nonpositive jumps. We study a discrete time approximation for the minimal solution to this class of BSDE when the time step goes to zero, which provides both an approximation for the value function and for an optimal control in feedback form. We obtained a convergence rate without any ellipticity condition on the controlled diffusion coefficient. Then, we provide an explicit implementable scheme based on Monte-Carlo simulations and empirical regressions, with a partial error analysis, and numerical experiments, notably on the problem of superreplication of option with uncertain volatilities and/or correlations.

5.1 Introduction

Let us consider the fully nonlinear generalized Hamilton-Jacobi-Bellman (HJB) equation:

$$\left\{ \begin{array}{rcll} \frac{\partial v}{\partial t} + \sup_{a\in A} \left[b(x,a).D_x v + \frac{1}{2}\mathrm{tr}(\sigma\sigma^\intercal(x,a)D_x^2 v) + f(x,a,v,\sigma^\intercal(x,a)D_x v)\right] & = & 0, & \text{on } [0,T)\times\mathbb{R}^d, \\ v(T,x) & = & g, & \text{on } \mathbb{R}^d. \end{array} \right. \tag{5.1.1}$$

In the particular case where $f(x,a)$ does not depend on v and $D_x v$, this partial differential equation (PDE) is the dynamic programming equation for the stochastic control problem:

$$v(t,x) \;=\; \sup_{\alpha} \mathbb{E}\Big[\int_t^T f(X_s^\alpha,\alpha_s)ds + g(X_T^\alpha)\Big| X_t^\alpha = x\Big], \tag{5.1.2}$$

with controlled diffusion in $\mathbb{R}^d$:

$$dX_t^\alpha \;=\; b(X_t^\alpha,\alpha_t)dt + \sigma(X_t^\alpha,\alpha_t)dW_t,$$

and where α is an adapted control process valued in a compact space A of $\mathbb{R}^q$. Numerical methods for parabolic partial differential equations (PDEs) are largely developed in the literature, but remain a big challenge for fully nonlinear PDEs, like the HJB equation (5.1.1), especially in high dimensional cases. We refer to the recent paper [49] for a review of some deterministic and probabilistic approaches.

In this chapter, we propose a new probabilistic numerical scheme for HJB equation, relying on the following Feynman-Kac formula for HJB equation obtained by randomization of the control process α. We consider the minimal solution (Y,Z,U,K) to the backward stochastic differential equation (BSDE) with nonpositive jumps:

$$\left\{ \begin{array}{rcl} Y_t & = & g(X_T) + \int_t^T f(X_s,I_s,Y_s,Z_s)ds + K_T - K_t \\ & & - \int_t^T Z_s dW_s - \int_t^T \int_A U_s(a)\tilde{\mu}(ds,da), \qquad 0\le t\le T, \\ U_t(a) & \le & 0, \end{array} \right. \tag{5.1.3}$$

with a forward Markov regime-switching diffusion process (X, I) valued in $\mathbb{R}^d \times A$ given by:

$$\begin{aligned} X_t &= X_0 + \int_0^t b(X_s, I_s)ds + \int_0^t \sigma(X_s, I_s)dW_s \\ I_t &= I_0 + \int_{(0,t]} \int_A (a - I_{s^-})\mu(ds, da). \end{aligned}$$

Here W is a standard Brownian motion, $\mu(dt, da)$ is a Poisson random measure on $[0, \infty) \times A$ with finite intensity measure $\lambda(da)$ of full topological support on A, and compensated measure $\tilde{\mu}(dt, da) = \mu(dt, da) - \lambda(da)dt$. Assumptions on the coefficients b, σ, f, g will be detailed in the next section, but we emphasize the important point that no degeneracy condition on the controlled diffusion coefficient σ is imposed. It is proved in [72] that the minimal solution to this class of BSDE is related to the HJB equation (5.1.1) through the relation $Y_t = v(t, X_t)$.

One purpose of this chapter is to provide and analyze a discrete-time approximation scheme for the minimal solution to (5.1.3), and thus an approximation scheme for the HJB equation (5.1.1). In the non-constrained jump case, approximations schemes for BSDE have been studied in the papers [82], [25], which extended works in [27], [109] for BSDEs in a Brownian framework. The issue is now to deal with the nonpositive jump constraint in (5.1.3), and we propose a discrete time approximation scheme of the form:

$$\left\{ \begin{array}{rcl} \bar{Y}_T^\pi = \bar{\mathcal{Y}}_T^\pi &=& g(\bar{X}_T^\pi) \\ \bar{\mathcal{Z}}_{t_k}^\pi &=& \mathbb{E}\Big[\bar{Y}_{t_{k+1}}^\pi \frac{W_{t_{k+1}} - W_{t_k}}{t_{k+1} - t_k} | \mathcal{F}_{t_k}\Big] \\ \bar{\mathcal{Y}}_{t_k}^\pi &=& \mathbb{E}\Big[\bar{Y}_{t_{k+1}}^\pi | \mathcal{F}_{t_k}\Big] + (t_{k+1} - t_k)\, f(\bar{X}_{t_k}^\pi, I_{t_k}, \bar{\mathcal{Y}}_{t_k}^\pi, \bar{\mathcal{Z}}_{t_k}^\pi) \\ \bar{Y}_{t_k}^\pi &=& \operatorname*{ess\,sup}_{a \in A} \mathbb{E}\Big[\bar{\mathcal{Y}}_{t_k}^\pi | \mathcal{F}_{t_k}, I_{t_k} = a\Big], \quad k = 0, \ldots, n-1, \end{array} \right. \tag{5.1.4}$$

where $\pi = \{t_0 = 0 < \ldots < t_k < \ldots < t_n = T\}$ is a partition of the time interval $[0, T]$, with modulus $|\pi|$, and $\bar{X}^\pi$ is the Euler scheme of X (notice that I is perfectly simulatable once we know how to simulate the distribution $\lambda(da)/\int_A \lambda(da)$ of the jump marks). The interpretation of this scheme is the following. The first three lines in (5.1.4) correspond to the standard scheme $(\bar{\mathcal{Y}}^\pi, \bar{\mathcal{Z}}^\pi)$ for a discretization of a BSDE with jumps (see [25]), where we omit here the computation of the jump component. The last line in (5.1.4) for computing the approximation $\bar{Y}^\pi$ of the minimal solution Y corresponds precisely to the minimality condition for the nonpositive jump constraint and should be understood as follows. By the Markov property of the forward process (X, I), the solution $(\mathcal{Y}, \mathcal{Z}, \mathcal{U})$ to the BSDE with jumps (without constraint) is in the form $\mathcal{Y}_t = \vartheta(t, X_t, I_t)$ for some deterministic function ϑ. Assuming that ϑ is a continuous function, the jump component of the BSDE, which is induced by a jump of the forward component I, is equal to $\mathcal{U}_t(a) = \vartheta(t, X_t, a) - \vartheta(t, X_t, I_{t^-})$. Therefore, the nonpositive jump constraint means that: $\vartheta(t, X_t, I_{t^-}) \geq \operatorname*{ess\,sup}_{a \in A} \vartheta(t, X_t, a)$. The minimality condition is thus written as:

$$Y_t = v(t, X_t) = \operatorname*{ess\,sup}_{a \in A} \vartheta(t, X_t, a) = \operatorname*{ess\,sup}_{a \in A} \mathbb{E}[\mathcal{Y}_t | X_t, I_t = a],$$

whose discrete time version is the last line in scheme (5.1.4).

In this work, we mainly consider the case where $f(x, a, y)$ does not depend on z, and our aim is to analyze the discrete time approximation error on Y, where we split the error between the positive and negative parts:

$$\mathrm{Err}_+^\pi(Y) := \Big(\max_{k \leq n-1} \mathbb{E}\Big[(Y_{t_k} - \bar{Y}_{t_k}^\pi)_+^2\Big]\Big)^{\frac{1}{2}}, \qquad \mathrm{Err}_-^\pi(Y) := \Big(\max_{k \leq n-1} \mathbb{E}\Big[(Y_{t_k} - \bar{Y}_{t_k}^\pi)_-^2\Big]\Big)^{\frac{1}{2}}.$$

We do not study directly the error on Z, and instead focus on the approximation of an optimal control for the HJB equation, which is more relevant in practice. It appears that the maximization step in the scheme (5.1.4) provides a control in feedback form $\{\hat{a}(t_k, \bar{X}^\pi_{t_k}),\, k \leq n-1\}$, which approximates the optimal control with an estimated error bound. The analysis of the error on Y proceeds as follows. We first introduce the solution $(Y^\pi, \mathcal{Y}^\pi, \mathcal{Z}^\pi, \mathcal{U}^\pi)$ of a discretely jump-constrained BSDE. This corresponds formally to BSDEs for which the nonpositive jump constraint operates only a finite set of times, and should be viewed as the analog of discretely reflected BSDEs defined in [7] and [24] in the context of the approximation for reflected BSDEs. By combining BSDE methods and PDE approach with comparison principles, and further with the shaking coefficients method of Krylov [78] and Barles, Jacobsen [10], we prove the monotone convergence of this discretely jump-constrained BSDE towards the minimal solution to the BSDE with nonpositive jump constraint, and obtained a convergence rate without any ellipticity condition on the diffusion coefficient σ. We next focus on the approximation error between the discrete time scheme in (5.1.4) and the discretely jump-constrained BSDE. The standard argument for studying rate of convergence of such error consists in getting an estimate of the error at time t_k: $\mathbb{E}[|Y^\pi_{t_k} - \bar{Y}^\pi_{t_k}|^2]$ in function of the same estimate at time t_{k+1}, and then conclude by induction together with classical estimates for the forward Euler scheme. However, due to the supremum in the conditional expectation in the scheme (5.1.4) for passing from $\bar{\mathcal{Y}}^\pi$ to $\bar{Y}^\pi$, which is a nonlinear operation violating the law of iterated conditional expectations, such argument does not work anymore. Instead, we consider the auxiliary error control at time t_k:

$$\mathcal{E}^\pi_k(\mathcal{Y}) \;:=\; \mathbb{E}\Big[\operatorname*{ess\,sup}_{a\in A} \mathbb{E}_{t_1,a}\big[\,\ldots \operatorname*{ess\,sup}_{a\in A} \mathbb{E}_{t_k,a}\big[|\mathcal{Y}^\pi_{t_k} - \bar{\mathcal{Y}}^\pi_{t_k}|^2\big]\ldots\big]\Big],$$

where $\mathbb{E}_{t_k,a}[.]$ denotes the conditional expectation $\mathbb{E}[.|\mathcal{F}_{t_k}, I_{t_k} = a]$, and we are able to express $\mathcal{E}^\pi_k(\mathcal{Y})$ in function of $\mathcal{E}^\pi_{k+1}(\mathcal{Y})$. We define similarly an error control $\mathcal{E}^\pi_k(X)$ for the forward Euler scheme, and prove that it converges to zero with a rate $|\pi|$. Proceeding by induction, we then obtain a rate of convergence $|\pi|$ for $\mathcal{E}^\pi_k(\mathcal{Y})$, and consequently for $\mathbb{E}[|Y^\pi_{t_k} - \bar{Y}^\pi_{t_k}|^2]$. This leads finally to a rate $|\pi|^{\frac{1}{2}}$ for $\mathrm{Err}^\pi_-(Y)$, $|\pi|^{\frac{1}{10}}$ for $\mathrm{Err}^\pi_+(Y)$, and so $|\pi|^{\frac{1}{10}}$ for the global error $\mathrm{Err}^\pi(Y) = \mathrm{Err}^\pi_+(Y) + \mathrm{Err}^\pi_-(Y)$. In fact, as noticed in Remark **??**, we believe that one can obtain a better rate at least of the order $|\pi|^{\frac{1}{6}}$. Anyway, our result improves the convergence rate of the mixed Monte-Carlo finite difference scheme proposed in [49], where the authors obtained a rate $|\pi|^{\frac{1}{4}}$ on one side and $|\pi|^{\frac{1}{10}}$ on the other side under a nondegeneracy condition.

Once this time dicretization is performed, the above discrete time scheme is still not yet directly implemented in practice, as it requires the estimation and computation of the conditional expectations together with the supremum. To do so, simulation-regression methods on basis functions defined on $\mathbb{R}^d \times A$ appear to be very efficient, and provide approximate optimal controls in feedback forms via the maximization operation in the last step of the scheme (5.1.4). We provide a partial analysis of the impact of this regression approximation, and the remaining obstacles towards a full analysis are highlighted. Notice that since it relies on the simulation of the forward process (X, I), our scheme does not suffer the curse of dimensionality encountered in finite difference scheme or controlled Markov chains methods (see [79], [21]), and takes advantage of the high-dimensional properties of Monte-Carlo methods.

Finally, we illustrate our scheme on several numerical tests, notably on the problem of superreplication of options under uncertain volatility and (for multi-dimensional claims) correlation, which is one of the major applications that takes advantage of the possibilities of our scheme. To our knowledge, the only other Monte Carlo scheme for HJB equations that can handle continuous controls as well as controlled volatility is described in [62], where they make use of another generalization of BSDEs, namely second-order BSDEs. Therefore we compare the performance of our scheme to the results provided in their paper.

To sum up, here is the outline of the chapter. The analysis of the time-discretization is performed in Section 5.2, and the analysis of the approximation of conditional expectations in Section 5.3. The numerical tests and comparisons are performed in Section 5.4. Finally, Section 5.5 concludes the chapter.

5.2 Time discretization

This long section is devoted to the analysis of the convergence of the solution of the discretized scheme (5.1.4) towards the minimal solution of the constrained BSDE (5.1.3). In Subsection 5.2.1, we state some useful auxiliary error estimate for the Euler scheme of the regime switching forward process. We introduce in Subsection 5.2.2 discretely jump-constrained BSDE and relate it to a system of integro-partial differential equations. Subsection 5.2.3 is devoted to the convergence of discretely jump-constrained BSDE to the minimal solution of BSDE with nonpositive jumps. We provide in Subsection 5.2.4 the approximation error for our discrete time scheme, and as a byproduct an estimate for the approximate optimal control in the case of classical HJB equation associated to stochastic control problem.

5.2.1 The forward regime switching process

Let $(\Omega, \mathcal{F}, \mathbb{P})$ be a probability space supporting d-dimensional Brownian motion W, and a Poisson random measure $\mu(dt, da)$ with intensity measure $\lambda(da)dt$ on $[0, \infty) \times A$, where A is a compact set of $\mathbb{R}^q$, endowed with its Borel tribe $\mathcal{B}(A)$, and λ is a finite measure on $(A, \mathcal{B}(A))$ with full topological support. We denote by $\mathbb{F} = (\mathcal{F}_t)_{t\geq 0}$ the completion of the natural filtration generated by (W, μ), and by $\mathcal{P}$ the σ-algebra of $\mathbb{F}$-predictable subsets of $\Omega \times \mathbb{R}_+$.

We fix a finite time horizon $T > 0$, and consider the solution (X, I) on $[0, T]$ of the regime-switching diffusion model:

$$\left\{ \begin{array}{rcl} X_t &=& X_0 + \int_0^t b(X_s, I_s)ds + \int_0^t \sigma(X_s, I_s)dW_s \\ I_t &=& I_0 + \int_{(0,t]} \int_A (a - I_{s^-})\mu(ds, da), \end{array} \right. \tag{5.2.1}$$

where $(X_0, I_0) \in \mathbb{R}^d \times A$, $b : \mathbb{R}^d \times A \to \mathbb{R}^d$ and $\sigma : \mathbb{R}^d \times A \to \mathbb{R}^{d\times d}$, are measurable functions, satisfying the Lipschitz condition:

(H1) There exists a constant L_1 such that

$$|b(x, a) - b(x', a')| + |\sigma(x, a) - \sigma(x', a')| \quad \leq \quad L_1(|x - x'| + |a - a'|),$$

for all $x, x' \in \mathbb{R}^d$ and $a, a' \in A$. The assumption **(H1)** stands in force throughout the chapter, and in this section, we shall denote by C_1 a generic positive constant which depends only on L_1, T, (X_0, I_0) and $\lambda(A) < \infty$, and may vary from lines to lines. Under **(H1)**, we have the existence and uniqueness of a solution to (5.2.1), and in the sequel, we shall denote by $(X^{t,x,a}, I^{t,a})$ the solution to (5.2.1) starting from (x, a) at time t.

Remark 5.2.1. We do not make any ellipticity assumption on σ. In particular, some lines and columns of σ may be equal to zero, and so there is no loss of generality by considering that the dimension d of X and W are equal.

We first study the discrete-time approximation of the forward process. Denoting by $(T_n, \iota_n)_n$ the jump times and marks associated to μ, we observe that I is explicitly written as:

$$I_t \;=\; I_0 \mathbb{1}_{[0,T_1)}(t) + \sum_{n\geq 1} \iota_n \mathbb{1}_{[T_n,T_{n+1})}(t), \quad 0 \leq t \leq T,$$

where the jump times $(T_n)_n$ evolve according to a Poisson distribution of parameter $\lambda := \int_A \lambda(da) < \infty$, and the i.i.d. marks $(\iota_n)_n$ follow a probability distribution $\bar{\lambda}(da) := \lambda(da)/\lambda$. Assuming that one can simulate the probability distribution $\bar{\lambda}$, we then see that the pure jump process I is perfectly simulated. Given a partition $\pi = \{t_0 = 0 < \ldots < t_k < \ldots t_n = T\}$ of $[0,T]$, we shall use the natural Euler scheme $\bar{X}^\pi$ for X, defined by:

$$\begin{aligned} \bar{X}^\pi_0 &= X_0 \\ \bar{X}^\pi_{t_{k+1}} &= \bar{X}^\pi_{t_k} + b(\bar{X}^\pi_{t_k}, I_{t_k})(t_{k+1} - t_k) + \sigma(\bar{X}^\pi_{t_k}, I_{t_k})(W_{t_{k+1}} - W_{t_k}), \end{aligned}$$

for $k = 0, \ldots, n-1$. We denote as usual by $|\pi| = \max_{k\leq n-1}(t_{k+1} - t_k)$ the modulus of π, and assume that $n|\pi|$ is bounded by a constant independent of n, which holds for instance when the grid is regular, i.e. $(t_{k+1} - t_k) = |\pi|$ for all $k \leq n-1$. We also define the continuous-time version of $\bar{X}^\pi$ by setting:

$$\bar{X}^\pi_t \;=\; \bar{X}^\pi_{t_k} + b(\bar{X}^\pi_{t_k}, I_{t_k})(t - t_k) + \sigma(\bar{X}^\pi_{t_k}, I_{t_k})(W_t - W_{t_k}), \;\; t \in [t_k, t_{k+1}], \; k < n.$$

By standard arguments, see e.g. [73], one can obtain under **(H1)** the L^2-error estimate for the above Euler scheme:

$$\mathbb{E}\Big[\sup_{t\in[t_k,t_{k+1}]} |X_t - \bar{X}^\pi_{t_k}|^2\Big] \;\leq\; C_1|\pi|, \quad k < n.$$

For our purpose, we shall need a stronger result, and introduce the following error control for the Euler scheme:

$$\mathcal{E}^\pi_k(X) \;:=\; \mathbb{E}\Big[\operatorname*{ess\,sup}_{a\in A} \mathbb{E}_{t_1,a}\big[\ldots \operatorname*{ess\,sup}_{a\in A} \mathbb{E}_{t_k,a}\big[\sup_{t\in[t_k,t_{k+1}]} |X_t - \bar{X}^\pi_{t_k}|^2\big]\ldots\big]\Big], \tag{5.2.2}$$

where $\mathbb{E}_{t_k,a}[.]$ denotes the conditional expectation $\mathbb{E}[.|\mathcal{F}_{t_k}, I_{t_k} = a]$. We also denote by $\mathbb{E}_{t_k}[.]$ the conditional expectation $\mathbb{E}[.|\mathcal{F}_{t_k}]$. Since I_{t_k} is $\mathcal{F}_{t_k}$-measurable, and by the law of iterated conditional expectations, we notice that

$$\mathbb{E}\Big[\sup_{t\in[t_k,t_{k+1}]} |X_t - \bar{X}^\pi_{t_k}|^2\Big] \;\leq\; \mathcal{E}^\pi_k(X), \quad k < n.$$

Lemma 5.2.1. *We have*

$$\max_{k<n} \mathcal{E}^\pi_k(X) \;\leq\; C_1|\pi|.$$

Proof. From the definition of the Euler scheme, and under the growth linear condition in **(H1)**, we easily see that

$$\mathbb{E}_{t_k}\Big[|\bar{X}^\pi_{t_{k+1}}|^2\Big] \;\leq\; C_1(1 + |\bar{X}^\pi_{t_k}|^2), \;\; k < n. \tag{5.2.3}$$

From the definition of the continuous-time Euler scheme, and by Burkholder-Davis-Gundy inequality, it is also clear that

$$\mathbb{E}_{t_k}\Big[\sup_{t\in[t_k,t_{k+1}]} |\bar{X}^\pi_t - \bar{X}^\pi_{t_k}|^2\Big] \;\leq\; C_1(1 + |\bar{X}^\pi_{t_k}|^2)|\pi|, \;\; k < n. \tag{5.2.4}$$

We also have the standard estimate for the pure jump process I (recall that A is assumed to be compact and $\lambda(A) < \infty$):

$$\mathbb{E}_{t_k}\Big[\sup_{t\in[t_k,t_{k+1}]}|I_s - I_{t_k}|^2\Big] \;\leq\; C_1|\pi|. \tag{5.2.5}$$

Let us denote by $\Delta X_t = X_t - \bar{X}^\pi_t$, and apply Itô's formula to $|\Delta X_t|^2$ so that for all $t \in [t_k, t_{k+1}]$:

$$\begin{aligned}
|\Delta X_t|^2 &= |\Delta X_{t_k}|^2 + \int_{t_k}^t 2(b(X_s,I_s) - b(\bar{X}^\pi_{t_k},I_{t_k})).\Delta X_s + |\sigma(X_s,I_s) - \sigma(\bar{X}^\pi_{t_k},I_{t_k})|^2 ds \\
&\quad + 2\int_{t_k}^t (\Delta X_s)'(\sigma(X_s,I_s) - \sigma(\bar{X}^\pi_{t_k},I_{t_k}))dW_s \\
&\leq |\Delta X_{t_k}|^2 + C_1\int_{t_k}^t |\Delta X_s|^2 + |\bar{X}^\pi_s - \bar{X}^\pi_{t_k}|^2 + |I_s - I_{t_k}|^2 ds \\
&\quad + 2\int_{t_k}^t (\Delta X_s)'(\sigma(X_s,I_s) - \sigma(\bar{X}^\pi_{t_k},I_{t_k}))dW_s,
\end{aligned}$$

from the Lipschitz condition on b, σ in **(H1)**. By taking conditional expectation in the above inequality, we then get:

$$\begin{aligned}
\mathbb{E}_{t_k}\Big[|\Delta X_t|^2\Big] &\leq |\Delta X_{t_k}|^2 + C_1\int_{t_k}^t \mathbb{E}_{t_k}\big[|\Delta X_s|^2 + |\bar{X}^\pi_s - \bar{X}^\pi_{t_k}|^2 + |I_s - I_{t_k}|^2\big]ds \\
&\leq |\Delta X_{t_k}|^2 + C_1(1 + |\bar{X}^\pi_{t_k}|^2)|\pi|^2 + C_1\int_{t_k}^t \mathbb{E}_{t_k}\big[|\Delta X_s|^2\big]ds, \quad t \in [t_k, t_{k+1}],
\end{aligned}$$

by (5.2.4)-(5.2.5). From Gronwall's lemma, we thus deduce that

$$\mathbb{E}_{t_k}\Big[|\Delta X_{t_{k+1}}|^2\Big] \;\leq\; e^{C_1|\pi|}|\Delta X_{t_k}|^2 + C_1(1 + |\bar{X}^\pi_{t_k}|^2)|\pi|^2, \quad k < n. \tag{5.2.6}$$

Since the right hand side of (5.2.6) does not depend on I_{t_k}, this shows that

$$\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{t_k,a}\Big[|\Delta X_{t_{k+1}}|^2\Big] \;\leq\; e^{C_1|\pi|}|\Delta X_{t_k}|^2 + C_1(1 + |\bar{X}^\pi_{t_k}|^2)|\pi|^2.$$

By taking conditional expectation w.r.t. $\mathcal{F}_{t_{k-1}}$ in the above inequality, using again estimate (5.2.6) together with (5.2.3) at step $k-1$, and iterating this backward procedure until the initial time $t_0 = 0$, we obtain:

$$\begin{aligned}
&\mathbb{E}\Big[\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{t_1,a}\big[\ldots\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{t_k,a}\big[|\Delta X_{t_{k+1}}|^2\big]\ldots\big]\Big] \\
&\leq\; e^{C_1 n|\pi|}|\Delta X_0|^2 + C_1(1 + |X_0|^2)|\pi|^2\frac{e^{C_1 n|\pi|} - 1}{e^{C_1|\pi|} - 1} \\
&\leq\; C_1|\pi|,
\end{aligned} \tag{5.2.7}$$

since $\Delta X_0 = 0$ and $n|\pi|$ is bounded.

Moreover, the process X satisfies the standard conditional estimate similarly as for the Euler scheme:

$$\begin{aligned}
\mathbb{E}_{t_k}\Big[|X_{t_{k+1}}|^2\Big] &\leq C_1(1 + |X_{t_k}|^2), \\
\mathbb{E}_{t_k}\Big[\sup_{t\in[t_k,t_{k+1}]}|X_t - X_{t_k}|^2\Big] &\leq C_1(1 + |X_{t_k}|^2)|\pi|, \quad k < n,
\end{aligned}$$

from which we deduce by backward induction on the conditional expectations:

$$\mathbb{E}\Big[\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{t_1,a}\big[\ldots\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{t_k,a}\big[\sup_{t\in[t_k,t_{k+1}]}|X_t-X_{t_k}|^2\big]\ldots\big]\Big] \;\leq\; C_1|\pi|. \tag{5.2.8}$$

Finally, by writing that $\sup_{t\in[t_k,t_{k+1}]}|X_t-\bar{X}^\pi_{t_k}|^2 \leq 2\sup_{t\in[t_k,t_{k+1}]}|X_t-X_{t_k}|^2+2\Delta X_{t_k}$, taking successive condition expectations w.r.t to $\mathcal{F}_{t_\ell}$ and essential supremum over $I_{t_\ell}=a$, for ℓ going recursively from k to 0, we get:

$$\begin{array}{rcl}\mathbb{E}_{t_k}\Big[\sup_{t\in[t_k,t_{k+1}]}|X_t-\bar{X}^\pi_{t_k}|^2\Big] &\leq& 2\mathbb{E}\Big[\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{t_1,a}\big[\ldots\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{t_k,a}\big[\sup_{t\in[t_k,t_{k+1}]}|X_t-X_{t_k}|^2\big]\ldots\big]\Big]\\ &&+\,2\mathbb{E}\Big[\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{t_1,a}\big[\ldots\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{t_{k-1},a}\big[|\Delta X_{t_k}|^2\big]\ldots\big]\Big]\\ &\leq& C_1|\pi|,\end{array}$$

by (5.2.7)-(5.2.8), which ends the proof. □

5.2.2 Discretely jump-constrained BSDE

Given the forward regime switching process (X,I) defined in the previous section, we consider the minimal quadruple solution (Y,Z,U,K) to the BSDE with nonpositive jumps:

$$\left\{\begin{array}{rcl} Y_t &=& g(X_T)+\int_t^T f(X_s,I_s,Y_s,Z_s)ds+K_T-K_t\\ &&-\int_t^T Z_s dW_s-\int_t^T\int_A U_s(a)\tilde{\mu}(ds,da), \quad 0\leq t\leq T.\\ U_t(a) &\leq& 0,\end{array}\right. \tag{5.2.9}$$

By solution to (5.2.9), we mean a quadruple $(Y,Z,U,K)\in\mathcal{S}^2\times L^2(W)\times L^2(\tilde{\mu})\times\mathcal{K}^2$, where $\mathcal{S}^2$ is the space of càd-làg or càg-làd $\mathbb{F}$-progressively measurable processes Y satisfying $\|Y\|^2 := \mathbb{E}[\sup_{t\in[0,T]}|Y_t|^2]<\infty$, $L^2(W)$ is the space of $\mathbb{R}^d$-valued $\mathcal{P}$-measurable processes such that $\|Z\|^2_{L^2(W)} := \mathbb{E}[\int_0^T|Z_t|^2dt]<\infty$, $L^2(\tilde{\mu})$ is the space of real-valued $\mathcal{P}\otimes\mathcal{B}(A)$-measurable processes U such that $\|U\|^2_{L^2(\tilde{\mu})} := \mathbb{E}[\int_0^T\int_A|U_t(a)|^2\lambda(da)\,dt]<\infty$, and $\mathcal{K}^2$ is the subspace of $\mathcal{S}^2$ consisting of nondecreasing predictable processes such that $K_0=0$, $\mathbb{P}$-a.s., and the equation in (5.2.9) holds $\mathbb{P}$-a.s., while the nonpositive jump constraint holds on $\Omega\times[0,T]\times A$ a.e. with respect to the measure $d\mathbb{P}\otimes dt\otimes\lambda(da)$. By minimal solution to the BSDE (5.1.3), we mean a quadruple solution $(Y,Z,U,K)\in\mathcal{S}^2\times L^2(W)\times L^2(\tilde{\mu})\times\mathcal{K}^2$ such that for any other solution (Y',Z',U',K') to the same BSDE, we have $\mathbb{P}$-a.s.: $Y_t\leq Y'_t$, $t\in[0,T]$.

In the rest of this chapter, we shall make the standing Lipschitz assumption on the functions $f:\mathbb{R}^d\times A\times\mathbb{R}\times\mathbb{R}^d\to\mathbb{R}$ and $g:\mathbb{R}^d\to\mathbb{R}$.

(H2) There exists a constant L_2 such that

$$|f(x,a,y,z)-f(x',a',y',z')|+|g(x)-g(x')| \;\leq\; L_2(|x-x'|+|a-a'|+|y-y'|+|z-z'|),$$

for all $x,x'\in\mathbb{R}^d$, $y,y'\in\mathbb{R}$, $z,z'\in\mathbb{R}^d$, $a,a'\in A$. In the sequel, we shall denote by C a generic positive constant which depends only on L_1, L_2, T, (X_0,I_0) and $\lambda(A)<\infty$, and may vary from lines to lines.

Under **(H1)-(H2)**, it is proved in [72] the existence and uniqueness of a minimal solution (Y, Z, U, K) to (5.2.9). Moreover, the minimal solution Y is in the form

$$Y_t = v(t, X_t), \quad 0 \le t \le T, \tag{5.2.10}$$

where $v : [0, T] \times \mathbb{R}^d \to \mathbb{R}$ is a viscosity solution with linear growth to the fully nonlinear HJB type equation:

$$\begin{cases} -\sup_{a \in A} \left[\mathcal{L}^a v + f(x, a, v, \sigma^\intercal(x, a) D_x v) \right] = 0, \text{ on } [0, T) \times \mathbb{R}^d, \\ v(T, x) = g, \text{ on } \mathbb{R}^d, \end{cases} \tag{5.2.11}$$

where

$$\mathcal{L}^a v = \frac{\partial v}{\partial t} + b(x, a).D_x v + \frac{1}{2} \mathrm{tr}(\sigma\sigma^\intercal(x, a) D_x^2 v).$$

We shall make the standing assumption that comparison principle holds for (5.2.11).

(HC) Let $\bar{w}$ (resp. $\underline{w}$) be a lower-semicontinuous (resp. upper-semicontinuous) viscosity supersolution (resp. subsolution) with linear growth condition to (5.2.11). Then, $\bar{w} \ge \underline{w}$.

When f does not depend on y, z, i.e. (5.2.11) is the usual HJB equation for a stochastic control problem, Assumption **(HC)** holds true, see [51] or [92]. In the general case, we refer to [44] for sufficient conditions to comparison principles. Under **(HC)**, the function v in (5.2.10) is the unique viscosity solution to (5.2.11), and is in particular continuous. Actually, we have the standard Hölder and Lipschitz property (see Appendix in [78] or [10]):

$$|v(t, x) - v(t', x')| \le C\big(|t - t'|^{\frac{1}{2}} + |x - x'|\big), \quad (t, t') \in [0, T], x, x' \in \mathbb{R}^d. \tag{5.2.12}$$

This implies that the process Y is continuous, and thus the jump component $U = 0$. In the sequel, we shall focus on the approximation of the remaining components Y and Z of the minimal solution to (5.2.9).

We introduce in this section discretely jump-constrained BSDE. The nonpositive jump constraint operates only at the times of the grid $\pi = \{t_0 = 0 < t_1 < \ldots < t_n = T\}$ of $[0, T]$, and we look for a quadruple $(Y^\pi, \mathcal{Y}^\pi, \mathcal{Z}^\pi, \mathcal{U}^\pi) \in \mathcal{S}^2 \times \mathcal{S}^2 \times L^2(W) \times L^2(\tilde{\mu})$ satisfying:

$$Y_T^\pi = \mathcal{Y}_T^\pi = g(X_T) \tag{5.2.13}$$

and

$$\begin{aligned} \mathcal{Y}_t^\pi &= Y_{t_{k+1}}^\pi + \int_t^{t_{k+1}} f(X_s, I_s, \mathcal{Y}_s^\pi, \mathcal{Z}_s^\pi) ds \\ &\quad - \int_t^{t_{k+1}} \mathcal{Z}_s^\pi dW_s - \int_t^{t_{k+1}} \int_A \mathcal{U}_s^\pi(a) \tilde{\mu}(ds, da) \,, \end{aligned} \tag{5.2.14}$$

$$Y_t^\pi = \mathcal{Y}_t^\pi \mathbb{1}_{(t_k, t_{k+1})}(t) + \operatorname*{ess\,sup}_{a \in A} \mathbb{E}\Big[\mathcal{Y}_t^\pi \big| X_t, I_t = a\Big] \mathbb{1}_{\{t_k\}}(t) \,, \tag{5.2.15}$$

for all $t \in [t_k, t_{k+1})$ and all $0 \le k \le n - 1$.

Notice that at each time t_k of the grid, the condition is not known a priori to be square integrable since it involves a supremum over A, and the well-posedness of the BSDE (5.2.13)-(5.2.14)-(5.2.15) is not a direct and standard issue. We shall use a PDE approach for proving the existence and uniqueness of a solution. Let us consider the system of integro-partial differential equations (IPDEs) for the functions v^π and ϑ^π defined recursively on $[0,T] \times \mathbb{R}^d \times A$ by:

• A terminal condition for v^π and ϑ^π:

$$v^\pi(T,x,a) \;=\; \vartheta^\pi(T,x,a) \;=\; g(x)\,, \quad (x,a) \in \mathbb{R}^d \times A\,, \tag{5.2.16}$$

• A sequence of IPDEs for ϑ^π

$$\left\{ \begin{array}{rl} -\mathcal{L}^a \vartheta^\pi - f\big(x,a,\vartheta^\pi,\sigma^\intercal(x,a) D_x \vartheta^\pi\big) & \\ -\int_A \big(\vartheta^\pi(t,x,a') - \vartheta^\pi(t,x,a)\big)\lambda(da') \;=\; 0, & (t,x,a) \in [t_k,t_{k+1}) \times \mathbb{R}^d \times A, \\ \vartheta^\pi(t_{k+1}^-,x,a) \;=\; \sup_{a' \in A} \vartheta^\pi(t_{k+1},x,a') & (x,a) \in \mathbb{R}^d \times A \end{array} \right. \tag{5.2.17}$$

for $k = 0 \ldots, n-1$,

• the relation between v^π and ϑ^π:

$$v^\pi(t,x,a) \;=\; \vartheta^\pi(t,x,a) 1\!\!1_{(t_k,t_{k+1})}(t) + \sup_{a' \in A} \vartheta^\pi(t,x,a') 1\!\!1_{\{t_k\}}(t)\,, \tag{5.2.18}$$

for all $t \in [t_k, t_{k+1})$ and $k = 0 \ldots, n-1$. The rest of this section is devoted to the proof of existence and uniqueness of a solution to (5.2.16)-(5.2.17)-(5.2.18), together with some uniform Lipschitz properties, and its connection to the discretely jump-constrained BSDE (5.2.13)-(5.2.14)-(5.2.15).

For any L-Lipschitz continuous function φ on $\mathbb{R}^d \times A$, and $k \leq n-1$, we denote:

$$\mathbb{T}_\pi^k[\varphi](t,x,a) \;:=\; w(t,x,a), \quad (t,x,a) \in [t_k,t_{k+1}) \times \mathbb{R}^d \times A, \tag{5.2.19}$$

where w is the unique continuous viscosity solution on $[t_k,t_{k+1}] \times \mathbb{R}^d \times A$ with linear growth condition in x to the integro partial differential equation (IPDE):

$$\left\{ \begin{array}{rl} -\mathcal{L}^a w - f\big(x,a,w,\sigma^\intercal D_x w\big) & \\ -\int_A \big(w(t,x,a') - w(t,x,a)\big)\lambda(da') \;=\; 0, & (t,x,a) \in [t_k,t_{k+1}) \times \mathbb{R}^d \times A, \\ w(t_{k+1}^-,x,a) \;=\; \varphi(x,a), & (x,a) \in \mathbb{R}^d \times A\,, \end{array} \right. \tag{5.2.20}$$

and we extend by continuity $\mathbb{T}_\pi^k[\varphi](t_{k+1},x,a) = \varphi(x,a)$. The existence and uniqueness of such a solution w to the semi linear IPDE (5.2.20), and its nonlinear Feynman-Kac representation in terms of BSDE with jumps, is obtained e.g. from Theorems 3.4 and 3.5 in [9].

Lemma 5.2.2. *There exists a constant C such that for any L-Lipschitz continuous function φ on $\mathbb{R}^d \times A$, and $k \leq n-1$, we have*

$$|\mathbb{T}_\pi^k[\varphi](t,x,a) - \mathbb{T}_\pi^k[\varphi](t,x',a')| \;\leq\; \max(L,1)\sqrt{1+|\pi|}e^{C|\pi|}(|x-x'| + |a-a'|)\,,$$

for all $t \in [t_k,t_{k+1})$, and $(x,a),(x',a') \in \mathbb{R}^d \times A$.

Proof. Fix $t \in [t_k, t_{k+1})$, $k \leq n-1$, $(x,a),(x',a') \in \mathbb{R}^d \times A$, and φ an L-Lipschitz continuous function on $\mathbb{R}^d \times A$. Let $(Y^\varphi, Z^\varphi, U^\varphi)$ and $(Y^{\varphi,'}, Z^{\varphi,'}, U^{\varphi,'})$ be the solutions on $[t, t_{k+1}]$ to the BSDEs

$$\begin{aligned}
Y_s^\varphi &= \varphi(X_{t_{k+1}}^{t,x,a}, I_{t_{k+1}}^{t,a}) + \int_s^{t_{k+1}} f(X_r^{t,x,a}, I_r^{t,a}, Y_r^\varphi, Z_r^\varphi)dr \\
&\quad - \int_s^{t_{k+1}} Z_r^\varphi dW_r - \int_s^{t_{k+1}} \int_A U_r^\varphi(e)\tilde{\mu}(dr,de), \qquad t \leq s \leq t_{k+1}, \\
Y_s^{\varphi,'} &= \varphi(X_{t_{k+1}}^{t,x',a'}, I_{t_{k+1}}^{t,a'}) + \int_s^{t_{k+1}} f(X_r^{t,x',a'}, I_r^{t,a'}, Y_r^{\varphi,'}, Z_r^{\varphi,'})dr \\
&\quad - \int_s^{t_{k+1}} Z_r^{\varphi,'} dW_r - \int_s^{t_{k+1}} \int_A U_r^{\varphi,'}(e)\tilde{\mu}(dr,de), \qquad t \leq s \leq t_{k+1}
\end{aligned}$$

From Theorems 3.4 and 3.5 in [9], we have the identification:

$$Y_t^\varphi = \mathbb{T}_\pi^k[\varphi](t,x,a) \quad \text{and} \quad Y_t^{\varphi,'} = \mathbb{T}_\pi^k[\varphi](t,x',a') \,. \tag{5.2.21}$$

We now estimate the difference between the processes $Y^\varphi$and $Y^{\varphi,'}$, and set $\delta Y^\varphi = Y^\varphi - Y^{\varphi,'}$, $\delta Z^\varphi = Z^\varphi - Z^{\varphi,'}$, $\delta X = X^{t,x,a} - X^{t,x',a'}$, $\delta I = I^{t,a} - I^{t,a'}$. By Itô's formula, the Lipschitz condition of f and φ, and Young inequality, we have

$$\begin{aligned}
\mathbb{E}\Big[|\delta Y_s^\varphi|^2\Big] + \mathbb{E}\Big[\int_s^{t_{k+1}} |\delta Z_s^\varphi|^2 ds\Big] &\leq L^2\mathbb{E}\Big[|\delta X_T|^2 + |\delta I_T|^2\Big] + C\int_s^{t_{k+1}} \mathbb{E}\Big[|\delta Y_r^\varphi|^2\Big]dr \\
&\quad + \frac{1}{2}\mathbb{E}\Big[\int_s^{t_{k+1}} (|\delta X_r|^2 + |\delta I_r|^2 + |\delta Z_r^\varphi|^2)dr\Big],
\end{aligned}$$

for any $s \in [t, t_{k+1}]$. Now, from classical estimates on jump-diffusion processes we have

$$\sup_{r\in[t,t_{k+1}]} \mathbb{E}\Big[|\delta X_r|^2 + |\delta I_r|^2\Big] \leq e^{C|\pi|}(|x-x'|^2 + |a-a'|^2),$$

and thus:

$$\mathbb{E}\Big[|\delta Y_s^\varphi|^2\Big] \leq (L^2 + |\pi|)e^{C|\pi|}(|x-x'|^2 + |a-a'|^2) + C\int_s^{t_{k+1}} \mathbb{E}\Big[|\delta Y_r^\varphi|^2\Big]dr \,,$$

for all $s \in [t, t_{k+1}]$. By Gronwall's Lemma, this yields

$$\sup_{s\in[t,t_{k+1}]} \mathbb{E}\Big[|\delta Y_s^\varphi|^2\Big] \leq (L^2 + |\pi|)e^{2C|\pi|}(|x-x'|^2 + |a-a'|^2),$$

which proves the required result from the identification (5.2.21):

$$\begin{aligned}
|\mathbb{T}_\pi^k[\varphi](t,x,a) - \mathbb{T}_\pi^k[\varphi](t,x',a')| &\leq \sqrt{L^2 + |\pi|}e^{C|\pi|}(|x-x'| + |a-a'|) \\
&\leq \max(L,1)\sqrt{1+|\pi|}e^{C|\pi|}(|x-x'| + |a-a'|).
\end{aligned}$$

□

Proposition 5.2.1. *There exists a unique viscosity solution ϑ^π with linear growth condition to the IPDE (5.2.16)-(5.2.17), and this solution satisfies:*

$$\begin{aligned} & |\vartheta^\pi(t,x,a) - \vartheta^\pi(t,x',a')| \\ \leq\ & \max(L_2,1)\sqrt{\Big(e^{2C|\pi|}(1+|\pi|)\Big)^{n-k}}(|x-x'|+|a-a'|), \end{aligned} \tag{5.2.22}$$

for all $k = 0,\dots,n-1$, $t \in [t_k, t_{k+1})$, $(x,a),(x',a') \in \mathbb{R}^d \times A$.

Proof. We prove by a backward induction on k that the IPDE (5.2.16)-(5.2.17) admits a unique solution on $[t_k,T]\times\mathbb{R}^d\times A$, which satisfies (5.2.22). • For $k = n-1$, we directly get the existence and uniqueness of ϑ^π on $[t_{n-1},T]\times\mathbb{R}^d\times A$ from Theorems 3.4 and 3.5 in [9], and we have $\vartheta^\pi = \mathbb{T}_\pi^{n-1}[g]$ on $[t_{n-1},T)\times\mathbb{R}^d\times A$. Moreover, we also get by Lemma 5.2.2:

$$|\vartheta^\pi(t,x,a) - \vartheta^\pi(t,x',a')| \ \leq\ \max(L_2,1)\sqrt{e^{2C|\pi|}(1+|\pi|)}(|x-x'|+|a-a'|)$$

for all $t \in [t_{n-1},t_n)$, $(x,a),(x',a') \in \mathbb{R}^d\times A$. • Suppose that the result holds true at step $k+1$ i.e. there exists a unique function ϑ^π on $[t_{k+1},T]\times\mathbb{R}^d\times A$ with linear growth and satisfying (5.2.16)-(5.2.17) and (5.2.22). It remains to prove that ϑ^π is uniquely determined by (5.2.17) on $[t_k,t_{k+1})\times\mathbb{R}^d\times A$ and that it satisfies (5.2.22) on $[t_k,t_{k+1})\times\mathbb{R}^d\times A$. Since ϑ^π satisfies (5.2.22) at time t_{k+1}, we deduce that the function

$$\psi_{k+1}(x) \ :=\ \sup_{a\in A}\vartheta^\pi(t_{k+1},x,a), \quad x\in\mathbb{R}^d,$$

is also Lipschitz continuous, and satisfies by the induction hypothesis:

$$|\psi_{k+1}(x) - \psi_{k+1}(x')| \ \leq\ \max(L_2,1)\sqrt{\Big(e^{2C|\pi|}(1+|\pi|)\Big)^{n-k-1}}|x-x'|, \tag{5.2.23}$$

for all $x,x' \in \mathbb{R}^d$. Under **(H1)** and **(H2)**, we can apply Theorems 3.4 and 3.5 in [9], and we get that ϑ^π is the unique viscosity solution with linear growth to (5.2.17) on $[t_k,t_{k+1})\times\mathbb{R}^d\times A$, with $\vartheta^\pi = \mathbb{T}_\pi^k[\psi_{k+1}]$. Thus it exists and is unique on $[t_k,T]\times\mathbb{R}^d\times A$. From Lemma 5.2.2 and (5.2.23), we then get

$$\begin{aligned} |\vartheta^\pi(t,x,a) - \vartheta^\pi(t,x',a')| &= |\mathbb{T}_\pi^k[\psi_{k+1}](t,x,a) - \mathbb{T}_\pi^k[\psi_{k+1}](t,x',a')| \\ &\leq \max(L_2,1)\sqrt{\Big(e^{2C|\pi|}(1+|\pi|)\Big)^{n-k-1}} \\ &\quad .\sqrt{(1+|\pi|)e^{2C|\pi|}}(|x-x'|+|a-a'|) \\ &\leq \max(L_2,1)\sqrt{\Big(e^{2C|\pi|}(1+|\pi|)\Big)^{n-k}}(|x-x'|+|a-a'|) \end{aligned}$$

for any $t \in [t_k,t_{k+1})$ and $(x,a),(x',a') \in \mathbb{R}^d\times A$, which proves the required induction inequality at step k. □

Remark 5.2.2. The function $a \to \vartheta^\pi(t,x,.)$ is continuous on A, for each (t,x), and so the function v^π is well-defined by (5.2.18). Moreover, the function ϑ^π may be written recursively as:

$$\begin{cases} \vartheta^\pi(T,.,.) &= g & \text{on } \mathbb{R}^d\times A, \\ \vartheta^\pi &= \mathbb{T}_\pi^k[v^\pi(t_{k+1},.)], & \text{on } [t_k,t_{k+1})\times\mathbb{R}^d\times A, \end{cases} \tag{5.2.24}$$

for $k = 0,\dots,n-1$. In particular, ϑ^π is continuous on $(t_k,t_{k+1})\times\mathbb{R}^d\times A$, $k \leq n-1$.

As a consequence of the above proposition, we obtain the uniform Lipschitz property of ϑ^π and v^π, with a Lipschitz constant independent of π.

Corollary 5.2.1. *There exists a constant C (independent of $|\pi|$) such that*

$$|\vartheta^\pi(t,x,a) - \vartheta^\pi(t,x',a')| + |v^\pi(t,x,a) - v^\pi(t,x',a')| \;\leq\; C(|x-x'| + |a-a'|),$$

for all $t \in [0,T]$, $x, x' \in \mathbb{R}^d$, $a, a' \in \mathbb{R}^d$.

Proof. Recalling that $n|\pi|$ is bounded, we see that the sequence appearing in (5.2.22): $\Big((e^{2C|\pi|}(1+|\pi|))^{n-k}\Big)_{0\leq k\leq n-1}$ is bounded uniformly in $|\pi|$ (or n), which shows the required Lipschitz property of ϑ^π. Since A is assumed to be compact, this shows in particular that the function v^π defined by the relation (5.2.18) is well-defined and finite. Moreover, by noting that

$$|\sup_{a\in A} \vartheta^\pi(t,x,a) - \sup_{a\in A} \vartheta^\pi(t,x',a)| \;\leq\; \sup_{a\in A} |\vartheta^\pi(t,x,a) - \vartheta^\pi(t,x',a)|$$

for all $(t,x) \in [0,T] \times \mathbb{R}^d$, we also obtain the required Lipschitz property for v^π.

□

We now turn to the existence of a solution to the discretely jump-constrained BSDE.

Proposition 5.2.2. *The BSDE (5.2.13)-(5.2.14)-(5.2.15) admits a unique solution $(Y^\pi, \mathcal{Y}^\pi, \mathcal{Z}^\pi, \mathcal{U}^\pi)$ in $\mathcal{S}^2 \times \mathcal{S}^2 \times L^2(W) \times L^2(\tilde{\mu})$. Moreover we have*

$$\mathcal{Y}^\pi_t \;=\; \vartheta^\pi(t, X_t, I_t), \quad \text{and} \quad Y^\pi_t \;=\; v^\pi(t, X_t, I_t) \tag{5.2.25}$$

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

Proof. We prove by backward induction on k that $(Y^\pi, \mathcal{Y}^\pi, \mathcal{Z}^\pi, \mathcal{U}^\pi)$ is well defined and satisfies (5.2.25) on $[t_k, T]$. • Suppose that $k = n-1$. From Corollary 2.3 in [9], we know that $(\mathcal{Y}^\pi, \mathcal{Z}^\pi, \mathcal{U}^\pi)$, exists and is unique on $[t_{n-1}, T]$. Moreover, from Theorems 3.4 and 3.5 in [9], we get $\mathcal{Y}^\pi_t = \mathbb{T}^k_\pi[g](t, X_t, I_t) = \vartheta^\pi(t, X_t, I_t)$ on $[t_{n-1}, T]$. By (5.2.15), we then have for all $t \in [t_{n-1}, T)$:

$$\begin{aligned} Y^\pi_t \;&=\; \mathbb{1}_{(t_{n-1},T)}(t)\, \vartheta^\pi(t, X_t, I_t) + \mathbb{1}_{t_{n-1}}(t)\, \operatorname*{ess\,sup}_{a\in A} \vartheta^\pi(t, X_t, a) \\ &=\; \mathbb{1}_{(t_{n-1},T)}(t)\, \vartheta^\pi(t, X_t, I_t) + \mathbb{1}_{t_{n-1}}(t)\, \sup_{a\in A} \vartheta^\pi(t, X_t, a) \;=\; v^\pi(t, X_t, I_t), \end{aligned}$$

since the essential supremum and supremum coincide by continuity of $a \to \vartheta^\pi(t, X_t, a)$ on the compact set A. • Suppose that the result holds true for some $k \leq n-1$. Then, we see that $(\mathcal{Y}^\pi, \mathcal{Z}^\pi, \mathcal{U}^\pi)$ is defined on $[t_{k-1}, t_k)$ as the solution to a BSDE driven by W and $\tilde{\mu}$ with a terminal condition $v^\pi(t_k, X_{t_k})$. Since v^π satisfies a linear growth condition, we know again by Corollary 2.3 in [9] that $(\mathcal{Y}^\pi, \mathcal{Z}^\pi, \mathcal{U}^\pi)$, thus also Y^π, exists and is unique on $[t_{k-1}, t_k)$. Moreover, using again Theorems 3.4 and 3.5 in [9], we get (5.2.25) on $[t_{k-1}, t_k)$. □

We end this section with a conditional regularity result for the discretely jump-constrained BSDE.

Proposition 5.2.3. *There exists some constant C such that*

$$\sup_{t\in[t_k,t_{k+1})} \mathbb{E}_{t_k}[|\mathcal{Y}_t^\pi - \mathcal{Y}_{t_k}^\pi|^2] + \sup_{t\in(t_k,t_{k+1}]} \mathbb{E}_{t_k}[|Y_t^\pi - Y_{t_{k+1}}^\pi|^2] \;\leq\; C(1+|X_{t_k}|^2)|\pi|,$$

for all $k = 0,\dots,n-1$.

Proof. Fix $k \leq n-1$. By Itô's formula, we have for all $t \in [t_k, t_{k+1})$:

$$\begin{aligned}
\mathbb{E}_{t_k}[|\mathcal{Y}_t^\pi - \mathcal{Y}_{t_k}^\pi|^2] &= 2\mathbb{E}_{t_k}\Big[\int_{t_k}^t f(X_s, I_s, \mathcal{Y}_s^\pi, \mathcal{Z}_s^\pi)(\mathcal{Y}_{t_k}^\pi - \mathcal{Y}_s^\pi)ds\Big] \\
&\quad + \mathbb{E}_{t_k}\Big[\int_{t_k}^t |\mathcal{Z}_s^\pi|^2\Big] + \mathbb{E}_{t_k}\Big[\int_{t_k}^t\int_A |\mathcal{U}_s^\pi(a)|^2\lambda(da)ds\Big] \\
&\leq \mathbb{E}_{t_k}\Big[\int_{t_k}^t |\mathcal{Y}_s^\pi - \mathcal{Y}_{t_k}^\pi|^2\Big] + C|\pi|\Big(1 + \mathbb{E}_{t_k}\Big[\sup_{s\in[t_k,t_{k+1}]}|X_s|^2\Big]\Big) \\
&\quad + C|\pi|\mathbb{E}_{t_k}\Big[\sup_{s\in[t_k,t_{k+1}]}\Big(|\mathcal{Y}_s^\pi|^2 + |\mathcal{Z}_s^\pi|^2 + \int_A |\mathcal{U}_s^\pi(a)|^2\lambda(da)\Big)\Big],
\end{aligned}$$

by the linear growth condition on f (recall also that A is compact), and Young inequality. Now, by standard estimate for X under growth linear condition on b and σ, we have:

$$\mathbb{E}_{t_k}\Big[\sup_{s\in[t_k,t_{k+1}]}|X_s|^2\Big] \;\leq\; C(1+|X_{t_k}|^2). \tag{5.2.26}$$

We also know from Proposition 4.2 in [25], under **(H1)** and **(H2)**, that there exists a constant C depending only on the Lipschitz constants of b, σ f and $v^\pi(t_{k+1},.)$ (which does not depend on π by Corollary 5.2.1), such that

$$\mathbb{E}_{t_k}\Big[\sup_{s\in[t_k,t_{k+1}]}\Big(|\mathcal{Y}_s^\pi|^2 + |\mathcal{Z}_s^\pi|^2 + \int_A |\mathcal{U}_s^\pi(a)|^2\lambda(da)\Big)\Big] \;\leq\; C(1+|X_{t_k}|^2). \tag{5.2.27}$$

We deduce that

$$\mathbb{E}_{t_k}[|\mathcal{Y}_t^\pi - \mathcal{Y}_{t_k}^\pi|^2] \;\leq\; \mathbb{E}_{t_k}\Big[\int_{t_k}^t |\mathcal{Y}_s^\pi - \mathcal{Y}_{t_k}^\pi|^2\Big] + C|\pi|(1+|X_{t_k}|^2),$$

and we conclude for the regularity of $\mathcal{Y}^\pi$ by Gronwall's lemma. Finally, from the definition (5.2.14)-(5.2.15) of Y^π and $\mathcal{Y}^\pi$, Itô isometry for stochastic integrals, and growth linear condition on f, we have for all $t \in (t_k, t_{k+1})$:

$$\begin{aligned}
\mathbb{E}_{t_k}[|Y_t^\pi - Y_{t_{k+1}}^\pi|^2] &= \mathbb{E}_{t_k}[|\mathcal{Y}_t^\pi - Y_{t_{k+1}}^\pi|^2] \\
&\leq 3\mathbb{E}_{t_k}\Big[\int_{t_k}^{t_{k+1}}\Big(|f(X_s, I_s, \mathcal{Y}_s^\pi, \mathcal{Z}_s^\pi)|^2 + |\mathcal{Z}_s^\pi|^2 + \int_A |\mathcal{U}_s^\pi(a)|^2\lambda(da)\Big)ds\Big] \\
&\leq C|\pi|\mathbb{E}_{t_k}\Big[1 + \sup_{s\in[t_k,t_{k+1}]}\Big(|X_s|^2 + |\mathcal{Y}_s^\pi|^2 + |\mathcal{Z}_s^\pi|^2 + \int_A |\mathcal{U}_s^\pi(a)|^2\lambda(da)\Big)\Big] \\
&\leq C|\pi|(1+|X_{t_k}|^2),
\end{aligned}$$

where we used again (5.2.26) and (5.2.27). This ends the proof. □

5.2.3 Convergence of discretely jump-constrained BSDE

This section is devoted to the convergence of the discretely jump-constrained BSDE towards the minimal solution to the BSDE with nonpositive jump.

5.2.3.1 Convergence result

Lemma 5.2.3. *We have the following assertions:*

1) The familly $(\vartheta^\pi)_\pi$ is nondecreasing and upper bounded by v: for any grids π and π' such that $\pi \subset \pi'$, we have

$$\vartheta^\pi(t,x,a) \;\leq\; \vartheta^{\pi'}(t,x,a) \;\leq\; v(t,x)\,, \quad (t,x,a) \in [0,T]\times\mathbb{R}^d\times A\,.$$

2) The familly $(\vartheta^\pi)_\pi$ satisfies a uniform linear growth condition: there exists a constant C such that

$$|\vartheta^\pi(t,x,a)| \;\leq\; C(1+|x|),$$

for any $(t,x,a) \in [0,T]\times\mathbb{R}^d\times A$ and any grid π.

Proof. 1) Let us first prove that $\vartheta^\pi \leq v$. Since v is a (continuous) viscosity solution to the HJB equation (5.2.11), and v does not depend on a, we see that v is a viscosity supersolution to the IPDE in (5.2.17) satisfied by ϑ^π on each interval $[t_k, t_{k+1})$. Now, since $v(T,x) = \vartheta^\pi(T,x,a)$, we deduce by comparison principle for this IPDE (see e.g. Theorem 3.4 in [9]) on $[t_{n-1},T)\times\mathbb{R}^d\times A$ that $v(t,x) \geq \vartheta^\pi(t,x,a)$ for all $t\in[t_{n-1},T]$, $(x,a)\in\mathbb{R}^d\times A$. In particular, $v(t_{n-1}^-,x) = v(t_{n-1},x)$ $\geq \sup_{a\in A}\vartheta^\pi(t_{n-1},x,a) = \vartheta^\pi(t_{n-1}^-,x,a)$. Again, by comparison principle for the IPDE (5.2.17) on $[t_{n-2},t_{n-1})\times\mathbb{R}^d\times A$, it follows that $v(t,x)\geq\vartheta^\pi(t,x,a)$ for all $t\in[t_{n-2},t_{n-1}]$, $(x,a)\in\mathbb{R}^d\times A$. By backward induction on time, we conclude that $v\geq\vartheta^\pi$ on $[0,T]\times\mathbb{R}^d\times A$.

Let us next consider two partitions $\pi = (t_k)_{0\leq k\leq n}$ and $\pi' = (t'_k)_{0\leq k\leq n'}$ of $[0,T]$ with $\pi\subset\pi'$, and denote by $m = \max\{k\leq n' : t'_m\notin\pi\}$. Thus, all the points of the grid π and π' coincide after time t'_m, and since ϑ^π and $\vartheta^{\pi'}$ are viscosity solution to the same IPDE (5.2.17) starting from the same terminal data g, we deduce by uniqueness that $\vartheta^\pi = \vartheta^{\pi'}$ on $[t'_m,T]\times\mathbb{R}^d\times A$. Then, we have $\vartheta^{\pi'}(t'^-_m,x,a) = \sup_{a\in A}\vartheta^\pi(t'_m,x,a) = \sup_{a\in A}\vartheta^\pi(t'_m,x,a)\geq\vartheta^\pi(t^-_m,x,a)$ since ϑ^π is continuous outside of the points of the grid π (recall Remark 5.2.2). Now, since ϑ^π and $\vartheta^{\pi'}$ are viscosity solution to the same IPDE (5.2.17) on $[t'_{m-1},t_m)$, we deduce by comparison principle that $\vartheta^{\pi'}\geq\vartheta^\pi$ on $[t'_{m-1},t'_m]\times\mathbb{R}^d\times A$. Proceeding by backward induction, we conclude that $\vartheta^{\pi'}$ $\geq\vartheta^\pi$ on $[0,T]\times\mathbb{R}^d\times A$. 2) Denote by $\pi_0 = \{t_0=0, t_1=T\}$ the trivial grid of $[0,T]$. Since $\vartheta^{\pi_0}\leq\vartheta^\pi\leq v$ and ϑ^{π_0} and v satisfy a linear growth condition, we get (recall that A is compact):

$$|\vartheta^\pi(t,x,a)| \;\leq\; |\vartheta^{\pi_0}(t,x,a)| + |v(t,x)| \;\leq\; C(1+|x|),$$

for any $(t,x,a)\in[0,T]\times\mathbb{R}^d\times A$ and any grid π. □

In the sequel, we denote by ϑ the increasing limit of the sequence $(\vartheta^\pi)_\pi$ when the grid increases by becoming finer, i.e. its modulus $|\pi|$ goes to zero. The next result shows that ϑ does not depend on the variable a in A.

Proposition 5.2.4. *The function ϑ is l.s.c. and does not depend on the variable $a\in A$:*

$$\vartheta(t,x,a) \;=\; \vartheta(t,x,a')\,, \quad (t,x)\in[0,T]\times\mathbb{R}^d,\; a,a'\in A\,.$$

To prove this result we use the following lemma. Observe by definition (5.2.18) of v^π that the function v^π does not depend on a on the grid times π, and we shall denote by misuse of notation: $v^\pi(t_k, x)$, for $k \leq n$, $x \in \mathbb{R}^d$.

Lemma 5.2.4. *There exists a constant C (not depending on π) such that*

$$|\vartheta^\pi(t,x,a) - v^\pi(t_{k+1},x)| \;\leq\; C(1+|x|)|\pi|^{\frac{1}{2}}$$

for all $k = 0,\ldots,n-1$, $t \in [t_k, t_{k+1})$, $(x,a) \in \mathbb{R}^d \times A$.

Proof. Fix $k = 0,\ldots,n-1$, $t \in [t_k, t_{k+1})$ and $(x,a) \in \mathbb{R}^d \times A$. Let $(\tilde{\mathcal{Y}}, \tilde{\mathcal{Z}}, \tilde{\mathcal{U}})$ be the solution to the BSDE

$$\begin{aligned}\tilde{\mathcal{Y}}_s &= v^\pi(t_{k+1}, X^{t,x,a}_{t_{k+1}}) + \int_s^{t_{k+1}} f(X^{t,x,a}_s, I^{t,a}_s, \tilde{\mathcal{Y}}_s, \tilde{\mathcal{Z}}_s)ds \\ &\quad - \int_s^{t_{k+1}} \tilde{\mathcal{Z}}_s dW_s - \int_s^{t_{k+1}} \int_A \tilde{\mathcal{U}}_s(a')\tilde{\mu}(ds,da')\,, \qquad s \in [t, t_{k+1}]\,.\end{aligned}$$

From Proposition 5.2.2, Markov property and uniqueness of a solution to the BSDE (5.2.13)-(5.2.14)-(5.2.15) we have: $\tilde{\mathcal{Y}}_s = \vartheta^\pi(s, X^{t,x,a}_s, I^{t,a}_s)$, for $s \in [t, t_{k+1}]$, and so:

$$\begin{aligned}|\vartheta^\pi(t,x,a) - v^\pi(t_{k+1},x)| &= |\tilde{\mathcal{Y}}_t - v^\pi(t_{k+1},x)| \\ &\leq \mathbb{E}|v^\pi(t_{k+1}, X^{t,x,a}_{t_{k+1}}) - v^\pi(t_{k+1},x)| \\ &\quad + \mathbb{E}\Big[\int_t^{t_{k+1}} |f(X^{t,x,a}_s, I^{t,a}_s, \tilde{\mathcal{Y}}_s, \tilde{\mathcal{Z}}_s)|ds\Big]. \qquad (5.2.28)\end{aligned}$$

From Corollary 5.2.1, we have

$$\mathbb{E}|v^\pi(t_{k+1}, X^{t,x,a}_{t_{k+1}}) - v^\pi(t_{k+1},x)| \;\leq\; C\sqrt{\mathbb{E}[|X^{t,x,a}_{t_{k+1}} - x|^2]} \;\leq\; C\sqrt{|\pi|}\,. \qquad (5.2.29)$$

Moreover, by the growth linear condition on f in **(H2)**, and on ϑ^π in Lemma 5.2.3, we have

$$\mathbb{E}\Big[\int_t^{t_{k+1}} |f(X_s, I_s, \tilde{\mathcal{Y}}_s, \tilde{\mathcal{Z}}_s)|ds\Big] \;\leq\; C\mathbb{E}\Big[\int_t^{t_{k+1}} \Big(1 + |X^{t,x,a}_s| + |\tilde{\mathcal{Z}}_s|\Big)ds\Big].$$

By classical estimates, we have

$$\sup_{s\in[t,T]} \mathbb{E}[|X^{t,x,a}_s|^2] \;\leq\; C(1+|x|^2).$$

Moreover, under **(H1)** and **(H2)**, we know from Proposition 4.2 in [25] that there exists a constant C depending only on the Lipschitz constants of b, σ f and $v^\pi(t_{k+1},.)$ such that

$$\mathbb{E}[\sup_{s\in[t_k,t_{k+1}]} |\tilde{\mathcal{Z}}_s|^2] \;\leq\; C(1+|x|^2).$$

This proves that

$$\mathbb{E}\Big[\int_t^{t_{k+1}} |f(X_s, I_s, \tilde{\mathcal{Y}}_s, \tilde{\mathcal{Z}}_s)|ds\Big] \;\leq\; C(1+|x|)|\pi|\,.$$

Combining this last estimate with (5.2.28) and (5.2.29), we get the result. □

Proof. of Proposition 5.2.4. The function ϑ is l.s.c. as the supremum of the l.s.c. functions ϑ^π. Fix $(t,x) \in [0,T) \times \mathbb{R}^d$ and $a, a' \in A$. Let $(\pi^p)_p$ be a sequence of subdivisions of $[0,T]$ such that $|\pi^p| \downarrow 0$ as $p \uparrow \infty$. We define the sequence $(t_p)_p$ of $[0,T]$ by

$$t_p \;=\; \min\{s \in \pi^p \;:\; s > t\}\,, \quad p \geq 0\,.$$

Since $|\pi^p| \to 0$ as $p \to \infty$ we get $t_p \to t$ as $p \to +\infty$. We then have from the previous lemma:

$$\begin{aligned} |\vartheta^{\pi^p}(t,x,a) - \vartheta^{\pi^p}(t,x,a')| &\leq |\vartheta^{\pi^p}(t,x,a) - v^{\pi^p}(t_p,x)| + |v^{\pi^p}(t_p,x) - \vartheta^{\pi^p}(t,x,a')| \\ &\leq 2C|\pi^p|^{\frac{1}{2}}\,. \end{aligned}$$

Sending p to ∞ we obtain that $\vartheta(t,x,a) = \vartheta(t,x,a')$. □

Corollary 5.2.2. *We have the identification: $\vartheta = v$, and the sequence $(v^\pi)_\pi$ also converges to v.*

Proof. We proceed in two steps.

Step 1. *The function ϑ is a supersolution to* (5.2.11). Since $\vartheta^{\pi_k}(T,.) = g$ for all $k \geq 1$, we first notice that $\vartheta(T,.) = g$. Next, since ϑ does not depend on the variable a, we have

$$\vartheta^\pi(t,x,a) \;\uparrow\; \vartheta(t,x) \quad \text{as} \quad |\pi| \;\downarrow\; 0$$

for any $(t,x,a) \in [0,T] \times \mathbb{R}^d \times A$. Moreover, since the function ϑ is l.s.c, we have

$$\vartheta \;=\; \vartheta_* \;=\; \liminf_{|\pi| \to 0}{}_*\,\vartheta^\pi, \tag{5.2.30}$$

where

$$\liminf_{|\pi| \to 0}{}_*\,\vartheta^\pi(t,x,a) \;:=\; \liminf_{\substack{|\pi| \to 0 \\ (t',x',a') \to (t,x,a) \\ t' < T}} \vartheta^\pi(t',x',a'), \qquad (t,x,a) \in [0,T] \times \mathbb{R}^d \times \mathbb{R}^q\,.$$

Fix now some $(t,x) \in [0,T] \times \mathbb{R}^d$ and $a \in A$ and $(p,q,M) \in \bar{J}^{2,-}\vartheta(t,x)$, the limiting parabolic subjet of ϑ at (t,x) (see definition in [44]). From standard stability results, there exists a sequence $(\pi_k, t_k, x_k, a_k, p_k, q_k, M_k)_k$ such that

$$(p_k, q_k, M_k) \;\in\; \bar{J}^{2,-}\vartheta^{\pi_k}(t_k, x_k, a_k)$$

for all $k \geq 1$ and

$$(t_k, x_k, a_k, \vartheta^{\pi_k}(t_k,x_k,a_k)) \;\longrightarrow\; (t,x,a,\vartheta(t,x,a)) \quad \text{as } k \to \infty,\ |\pi_k| \to 0.$$

From the viscosity supersolution property of ϑ^{π_k} to (5.2.17) in terms of subjets, we have

$$\begin{aligned} -p_k - b(x_k,a_k).q_k - \frac{1}{2}\mathrm{tr}(\sigma\sigma^\intercal(x_k,a_k)M_k) - f\big(x_k,a_k,\vartheta^{\pi_k}(t_k,x_k,a_k),\sigma^\intercal(x_k,a_k)q_k\big) \\ - \int_A \big(\vartheta^{\pi_k}(t_k,x_k,a') - \vartheta^{\pi_k}(t_k,x_k,a_k)\big)\lambda(da') \;\geq\; 0 \end{aligned}$$

for all $k \geq 1$. Sending k to infinity and using (5.2.30), we get

$$-p - b(x,a).q - \frac{1}{2}\mathrm{tr}(\sigma\sigma^\intercal(x,a)M) - f\big(x,a,\vartheta(t,x),\sigma^\intercal(x,a)q\big) \;\geq\; 0.$$

Since a is arbitrary in A, this shows

$$-p - \sup_{a\in A}\big[b(x,a).q + \frac{1}{2}\mathrm{tr}(\sigma\sigma^\intercal(x,a)M) + f\big(x,a,\vartheta(t,x),\sigma^\intercal(x,a)q\big)\big] \quad \geq \quad 0,$$

i.e. the viscosity supersolution property of ϑ to (5.2.11). **Step 2.** *Comparison.* Since the PDE (5.2.11) satisfies a comparison principle, we have from the previous step $\vartheta \geq v$, and we conclude with Lemma 5.2.3 that $\vartheta = v$. Finally, by definition (5.2.18) of v^π and from Lemma 5.2.3, we clearly have $\vartheta^\pi \leq v^\pi \leq v$, which also proves that $(v^\pi)_\pi$ converges to v. □

In terms of the discretely jump-constrained BSDE, the convergence result is formulated as follows:

Proposition 5.2.5. *We have $\mathcal{Y}_t^\pi \leq Y_t^\pi \leq Y_t$, $0 \leq t \leq T$, and*

$$\mathbb{E}\Big[\sup_{t\in[0,T]}|Y_t - \mathcal{Y}_t^\pi|^2\Big] + \mathbb{E}\Big[\sup_{t\in[0,T]}|Y_t - Y_t^\pi|^2\Big] + \mathbb{E}\Big[\int_0^T |Z_t - \mathcal{Z}_t^\pi|^2 dt\Big] \quad \to \quad 0,$$

as $|\pi|$ goes to zero.

Proof. Recall from (5.2.10) and (5.2.25) that we have the representation:

$$Y_t \;=\; v(t,X_t), \quad Y_t^\pi \;=\; v^\pi(t,X_t,I_t), \quad \mathcal{Y}_t^\pi \;=\; \vartheta(t,X_t,I_t), \tag{5.2.31}$$

and the first assertion of Lemma 5.2.3, we clearly have: $\mathcal{Y}_t^\pi \leq Y_t^\pi \leq Y_t$, $0 \leq t \leq T$. The convergence in $\mathcal{S}^2$ for $\mathcal{Y}^\pi$ to Y and Y^π to Y comes from the above representation (5.2.31), the pointwise convergence of ϑ^π and v^π to v in Corollary 5.2.2, and the dominated convergence theorem by recalling that $0 \leq (v - v^\pi)(t,x,a) \leq (v - \vartheta^\pi)(t,x,a) \leq v(t,x) \leq C(1+|x|)$. Let us now turn to the component Z. By definition (5.2.13)-(5.2.14)-(5.2.15) of the discretely jump-constrained BSDE we notice that $\mathcal{Y}^\pi$ can be written on $[0,T]$ as:

$$\mathcal{Y}_t^\pi \;=\; g(X_T) + \int_t^T f(X_s,I_s,\mathcal{Y}_s^\pi,\mathcal{Z}_s^\pi) - \int_t^T \mathcal{Z}_s^\pi dW_s - \int_t^T\int_A \mathcal{U}_s^\pi(a)\tilde{\mu}(ds,da) + \mathcal{K}_T^\pi - \mathcal{K}_t^\pi,$$

where $\mathcal{K}^\pi$ is the nondecreasing process defined by: $\mathcal{K}_t^\pi = \sum_{t_k\leq t}(Y_{t_k}^\pi - \mathcal{Y}_{t_k}^\pi)$, for $t \in [0,T]$. Denote by $\delta Y = Y - \mathcal{Y}^\pi$, $\delta Z = Z - \mathcal{Z}^\pi$, $\delta U = U - \mathcal{U}^\pi$ and $\delta K = K - \mathcal{K}^\pi$. From Itô's formula, Young Inequality and **(H2)**, there exists a constant C such that

$$\begin{aligned} &\mathbb{E}\Big[|\delta Y_t|^2\Big] + \frac{1}{2}\mathbb{E}\Big[\int_t^T |\delta Z_s|^2 ds\Big] + \frac{1}{2}\mathbb{E}\Big[\int_t^T |\delta U_s(a)|^2\lambda(da)ds\Big] \\ \leq\; & C\int_t^T \mathbb{E}\Big[|\delta Y_s|^2\Big]ds + \frac{1}{\varepsilon}\mathbb{E}\Big[\sup_{s\in[0,T]}|\delta Y_s|^2\Big] + \varepsilon\mathbb{E}\Big[\big|\delta K_T - \delta K_t\big|^2\Big] \end{aligned} \tag{5.2.32}$$

for all $t \in [0,T]$, with ε a constant to be chosen later. From the definition of δK we have

$$\begin{aligned} \delta K_T - \delta K_t \;=\; & \delta Y_t - \int_t^T \big(f(X_s,I_s,Y_s,Z_s) - f(X_s,I_s,\mathcal{Y}_s^\pi,\mathcal{Z}_s^\pi)\big)ds \\ & + \int_0^T \delta Z_s dW_s + \int_t^T\int_A \delta U_s(a)\tilde{\mu}(ds,da)\,. \end{aligned}$$

Therefore, by (H2), we get the existence of a constant C' such that

$$\mathbb{E}\Big[|\delta K_T - \delta K_t|^2\Big] \;\leq\; C'\Big(\mathbb{E}\Big[\sup_{s\in[0,T]}|\delta Y_s|^2\Big] + \mathbb{E}\Big[\int_t^T |\delta Z_s|^2 ds\Big] + \mathbb{E}\Big[\int_t^T |\delta U_s(a)|^2 \lambda(da)ds\Big]\Big)$$

Taking $\varepsilon = \frac{C'}{4}$ and plugging this last inequality in (5.2.32), we get the existence of a constant C'' such that

$$\mathbb{E}\Big[\int_t^T |\delta Z_s|^2 ds\Big] + \mathbb{E}\Big[\int_t^T |\delta U_s(a)|^2 \lambda(da)ds\Big] \;\leq\; C''\Big(\mathbb{E}\Big[\sup_{s\in[0,T]}|\delta Y_s|^2\Big]\Big), \qquad (5.2.33)$$

which shows the $L^2(W)$ convergence of $\mathcal{Z}^\pi$ to Z from the $\mathcal{S}^2$ convergence of $\mathcal{Y}^\pi$ to Y. □

5.2.3.2 Rate of convergence

We next provide an error estimate for the convergence of the discretely jump-constrained BSDE. We shall combine BSDE methods and PDE arguments adapted from the shaking coefficients approach of Krylov [78] and switching systems approximation of Barles, Jacobsen [10]. We make further assumptions:

(H1') The functions b and σ are uniformly bounded:

$$\sup_{x\in\mathbb{R}^d, a\in A} |b(x,a)| + |\sigma(x,a)| \;<\; \infty.$$

(H2') The function f does not depend on z: $f(x,a,y,z) = f(x,a,y)$ for all $(x,a,y,z) \in \mathbb{R}^d \times A \times \mathbb{R} \times \mathbb{R}^d$ and

[(i)]

1. the functions $f(.,.,0)$ and g are uniformly bounded:

$$\sup_{x\in\mathbb{R}^d, a\in A} |f(x,a,0)| + |g(x)| \;<\; \infty,$$

2. for all $(x,a) \in \mathbb{R}^d \times A$, $y \mapsto f(x,a,y)$ is convex.

Under these assumptions, we obtain the rate of convergence for v^π and ϑ^π towards v.

Theorem 5.2.1. *Under **(H1')** and **(H2')**, there exists a constant C such that*

$$0 \;\leq\; v(t,x) - v^\pi(t,x,a) \;\leq\; v(t,x) - \vartheta^\pi(t,x,a) \;\leq\; C|\pi|^{\frac{1}{10}}$$

for all $(t,x,a) \in [0,T] \times \mathbb{R}^d \times A$ and all grid π with $|\pi| \leq 1$. Moreover, when $f(x,a)$ does not depend on y, the rate of convergence is improved to $|\pi|^{\frac{1}{6}}$.

Before proving this result, we give as corollary the rate of convergence for the discretely jump-constrained BSDE.

Corollary 5.2.3. *Under* ***(H1')*** *and* ***(H2')****, there exists a constant* C *such that*

$$\mathbb{E}\Big[\sup_{t\in[0,T]}|Y_t-\mathcal{Y}_t^\pi|^2\Big]+\mathbb{E}\Big[\sup_{t\in[0,T]}|Y_t-Y_t^\pi|^2\Big]+\mathbb{E}\Big[\int_0^T|Z_t-\mathcal{Z}_t^\pi|^2dt\Big] \;\leq\; C|\pi|^{\frac{1}{5}}.$$

for all grid π *with* $|\pi|\leq 1$*, and the above rate is improved to* $|\pi|^{\frac{1}{3}}$ *when* $f(x,a)$ *does not depend on* y*.*

Proof. From the representation 5.2.31, and Theorem 5.2.1, we immediately have

$$\mathbb{E}\Big[\sup_{t\in[0,T]}|Y_t-\mathcal{Y}_t^\pi|^2\Big]+\mathbb{E}\Big[\sup_{t\in[0,T]}|Y_t-Y_t^\pi|^2\Big] \;\leq\; C|\pi|^{\frac{1}{5}}. \tag{5.2.34}$$

(resp. $|\pi|^{\frac{1}{3}}$ when $f(x,a)$ does not depend on y). Finally, the convergence rate for Z follows from the inequality (5.2.33). □

Remark 5.2.3. The above convergence rate $|\pi|^{\frac{1}{10}}$ is the optimal rate that one can prove in our generalized stochastic control context with fully nonlinear HJB equation by PDE approach and shaking coefficients technique, see [78], [10], [49] or [100]. However, this rate may not be the sharpest one. In the case of continuously reflected BSDEs, i.e. BSDEs with upper or lower constraint on Y, it is known that Y can be approximated by discretely reflected BSDEs, i.e. BSDEs where reflection on Y operates a finite set of times on the grid π, with a rate $|\pi|^{\frac{1}{2}}$ (see [7]). The standard arguments for proving this rate is based on the representation of the continuously (resp. discretely) reflected BSDE as optimal stopping problems where stopping is possible over the whole interval time (resp. only on the grid times). In our jump-constrained case, we know from [72] that the minimal solution to the BSDE with nonpositive jumps has the stochastic control representation 5.1.2 when $f(x,a)$ does not depend on y and z. We shall prove an analog representation for discretely jump-constrained BSDEs, and this helps to improve the rate of convergence from $|\pi|^{\frac{1}{10}}$ to $|\pi|^{\frac{1}{6}}$.

The rest of this section is devoted to the proof of Theorem 5.2.1. We first consider the special case where $f(x,a)$ does not depend on y, and then address the case $f(x,a,y)$.

Proof. **Case $f(x,a)$.**

In the case where $f(x,a)$ does not depend on y, z, by (linear) Feynman-Kac formula for ϑ^π solution to (5.2.17), and by definition of v^π in (5.2.18), we have:

$$v^\pi(t_k,x)=\sup_{a\in A}\mathbb{E}\left[\int_{t_k}^{t_{k+1}}f\left(X_t^{t_k,x,a},I_t^{t_k,a}\right)dt+v^\pi\left(t_{k+1},X_{t_{k+1}}^{t_k,x,a}\right)\right],\; k\leq n-1,\; x\in\mathbb{R}^d.$$

By induction, this dynamic programming relation leads to the following stochastic control problem with discrete time policies:

$$v^\pi(t_k,x)=\sup_{a\in\mathcal{A}_{\mathbb{F}}^\pi}\mathbb{E}\left[\int_{t_k}^{T}f\left(\bar{X}_t^{t_k,x,a},\bar{I}_t^{\alpha}\right)dt+g\left(\bar{X}_T^{t_k,x,a}\right)\right],$$

where $\mathcal{A}_{\mathbb{F}}^\pi$ is the set of discrete time processes $\alpha=\left(\alpha_{t_j}\right)_{j\leq n-1}$ with α_{t_j} $\mathcal{F}_{t_j}$-measurable, valued in A, and

$$\begin{aligned}\bar{X}_t^{t_k,x,\alpha}=&x+\int_{t_k}^t b\left(\bar{X}_s^{t,x,\alpha},\bar{I}_s^{\alpha}\right)ds+\int_{t_k}^t\sigma\left(\bar{X}_s^{t,x,\alpha},\bar{I}_s^{\alpha}\right)dW_s,\;\; t_k\leq t\leq T,\\ \bar{I}_t^{\alpha}=&\alpha_{t_j}+\int_{(t_j,t]}\int_A\left(a-\bar{I}_{s^-}^{\alpha}\right)\mu(ds,da)\;,\;\; t_k\leq t<t_{j+1},\;\; j\leq n-1.\end{aligned}$$

In other words, $v^\pi(t_k,x)$ corresponds to the value function for a stochastic control problem where the controller can act only at the dates t_j of the grid π, and then let the regime of the coefficients of the diffusion evolve according to the Poisson random measure μ. Let us introduce the following stochastic control problem with piece-wise constant control policies:

$$\tilde{v}^\pi(t_k,x) = \sup_{a\in\mathcal{A}_{\mathbb{F}}^\pi} \mathbb{E}\left[\int_{t_k}^T f\left(\tilde{X}_t^{t_k,x,a},\tilde{I}_t^\alpha\right)dt + g\left(\tilde{X}_T^{t_k,x,a}\right)\right],$$

where for $\alpha = \left(\alpha_{t_j}\right)_{j\le n-1} \in \mathcal{A}_{\mathbb{F}}^\pi$:

$$\begin{aligned}\tilde{X}_t^{t_k,x,\alpha} =& x + \int_{t_k}^t b\left(\tilde{X}_s^{t,x,\alpha},\tilde{I}_s^\alpha\right)ds + \int_{t_k}^t \sigma\left(\tilde{X}_s^{t,x,\alpha},\tilde{I}_s^\alpha\right)dW_s\,,\ t_k\le t\le T\,,\\ \tilde{I}_t^\alpha =& \alpha_{t_j}\,,\ t_j\le t<t_{j+1}\,,\ j\le n-1\,.\end{aligned}$$

It is shown in [77] that $\tilde{v}^\pi$ approximates the value function v for the controlled diffusion problem (5.1.2), solution to the HJB equation (5.2.11), with a rate $|\pi|^{\frac{1}{6}}$:

$$0\le v(t_k,x)-\tilde{v}^\pi(t_k,x)\le C|\pi|^{\frac{1}{6}}\,, \tag{5.2.35}$$

for all $t_k\in\pi$, $x\in\mathbb{R}^d$. Now, recalling that A is compact and $\lambda(A)<\infty$, it is clear that there exists some positive constant C such that for all $\alpha\in\mathcal{A}_{\mathbb{F}}^\pi$, $j\le n-1$:

$$\mathbb{E}\left[\sup_{t\in[t_j,t_{j+1})}\left|\bar{I}_t^\alpha-\tilde{I}_t^\alpha\right|^2\right]\le C|\pi|\,,$$

and then by standard arguments under Lipschitz condition on b, σ:

$$\mathbb{E}\left[\sup_{t\in[t_j,t_{j+1}]}\left|\bar{X}_t^{t_k,x,\alpha}-\tilde{X}_t^{t_k,x,\alpha}\right|^2\right]\le C|\pi|\,,\ k\le j\le n-1\,,\ x\in\mathbb{R}^d\,.$$

By the Lipschitz conditions on f and g, it follows that

$$|v^\pi(t_k,x)-\tilde{v}^\pi(t_k,x)|\le C|\pi|^{\frac{1}{2}}\,,$$

and thus with (5.2.35):

$$0\le\sup_{x\in\mathbb{R}^d}(v-v^\pi)(t_k,x)\le C|\pi|^{\frac{1}{6}}\,.$$

Finally, by combining with the estimate in Lemma 5.2.4, which gives actually under **(H2')**(i):

$$|\vartheta^\pi(t,x,a)-v^\pi(t_{k+1},x)|\le C|\pi|^{\frac{1}{2}}\,,\ t\in[t_k,t_{k+1})\,,\ (x,a)\in\mathbb{R}^d\times A\,,$$

together with the 1/2-Hölder property of v in time (see (5.2.12)), we obtain:

$$\sup_{(t,x,a)\in[0,T]\times\mathbb{R}^d\times A}(v-\vartheta^\pi)(t,x,a)\le C\left(|\pi|^{\frac{1}{6}}+|\pi|^{\frac{1}{2}}\right)\le C|\pi|^{\frac{1}{6}}\,,$$

for $|\pi|\le 1$. This ends the proof in the case $f(x,a)$.

Case $f(x,a,y)$.

Let us now turn to the case where $f(x,a,y)$ may also depend on y. We cannot rely anymore on the convergence rate result in [77]. Instead, recalling that A is compact and since σ, b and f are Lipschitz in (x,a), we are allowed to apply the switching system method of Barles and Jacobsen [10], which is a variation of the shaken coefficients method and smoothing technique

used in Krylov [78], in order to obtain approximate smooth subsolution to (5.2.11). By Lemmas 3.3 and 3.4 in [10], one can find a family of smooth functions $(w_\varepsilon)_{0<\varepsilon\leq 1}$ on $[0,T]\times\mathbb{R}^d$ such that:

$$\sup_{[0,T]\times\mathbb{R}^d}|w_\varepsilon| \;\leq\; C, \tag{5.2.36}$$

$$\sup_{[0,T]\times\mathbb{R}^d}|w_\varepsilon - w| \;\leq\; C\varepsilon^{\frac{1}{3}}, \tag{5.2.37}$$

$$\sup_{[0,T]\times\mathbb{R}^d}|\partial_t^{\beta_0}D^\beta w_\varepsilon| \;\leq\; C\varepsilon^{1-2\beta_0-\sum_{i=1}^d\beta^i},\;\; \beta_0\in\mathbb{N},\; \beta=(\beta^1,\ldots,\beta^d)\in\mathbb{N}^d, \tag{5.2.38}$$

for some positive constant C independent of ε, and by convexity of f in **(H2')**(ii), for any $\varepsilon\in(0,1]$, $(t,x)\in[0,T]\times\mathbb{R}^d$, there exists $a_{t,x,\varepsilon}\in A$ such that:

$$-\mathcal{L}^{a_{t,x,\varepsilon}}w_\varepsilon(t,x) - f(x,a_{t,x,\varepsilon},w_\varepsilon(t,x)) \;\geq\; 0. \tag{5.2.39}$$

□

Recalling the definition of the operator $\mathbb{T}_\pi^k$ in (5.2.19), we define for any function φ on $[0,T]\times\mathbb{R}^d\times A$, Lipschitz in (x,a):

$$\mathbb{T}_\pi[\varphi](t,x,a) \;:=\; \mathbb{T}_\pi^k[\varphi(t_{k+1},.,.)](t,x,a),\quad t\in[t_k,t_{k+1}),\;(x,a)\in\mathbb{R}^d\times A,$$

for $k=0,\ldots,n-1$, and

$$\begin{aligned}\mathbb{S}_\pi[\varphi](t,x,a) \;:=\; \frac{1}{|\pi|}\Big[&\varphi(t,x)-\mathbb{T}_\pi[\varphi](t,x,a)\\ &+(t_{k+1}-t)\big(\mathcal{L}^a\varphi(t,x)+f(x,a,\varphi(t,x))\big)\Big],\end{aligned}$$

for $(t,x,a)\in[t_k,t_{k+1})\times\mathbb{R}^d\times A$, $k\leq n-1$.

We have the following key error bound on $\mathbb{S}_\pi$.

Lemma 5.2.5. *Let **(H1')** and **(H2')**(i) hold. There exists a constant C such that*

$$|\mathbb{S}_\pi[\varphi_\varepsilon](t,x,a)| \;\leq\; C\Big(|\pi|^{\frac{1}{2}}(1+\varepsilon^{-1})+|\pi|\varepsilon^{-3}\Big),\quad (t,x,a)\in[0,T]\times\mathbb{R}^d\times A,$$

for any family $(\varphi_\varepsilon)_\varepsilon$ of smooth functions on $[0,T]\times\mathbb{R}^d$ satisfying (5.2.36) and (5.2.38).

Proof. Fix $(t,x,a)\in[0,T]\times\mathbb{R}^d\times A$. If $t=T$, we have $|\mathbb{S}_\pi[\varphi_\varepsilon](t,x,a)|=0$. Suppose that $t<T$ and fix $k\leq n$ such that $t\in[t_k,t_{k+1})$. Given a smooth function φ_ε satisfying (5.2.36) and (5.2.38), we split:

$$|\mathbb{S}_\pi[\varphi_\varepsilon](t,x,a)| \;\leq\; A_\varepsilon(t,x,a)+B_\varepsilon(t,x,a),$$

where

$$A_\varepsilon(t,x,a) \;:=\; \frac{1}{|\pi|}\Big|\mathbb{T}_\pi[\varphi_\varepsilon](t,x,a)-\mathbb{E}[\varphi_\varepsilon(t_{k+1},X_{t_{k+1}}^{t,x,a})]-(t_{k+1}-t)f(x,a,\varphi_\varepsilon(t,x))\Big|,$$

and

$$B_\varepsilon(t,x,a) \;:=\; \frac{1}{|\pi|}\Big|\mathbb{E}[\varphi_\varepsilon(t_{k+1},X_{t_{k+1}}^{t,x,a})]-\varphi_\varepsilon(t,x)-(t_{k+1}-t)\mathcal{L}^a\varphi_\varepsilon(t,x)\Big|,$$

and we study each term A_ε and B_ε separately. **1.** *Estimate on* $A_\varepsilon(t,x,a)$.

Define $(Y^{\varphi_\varepsilon}, Z^{\varphi_\varepsilon}, U^{\varphi_\varepsilon})$ as the solution to the BSDE on $[t, t_{k+1}]$:

$$\begin{aligned} Y_s^{\varphi_\varepsilon} &= \varphi_\varepsilon(t_{k+1}, X_{t_{k+1}}^{t,x,a}) + \int_s^{t_{k+1}} f(X_r^{t,x,a}, I_r^{t,a}, Y_r^{\varphi_\varepsilon}) dr \\ &\quad - \int_s^{t_{k+1}} Z_r^{\varphi_\varepsilon} dW_r - \int_s^{t_{k+1}} \int_A U_r^{\varphi_\varepsilon}(a) \tilde{\mu}(dr, da)\ , \quad s \in [t, t_{k+1}]. \end{aligned} \tag{5.2.40}$$

From Theorems 3.4 and 3.5 in [9], we have $Y_t^{\varphi_\varepsilon} = \mathbb{T}_\pi[\varphi_\varepsilon](t,x,a)$, and by taking expectation in (5.2.40), we thus get:

$$Y_t^{\varphi_\varepsilon} = \mathbb{T}_\pi[\varphi_\varepsilon](t,x,a) = \mathbb{E}[\varphi_\varepsilon(t_{k+1}, X_{t_{k+1}}^{t,x,a})] + \mathbb{E}\Big[\int_t^{t_{k+1}} f(X_s^{t,x,a}, I_s^{t,a}, Y_s^{\varphi_\varepsilon}) ds\Big]$$

and so:

$$\begin{aligned} A_\varepsilon(t,x,a) &\le \frac{1}{|\pi|}\mathbb{E}\Big[\int_t^{t_{k+1}} |f(X_s^{t,x,a}, I_s^{t,a}, Y_s^{\varphi_\varepsilon}) - f(x,a,\varphi_\varepsilon(t,x))| ds\Big] \\ &\le C\Big(\mathbb{E}[\sup_{s\in[t,t_{k+1}]} |X_s^{t,x,a} - x| + |I_s^{t,a} - a|] + \mathbb{E}[\sup_{s\in[t,t_{k+1}]} |Y_s^{\varphi_\varepsilon} - \varphi_\varepsilon(t,x)|]\Big), \end{aligned}$$

by the Lipschitz continuity of f. From standard estimate for SDE, we have (recall that the coefficients b and σ are bounded under **(H1')** and A is compact):

$$\mathbb{E}[\sup_{s\in[t,t_{k+1}]} |X_s^{t,x,a} - x| + |I_s^{t,a} - a|] \le C|\pi|^{\frac{1}{2}}. \tag{5.2.41}$$

Moreover, by (5.2.40), the boundedness condition in **(H2')**(i) together with the Lipschitz condition of f, and Burkholder-Davis-Gundy inequality, we have:

$$\begin{aligned} \mathbb{E}\Big[\sup_{s\in[t,t_{k+1}]} |Y_s^{\varphi_\varepsilon} - \varphi_\varepsilon(t,x)|\Big] &\le \mathbb{E}[|\varphi_\varepsilon(t_{k+1}, X_{t_{k+1}}^{t,x,a}) - \varphi_\varepsilon(t,x)|] \\ &\quad + C|\pi|\mathbb{E}[1 + \sup_{s\in[t,t_{k+1}]} |Y_s^{\varphi_\varepsilon}|] \\ &\quad + C|\pi|\Big(\mathbb{E}[\sup_{s\in[t,t_{k+1}]} |Z_s^{\varphi_\varepsilon}|^2] + \mathbb{E}[\sup_{s\in[t,t_{k+1}]} \int_A |U_s^{\varphi_\varepsilon}(a)|^2 \lambda(da)]\Big). \end{aligned}$$

From standard estimate for the BSDE (5.2.40), we have:

$$\mathbb{E}[\sup_{s\in[t,t_{k+1}]} |Y_s^{\varphi_\varepsilon}|^2] \le C,$$

for some positive constant C depending only on the Lipschitz constant of f, the upper bound of $|f(x,a,0,0)|$ in **(H2')**(i), and the upper bound of $|\varphi_\varepsilon|$ in (5.2.36). Moreover, from the estimate in Proposition 4.2 in [25] about the coefficients Z^{φ_ε} and U^{φ_ε} of the BSDE with jumps (5.2.40), there exists some constant C depending only on the Lipschitz constant of b, σ, f, and of the Lipschitz constant of $\varphi_\varepsilon(t_{k+1}, .)$ (which does not depend on ε by (5.2.38)), such that:

$$\mathbb{E}[\sup_{s\in[t,t_{k+1}]} |Z_s^{\varphi_\varepsilon}|^2] + \mathbb{E}[\sup_{s\in[t,t_{k+1}]} \int_A |U_s^{\varphi_\varepsilon}(a)|^2 \lambda(da)] \le C.$$

From (5.2.38), we then have:

$$\begin{aligned} \mathbb{E}\Big[\sup_{s\in[t,t_{k+1}]} |Y_s^{\varphi_\varepsilon} - \varphi_\varepsilon(t,x)|\Big] &\le C(|t_{k+1} - t|\varepsilon^{-1} + \mathbb{E}[|X_{t_{k+1}}^{t,x,a} - x|] + |\pi|) \\ &\le C|\pi|^{\frac{1}{2}}(1 + \varepsilon^{-1}), \end{aligned}$$

by (5.2.41). This leads to the error bound for $A_\varepsilon(t,x,a)$:

$$A_\varepsilon(t,x,a) \quad \leq \quad C|\pi|^{\frac{1}{2}}(1+\varepsilon^{-1}).$$

2. *Estimate on* $B_\varepsilon(t,x,a)$.

From Itô's formula we have

$$\begin{aligned} B_\varepsilon(t,x,a) &= \frac{1}{|\pi|}\Big|\mathbb{E}\Big[\int_t^{t_{k+1}} (\mathcal{L}^{I_s^{t,a}}\varphi_\varepsilon(s,X_s^{t,x,a}) - \mathcal{L}^a\varphi_\varepsilon(t,x))ds\Big]\Big| \\ &\leq B_\varepsilon^1(t,x,a) + B_\varepsilon^2(t,x,a) \end{aligned}$$

where

$$\begin{aligned} B_\varepsilon^1(t,x,a) &= \frac{1}{|\pi|}\mathbb{E}\Big[\int_t^{t_{k+1}} \Big|(b(X_s^{t,x,a},I_s^{t,a}) - b(x,a)).D_x\varphi_\varepsilon(s,X_s^{t,x,a}) \\ &\qquad + \frac{1}{2}\mathrm{tr}\big[(\sigma\sigma^\intercal(X_s^{t,x,a},I_s^{t,a}) - \sigma\sigma^\intercal(x,a))D_x^2\varphi_\varepsilon(t,x)\big]\Big|ds\Big] \end{aligned}$$

and

$$B_\varepsilon^2(t,x,a) \quad = \quad \frac{1}{|\pi|}\mathbb{E}\Big[\int_t^{t_{k+1}} |\tilde{\mathcal{L}}_{t,x}^a\varphi_\varepsilon(s,X_s^{t,x,a}) - \tilde{\mathcal{L}}_{t,x}^a\varphi_\varepsilon(t,x)|ds\Big]\,,$$

with $\tilde{\mathcal{L}}_{t,x}^a$ defined by

$$\tilde{\mathcal{L}}_{t,x}^a\varphi_\varepsilon(t',x') \quad = \quad \frac{\partial\varphi_\varepsilon}{\partial t}(t',x') + b(x,a).D_x\varphi_\varepsilon(t',x') + \frac{1}{2}\mathrm{tr}(\sigma\sigma^\intercal(x,a)D_x^2\varphi_\varepsilon(t',x')).$$

Under **(H1)**, **(H1')**, and by (5.2.38), we have

$$\begin{aligned} B_\varepsilon^1(t,x,a) &\leq C(1+\varepsilon^{-1})\mathbb{E}\Big[\sup_{s\in[t,t_{k+1}]}|X_s^{t,x,a}-x| + |I_s^{t,a}-a|\Big] \\ &\leq C(1+\varepsilon^{-1})|\pi|^{\frac{1}{2}}, \end{aligned}$$

where we used again (5.2.41). On the other hand, since φ_ε is smooth, we have from Itô's formula

$$B_\varepsilon^2(t,x,a) \quad = \quad \frac{1}{|\pi|}\mathbb{E}\Big[\int_t^{t_{k+1}} \Big|\int_t^s \mathcal{L}^{I_r^{t,a}}\tilde{\mathcal{L}}_{t,x}^a\phi(r,X_r^{t,x,a})dr\Big|ds\Big]\,.$$

Under **(H1')** and by (5.2.38), we then see that

$$B_\varepsilon^2(t,x,a) \quad \leq \quad C|\pi|\varepsilon^{-3},$$

and so:

$$B_\varepsilon(t,x,a) \quad \leq \quad C\Big(|\pi|^{\frac{1}{2}}(1+\varepsilon^{-1}) + |\pi|\varepsilon^{-3}\Big).$$

Together with the estimate for $A_\varepsilon(t,x,a)$, this proves the error bound for $|\mathbb{S}_\pi[\varphi_\varepsilon](t,x,a)|$. □

We next state a maximum principle type result for the operator $\mathbb{T}_\pi$.

Lemma 5.2.6. *Let φ and ψ be two functions on $[0,T]\times\mathbb{R}^d\times A$, Lipschitz in (x,a). Then, there exists some positive constant C independent of π such that*

$$\sup_{(t,x,a)\in[t_k,t_{k+1}]\times\mathbb{R}^d\times A}(\mathbb{T}_\pi[\varphi]-\mathbb{T}_\pi[\psi])(t,x,a) \;\le\; e^{C|\pi|}\sup_{(x,a)\in\mathbb{R}^d\times A}(\varphi-\psi)(t_{k+1},x,a)\,,$$

for all $k=0,\dots,n-1$.

Proof. Fix $k\le n-1$, and set

$$M \;:=\; \sup_{(x,a)\in\mathbb{R}^d\times A}(\varphi-\psi)(t_{k+1},x,a).$$

We can assume w.l.o.g. that $M<\infty$ since otherwise the required inequality is trivial. Let us denote by Δv the function

$$\Delta v(t,x,a) \;=\; \mathbb{T}_\pi[\varphi](t,x,a)-\mathbb{T}_\pi[\psi](t,x,a),$$

for all $(t,x,a)\in[t_k,t_{k+1}]\times\mathbb{R}^d\times A$. By definition of $\mathbb{T}_\pi$, and from the Lipschitz condition of f, we see that Δv is a viscosity subsolution to

$$\begin{cases} -\mathcal{L}^a\Delta v(t,x,a)-C\big(|\Delta v(t,x,a)|+|D\Delta v(t,x,a)|\big) \\ -\int_A(\Delta v(t,x,a')-\Delta v(t,x,a))\lambda(da')=0, \quad \text{for } (t,x,a)\in[t_k,t_{k+1})\times\mathbb{R}^d\times A, \\ \Delta v(t_{k+1},x,a)\le M\,, \qquad \text{for } (x,a)\in\mathbb{R}^d\times A\,. \end{cases} \tag{5.2.42}$$

Then, we easily check that the function Φ defined by

$$\Phi(t,x,a) \;=\; Me^{C(t_{k+1}-t)}\,, \quad (t,x,a)\in[t_k,t_{k+1}]\times\mathbb{R}^d\times A\,,$$

is a solution to

$$\begin{cases} -\mathcal{L}^a\Phi(t,x,a)-C\big(|\Phi(t,x,a)|+|D\Phi(t,x,a)|\big) \\ -\int_A(\Phi(t,x,a')-\Phi(t,x,a))\lambda(da') \;=\; 0, \quad \text{for } (t,x,a)\in[t_k,t_{k+1})\times\mathbb{R}^d\times A, \\ \Phi(t_{k+1},x,a) \;=\; M\,, \qquad \text{for } (x,a)\in\mathbb{R}^d\times A\,. \end{cases} \tag{5.2.43}$$

From the comparison theorem in [9] for viscosity solutions of semi-linear IPDEs, we get that $\Delta v\le\Phi$ on $[t_k,t_{k+1}]\times\mathbb{R}^d\times A$, which proves the required inequality. □

Proof. of Theorem 5.2.1. By (5.2.18) and (5.2.24), we observe that v^π is a fixed point of $\mathbb{T}_\pi$, i.e.

$$\mathbb{T}_\pi[v^\pi] \;=\; v^\pi.$$

On the other hand, by (5.2.39), and the estimate of Lemma 5.2.5 applied to w_ε, we have:

$$w_\varepsilon(t,x)-\mathbb{T}_\pi[w_\varepsilon](t,x,a_{t,x,\varepsilon}) \;\le\; |\pi|\mathbb{S}_\pi[w_\varepsilon](t,x,a_{t,x,\varepsilon}) \;\le\; C|\pi|\bar{S}(\pi,\varepsilon)$$

where we set: $\bar{S}(\pi,\varepsilon)=(|\pi|^{\frac{3}{2}}(1+\varepsilon^{-1})+|\pi|^2\varepsilon^{-3})$. Fix $k\le n-1$. By Lemma 5.2.6, we then have for all $t\in[t_k,t_{k+1}]$, $x\in\mathbb{R}^d$:

$$\begin{aligned} w_\varepsilon(t,x)-v^\pi(t,x,a_{t,x,\varepsilon}) &= w_\varepsilon(t,x)-\mathbb{T}_\pi[w_\varepsilon](t,x,a_{t,x,\varepsilon})+(\mathbb{T}_\pi[w_\varepsilon]-\mathbb{T}_\pi[v^\pi])(t,x,a_{t,x,\varepsilon}) \\ &\le C|\pi|\bar{S}(\pi,\varepsilon)+e^{C|\pi|}\sup_{(x,a)\in\mathbb{R}^d\times A}(w_\varepsilon-v^\pi)(t_{k+1},x,a). \end{aligned} \tag{5.2.44}$$

Recalling by its very definition that v^π does not depend on $a \in A$ on the grid times of π, and denoting then $M_k := \sup_{x\in\mathbb{R}^d}(w_\varepsilon - v^\pi)(t_k,x)$, we have by (5.2.44) the relation:

$$M_k \;\leq\; C|\pi|\bar{S}(\pi,\varepsilon) + e^{C|\pi|}M_{k+1}.$$

By induction, this yields:

$$\begin{aligned}\sup_{x\in\mathbb{R}^d}(w_\varepsilon - v^\pi)(t_k,x) &\leq C\frac{e^{Cn|\pi|}-1}{e^{C|\pi|}-1}|\pi|\bar{S}(\pi,\varepsilon) + e^{Cn|\pi|}\sup_{x\in\mathbb{R}^d}(w_\varepsilon - v^\pi)(T,x)\\ &\leq C\bar{S}(\pi,\varepsilon) + C\sup_{x\in\mathbb{R}^d}(w_\varepsilon - v)(T,x),\end{aligned}$$

since $n|\pi|$ is bounded and $v(T,x) = v^\pi(T,x)$ $(= g(x))$. From (5.2.37), we then get:

$$\sup_{x\in\mathbb{R}^d}(v - v^\pi)(t_k,x) \;\leq\; C(\varepsilon^{\frac{1}{3}} + |\pi|^{\frac{1}{2}}(1+\varepsilon^{-1}) + |\pi|\varepsilon^{-3}).$$

By minimizing the r.h.s of this estimate with respect to ε, this leads to the error bound when taking $\varepsilon = |\pi|^{\frac{3}{10}} \leq 1$:

$$\sup_{x\in\mathbb{R}^d}(v - v^\pi)(t_k,x) \;\leq\; C|\pi|^{\frac{1}{10}}.$$

Finally, by combining with the estimate in Lemma 5.2.4, which gives actually under **(H2')**(i):

$$|\vartheta^\pi(t,x,a) - v^\pi(t_{k+1},x)| \;\leq\; C|\pi|^{\frac{1}{2}}, \quad t\in[t_k,t_{k+1}),(x,a)\in\mathbb{R}^d\times A,$$

together with the 1/2-Hölder property of v in time (see (5.2.12)), we obtain:

$$\sup_{(t,x,a)\in[0,T]\times\mathbb{R}^d\times A}(v - \vartheta^\pi)(t,x,a) \;\leq\; C(|\pi|^{\frac{1}{10}} + |\pi|^{\frac{1}{2}}) \;\leq\; C|\pi|^{\frac{1}{10}}.$$

This ends the proof. □

5.2.4 Approximation scheme for jump-constrained BSDE and stochastic control problem

We consider the discrete time approximation of the discretely jump-constrained BSDE in the case where $f(x,a,y)$ does not depend on z, and define the scheme $(\bar{Y}^\pi, \bar{\mathcal{Y}}^\pi, \bar{\mathcal{Z}}^\pi)$ by induction on the grid $\pi = \{t_0 = 0 < \ldots < t_k < \ldots < t_n = T\}$ by:

$$\left\{\begin{array}{rcl} \bar{Y}^\pi_T = \bar{\mathcal{Y}}^\pi_T &=& g(\bar{X}^\pi_T)\\ \bar{\mathcal{Y}}^\pi_{t_k} &=& \mathbb{E}_{t_k}[\bar{Y}^\pi_{t_{k+1}}] + f(\bar{X}^\pi_{t_k}, I_{t_k}, \bar{\mathcal{Y}}^\pi_{t_k})\Delta t_k\\ \bar{Y}^\pi_{t_k} &=& \operatorname{ess\,sup}_{a\in A}\mathbb{E}_{t_k,a}[\bar{\mathcal{Y}}^\pi_{t_k}], \quad k = 0,\ldots,n-1,\end{array}\right. \tag{5.2.45}$$

where $\Delta t_k = t_{k+1} - t_k$, $\mathbb{E}_{t_k}[.]$ stands for $\mathbb{E}[.|\mathcal{F}_{t_k}]$, and $\mathbb{E}_{t_k,a}[.]$ for $\mathbb{E}[.|\mathcal{F}_{t_k}, I_{t_k} = a]$.

By induction argument, we easily see that $\bar{\mathcal{Y}}^\pi_{t_k}$ is a deterministic function of $(\bar{X}^\pi_{t_k}, I_{t_k})$, while $\bar{Y}^\pi_{t_k}$ is a deterministic function of $\bar{X}^\pi_{t_k}$, for $k = 0,\ldots,n$, and by the Markov property of the process $(\bar{X}^\pi, I)$, the conditional expectations in (5.2.45) can be replaced by the corresponding regressions:

$$\mathbb{E}_{t_k}[\bar{Y}^\pi_{t_{k+1}}] \;=\; \mathbb{E}[\bar{Y}^\pi_{t_{k+1}}|\bar{X}^\pi_{t_k}, I_{t_k}] \quad \text{and} \quad \mathbb{E}_{t_k,a}[\bar{\mathcal{Y}}^\pi_{t_k}] \;=\; \mathbb{E}[\bar{\mathcal{Y}}^\pi_{t_k}|\bar{X}^\pi_{t_k}, I_{t_k} = a].$$

We then have:

$$\bar{\mathcal{Y}}^\pi_{t_k} \;=\; \bar{\vartheta}^\pi_k(\bar{X}^\pi_{t_k}, I_{t_k}), \qquad Y^\pi_{t_k} \;=\; \bar{v}^\pi_k(\bar{X}^\pi_{t_k}),$$

for some sequence of functions $(\bar{\vartheta}^\pi_k)_k$ and $(\bar{v}^\pi_k)_k$ defined respectively on $\mathbb{R}^d \times A$ and $\mathbb{R}^d$ by backward induction:

$$\left\{ \begin{array}{rcl} \bar{v}^\pi_n(x,a) \;=\; \bar{\vartheta}^\pi_n(x) & = & g(x) \\ \bar{\vartheta}^\pi_k(x,a) & = & \mathbb{E}\Big[\bar{v}^\pi_{k+1}(\bar{X}^\pi_{t_{k+1}}, I_{t_{k+1}}) | (\bar{X}^\pi_{t_k}, I_{t_k}) = (x,a)\Big] + f(x,a,\bar{\vartheta}^\pi_k(x,a))\Delta t_k \\ \bar{v}^\pi_k(x) & = & \sup_{a\in A} \bar{\vartheta}^\pi_k(x,a), \quad k = 0,\ldots,n-1. \end{array} \right. \tag{5.2.46}$$

There are well-known different methods (Longstaff-Schwartz least square regression, quantization, Malliavin integration by parts, see e.g. [7], [82], [27]) for computing the above conditional expectations, and so the functions $\bar{\vartheta}^\pi_k$ and $\bar{v}^\pi_k$. It appears that in our context, the simulation-regression method on basis functions defined on $\mathbb{R}^d \times A$, is quite suitable since it allows us to derive at each time step $k \leq n-1$, a functional form $\hat{a}_k(x)$, which attains the supremum over A in $\bar{\vartheta}^\pi_k(x,a)$. We shall see later in this section that the feedback control $(\hat{a}_k(x))_k$ provides an approximation of the optimal control for the HJB equation associated to a stochastic control problem when $f(x,a)$ does not depend on y. The practical details about the computation of functions $\bar{\vartheta}^\pi_k$, $\bar{v}^\pi_k$, $\hat{a}_k$ by simulation-regression methods will be provided in Section 5.3.

5.2.4.1 Error estimate for the discrete time scheme

The main result of this section is to state an error bound between the component Y^π of the discretely jump-constrained BSDE and the solution $(\bar{Y}^\pi, \bar{\mathcal{Y}}^\pi)$ to the above discrete time scheme.

Theorem 5.2.2. *There exists some constant C such that:*

$$\mathbb{E}\Big[|Y^\pi_{t_k} - \bar{Y}^\pi_{t_k}|^2\Big] + \sup_{t\in(t_k,t_{k+1}]} \mathbb{E}\Big[|Y^\pi_t - \bar{Y}^\pi_{t_{k+1}}|^2\Big] + \sup_{t\in[t_k,t_{k+1})} \mathbb{E}\Big[|Y^\pi_t - \bar{\mathcal{Y}}^\pi_{t_k}|^2\Big] \;\leq\; C|\pi|,$$

for all $k = 0,\ldots,n-1$.

The above convergence rate $|\pi|^{\frac{1}{2}}$ in the L^2-norm for the discretization of the discretely jump-constrained BSDE is the same as for standard BSDE, see [27], [109]. By combining with the convergence result in Section 5.2.3, we finally obtain an estimate on the error due to the discrete time approximation of the minimal solution Y to the BSDE with nonpositive jumps. We split the error between the positive and negative parts:

$$\mathrm{Err}^\pi_+(Y) \;:=\; \max_{k\leq n-1}\Big(\mathbb{E}\Big[(Y_{t_k} - \bar{Y}^\pi_{t_k})^2_+\Big] + \sup_{t\in(t_k,t_{k+1}]} \mathbb{E}\Big[(Y_t - \bar{Y}^\pi_{t_{k+1}})^2_+\Big] + \sup_{t\in[t_k,t_{k+1})} \mathbb{E}\Big[(Y_t - \bar{\mathcal{Y}}^\pi_{t_k})^2_+\Big]\Big)^{\frac{1}{2}}$$

$$\mathrm{Err}^\pi_-(Y) \;:=\; \max_{k\leq n-1}\Big(\mathbb{E}\Big[(Y_{t_k} - \bar{Y}^\pi_{t_k})^2_-\Big] + \sup_{t\in(t_k,t_{k+1}]} \mathbb{E}\Big[(Y_t - \bar{Y}^\pi_{t_{k+1}})^2_-\Big] + \sup_{t\in[t_k,t_{k+1})} \mathbb{E}\Big[(Y_t - \bar{\mathcal{Y}}^\pi_{t_k})^2_-\Big]\Big)^{\frac{1}{2}}.$$

Corollary 5.2.4. *We have:*

$$\mathrm{Err}^\pi_-(Y) \;\leq\; C|\pi|^{\frac{1}{2}}.$$

*Moreover, under **(H1')** and **(H2')**,*

$$\mathrm{Err}^\pi_+(Y) \;\leq\; C|\pi|^{\frac{1}{10}},$$

and when $f\,(x,a)$ does not depend on y:

$$\mathrm{Err}^\pi_+(Y) \;\leq\; C|\pi|^{\frac{1}{6}}.$$

Proof. Recall from Proposition 5.2.5 that $\mathcal{Y}_t^\pi \le Y_t^\pi \le Y_t$, $0 \le t \le T$. Then, we have: $(Y_{t_k} - \bar{Y}_{t_k}^\pi)_- \le |Y_{t_k}^\pi - \bar{Y}_{t_k}^\pi|$, $(Y_t - \bar{Y}_{t_{k+1}}^\pi)_- \le |Y_t^\pi - \bar{Y}_{t_{k+1}}^\pi|$, and $(Y_{t_k} - \bar{\mathcal{Y}}_{t_k}^\pi)_- \le |Y_{t_k}^\pi - \bar{\mathcal{Y}}_{t_k}^\pi|$, for all $k \le n-1$, and $t \in [0,T]$. The error bound on $\mathrm{Err}_-^\pi(Y)$ follows then from the estimation in Theorem 5.2.2. The error bound on $\mathrm{Err}_-^\pi(Y)$ follows from Corollary 5.2.3 and Theorem 5.2.2. □

Remark 5.2.4. In the particular case where f depends only on (x,a), our discrete time approximation scheme is a probabilistic scheme for the fully nonlinear HJB equation associated with the stochastic control problem (5.1.2). As in [78], [10] or [49], we have non symmetric bounds on the rate of convergence. For instance, in [49], the authors obtained a convergence rate $|\pi|^{\frac{1}{4}}$ on one side and $|\pi|^{\frac{1}{10}}$ on the other side, while we improve the rate to $|\pi|^{\frac{1}{2}}$ for one side, and $|\pi|^{\frac{1}{6}}$ on the other side. This induces a global error $\mathrm{Err}^\pi(Y) = \mathrm{Err}_+^\pi(Y) + \mathrm{Err}_-^\pi(Y)$ of order $|\pi|^{\frac{1}{6}}$, which is derived without any non degeneracy condition on the controlled diffusion coefficient.

Proof. of Theorem 5.2.2. Let us introduce the continuous time version of (5.2.45). By the martingale representation theorem, there exists $\tilde{\mathcal{Z}}^\pi \in L^2(W)$ and $\tilde{\mathcal{U}}^\pi \in L^2(\tilde{\mu})$ such that

$$\bar{Y}_{t_{k+1}}^\pi = \mathbb{E}_{t_k}[\bar{Y}_{t_{k+1}}^\pi] + \int_{t_k}^{t_{k+1}} \tilde{\mathcal{Z}}_t^\pi dW_t + \int_{t_k}^{t_{k+1}} \int_A \tilde{\mathcal{U}}_t^\pi(a)\tilde{\mu}(dt,da), \quad k < n,$$

and we can then define the continuous-time processes $\bar{Y}^\pi$ and $\bar{\mathcal{Y}}^\pi$ by:

$$\begin{aligned} \bar{\mathcal{Y}}_t^\pi &= \bar{Y}_{t_{k+1}}^\pi + (t_{k+1}-t) f(\bar{X}_{t_k}^\pi, I_{t_k}, \bar{\mathcal{Y}}_{t_k}^\pi) \\ &\quad - \int_t^{t_{k+1}} \tilde{\mathcal{Z}}_t^\pi dW_t - \int_t^{t_{k+1}} \int_A \tilde{\mathcal{U}}_t^\pi(a)\tilde{\mu}(dt,da), \qquad t \in [t_k, t_{k+1}), \end{aligned} \tag{5.2.47}$$

$$\begin{aligned} \bar{Y}_t^\pi &= \bar{Y}_{t_{k+1}}^\pi + (t_{k+1}-t) f(\bar{X}_{t_k}^\pi, I_{t_k}, \bar{\mathcal{Y}}_{t_k}^\pi) \\ &\quad - \int_t^{t_{k+1}} \tilde{\mathcal{Z}}_t^\pi dW_t - \int_t^{t_{k+1}} \int_A \tilde{\mathcal{U}}_t^\pi(a)\tilde{\mu}(dt,da), \qquad t \in (t_k, t_{k+1}], \end{aligned} \tag{5.2.48}$$

for $k = 0, \dots, n-1$. Denote by $\delta Y_t^\pi = Y_t^\pi - \bar{Y}_t^\pi$, $\delta \mathcal{Y}_t^\pi = \mathcal{Y}_t^\pi - \bar{\mathcal{Y}}_t^\pi$, $\delta \mathcal{Z}_t^\pi = \mathcal{Z}_t^\pi - \tilde{\mathcal{Z}}_t^\pi$, $\delta \mathcal{U}_t^\pi = \mathcal{U}_t^\pi - \tilde{\mathcal{U}}_t^\pi$ and $\delta f_t = f(X_t, I_t, \mathcal{Y}_t^\pi) - f(\bar{X}_{t_k}^\pi, I_{t_k}, \bar{\mathcal{Y}}_{t_k}^\pi)$ for $t \in [t_k, t_{k+1})$. Recalling (5.2.14) and (5.2.47), we have by Itô's formula:

$$\begin{aligned} \Delta_t &:= \mathbb{E}_{t_k}\Big[|\delta \mathcal{Y}_t^\pi|^2 + \int_t^{t_{k+1}} |\delta \mathcal{Z}_s^\pi|^2 ds + \int_t^{t_{k+1}} \int_A |\delta \mathcal{U}_s^\pi(a)|^2 \lambda(da) ds\Big] \\ &= \mathbb{E}_{t_k}[|\delta Y_{t_{k+1}}^\pi|^2|] + \mathbb{E}_{t_k}\Big[\int_t^{t_{k+1}} 2\delta \mathcal{Y}_s^\pi \delta f_s\Big] ds \end{aligned}$$

for all $t \in [t_k, t_{k+1})$. By the Lipschitz continuity of f in **(H2)** and Young inequality, we then have:

$$\begin{aligned} \Delta_t &\le \mathbb{E}_{t_k}[|\delta Y_{t_{k+1}}^\pi|^2|] + \mathbb{E}_{t_k}\Big[\int_t^{t_{k+1}} \eta |\delta \mathcal{Y}_s^\pi|^2 ds + \frac{C}{\eta}\pi |\delta \mathcal{Y}_{t_k}^\pi|^2\Big] \\ &\quad + \frac{C}{\eta}\mathbb{E}_{t_k}\Big[\int_t^{t_{k+1}} (|X_s - \bar{X}_{t_k}^\pi|^2 + |I_s - I_{t_k}|^2 + |\mathcal{Y}_s^\pi - \mathcal{Y}_{t_k}^\pi|^2) ds\Big]. \end{aligned}$$

From Gronwall's lemma, and by taking η large enough, this yields for all $k \le n-1$:

$$\mathbb{E}_{t_k}\Big[|\delta \mathcal{Y}_{t_k}^\pi|^2\Big] \le e^{C|\pi|}\mathbb{E}_{t_k}[|\delta Y_{t_{k+1}}^\pi|^2|] + CB_k \tag{5.2.49}$$

where

$$\begin{aligned} B_k &= \mathbb{E}_{t_k}\Big[\int_{t_k}^{t_{k+1}} (|X_s - \bar{X}^\pi_{t_k}|^2 + |I_s - I_{t_k}|^2 + |\mathcal{Y}^\pi_s - \mathcal{Y}^\pi_{t_k}|^2)ds\Big] \\ &\le C|\pi|\Big(\mathbb{E}_{t_k}[\sup_{s\in[t_k,t_{k+1}]} |X_s - \bar{X}^\pi_{t_k}|^2] + |\pi|(1+|X_{t_k}|)\Big), \end{aligned} \tag{5.2.50}$$

by (5.2.5) and Proposition 5.2.3. Now, by definition of $Y^\pi_{t_{k+1}}$ and $\bar{Y}^\pi_{t_{k+1}}$, we have

$$|\delta Y^\pi_{t_{k+1}}|^2 \le \operatorname*{ess\,sup}_{a\in A} \mathbb{E}_{t_{k+1},a}[|\delta\mathcal{Y}^\pi_{t_{k+1}}|^2]. \tag{5.2.51}$$

By plugging (5.2.50), (5.2.51) into (5.2.49), taking conditional expectation with respect to $I_{t_k} = a$, and taking essential supremum over a, we obtain:

$$\begin{aligned} \operatorname*{ess\,sup}_{a\in A} \mathbb{E}_{t_k,a}\Big[|\delta\mathcal{Y}^\pi_{t_k}|^2\Big] &\le e^{C|\pi|}\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{t_k,a}\big[\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{t_{k+1},a}[|\delta\mathcal{Y}^\pi_{t_{k+1}}|^2]\big] \\ &\quad + C|\pi|\Big(\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{t_k,a}[\sup_{s\in[t_k,t_{k+1}]}|X_s - \bar{X}^\pi_{t_k}|^2] + |\pi|(1+|X_{t_k}|)\Big). \end{aligned}$$

By taking conditional expectation with respect to $\mathcal{F}_{t_{k-1}}$, and $I_{t_{k-1}} = a$, taking essential supremum over a in the above inequality, and iterating this backward procedure until time $t_0 = 0$, we obtain:

$$\begin{aligned} \mathcal{E}^\pi_k(\mathcal{Y}) &\le e^{C|\pi|}\mathcal{E}^\pi_{k+1}(\mathcal{Y}) + C|\pi|(\mathcal{E}^\pi_k(X) + |\pi|(1+\mathbb{E}[|X_{t_k}|])) \\ &\le e^{C|\pi|}\mathcal{E}^\pi_{k+1}(\mathcal{Y}) + C|\pi|^2, \quad k \le n-1, \end{aligned} \tag{5.2.52}$$

where we recall the auxiliary error control $\mathcal{E}^\pi_k(X)$ on X in (5.2.2) and its estimate in Lemma 5.2.1, and set:

$$\mathcal{E}^\pi_k(\mathcal{Y}) := \mathbb{E}\Big[\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{t_1,a}[\dots\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{t_k,a}[|\delta\mathcal{Y}^\pi_{t_k}|^2]\dots]\Big].$$

By a direct induction on (5.2.52), and recalling that $n|\pi|$ is bounded, we get

$$\begin{aligned} \mathcal{E}^\pi_k(\mathcal{Y}) &\le C(\mathcal{E}^\pi_n(\mathcal{Y}) + |\pi|) \\ &\le C(\mathcal{E}^\pi_n(X) + |\pi|) \le C|\pi|, \end{aligned}$$

since g is Lipschitz, and using again the estimate in Lemma 5.2.1. Observing that $\mathbb{E}[|\delta Y^\pi_{t_k}|^2]$, $\mathbb{E}[|\delta\mathcal{Y}^\pi_{t_k}|^2] \le \mathcal{E}^\pi_k(\mathcal{Y})$, we get the estimate:

$$\max_{k\le n}\mathbb{E}[|Y^\pi_{t_k} - \bar{Y}^\pi_{t_k}|^2] + \mathbb{E}[|\mathcal{Y}^\pi_{t_k} - \bar{\mathcal{Y}}^\pi_{t_k}|^2] \le C|\pi|.$$

Moreover, by Proposition 5.2.3, we have

$$\begin{aligned} \sup_{t\in[t_k,t_{k+1})}\mathbb{E}\Big[|\mathcal{Y}^\pi_t - \mathcal{Y}^\pi_{t_k}|^2\Big] + \sup_{t\in(t_k,t_{k+1}]}\mathbb{E}\Big[|Y^\pi_t - Y^\pi_{t_{k+1}}|^2\Big] &\le C(1+\mathbb{E}[|X_{t_k}|])|\pi| \\ &\le C(1+|X_0|)|\pi|. \end{aligned}$$

This implies finally that:

$$\begin{aligned} \sup_{s\in(t_k,t_{k+1}]}\mathbb{E}\Big[|Y^\pi_t - \bar{Y}^\pi_{t_{k+1}}|^2\Big] &\le 2\sup_{s\in(t_k,t_{k+1}]}\mathbb{E}\Big[|Y^\pi_t - Y^\pi_{t_{k+1}}|^2\Big] + 2\mathbb{E}\Big[|Y^\pi_{t_{k+1}} - \bar{Y}^\pi_{t_{k+1}}|^2\Big] \\ &\le C|\pi|, \end{aligned}$$

as well as

$$\begin{aligned} \sup_{s\in[t_k,t_{k+1})}\mathbb{E}\Big[|Y^\pi_t - \bar{\mathcal{Y}}^\pi_{t_k}|^2\Big] &\le 2\sup_{s\in[t_k,t_{k+1})}\mathbb{E}\Big[|Y^\pi_t - \mathcal{Y}^\pi_{t_k}|^2\Big] + 2\mathbb{E}\Big[|\mathcal{Y}^\pi_{t_k} - \bar{\mathcal{Y}}^\pi_{t_k}|^2\Big] \\ &\le C|\pi|. \end{aligned}$$

□

5.2.4.2 Approximate optimal control

We now consider the special case where $f(x,a)$ does not depend on y, so that the discrete time scheme (5.1.4) is an approximation for the value function of the stochastic control problem:

$$\begin{aligned} V_0 &:= \sup_{\alpha\in\mathcal{A}} J(\alpha) = Y_0, \\ J(\alpha) &= \mathbb{E}\Big[\int_0^T f(X_t^\alpha,\alpha_t)dt + g(X_T^\alpha)\Big], \end{aligned} \tag{5.2.53}$$

where $\mathcal{A}$ is the set of $\mathbb{G}$-adapted control processes α valued in A, and X^α is the controlled diffusion in $\mathbb{R}^d$:

$$X_t^\alpha = X_0 + \int_0^t b(X_s^\alpha,\alpha_s)ds + \int_0^t \sigma(X_s^\alpha,\alpha_s)dW_s, \quad 0\le t\le T.$$

(Here $\mathbb{G} = (\mathcal{G}_t)_{0\le t\le T}$ denotes some filtration under which W is a standard Brownian motion). Let us now define the discrete time version of (5.2.53). We introduce the set $\mathcal{A}^\pi$ of discrete time processes $\alpha = (\alpha_{t_k})_k$ with α_{t_k} $\mathcal{G}_{t_k}$-measurable, and valued in A. For each $\alpha\in\mathcal{A}^\pi$, we consider the controlled discrete time process $(X_{t_k}^{\pi,\alpha})_k$ of Euler type defined by:

$$X_{t_k}^{\pi,\alpha} = X_0 + \sum_{j=0}^k b(X_{t_j}^{\pi,\alpha},\alpha_{t_j})\Delta t_j + \sum_{j=0}^k \sigma(X_{t_j}^{\pi,\alpha},\alpha_{t_j})\Delta W_{t_j}, \quad k\le n,$$

where $\Delta W_{t_j} = W_{t_{j+1}} - W_{t_j}$, and the gain functional:

$$J^\pi(\alpha) = \mathbb{E}\Big[\sum_{k=0}^{n-1} f(X_{t_k}^{\pi,\alpha},\alpha_{t_k})\Delta t_k + g(X_{t_n}^{\pi,\alpha})\Big].$$

Given any $\alpha\in\mathcal{A}^\pi$, we define its continuous time piecewise-constant interpolation $\alpha\in\mathcal{A}$ by setting: $\alpha_t = \alpha_{t_k}$, for $t\in[t_k,t_{k+1})$ (by misuse of notation, we keep the same notation α for the discrete time and continuous time interpolation). By standard arguments similar to those for Euler scheme of SDE, there exists some positive constant C such that for all $\alpha\in\mathcal{A}^\pi$, $k\le n-1$:

$$\mathbb{E}\Big[\sup_{t\in[t_k,t_{k+1}]} |X_t^\alpha - X_{t_k}^{\pi,\alpha}|^2\Big] \le C|\pi|,$$

from which we easily deduce by Lipschitz property of f and g:

$$|J(\alpha) - J^\pi(\alpha)| \le C|\pi|^{\frac{1}{2}}, \quad \forall\alpha\in\mathcal{A}^\pi. \tag{5.2.54}$$

Let us now consider at each time step $k\le n-1$, the function $\hat{a}_k(x)$ which attains the supremum over $a\in A$ of $\bar{\vartheta}_k^\pi(x,a)$ in the scheme (5.2.46), so that:

$$\bar{v}_k^\pi(x) = \bar{\vartheta}_k^\pi(x,\hat{a}_k(x)), \quad k=0,\dots,n-1.$$

Let us define the process $(\hat{X}_{t_k}^\pi)_k$ by: $\hat{X}_0^\pi = X_0$,

$$\hat{X}_{t_{k+1}}^\pi = \hat{X}_{t_k}^\pi + b(\hat{X}_{t_k}^\pi,\hat{a}_k(\hat{X}_{t_k}^\pi))\Delta t_k + \sigma(\hat{X}_{t_k}^\pi,\hat{a}_k(\hat{X}_{t_k}^\pi))\Delta W_{t_k}, \quad k\le n,$$

and notice that $\hat{X}^\pi = X^{\pi,\hat{\alpha}}$, where $\hat{\alpha}\in\mathcal{A}^\pi$ is a feedback control defined by:

$$\hat{\alpha}_{t_k} = \hat{a}_k(\hat{X}_{t_k}^\pi) = \hat{a}_k(X_{t_k}^{\pi,\hat{\alpha}}), \quad k=0,\dots,n.$$

Next, we observe that the conditional law of $\bar{X}^\pi_{t_{k+1}}$ given $(\bar{X}^\pi_{t_k} = x, I_{t_k} = \hat{a}_k(\bar{X}^\pi_{t_k}) = \hat{a}_k(x))$ is the same than the conditional law of $X^{\pi,\hat{\alpha}}_{t_{k+1}}$ given $X^{\pi,\hat{\alpha}}_{t_k} = x$, for $k \leq n-1$, and thus the induction step in the scheme (5.2.45) or (5.2.46) reads as:

$$\bar{v}^\pi_k(X^{\pi,\hat{\alpha}}_{t_k}) \;=\; \mathbb{E}\Big[\bar{v}^\pi_{k+1}(X^{\pi,\hat{\alpha}}_{t_{k+1}})\big|X^{\pi,\hat{\alpha}}_{t_k}\Big] + f(X^{\pi,\hat{\alpha}}_{t_k}, \hat{\alpha}_{t_k})\Delta t_k, \quad k \leq n-1.$$

By induction, and law of iterated conditional expectations, we then get:

$$\bar{Y}^\pi_0 \;=\; \bar{v}^\pi_0(X_0) \;=\; J^\pi(\hat{\alpha}). \tag{5.2.55}$$

Consider the continuous time piecewise-constant interpolation $\hat{\alpha} \in \mathcal{A}$ defined by: $\hat{\alpha}_t = \hat{\alpha}_{t_k}$, for $t \in [t_k, t_{k+1})$. By (5.2.53), (5.2.55), (5.2.54), and Corollary 5.2.4, we finally obtain:

$$\begin{aligned} 0 \;\leq\; V_0 - J(\hat{\alpha}) \;&=\; Y_0 - \bar{Y}^\pi_0 + J^\pi(\hat{\alpha}) - J(\hat{\alpha}) \\ &\leq\; C|\pi|^{\frac{1}{6}} + C|\pi|^{\frac{1}{2}} \;\leq\; C|\pi|^{\frac{1}{6}}, \end{aligned}$$

for $|\pi| \leq 1$. In other words, for any small $\varepsilon > 0$, we obtain an ε-approximate control $\hat{\alpha}$ for the stochastic control problem (5.2.53) by taking $|\pi|$ of order ε^6.

5.3 Approximation of conditional expectations

So far, we proposed and analyzed the discrete-time scheme (5.1.4) to approximate the minimal solution of the BSDE with nonpositive jumps (5.1.3), and thus an approximation scheme for the HJB equation (5.1.1), under the assumptions **(H1)** (b and σ Lipschitz) and **(H2)** (f and g Lipschitz). Now, the discrete scheme (5.1.4) is in general not readily implementable because it involves conditional expectations that cannot be computed explicitly. It is thus necessary in practice to approximate these conditional expectations. Here we follow the empirical regression approach ([82, 28, 58, 108, 2]). In our context, apart from being easy to implement, the strong advantage of this choice is that, unlike other standard methods, it provides as a by-product a parametric feedback estimate $\hat{\alpha}(t, X_t)$ for the optimal control. The idea is to replace the conditional expectations from (5.1.4) by empirical regressions. This section is devoted to the analysis of the error generated by this replacement.

5.3.1 Localizations

The first step is to localize the discrete BSDE (5.1.4), i.e. to truncate it so that it admits a.s. deterministic bounds. Introduce $R_X \in \mathbb{R}^d_+$ and $R_w \in \mathbb{R}_+$ and define the following truncations of X_i and ΔW_i:

$$[X_i]_X := -R_X \vee X_i \wedge R_X = \{-R_{1,X} \vee X_{1,i} \wedge R_{1,X}, \ldots, -R_{d,X} \vee X_{d,i} \wedge R_{d,X}\}^\top \tag{5.3.1}$$

$$[\Delta W_i]_w := -R_w\sqrt{\Delta_i} \vee \Delta W_i \wedge R_w\sqrt{\Delta_i} = \Big\{-R_w\sqrt{\Delta_i} \vee \Delta W_{1,i} \wedge R_w\sqrt{\Delta_i}, \ldots, -R_w\sqrt{\Delta_i} \vee \Delta W_{q,i} \wedge R_w\sqrt{\Delta_i}\Big\}^\top \tag{5.3.2}$$

Define $R = \{R_X, R_w\}$ and define the localized version of the discrete BSDE (5.1.4), using the truncations (5.3.1) and (5.3.2).

$$\begin{cases} Y^R_N & = g\left([X_N]_X\right) \\ \Delta_i \mathcal{Z}^R_i & = \mathbb{E}_i\left[Y^R_{i+1}\left[\Delta W_i^\top\right]_w\right] \\ \mathcal{Y}^R_i & = \mathbb{E}_i\left[Y^R_{i+1} + f\left([X_i]_X, I_i, Y^R_{i+1}, \mathcal{Z}^R_i\right)\Delta_i\right] \\ Y^R_i & = \operatorname{ess\,sup}_{a\in A} \mathbb{E}_{i,a}\left[\mathcal{Y}^R_i\right] \end{cases} \tag{5.3.3}$$

First, we check that this localized BSDE does admit a.s. bounds.

Lemma 5.3.1. *[almost sure bounds] For every $R = \{R_X, R_w\} \in [0,\infty)^d \times [0,\infty]$ and every $1 \leq i \leq N$, the following uniform bounds hold a.s.:*

$$\begin{aligned}\left|\mathcal{Y}_i^R\right|, \left|Y_i^R\right| &\leq C_y = C_y(R_X) := e^{\frac{C}{2}T}\sqrt{C_g^2(R_X) + \frac{e^{C|\pi|}}{L_f^2}C_f^2(R_X)} \\ \left|\mathcal{Z}_i^R\right|, \left|Z_i^R\right| &\leq C_z = C_y(R_X) := \frac{\sqrt{q}}{\sqrt{\Delta_i}}C_y\end{aligned}$$

where $C := 3L_f^2(q+|\pi|)+\frac{1}{q}$, $C_g(R_X) := \max_{-R_X\leq x\leq R_X}|g(x)|$ and $C_f(R_X) := L_f\left(|R_X| + \left|\bar{A}\right|\right) + f(0,0,0,0)$ ($\bar{A}$ is an upper bound of the compact set A, L_f is the Lipschitz constant of f and L_g is the Lipschitz constant of g).

Proof. First, as g is continuous, there exists $C_g = C_g(R_X) > 0$ such that for all $-R_X \leq x \leq R_X$, $|g(x)| \leq C_g(R_X)$. Hence

$$\left(Y_N^R\right)^2 = g^2\left([X_N]_X\right) \leq C_g^2(R_X) \tag{5.3.4}$$

Next,

$$\Delta_i\mathcal{Z}_i^R = \mathbb{E}_i\left[Y_{i+1}^R[\Delta W_i]_w\right] = \mathbb{E}_i\left[\left(Y_{i+1}^R - \mathbb{E}_i\left[Y_{i+1}^R\right]\right)[\Delta W_i]_w\right]$$

Thus, using the Cauchy-Schwarz inequality and dividing by Δ_i:

$$\Delta_i\left(\mathcal{Z}_i^R\right)^2 \leq q\left(\mathbb{E}_i\left[\left(Y_{i+1}^R\right)^2\right] - \mathbb{E}_i\left[Y_{i+1}^R\right]^2\right)$$

Now, using Young's inequality $(a+b)^2 \leq (1+\gamma\Delta_i)a^2 + \left(1+\frac{1}{\gamma\Delta_i}\right)b^2$ with $\gamma > 0$:

$$\left(\mathcal{Y}_i^R\right)^2 \leq (1+\gamma\Delta_i)\,\mathbb{E}_i\left[Y_{i+1}^R\right]^2 + \left(1+\frac{1}{\gamma\Delta_i}\right)\Delta_i^2\mathbb{E}_i\left[f^2\left([X_i]_X, I_i, Y_{i+1}^R, \mathcal{Z}_i^R\right)\right]$$

Remark that

$$\begin{aligned}\left|f\left([X_i]_X, I_i, Y_{i+1}^R, \mathcal{Z}_i^R\right)\right| &\leq \left|f\left([X_i]_X, I_i, Y_{i+1}^R, \mathcal{Z}_i^R\right) - f(0,0,0,0)\right| + |f(0,0,0,0)| \\ &\leq L_f\left(|[X_i]_X| + |I_i| + \left|Y_{i+1}^R\right| + \left|\mathcal{Z}_i^R\right|\right) + |f(0,0,0,0)| \\ &\leq C_f(R_X) + L_f\left(\left|Y_{i+1}^R\right| + \left|\mathcal{Z}_i^R\right|\right)\end{aligned}$$

where $C_f(R_X) := L_f\left(|R_X| + \left|\bar{A}\right|\right) + |f(0,0,0,0)|$. Hence

$$\begin{aligned}\left(\mathcal{Y}_i^R\right)^2 &\leq (1+\gamma\Delta_i)\,\mathbb{E}_i\left[Y_{i+1}^R\right]^2 + 3\left(\Delta_i + \frac{1}{\gamma}\right)\Delta_i\left(C_f^2(R_X) + L_f^2\mathbb{E}_i\left[\left(Y_{i+1}^R\right)^2\right] + L_f^2\mathbb{E}_i\left[\left|\mathcal{Z}_i^R\right|\right]\right) \\ &\leq \left(\Delta_i + \frac{1}{\gamma}\right)\left(\mathbb{E}_i\left[Y_{i+1}^R\right]^2\left(\gamma - 3qL_f^2\right) + \mathbb{E}_i\left[\left(Y_{i+1}^R\right)^2\right]3L_f^2(\Delta_i - q) + 3C_f^2(R_X)\Delta_i\right)\end{aligned}$$

Thus, for every $\gamma \geq 3qL_f^2$, one can group together the terms involving $\mathbb{E}_i\left[Y_{i+1}^R\right]^2$ and $\mathbb{E}_i\left[\left(Y_{i+1}^R\right)^2\right]$ using Jensen's inequality:

$$\left(\mathcal{Y}_i^R\right)^2 \leq (1+\theta(3,\gamma)\Delta_i)\,\mathbb{E}_i\left[\left(Y_{i+1}^R\right)^2\right] + 3\left(|\pi| + \frac{1}{\gamma}\right)C_f^2(R_X)\Delta_i$$

where $\theta(c,\gamma) := \gamma + cL_f^2\left(|\pi| + \frac{1}{\gamma}\right)$. Finally:

$$\left(Y_i^R\right)^2 \leq \operatorname*{ess\,sup}_{a\in A} \mathbb{E}_{i,a}\left[\left(Y_{i+1}^R\right)^2\right](1+\theta(3,\gamma)\Delta_i) + 3\left(|\pi| + \frac{1}{\gamma}\right) C_f^2(R_X)\Delta_i \tag{5.3.5}$$

Using equations (5.3.4) and (5.3.5), one obtains by induction that:

$$\left(Y_i^R\right)^2 \leq \Gamma_i^{N-1}(3,\gamma) C_g^2(R_X) + 3\left(|\pi| + \frac{1}{\gamma}\right) C_f^2(R_X) \sum_{k=i}^{N-1} \Gamma_i^k(3,\gamma)\Delta_k \tag{5.3.6}$$

where $\Gamma_i^j(c,\gamma) := \Pi_{k=i}^j(1+\theta(c,\gamma)\Delta_k)$. Finally remark that $\forall c,\gamma > 0$

$$\ln\left(\Gamma_i^j(c,\gamma)\right) = \sum_{k=i}^j \ln(1+\theta(c,\gamma)\Delta_k) \leq \sum_{k=i}^j \theta(c,\gamma)\Delta_k = \theta(c,\gamma)(t_{j+1}-t_i)$$

Thus

$$\Gamma_i^j(c,\gamma) \leq \exp(\theta(c,\gamma)(t_{j+1}-t_i)) \tag{5.3.7}$$

And

$$\begin{aligned}\sum_{k=i}^{N-1}\Gamma_i^k(c,\gamma)\Delta_k &\leq \sum_{k=i}^{N-1} e^{\theta(c,\gamma)(t_{j+1}-t_i)}\Delta_k \leq e^{\theta(c,\gamma)|\pi|}\sum_{k=i}^{N-1} e^{\theta(c,\gamma)(t_j-t_i)}\Delta_k \\ &\leq e^{\theta(c,\gamma)|\pi|}\int_{t_i}^{t_N} e^{\theta(c,\gamma)(t-t_i)}dt = \frac{e^{\theta(c,\gamma)|\pi|}}{\theta(c,\gamma)}\left(e^{\theta(c,\gamma)(t_N-t_i)}-1\right)\end{aligned} \tag{5.3.8}$$

Finally, combine equations (5.3.6), (5.3.7) and (5.3.8) with $c=3$ and $\gamma \geq 3qL_f^2$ to obtain the following a.s. bound for Y_i^R :

$$\begin{aligned}\left(Y_i^R\right)^2 &\leq e^{\theta(3,\gamma)(t_N-t_i)}C_g^2(R_X) + 3\left(|\pi|+\frac{1}{\gamma}\right)C_f^2(R_X)\frac{e^{\theta(3,\gamma)|\pi|}}{\theta(3,\gamma)}\left(e^{\theta(3,\gamma)(t_N-t_i)}-1\right) \\ &\leq e^{\theta(3,\gamma)T}\left\{C_g^2(R_X) + 3\frac{e^{\theta(3,\gamma)|\pi|}}{\theta(3,\gamma)}\left(|\pi|+\frac{1}{\gamma}\right)C_f^2(R_X)\right\}\end{aligned}$$

In particular, for $c=3$ and $\gamma = 3qL_f^2$:

$$\left(Y_i^R\right)^2 \leq e^{CT}\left\{C_g^2(R_X) + \frac{e^{C|\pi|}}{L_f^2}C_f^2(R_X)\right\} =: C_y^2$$

where $C := 3L_f^2(q+|\pi|) + \frac{1}{q}$. The same inequality holds for $\left(\mathcal{Y}_i^R\right)^2$. For $\mathcal{Z}_i^R$, use the Cauchy-Schwarz inequality to obtain:

$$\left(\mathcal{Z}_i^R\right)^2 \leq \frac{q}{\Delta_i}\mathbb{E}_i\left[\left(Y_{i+1}^R\right)^2\right] \leq \frac{q}{\Delta_i}C_y^2 =: C_z^2$$

and the same inequality holds for $\left(Z_i^R\right)^2$. □

Lemma 5.3.2. *For $R>0$, define $\mathcal{T}_R = \mathbb{E}\left[(\mathcal{N} - (-R)\vee\mathcal{N}\wedge R)^2\right]$ where $\mathcal{N}$ is a Gaussian random variable with mean 0 and variance 1. Then:*

$$\mathcal{T}_R \leq \sqrt{\frac{2}{\pi}}\frac{1}{R}e^{-\frac{R^2}{2}}$$

Proof. Developing the square yields

$$\mathcal{T}_R = 2R^2\mathbb{P}\left(\mathcal{N} > R\right) - 4R\mathbb{E}\left[\mathcal{N}\mathbf{1}\left\{\mathcal{N} > R\right\}\right] + 2\mathbb{E}\left[\mathcal{N}^2\mathbf{1}\left\{\mathcal{N} > R\right\}\right]$$

Then the two expectations can be explicited as follows

$$\begin{aligned}\mathbb{E}\left[\mathcal{N}\mathbf{1}\left\{\mathcal{N} > R\right\}\right] &= \frac{e^{-\frac{R^2}{2}}}{\sqrt{2\pi}} \\ \mathbb{E}\left[\mathcal{N}^2\mathbf{1}\left\{\mathcal{N} > R\right\}\right] &= \frac{R}{\sqrt{2\pi}}e^{-\frac{R^2}{2}} + \mathbb{P}\left(\mathcal{N} > R\right)\end{aligned}$$

Finally, the use of Mill's ratio inequality $\mathbb{P}\left(\mathcal{N} > R\right) < \frac{1}{R}\frac{e^{-\frac{R^2}{2}}}{\sqrt{2\pi}}$ concludes the proof. □

Then, we can estimate bounds between the BSDEs (5.1.4) and (5.3.3).

Proposition 5.3.1. *The following bounds hold:*

$$\left(Y_i - Y_i^R\right)^2 \leq e^{CT}\left\{L_g\left(|\Delta X_N|^2\right)^* + C\sum_{k=i}^{N-1}\Delta_k\left(|\Delta X_k|^2\right)^* + 2qCTC_y^2\mathcal{T}_{R_w}\right\}$$

where $C := 3L_f^2\left(2q + |\pi|\right) + \frac{1}{2q}$, and $\left(|\Delta X_k|^2\right)^*$, $k \geq i$, is the solution of the following linear constrained BSDE:

$$\begin{cases} Y_k &= \left(X_k - [X_k]_X\right)^2 \\ Y_j &= \operatorname{ess\,sup}_{a\in A}\mathbb{E}_{j,a}\left[Y_{j+1}\right], \ j = k-1,\dots,i \end{cases}$$

Proof. Define $\Delta X_i = X_i - [X_i]_X$, $\Delta Y_i = Y_i - Y_i^R$, $\Delta\mathcal{Y}_i = \mathcal{Y}_i - \mathcal{Y}_i^R$, $\Delta Z_i = Z_i - Z_i^R$ and $\Delta\mathcal{Z}_i = \mathcal{Z}_i - \mathcal{Z}_i^R$. First

$$|\Delta Y_N| = |g\left(X_N\right) - g\left([X_N]_X\right)| \leq L_g\left|\Delta X_N^p\right|$$

Then

$$\begin{aligned}\Delta_i\Delta\mathcal{Z}_i &= \mathbb{E}_i\left[Y_{i+1}\Delta W_i^\top - Y_{i+1}^R\left[\Delta W_i^\top\right]_w\right] \\ &= \mathbb{E}_i\left[\Delta Y_{i+1}\Delta W_i^\top + Y_{i+1}^R\left\{\Delta W_i - [\Delta W_i]_w\right\}^\top\right] \\ &\quad \mathbb{E}_i\left[\left(\Delta Y_{i+1} - \mathbb{E}_i\left[\Delta Y_{i+1}\right]\right)\Delta W_i^\top\right] + \mathbb{E}_i\left[Y_{i+1}^R\left\{\Delta W_i - [\Delta W_i]_w\right\}^\top\right]\end{aligned}$$

Hence

$$\Delta_i\left(\Delta\mathcal{Z}_i\right)^2 \leq 2q\left(\mathbb{E}_i\left[\left(\Delta Y_{i+1}\right)^2\right] - \mathbb{E}_i\left[\Delta Y_{i+1}\right]^2\right) + 2qC_y^2\mathcal{T}_{R_w}$$

Then

$$\Delta\mathcal{Y}_i = \mathbb{E}_i\left[\Delta Y_{i+1} + \left\{f\left(X_i, I_i, Y_{i+1}, \mathcal{Z}_i\right) - f\left([X_i]_X, I_i, Y_{i+1}^R, \mathcal{Z}_i^R\right)\right\}\Delta_i\right]$$

Using Jensen's inequality and Young's inequality with parameter $\gamma\Delta_i$, $\gamma > 0$:

$$\begin{aligned}\left(\Delta\mathcal{Y}_i\right)^2 &\leq \left(1 + \gamma\Delta_i\right)\mathbb{E}_i\left[\Delta Y_{i+1}\right]^2 + \left(1 + \frac{1}{\gamma\Delta_i}\right)\Delta_i^2 3L_f^2\mathbb{E}_i\left[\left(\Delta X_i\right)^2 + \left(\Delta Y_{i+1}\right)^2 + \left(\Delta\mathcal{Z}_i\right)^2\right] \\ &\leq \mathbb{E}_i\left[\Delta Y_{i+1}\right]^2\left(\Delta_i + \frac{1}{\gamma}\right)\left(\gamma - 6qL_f^2\right) + \mathbb{E}_i\left[\left(\Delta Y_{i+1}\right)^2\right]\left(\Delta_i + \frac{1}{\gamma}\right)3L_f^2\left(\Delta_i + 2q\right) \\ &\quad + \left(\Delta_i + \frac{1}{\gamma}\right)\Delta_i 3L_f^2\left\{\left(\Delta X_i\right)^2 + 2qC_y^2\mathcal{T}_{R_w}\right\}\end{aligned}$$

Now, for any $\gamma \geq 6qL_f^2$, one can group together the terms in $\mathbb{E}_i\left[\Delta Y_{i+1}\right]^2$ and $\mathbb{E}_i\left[\left(\Delta Y_{i+1}\right)^2\right]$ using Jensen's inequality:

$$\left(\Delta \mathcal{Y}_i\right)^2 \leq \mathbb{E}_i\left[\left(\Delta Y_{i+1}\right)^2\right]\left\{1+\theta\left(3,\gamma\right)\Delta_i\right\} + 3L_f^2\left(\left|\pi\right|+\frac{1}{\gamma}\right)\Delta_i\left\{\left(\Delta X_i\right)^2 + 2qC_y^2\mathcal{T}_{R_w}\right\}$$

where, as in Lemma 5.3.1, $\theta\left(c,\gamma\right) := \gamma + cL_f^2\left(\left|\pi\right|+\frac{1}{\gamma}\right)$. Hence, using that for any random variables Θ and Θ',

$$\left(\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{i,a}\left[\Theta\right] - \sup_{a\in A}\mathbb{E}_{i,a}\left[\Theta'\right]\right)^2 \leq \operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{i,a}\left[\left(\Theta-\Theta'\right)^2\right],$$

the following holds:

$$\left(\Delta Y_i\right)^2 \leq \left\{1+\theta\left(3,\gamma\right)\Delta_i\right\}\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{i,a}\left[\left(\Delta Y_{i+1}\right)^2\right] + 3L_f^2\left(\left|\pi\right|+\frac{1}{\gamma}\right)\Delta_i\left\{\left(\Delta X_i\right)^2 + 2qC_y^2\mathcal{T}_{R_w}\right\}$$

By induction

$$\begin{aligned}\left(\Delta Y_i\right)^2 \quad &\leq \quad L_g\Gamma_i^{N-1}\left(3,\gamma\right)\left(\left|\Delta X_N\right|^2\right)^* \\ &\quad +3L_f^2\left(\left|\pi\right|+\frac{1}{\gamma}\right)\sum_{k=i}^{N-1}\Delta_k\Gamma_i^k\left(3,\gamma\right)\left\{\left(\left|\Delta X_k\right|^2\right)^* + 2qC_y^2\mathcal{T}_{R_w}\right\}\end{aligned}$$

where, as in Lemma 5.3.1, $\Gamma_i^j\left(c,\gamma\right) := \Pi_{k=i}^j\left(1+\theta\left(c,\gamma\right)\Delta_k\right) \leq \exp\left(\theta\left(c,\gamma\right)\left(t_{j+1}-t_i\right)\right)$. Finally, take $\gamma = 6qL_f^2$ to obtain the desired bound. □

5.3.2 Projections

In its current form, the scheme (5.3.3) is not readily implementable, because its conditional expectations cannot be computed in general. Therefore, there is a need to approximate these conditional expectations. For handiness and efficiency, we choose, in the spirit of [82] and [58], to approximate them by empirical least-squares regression.

First, we will study the impact of the replacement of the conditional expectations by theoretical least-squares regressions. We will see that the resulting scheme is not easy to analyze. Therefore, we will study a stronger version of it, and discuss their practical differences. As it is already a daunting task for standard BSDEs (cf. [82]), and in view of the difficulties already raised at theoretical regression level, we leave the study of the final replacement of these regressions by their empirical counterparts for further research.

Hence, for each $i \in \{0,\ldots,N-1\}$, consider $\mathcal{S}_i^Y$ and $\mathcal{S}_i^Z = \left\{\mathcal{S}_i^{Z,1},\ldots,\mathcal{S}_i^{Z,q}\right\}$ that are non-empty closed convex subsets of $\mathbf{L}_2\left(\mathcal{F}_{t_i},\mathbb{P}\right)$, as well as the corresponding projection operators $\mathcal{P}_i^Y$ and $\mathcal{P}_i^Z = \left\{\mathcal{P}_i^{Z,1},\ldots,\mathcal{P}_i^{Z,q}\right\}$. Using the above projection operators in lieu of the conditional expectations in (5.3.3), we obtain the following approximation scheme:

$$\begin{cases}\tilde{Y}_N^R & = g\left(\left[X_N\right]_X\right) \\ \Delta_i\tilde{\mathcal{Z}}_i^R & = \left[\mathcal{P}_i^Z\left(\tilde{Y}_{i+1}^R\left[\Delta W_i^\top\right]_w\right)\right]_{i,z} \\ \tilde{\mathcal{Y}}_i^R & = \left[\mathcal{P}_i^Y\left(\tilde{Y}_{i+1}^R + f\left(\left[X_i\right]_X, I_i, \tilde{Y}_{i+1}^R, \tilde{\mathcal{Z}}_i^R\right)\Delta_i\right)\right]_y \\ \tilde{Y}_i^R & = \operatorname{ess\,sup}_{a\in A}\mathbb{E}_{i,a}\left[\tilde{\mathcal{Y}}_i^R\right]\end{cases} \tag{5.3.9}$$

where $[.]_{i,z} := -\Delta_i C_z \wedge . \vee \Delta_i C_z$ and $[.]_y := -C_y \wedge . \vee C_y$ are truncation operators that ensure that the a.s. upper bounds for $\left(Y^R, Z^R\right)$ from Lemma 5.3.1 will also hold for $\left(\tilde{Y}^R, \tilde{Z}^R\right)$.

To be more specific, choose the subsets $\mathcal{S}_i^Y$ and $\mathcal{S}_i^Z$ as follows:

$$\begin{aligned}\mathcal{S}_i^Y &= \left\{\lambda.p_i^Y\left(X_i, I_i\right) \,;\, \lambda \in \mathbb{R}^{B_i^Y}\right\} \\ \mathcal{S}_i^{Z,k} &= \left\{\lambda.p_i^{Z,k}\left(X_i, I_i\right) \,;\, \lambda \in \mathbb{R}^{B_i^{Z,k}}\right\},\ k = 1, \ldots, q\end{aligned}$$

where $p_i^Y = \left(p_{i,1}^Y, \ldots, p_{i,B_i^Y}^Y\right)^\top$, $B_i^Y \geq 1$, and $p_i^{Z,k} = \left(p_{i,1}^{Z,k}, \ldots, p_{i,B_i^{Z,k}}^{,Z,k}\right)^\top$, $B_i^{Z,k} \geq 1$, are predefined sets of deterministic functions from $\mathbb{R}^d \times \mathbb{R}^q$ into $\mathbb{R}$. Hence, for any random variable U in $\mathbf{L}_2\left(\mathcal{F}_T, \mathbb{P}\right)$, $\mathcal{P}_i^Y\left(U\right)$ is defined as follows:

$$\begin{aligned}\hat{\lambda}_i^Y\left(U\right) &:= \arg \inf_{\lambda \in \mathbb{R}^{B_i^Y}} \mathbb{E}\left[\left(\lambda.p_i^Y\left(X_i, I_i\right) - U\right)^2\right] \\ \mathcal{P}_i^Y\left(U\right) &:= \hat{\lambda}_i^Y\left(U\right).p_i^Y\left(X_i, I_i\right)\end{aligned} \tag{5.3.10}$$

and $\mathcal{P}_i^Z\left(U\right)$ is defined in a similar manner. With these notations, the scheme (5.3.9) can be explicited further as follows:

$$\begin{cases}\tilde{Y}_N^R &= g\left(\left[X_N\right]_X\right) \\ \Delta_i \tilde{\mathcal{Z}}_i^R &= \left[\hat{\lambda}_i^Z\left(\tilde{Y}_{i+1}^R\left[\Delta W_i^\top\right]_w\right).p_i^Z\left(X_i, I_i\right)\right]_{i,z} \\ \tilde{Y}_i^R &= \operatorname{ess\,sup}_{a \in \mathcal{A}_i}\left[\hat{\lambda}_i^Y\left(\tilde{Y}_{i+1}^R + f\left(\left[X_i\right]_X, I_i, \tilde{Y}_{i+1}^R, \tilde{\mathcal{Z}}_i^R\right)\Delta_i\right).p_i^Y\left(X_i, a\right)\right]_y,\end{cases} \tag{5.3.11}$$

where $\mathcal{A}_i$ is the set of $\sigma\left(X_i\right)$-measurable random variables taking values in A. Now, we would like to analyze the error between $\left(Y^R, Z^R\right)$ and $\left(\tilde{Y}^R, \tilde{Z}^R\right)$. Unfortunately, in spite of the simplicity of the scheme (5.3.11), this analysis is made strenuous by the fact that $\tilde{Y}_i^R$ is not itself a projection, as it combines regression coefficients computed using the random variable I_i and regression functions valued at another random variable a. This prevents the analysis from taking advantage of standard tools to deal with least-squares regressions. For comparison, consider the following alternative scheme:

$$\begin{cases}\hat{Y}_N^R &= g\left(\left[X_N\right]_X\right) \\ \Delta_i \hat{\mathcal{Z}}_{i,a}^R &= \left[\hat{\lambda}_{i,a}^Z\left(\hat{Y}_{i+1}^R\left[\Delta W_i^\top\right]_w\right).p_i^Z\left(X_i, I_i\right)\right]_{i,z},\ a \in \mathcal{A}_i \\ \hat{Y}_i^R &= \operatorname{ess\,sup}_{a \in \mathcal{A}_i}\left[\hat{\lambda}_{i,a}^Y\left(\hat{Y}_{i+1}^R + f\left(\left[X_i\right]_X, I_i, \hat{Y}_{i+1}^R, \hat{\mathcal{Z}}_{i,a}^R\right)\Delta_i\right).p_i^Y\left(X_i, a\right)\right]_y\end{cases} \tag{5.3.12}$$

where, unlike equation (5.3.10), the regression coefficients $\hat{\lambda}_{i,a}^Y$ are computed as follows:

$$\begin{aligned}\hat{\lambda}_{i,a}^Y\left(U\right) &:= \arg \inf_{\lambda \in \mathbb{R}^{B_i^Y}} \mathbb{E}\left[\left(\lambda.p_i^Y\left(X_i, a\right) - U_a\right)^2\right] \\ \mathcal{P}_{i,a}^Y\left(U\right) &:= \hat{\lambda}_{i,a}^Y\left(U\right).p_i^Y\left(X_i, a\right)\end{aligned} \tag{5.3.13}$$

for every $U \in \mathbf{L}_2\left(\mathcal{F}_T, \mathbb{P}\right)$ and $a \in \mathcal{A}_i$, where U_a corresponds to the conditional random variable $U\,|\{I_i = a\}$. $\mathcal{P}_{i,a}^Z\left(U\right)$ is defined in a similar manner. Remark that $\mathcal{P}_{i,I_i}\left(U\right) = \mathcal{P}_i\left(U\right)$, and that $\mathcal{P}_{i,a}\left(U_a\right) = \mathcal{P}_{i,a}\left(U\right)$. With this new scheme, the estimated regression coefficients are changed along with the strategy a when computing the optimal strategy. Therefore, compared with the scheme (5.3.11), the implementation of an empirical version of the scheme (5.3.12) is much more involved, as it may require, for the same time step, many regressions involving several random

variables a different from I_i (which is used to simulate the forward process). However, these modifications ease considerably the analysis of the impact of the projections compared with $\left(Y^R, Z^R\right)$ as shown below in the remaining of this subsection.

First, the scheme (5.3.12) can be written as follows:

$$\begin{cases} \hat{Y}_N^R & = g\left([X_N]_X\right) \\ \Delta_i \hat{\mathcal{Z}}_{i,a}^R & = \left[\mathcal{P}_{i,a}^Z\left(\hat{Y}_{i+1}^R\left[\Delta W_i^\top\right]_w\right)\right]_{i,z} \\ \hat{\mathcal{Y}}_{i,a}^R & = \left[\mathcal{P}_{i,a}^Y\left(\hat{Y}_{i+1}^R + f\left([X_i]_X, I_i, \hat{Y}_{i+1}^R, \hat{\mathcal{Z}}_{i,a}^R\right)\Delta_i\right)\right]_y \\ \hat{Y}_i^R & = \operatorname{ess\,sup}_{a\in A}\hat{\mathcal{Y}}_{i,a}^R \end{cases} \tag{5.3.14}$$

Then, we recall below some useful properties of the projection operators $\mathcal{P}_{i,a}^{\cdot}$.

Lemma 5.3.3. *For any fixed $a \in \mathcal{A}_i$:*

$$\mathcal{P}_{i,a}^{\cdot}(U) = \mathcal{P}_{i,a}^{\cdot}\left(\mathbb{E}_{i,a}[U]\right),\ \forall U \in \mathbf{L}_2\left(\mathcal{F}_{t_i}, \mathbb{P}\right) \tag{5.3.15}$$

$$\mathbb{E}\left[\left(\mathcal{P}_{i,a}^{\cdot}(U) - \mathcal{P}_{i,a}^{\cdot}(V)\right)^2\right] \leq \mathbb{E}\left[(U_a - V_a)^2\right],\ \forall U, V \text{ in } \mathbf{L}_2\left(\mathcal{F}_T, \mathbb{P}\right). \tag{5.3.16}$$

Proof. The proof can be found in [58]. □

We now assess the error between $\left(Y^R, Z^R\right)$ and $\left(\hat{Y}^R, \hat{Z}^R\right)$.

Proposition 5.3.2. *[projection error] The following bound holds:*

$$\mathbb{E}\left[\left|Y_i^R - \hat{Y}_i^R\right|^2\right],\ \Delta_i\mathbb{E}\left[\left|Z_i^R - \hat{Z}_i^R\right|^2\right] \leq e^{C(T-t_i)}\sum_{k=i}^{N-1}\left\{\mathbb{E}\left[\left(\left|\Delta\mathcal{P}\mathcal{Y}_k\right|^2\right)^*\right] + C\Delta_k\mathbb{E}\left[\left(\left|\Delta\mathcal{P}\mathcal{Z}_k\right|^2\right)^*\right]\right\}$$

where $C := 2L_f^2\left(|\pi| + q\right) + \frac{1}{q}$, and $\left(\left|\Delta\mathcal{P}\mathcal{Y}_k\right|^2\right)^*$ (resp. $\left(\left|\Delta\mathcal{P}\mathcal{Z}_k\right|^2\right)^*$), $k \geq i$, is solution of the linear constrained BSDE:

$$\begin{cases} Y_k & = \operatorname{ess\,sup}_{a\in A}\mathbb{E}_{k,a}\left[\left|\mathcal{Y}_k^R - \mathcal{P}_k^Y\left(\mathcal{Y}_k^R\right)\right|^2\right],\ \left(\text{resp. } \operatorname{ess\,sup}_{a\in A}\mathbb{E}_{k,a}\left[\left|\mathcal{Z}_k^R - \mathcal{P}_k^Z\left(\mathcal{Z}_k^R\right)\right|^2\right]\right) \\ Y_j & = \operatorname{ess\,sup}_{a\in A}\mathbb{E}_{j,a}\left[Y_{j+1}\right],\ j = k-1, \ldots, i \end{cases}$$

Moreover, the same upper bound holds for $\mathbb{E}\left[\operatorname{ess\,sup}_{a\in A}\left|\mathcal{Y}_{i,a}^R - \hat{\mathcal{Y}}_{i,a}^R\right|^2\right]$ and $\Delta_i\mathbb{E}\left[\operatorname{ess\,sup}_{a\in A}\left|\mathcal{Z}_{i,a}^R - \hat{\mathcal{Z}}_{i,a}^R\right|^2\right]$.

Proof. Fix $a \in \mathcal{A}_i$. Define $\Delta Y_i^R = Y_i^R - \hat{Y}_i^R$, $\Delta\mathcal{Y}_{i,a}^R = \mathcal{Y}_{i,a}^R - \hat{\mathcal{Y}}_{i,a}^R$, $\Delta Z_i^R = Z_i^R - \hat{Z}_i^R$ and $\Delta\mathcal{Z}_{i,a}^R = \mathcal{Z}_{i,a}^R - \hat{\mathcal{Z}}_{i,a}^R$, where, as in equation (5.3.13), $\mathcal{Y}_{i,a}^R$ (resp. $\mathcal{Z}_{i,a}^R$) stands for the conditional variable $\mathcal{Y}_i^R\,|\{I_i = a\}$ (resp. $\mathcal{Z}_i^R\,|\{I_i = a\}$).

First, using that $\Delta_i\mathcal{Z}_{i,a}^R = \left[\Delta_i\mathcal{Z}_{i,a}^R\right]_{i,z}$ and the 1-Lipschitz property of $[.]_{i,z}$:

$$\left|\Delta_i\Delta\mathcal{Z}_{i,a}^R\right|^2 \leq \left|\Delta_i\mathcal{Z}_{i,a}^R - \mathcal{P}_{i,a}^Z\left(\hat{Y}_{i+1}^R\left[\Delta W_i^\top\right]_w\right)\right|^2$$

Using Pythagoras' theorem:

$$\mathbb{E}\left[\left|\Delta_i\Delta\mathcal{Z}_{i,a}^R\right|^2\right] = \mathbb{E}\left[\left|\Delta_i\mathcal{Z}_{i,a}^R - \mathcal{P}_{i,a}^Z\left(\Delta_i\mathcal{Z}_{i,a}^R\right)\right|^2\right] + \mathbb{E}\left[\left|\mathcal{P}_{i,a}^Z\left(\Delta_i\mathcal{Z}_{i,a}^R\right) - \mathcal{P}_{i,a}^Z\left(\hat{Y}_{i+1}^R\left[\Delta W_i^\top\right]_w\right)\right|^2\right]$$

where, using equation (5.3.15):

$$\begin{aligned}&\mathcal{P}_{i,a}^{Z}\left(\hat{Y}_{i+1}^{R}\left[\Delta W_i^{\top}\right]_w\right)=\mathcal{P}_{i,a}^{Z}\left(\mathbb{E}_{i,a}\left[\hat{Y}_{i+1}^{R}\left[\Delta W_i^{\top}\right]_w\right]\right)\\=&\mathcal{P}_{i,a}^{Z}\left(\mathbb{E}_{i,a}\left[\left(\hat{Y}_{i+1}^{R}-\mathbb{E}_{i,a}\left[\hat{Y}_{i+1}^{R}\right]\right)\left[\Delta W_i^{\top}\right]_w\right]\right)\end{aligned}$$

Then, using equation (5.3.16):

$$\begin{aligned}\mathbb{E}\left[\left|\mathcal{P}_{i,a}^{Z}\left(\Delta_i\mathcal{Z}_{i,a}^{R}\right)-\mathcal{P}_{i,a}^{Z}\left(\hat{Y}_{i+1}^{R}\left[\Delta W_i^{\top}\right]_w\right)\right|^2\right]&\leq\mathbb{E}\left[\left|\mathbb{E}_{i,a}\left[\Delta_i\mathcal{Z}_{i,a}^{R}\right]-\mathbb{E}_{i,a}\left[\left(\hat{Y}_{i+1}^{R}-\mathbb{E}_{i,a}\left[\hat{Y}_{i+1}^{R}\right]\right)\left[\Delta W_i^{\top}\right]_w\right]\right|^2\right]\\&=\mathbb{E}\left[\left|\mathbb{E}_{i,a}\left[\left(\Delta Y_{i+1}^{R}-\mathbb{E}_{i,a}\left[\Delta Y_{i+1}^{R}\right]\right)\left[\Delta W_i^{\top}\right]_w\right]\right|^2\right]\\&\leq q\Delta_i\mathbb{E}\left[\mathbb{E}_{i,a}\left[\left(\Delta Y_{i+1}^{R}\right)^2\right]-\mathbb{E}_{i,a}\left[\Delta Y_{i+1}^{R}\right]^2\right]\end{aligned}$$

To sum up for the Z component:

$$\Delta_i\mathbb{E}\left[\left|\Delta\mathcal{Z}_{i,a}^{R}\right|^2\right]\leq\Delta_i\mathbb{E}\left[\left|\mathcal{Z}_{i,a}^{R}-\mathcal{P}_{i,a}^{Z}\left(\mathcal{Z}_{i,a}^{R}\right)\right|^2\right]+q\mathbb{E}\left[\mathbb{E}_{i,a}\left[\left(\Delta Y_{i+1}^{R}\right)^2\right]-\mathbb{E}_{i,a}\left[\Delta Y_{i+1}^{R}\right]^2\right]$$

For the Y component, start similarly by using the 1-Lipschitz property of $[.]_y$ and Pythagoras' theorem:

$$\mathbb{E}\left[\left|\Delta\mathcal{Y}_{i,a}^{R}\right|^2\right]=\mathbb{E}\left[\left(\mathcal{Y}_{i,a}^{R}-\mathcal{P}_{i,a}^{Y}\left(\mathcal{Y}_{i,a}^{R}\right)\right)^2\right]+\mathbb{E}\left[\left(\mathcal{P}_{i,a}^{Y}\left(\mathcal{Y}_{i,a}^{R}\right)-\mathcal{P}_{i,a}^{Y}\left(\hat{Y}_{i+1}^{R}+f\left(\left[X_i\right]_X,I_i,\hat{Y}_{i+1}^{R},\hat{\mathcal{Z}}_{i,a}^{R}\right)\Delta_i\right)\right)^2\right]$$

And then, using again equations (5.3.15), (5.3.16), Jensen's inequality and Young's inequality with parameter $\gamma\Delta_i$, $\gamma>0$:

$$\begin{aligned}&\mathbb{E}\left[\left(\mathcal{P}_{i,a}^{Y}\left(\mathcal{Y}_{i,a}^{R}\right)-\mathcal{P}_{i,a}^{Y}\left(\hat{Y}_{i+1}^{R}+f\left(\left[X_i\right]_X,I_i,\hat{Y}_{i+1}^{R},\hat{\mathcal{Z}}_{i,a}^{R}\right)\Delta_i\right)\right)^2\right]\\&\leq\mathbb{E}\left[\left(\mathbb{E}_{i,a}\left[\Delta Y_{i+1}^{R}+L_f\left(\left|\Delta Y_{i+1}^{R}\right|+\left|\Delta\mathcal{Z}_{i,a}^{R}\right|\right)\Delta_i\right]\right)^2\right]\\&\leq\mathbb{E}\left[\left(1+\gamma\Delta_i\right)\mathbb{E}_{i,a}\left[\Delta Y_{i+1}^{R}\right]^2+\left(1+\frac{1}{\gamma\Delta_i}\right)\Delta_i^2L_f^2 2\left\{\mathbb{E}_{i,a}\left[\left(\Delta Y_{i+1}^{R}\right)^2\right]+\mathbb{E}_{i,a}\left[\left|\Delta\mathcal{Z}_{i,a}^{R}\right|^2\right]\right\}\right]\\&\leq\left(\Delta_i+\frac{1}{\gamma}\right)\mathbb{E}\left[\left(\gamma-2qL_f^2\right)\mathbb{E}_{i,a}\left[\Delta Y_{i+1}^{R}\right]^2+2L_f^2\left(\Delta_i+q\right)\mathbb{E}_{i,a}\left[\left(\Delta Y_{i+1}^{R}\right)^2\right]+2L_f^2\Delta_i\left|\mathcal{Z}_{i,a}^{R}-\mathcal{P}_{i,a}^{Z}\left(\mathcal{Z}_{i,a}^{R}\right)\right|^2\right]\end{aligned}$$

For all $\gamma\geq 2qL_f^2$, one can group together the terms involving $\mathbb{E}_{i,a}\left[Y_{i+1}^{R}\right]^2$ and $\mathbb{E}_{i,a}\left[\left(Y_{i+1}^{R}\right)^2\right]$ using Jensen's inequality:

$$\begin{aligned}\mathbb{E}\left[\left|\Delta\mathcal{Y}_{i,a}^{R}\right|^2\right]\leq&\mathbb{E}\left[\left(\mathcal{Y}_{i,a}^{R}-\mathcal{P}_{i,a}^{Y}\left(\mathcal{Y}_{i,a}^{R}\right)\right)^2\right]+\left(1+\theta\left(2,\gamma\right)\Delta_i\right)\mathbb{E}\left[\mathbb{E}_{i,a}\left[\left(\Delta Y_{i+1}^{R}\right)^2\right]\right]\\&+2L_f^2\left(|\pi|+\frac{1}{\gamma}\right)\Delta_i\mathbb{E}\left[\left|\mathcal{Z}_{i,a}^{R}-\mathcal{P}_{i,a}^{Z}\left(\mathcal{Z}_{i,a}^{R}\right)\right|^2\right]\\\leq&\mathbb{E}\left[\operatorname*{ess\,sup}_{a\in A}\left(\mathcal{Y}_{i,a}^{R}-\mathcal{P}_{i,a}^{Y}\left(\mathcal{Y}_{i,a}^{R}\right)\right)^2\right]+\left(1+\theta\left(2,\gamma\right)\Delta_i\right)\mathbb{E}\left[\operatorname*{ess\,sup}_{a\in A}\mathbb{E}_{i,a}\left[\left(\Delta Y_{i+1}^{R}\right)^2\right]\right]\\&+2L_f^2\left(|\pi|+\frac{1}{\gamma}\right)\Delta_i\mathbb{E}\left[\operatorname*{ess\,sup}_{a\in A}\left|\mathcal{Z}_{i,a}^{R}-\mathcal{P}_{i,a}^{Z}\left(\mathcal{Z}_{i,a}^{R}\right)\right|^2\right]\end{aligned}\tag{5.3.17}$$

where $\theta\left(c,\gamma\right)=\gamma+cL_f^2\left(|\pi|+\frac{1}{\gamma}\right)$.

Therefore, as equation (5.3.17) is true for every $a \in \mathcal{A}_i$ on its left-hand side, and as $\left|\Delta Y_{i+1}^R\right|^2 \leq \operatorname{ess\,sup}_{a\in A}\left|\Delta \mathcal{Y}_{i+1,a}^R\right|^2$, the following holds by induction:

$$\mathbb{E}\left[\operatorname*{ess\,sup}_{a\in A}\left|\Delta \mathcal{Y}_{i,a}^R\right|^2\right] \leq \sum_{k=i}^{N-1} \Gamma_i^k(2,\gamma)\left\{\mathbb{E}\left[\left(\left|\Delta \mathcal{P}\mathcal{Y}_k\right|^2\right)^*\right] + 2L_f^2\left(|\pi| + \frac{1}{\gamma}\right)\Delta_k \mathbb{E}\left[\left(\left|\Delta \mathcal{P}\mathcal{Z}_k\right|^2\right)^*\right]\right\}$$

where $\Gamma_i^j(c,\gamma) = \Pi_{k=i}^j\left(1+\theta(c,\gamma)\Delta_k\right) \leq \exp\left(\theta(c,\gamma)\left(t_{j+1}-t_i\right)\right)$. Finally, take $\gamma = 2qL_f^2$ to obtain the desired bound for $\left|\Delta \mathcal{Y}_{i,a}^R\right|^2$. Moreover, as $\left|\Delta Y_i^R\right|^2 \leq \sup_{a\in A}\left|\Delta \mathcal{Y}_{i,a}^R\right|^2$, the same bound holds for $\left|\Delta Y_i^R\right|^2$. For the bound on $\left|\Delta Z_i^R\right|^2$, use that:

$$\Delta_i \mathbb{E}\left[\left|\Delta Z_i^R\right|^2\right] \leq \Delta_i \mathbb{E}\left[\operatorname*{ess\,sup}_{a\in A}\left|\Delta \mathcal{Z}_{i,a}^R\right|^2\right] \leq \Delta_i \mathbb{E}\left[\operatorname*{ess\,sup}_{a\in A}\left|\mathcal{Z}_{i,a}^R - \mathcal{P}_{i,a}^Z\left(\mathcal{Z}_{i,a}^R\right)\right|^2\right] + \mathbb{E}\left[\operatorname*{ess\,sup}_{a\in A}\left|\Delta \mathcal{Y}_{i+1,a}^R\right|^2\right]$$

□

5.4 Applications

In this section, we illustrate our numerical scheme on various examples.

5.4.1 Linear Quadratic stochastic control problem

The first application is an example of a linear-quadratic stochastic control problem. We consider the following problem:

$$v(t,x) = \sup_{\alpha\in\mathcal{A}} \mathbb{E}\left[-\lambda_0 \int_t^T (\alpha_s)^2\, ds - \lambda_1 \left(X_T^\alpha\right)^2\right] \tag{5.4.1}$$

$$dX_s^\alpha = \left(-\mu_0 X_s^\alpha + \mu_1 \alpha_s\right) dt + \left(\sigma_0 + \sigma_1 \alpha_s\right) dW_s\,,\ X_0^\alpha = 0 \tag{5.4.2}$$

where $\lambda_i, \mu_i, \sigma_i > 0$, $i = 1, 2$. It is called linear-quadratic because the drift and the volatility of X^α are linear in α and X^α, while the terms in the objective function v are quadratic in α and X^α. We choose this example as a first, simple application for our numerical scheme because there exists analytical solutions to this class of stochastic control problem (cf. [107]) to which our results can be compared in order to assess the accuracy of our method.

Now, let us look closer to this specific example. As can be seen from equation (5.4.1), the objective function v penalizes the terminal value X_T^α of the controlled diffusion if it is away from zero (with the $-\lambda_1 \left(X_T^\alpha\right)^2$ term). Hence, X^α, which starts from zero, has to be controlled carefully over time so as not divert too much from this initial value. This can be achieved through the control α in the drift term $\left(-\mu_0 X_s^\alpha + \mu_1 \alpha_s\right)$, which can reinforce the default mean-reversion speed μ_0. However, this control also impacts the volatility $\left(\sigma_0 + \sigma_1 \alpha_s\right)$, which makes it easier to decrease X^α than to increase it. Moreover, the controls are penalized over time $\left(-\lambda_0 \int_t^T (\alpha_s)^2\, ds\right)$, meaning that they must be exerted parsimoniously.

We test our numerical scheme on this specific problem. We set the parameters to the following values:

λ_0	λ_1	μ_0	μ_1	σ_0	σ_1	T
20	200	0.02	0.5	0.2	0.1	2

For the numerical parameters, we use $n = 52$ time-discretization steps, and a sample of $M = 10^6$ Monte Carlo simulations. For the regressions, we use a basis function of global polynomial of degree two:

$$\phi(t, x, \alpha) = \beta_0 + \beta_1 x + \beta_2 \alpha + \beta_3 x\alpha + \beta_4 x^2 + \beta_5 \alpha^2 .$$

In particular, assuming $\beta_5 < 0$, the optimal control will be linear w.r.t. x:

$$\alpha^* = \alpha^*(t, x) := \arg\max_{\alpha} \phi(t, x, \alpha) = A(t)\, x + B(t)$$

$$A(t) := -\frac{\beta_3}{2\beta_5} \quad , \quad B(t) := -\frac{\beta_2}{2\beta_5}$$

This behaviour is illustrated on Figure 5.4.1 below.

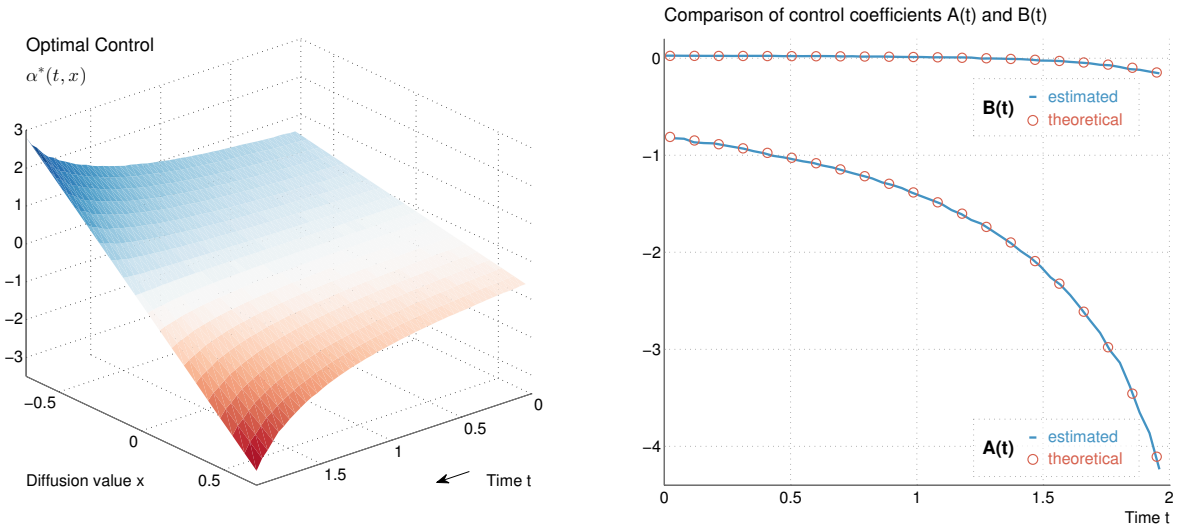

(a) Shape

(b) Optimal coefficients vs. theoretical values

Figure 5.4.1: Optimal control

Figure 5.4.1a displays the shape of the optimal control $\alpha^*(t, x)$.

First, as expected from the drift term in the dynamics of X^α (equation (5.4.2)), α^* is a decreasing function of x ($A(t) \leq 0$):

- If X_t^α takes a large positive value, then $\alpha^*(t, X_t)$ will take a large negative value so as to push it back more quickly to zero (recall the drift term $-\mu_0 X_s^\alpha + \mu_1 \alpha_s$).

- Conversely, if X_t^α takes a large negative value, then $\alpha^*(t, X_t)$ will take a large positive value for the same reason.

Second, the strength of the control increases as time reaches maturity (i.e. $A(t)$ decreases with t). Indeed, the penalization of the control becomes relatively cheaper compared with the penalization of the final value when time is close to maturity.

The strengthening of the control can also be assessed on Figure 5.4.1b, which displays the time evolution of the estimated coefficients A and B ($\alpha^*(t, x) = A(t)\, x + B(t)$). Moreover, one can see that the coefficient B is slightly negative close to maturity. This creates an asymmetry in the control (as $\alpha^*(t, 0) = B(t) \neq 0$), which comes from the asymmetric effect of the control on the volatility of X^α.

The effect of the optimal control α^* is clearly visible on Figure 5.4.2 below, which compares the distribution of X^α without control (Figure 5.4.2a) and when the optimal control is used

(Figure5.4.2b). The strengthening of the control at the end of the time period, as well as the slightly asymmetric shape of the distribution are prominent.

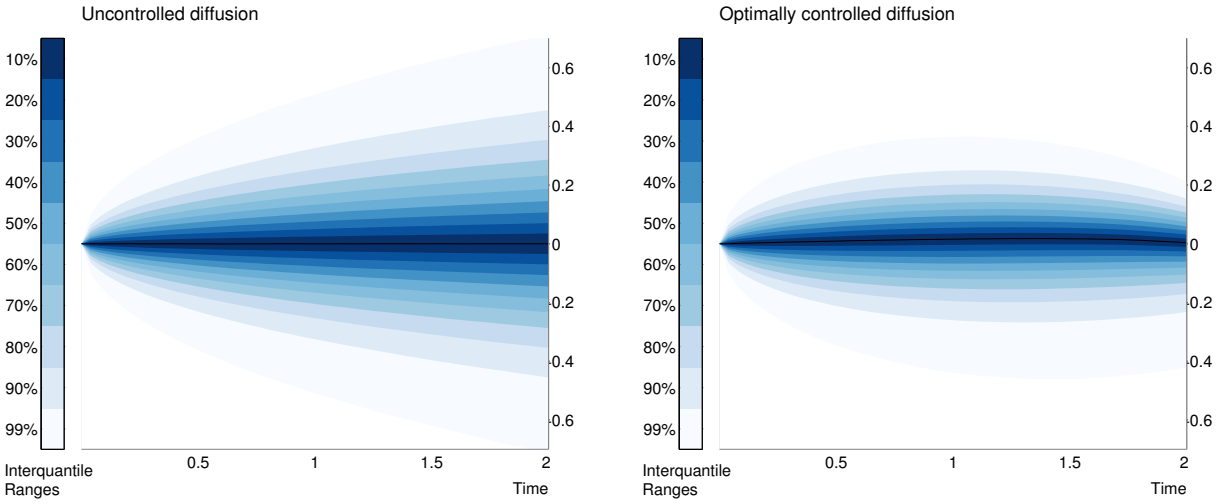

(a) Without control

(b) With control

Figure 5.4.2: Time-evolution of the distribution of the diffusion

Finally, regarding the accuracy of the method, the comparison between the estimated coefficients and their theoretical values is reported on Figure 5.4.1b. Indeed an analytical characterization of the solution of linear quadratic stochastic control problems is available using ordinary differential equations (cf. [107]). On our one-dimensional example (5.4.1), it is given by:

$$\begin{aligned} \alpha^{*}(t, X_t) &= A(t) X_t + B(t) \\ A(t) &= -\frac{\mu_1 P(t)}{2\lambda_0 + \sigma_1^2 P(t)} \\ B(t) &= -\frac{\mu_1}{2\lambda_0} Q(t) + A(t) \left(\frac{\sigma_0 \sigma_1}{\mu_1} - \frac{\sigma_1^2}{2\lambda_0} Q(t) \right) \end{aligned}$$

where $P(t)$ and $Q(t)$ are the solutions of the following ordinary differential equations:

$$\begin{aligned} P'(t) &= 2\mu_0 P(t) + \frac{\mu_1^2 P^2(t)}{2\lambda_0 + \sigma_1^2 P(t)} \\ P(T) &= 2\lambda_1 \\ Q'(t) &= \left(\mu_0 + \frac{\mu_1^2 P(t)}{2\lambda_0 + \sigma_1^2 P(t)} \right) Q(t) + \frac{\sigma_0 \sigma_1 \mu_1 P^2(t)}{2\lambda_0 + \sigma_1^2 P(t)} \\ Q(T) &= 0 \end{aligned}$$

As can be seen from the comparison on Figure 5.4.1b, our estimates of the control coefficients are very accurate. Regarding the value function, our method provides the estimate $\hat{v}(0,0) = -5.761$. The theoretical value being equal to -5.705, this means a relative error of 1%.

5.4.2 Uncertain volatility/correlation model

The second application is the problem of pricing and hedging an option under uncertain volatility.

Instead of specifying the parameters of the dynamics of an underlying process, one can, for robustness, consider them uncertain. To some extent, this parameter uncertainty provides hedging strategies that are more robust to model risk (cf. [99]). To handle these uncertain parameters, the usual approach is to resort to superhedging strategies, that is, to find the smallest amount of money from which it is possible to superreplicate the option, i.e. to build a strategy that will almost surely provide an amount greater than (or equal to) the payoff at the maturity of the option.

To compute these prices in practice, the most common approach is to resort to numerical methods for partial differential equations. For instance, [86] computes the superhedging price under uncertain correlation of a digital outperformance option using a finite differences sheme. Unfortunately, these PDE methods suffer from the curse of dimensionality, which means that they cannot handle many state variables (no more than three in practice).

This is why a few authors tried recently to resort to Monte Carlo techniques to solve this problem of pricing and hedging options under uncertain volatility and/or correlation.

To our knowledge, the first attempt to do so was made in [88]. In this thesis, along the usual backward induction, the conditional expectation are computed using the Malliavin calculus approach. This approach uses the representation of conditional expectations in terms of a suitable ratio of unconditional expectations. Then, to find the optimal covariance matrix at each time step, an exhaustive comparison is performed. Of course, this methodology works only if the set of possible matrices is finite, which is the case when the optimal control is of bang-bang type. For instance, it includes the case of unknown correlations with known volatilies, but not the case when both volatilities and correlations are unknown, a shortcoming that is acknowledged in [88]. This means that this methodology can only deal with optimal switching problems, for which the control set is finite.

To overcome this limitation, [62] propose to restrict the maximization domain to a parameterized set of relevant functions, indexed by a low-dimensional parameter. They then perform this much simpler optimization inductively at each time step, by the downhill simplex method (when the optimum is not of bang-bang type). Once it is done, say, at time t_i, they immediately use these estimated volatilities and correlations (along with those from $t_j > t_i$) to resample the whole Monte Carlo set from t_i to T (and idea also used in the Multiple Step Forward scheme from [58]). Remark that this parameterization avoids the computation of conditional expectations for each point and time step.

In [62], a second Monte Carlo scheme is proposed. It is a Monte Carlo scheme for 2-BSDEs, very similar to the schemes [41] and [49], but fine-tuned for the uncertain volatility problem under log-normal processes. The conditional expectations are computed by parametric regression (non-parametric regression in dimension 1). Then for each point and each time step, a deterministic optimization procedure has to be performed to find the optimal covariance matrix. However, unlike their previous algorithm, there is no resampling of the underlying diffusion using the newly computed covariances, which means that ensuring a proper simulation of the forward process becomes an issue.

Finally, we would like to draw attention to the work [89], which is not devoted to the uncertain volatility problem (it deals with the partial hedging of power futures with others futures with larger delivery period), but the probabilistic numerical scheme they propose can deal with a control in the volatility. Their specific application allows to retrieve the optimal control by a fixed point argument, within a backward scheme. However, as in the previous algorithm, an a priori control has to be used to simulate the forward process.

Here, our numerical scheme provides an alternative numerical sheme for dealing with the problem of pricing and hedging an option under uncertain volatility. To illustrate this, we implement it

below on a simple example.

Consider two underlyings driven by the following dynamics:

$$dS_i(t) = \sigma_i S_i(t)\, dW_i(t) \ , \ i = 1,2 \tag{5.4.3}$$

$$\langle dW_1(t), dW_2(t) \rangle = \rho(t, S_1(t), S_2(t))\, dt \tag{5.4.4}$$

where $\sigma_1, \sigma_2 > 0$, W_1 and W_2 are two correlated brownian motions. We consider no drift and no interest rate for simplicity. We instead focus our attention on the following crucial feature: we consider the correlation ρ to be *uncertain*. We only assume that ρ always lies between two known bounds $-1 \leq \rho_{\min} \leq \rho_{\max} \leq 1$:

$$\rho_{\min} \leq \rho \leq \rho_{\max} \tag{5.4.5}$$

Notice that when $\rho_{\min} = -1$ or $\rho_{\max} = 1$, the diffusion matrix of (S_1, S_2) can be degenerate.

We could also consider the two volatilities to be uncertain as well, but for illustration purposes, we focus on the uncertainty of the correlation parameter.

Finally, consider a payoff function $\Phi = \Phi(T, S_1(T), S_2(T))$ at a time horizon $T > 0$.

Now, the problem is to estimate the price of an option that delivers the payoff Φ at time T, and, if possible, to build a hedging strategy for this option.

Given that ρ is uncertain, the model is incomplete, i.e. it is not possible to construct a hedging strategy that replicates perfectly the payoff Φ from any given amount of money. We thus look for superhedging strategies instead.

Hence, consider the class $\mathbf{Q}$ of all probability measures $\mathbb{Q}$ on the sets of paths $\{S_i(t)\}_{0 \leq t \leq T}^{i=1,2}$ such that equations (5.4.4) and (5.4.5) hold for a particular $\rho^{\mathbb{Q}}$. The superhedging price is thus given by:

$$P_0^+ := \sup_{\mathbb{Q} \in \mathbf{Q}} \mathbb{E}^{\mathbb{Q}} \left[\Phi(T, S_1(T), S_2(T)) \right] \tag{5.4.6}$$

and the superhedging strategy is simply given by the usual delta-hedging strategy with ρ equal to the correlation that attains the supremum in equation (5.4.6). In particular it provides an upper arbitrage bound to the price of the option. Symmetrically, a lower bound is provided by the subreplication price:

$$P_0^- := \inf_{\mathbb{Q} \in \mathbf{Q}} \mathbb{E}^{\mathbb{Q}} \left[\Phi(T, S_1(T), S_2(T)) \right] \tag{5.4.7}$$

The practical computation of P_0^+ and P_0^- falls within the scope of our numerical scheme.

We thus test our numerical scheme on this specific problem. We consider the example of a call spread on the spread $S_1(T) - S_2(T)$, i.e.:

$$\Phi = (S_1(T) - S_2(T) - K_1)^+ - (S_1(T) - S_2(T) - K_2)^+$$

where $K_1 < K_2$. Unless stated otherwise, the parameters of the model are fixed to the following values:

$S_1(0)$	$S_2(0)$	σ_1	σ_2	$\rho_{\min}$	$\rho_{\max}$	K_1	K_2	T
50	50	0.4	0.3	-0.8	0.8	-5	5	0.25

For the numerical parameters, we use $n = 26$ time-discretization steps, and a sample of $M = 10^6$ Monte Carlo simulations. For the regressions, we use a basis function of sigmoid transforms of polynomial of degree two:

$$\begin{aligned} \phi(t, s_1, s_2, \rho) &:= (K_2 - K_1) \times \mathcal{S}(\beta_0 + \beta_1 s_1 + \beta_2 s_2 + \beta_3 \rho + \beta_4 \rho s_1 + \beta_5 \rho s_2) \\ \mathcal{S}(u) &:= \frac{1}{1 + e^{-u}} \end{aligned}$$

We chose the sigmoid function for its resemblance to the call spread payoff, and the terms inside the sigmoid according to their statistical significance. With this choice of basis, the optimal control will be bang-bang:

$$\rho^* = \rho^*(t, s_1, s_2) := \arg\max_{\rho} \phi(t, s_1, s_2, \rho) = \rho_{\max}\mathbf{1}\{\beta_3+\beta_4 s_1+\beta_5 s_2 \geq 0\}+\rho_{\min}\mathbf{1}\{\beta_3+\beta_4 s_1+\beta_5 s_2 < 0\}$$

Figure 5.4.3 below reports our results.

Figure 5.4.3a reports the superhedging and subhedging prices of the option, for different values of the moneyness ($S_2(0) = 50$ is kept fixed and different values of $S_1(0) = 50 +$ Moneyness are tested). One can clearly see the range of non-arbitrage prices that they define. For comparison, the prices obtained when ρ is constant are reported on the same graph for different values ($\rho_{\min}$, 0 and $\rho_{\max}$). One can see that, even though these prices belong to the non-arbitrage range, they do not cover the whole range, especially close to the money. This clearly indicates that, as already observed in [86] for instance, the practice of pricing under the hypothesis of constant parameters, and then testing different values for the parameters can be a very deceptive assessment of risk (as "uncertain" is not the same as "uncertain but constant").

Figure 5.4.3b illustrates the impact of the size of the correlation range $[\rho_{\min}, \rho_{\max}]$. Naturally, the wider the correlation range, the wider the price range. On average, an increase of 0.1 of the correlation range increases the price range by 0.135.

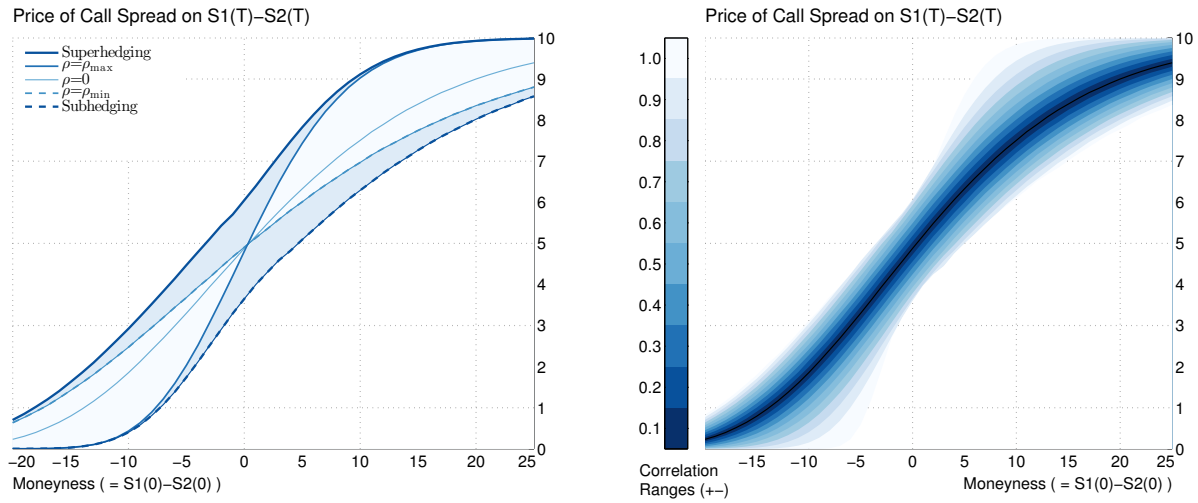

(a) Price of Call Spread

(b) Influence of the correlation range

Figure 5.4.3: Prices under uncertain correlation

5.4.3 Comparisons with [62]

Finally, we test our algorithm on several payoffs proposed in [62], and compare the behaviour of our method to their results. To be more specific, we will not focus our comparison of algorithms to their parametric approach[1], but to their second-order BSDE approach, as both algorithms are similar in nature (forward-backward schemes involving simulations and regressions).

[1]For comprehensiveness, here are the main pros and cons of the parametric approach: it is very accurate (especially when the optimal control belongs to the chosen parametric class) but requires $\mathcal{O}\left(N^2 \times M\right)$ operations, as at each time step t_i the simulations of the forward process are recomputed between t_i and t_N using the newly estimated optimal controls.

Actually, we are going to implement and compare two different versions of our scheme. The first one corresponds to the empirical version of the scheme studied in Section5.2:

$$\begin{aligned}\hat{Y}_N &= g\left(X_N\right)\\ \hat{\mathcal{Y}}_i &= \hat{\mathbb{E}}_i\left[\hat{Y}_{i+1} + f\left(X_i, I_i\right)\Delta_i\right]\\ \hat{Y}_i &= \operatorname*{ess\,sup}_{a\in A} \mathbb{E}_{i,a}\left[\hat{\mathcal{Y}}_i\right]\end{aligned} \tag{5.4.8}$$

where $\hat{\mathbb{E}}_i$ corresponds to an empirical least-squares regression which approximates the true conditional expectation $\mathbb{E}_i$. In the simpler context of American option pricing, this scheme would correspond to the Tsitsiklis-van Roy algorithm ([102]).

The second one makes use of the estimated optimal policies computed by the first algorithm, which are then directly plugged into the stochastic control problem under consideration:

$$\begin{aligned}\hat{\alpha}_i &= \arg\operatorname*{ess\,sup}_{a\in A} \mathbb{E}_{i,a}\left[\hat{\mathcal{Y}}_i\right]\\ \hat{X}_{i+1} &= b(\hat{X}_i, \hat{\alpha}_i)\Delta_i + \sigma(\hat{X}_i, \hat{\alpha}_i)\Delta W_i\\ \hat{v}\left(t_0, x_0\right) &= \frac{1}{M}\sum_{m=1}^{M}\left[\sum_{i=1}^{N} f(\hat{X}_{i+1}, \hat{\alpha}_i)\Delta_i + g(\hat{X}_N)\right]\end{aligned} \tag{5.4.9}$$

In the context of American option pricing, this scheme would correspond to the Longstaff-Schwarz algorithm ([83]).

We compute both prices as they are somehow complementary. Indeed, as noticed in [28] and detailed in [2], the first algorithm tend to be upward biased (up to the Monte Carlo error and the regression bias) compared with the discretized price, while the second one tend to be downward biased (up to the Monte Carlo error). Therefore, computing both prices provides a kind of empirical confidence interval, with the length of the interval being due to the choice of regression basis, thus providing an empirical assessment of the quality of the chosen regression basis.

Call Spread Let S be a geometric brownian motion with $S\left(0\right) = 100$ and with uncertain volatility σ taking values in $[0.1, 0.2]$.

Consider a call spread option, with payoff $\left(S\left(T\right) - K_1\right)^+ - \left(S\left(T\right) - K_2\right)^+$ and time horizon $T = 1$, with $K1 = 90$ and $K2 = 110$. The true price of the option (as estimated by PDE methods in [62]) is $\mathcal{C}_{PDE} = 11.20$, and the Black-Scholes price with constant volatility $\sigma_{\text{mid}} = 0.15$ is $\mathcal{C}_{BS} = 9.52$. We implement our scheme using the following set of basis functions:

$$\phi\left(t, s, \sigma\right) = \left(K_2 - K_1\right) \times \mathcal{S}\left(\beta_0 + \beta_1 s + \beta_2 s^2 + \beta_3 \sigma + \beta_4 \sigma s + \beta_5 \sigma s^2\right)$$

where, as in Subsection 5.4.2, $\mathcal{S}$ denotes the sigmoid function.

Figure 5.4.4 describes the estimates obtained with both algorithms (5.4.8) and (5.4.9), for various values of the number M of Monte Carlo simulations, and of the length of the constant discretization time step. For comparison, the red line corresponds to the price $\mathcal{C}_{PDE}$ of the option.

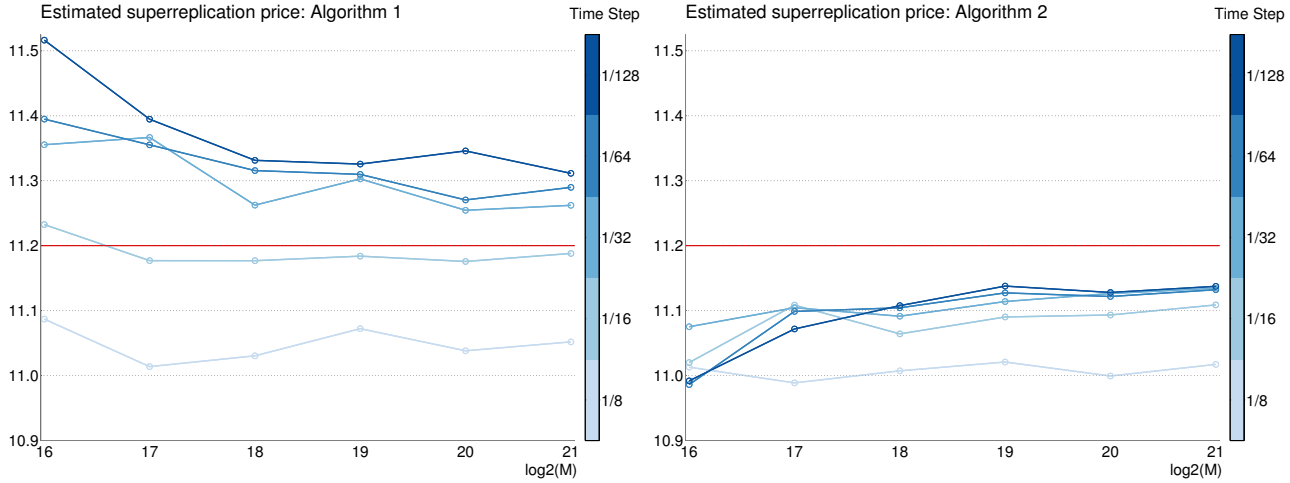

(a) First Algorithm

(b) Second Algorithm

Figure 5.4.4: Price of Call Spread

The following general observations can be made.

First, for a small enough time step, the prices computed using the first algorithm (5.4.8) (Figure 5.4.4a) tend as expected to be above the true price, while the second algorithm (5.4.9) (Figure 5.4.4b) tend to be below it.

Our best estimate here ($M = 2^{21}$, $\Delta_t = 1/128$) is 11.31 with the first algorithm (+1% compared with the true price) and 11.14 with the second one (−0.6%). The true price lies indeed between those two bounds, and their average (11.22) is even closer to the true price than any of the two estimates (+0.2%).

The prices computed with the first algorithm always lie above the prices computed with the second algorithm. As these prices are expected to surround the true discretized price (as would be computed by the scheme (5.4.8) with $\mathbb{E}_i$ instead of $\hat{\mathbb{E}}_i$), the fact that for large discretization steps ($\Delta_t = 1/8$ or $1/16$) the prices computed using the first algorithm are below the true price 11.20 simply means that, for such discretization steps, the true discretized price lies below the true price (in other words the time discretization generates here a negative bias).

Finally, increasing the number of Monte Carlo simulations tends as expected to improve the price estimates. However, the Monte Carlo error can be negligible compared with the discretization error for small time steps, which is why both a large number of Monte Carlo simulations and a small discretization time step are required to obtain accurate estimates.

In [62], the algorithm based on second-order BSDEs produces the estimates 11.04 for $(1/\Delta_t, \log_2(M)) = (8, 16)$ and 11.11 for $(1/\Delta_t, \log_2(M)) = (8, 17)$. This is close to our estimates for similar parameters. However, a more accurate comparison would require to test their algorithm with smaller time steps and more Monte Carlo simulations (they only consider parameters $(1/\Delta_t, \log_2(M))$ within $[2, 8] \times [12, 17]$, whereas we consider here the range $[8, 128] \times [16, 21]$, as it provides much greater accuracy of the estimates, providing a sound basis for the analysis of the results).

Digital option: Consider a digital option, with payoff $100 \times \mathbf{1}\{S(T) \geq K\}$ and $T = 1$ on the samee asset, with $K = 100$. The true (PDE) price is $\mathcal{C}_{PDE} = 63.33$, and the Black-Scholes price

with mid-volatility is $\mathcal{C}_{BS} = 46.54$. We use the following set of basis functions:

$$\phi(t,s,\sigma) = 100 \times \mathcal{S}\left(\beta_0 + \beta_1 s + \beta_2 s^2 + \beta_3 \sigma + \beta_4 \sigma s + \beta_5 \sigma s^2\right)$$

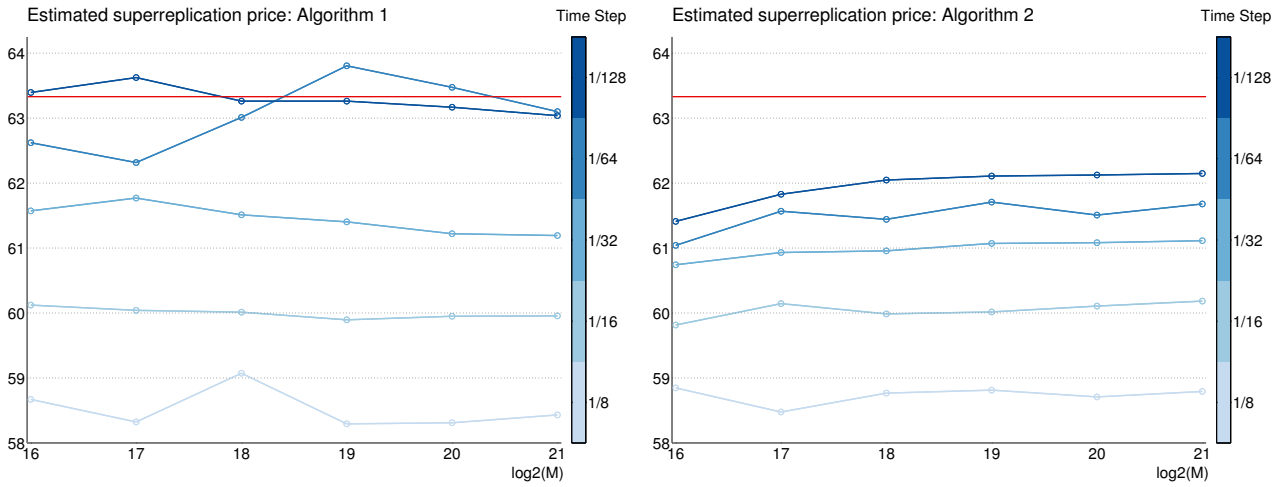

(a) First Algorithm

(b) Second Algorithm

Figure 5.4.5: Price of Digital Option

As can be seen on Figure 5.4.5, the time discretization error is much more pronounced with this discontinuous payoff, compared with the previous call spread example. We manage to reach estimates of 63.04 (-0.5%) and 62.15 (-1.9%), even though smaller time steps would be required for better accuracy.

For small parameters $((1/\Delta_t, \log_2(M)) = (8,16))$, the accuracy is better in [62] (60.53), even though shortening the time step tends to degrade the results in their case.

Outperformer Option: Consider now two geometric Brownian motions S_1 and S_2, starting from 100 at time 0, with uncertain volatilities σ_1 and σ_2 taking values in $[0.1, 0.2]$. For the moment, suppose that the correlation ρ between the two underlying Brownian motions is zero.

Consider an outperformer option, with payoff $(S_1(T) - S_2(T))^+$ and time horizon $T = 1$. The true price is $\mathcal{C} = 11.25$. We use the following set of basis functions:

$$\begin{aligned}\phi(t,s_1,s_2,\sigma_1,\sigma_2) =& 100 \times \left(\beta_0 + \beta_1 s_1 + \beta_2 s_1^2 + \beta_3 s_2 + \beta_4 s_2^2 + \beta_5 s_1 s_2 + \beta_6 \sigma_1 + \beta_7 \sigma_1 s_1 + \beta_8 \sigma_1 s_1^2\right. \\ &\left. + \beta_9 \sigma_1 s_2 + \beta_{10} \sigma_1 s_2^2 + \beta_{11} \sigma_2 + \beta_{12} \sigma_2 s_1 + \beta_{13} \sigma_2 s_1^2 + + \beta_{14} \sigma_2 s_2 + \beta_{15} \sigma_2 s_2^2\right)\end{aligned}$$

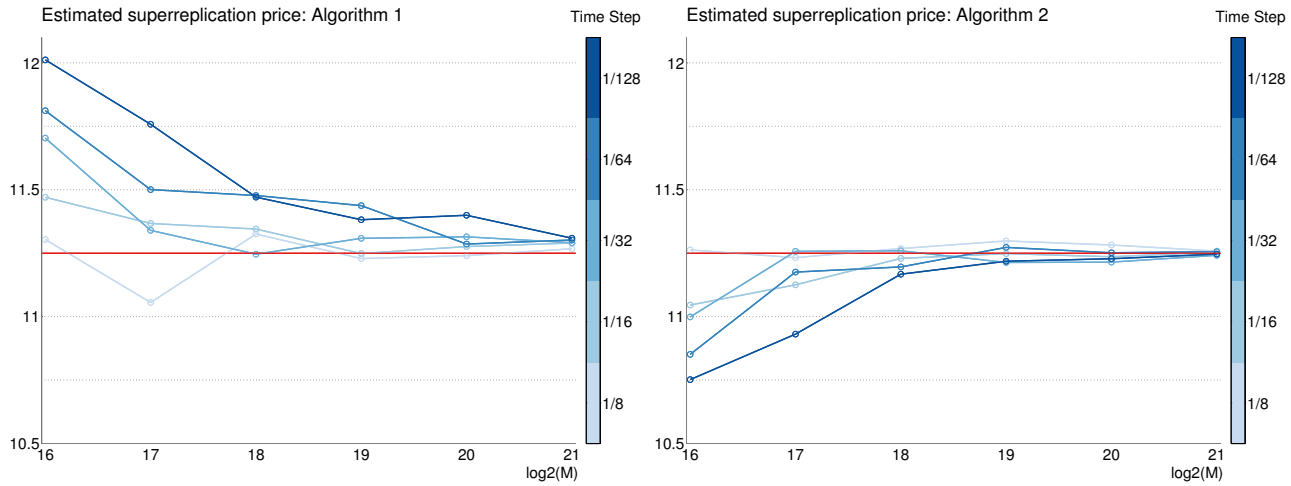

(a) First Algorithm

(b) Second Algorithm

Figure 5.4.6: Price of Outperformer Option ($\rho = 0$)

Here, in contrast with the previous examples, the bulk of the error comes from the Monte Carlo simulations, and not from the time discretization. Moreover, both algorithms provide very accurate estimates. Indeed, this convex option is easy to price under the uncertain volatility model, as it is given by the price obtained with the maximum volatilities. With our choice of regression basis, the algorithm correctly detects that the maximum volatilities are to be used, leading to these very accurate estimates 11.31 (+0.5%) and 11.25 (−0%). For the same reason, the estimates from [62] are accurate too.

Figure 5.4.7 below depicts the estimated price of the same option but now with a negative constant correlation $\rho = -0.5$. Its true price is $\mathcal{C} = 13.75$.

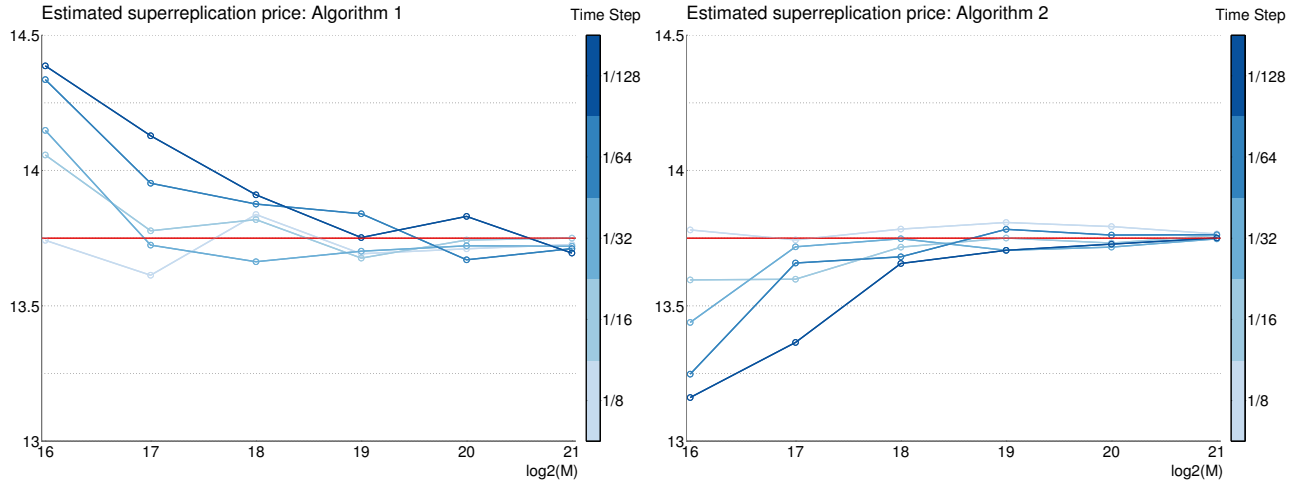

(a) First Algorithm

(b) Second Algorithm

Figure 5.4.7: Price of Outperformer Option ($\rho = -0.5$)

The same behaviour can be observed. Both algorithms are accurate here (13.69 (−0.4%) and 13.75 (−0%)).

As the estimate from the first algorithm happens to lie below the true price, we take advantage of this result to recall from the introduction of this subsection that the bias of the first algorithm bears one more source of error (the regression bias) than the bias of the second algorithm. This means that in general the sign of the bias wrt. the true discretized price is more reliable with the second algorithm. With this observation in mind, we propose, from the two estimates P_1 and P_2 computed by the two algorithms, to consider the following general estimate P:

$$P := \max\left(P_2, \frac{P_1 + P_2}{2}\right)$$

Indeed, if $P_1 \geq P_2$ (which is the expected behaviour), then $P := \frac{P_1+P_2}{2}$ may provide a better estimate than both P_1 and P_2 separately (as is the case for the call spread example from Figure 5.4.4). However, when $P_1 < P_2$ (which is not expected), then, recalling that P_2 may be more accurate than P_1, it is better to consider $P := P_2$ (as is the case here of this outperformer option with $\rho = -0.5$). In the following, we will call P the mid-estimate (with a slight abuse of terminology, as P is usually but not always the average between P_1 and P_2).

Outperformer spread option: We now analyze a more complex payoff. Consider an outperformer spread option, with payoff $(S_2(T) - K_1 S_1(T))^+ - (S_2(T) - K_2 S_1(T))^+$, time horizon $T = 1$ and constant correlation $\rho = -0.5$. The true (PDE) price is $\mathcal{C}_{PDE} = 11.41$, and the Black-Scholes price with mid-volatility is $\mathcal{C}_{BS} = 9.04$. We use the following set of basis functions:

$$\phi(t, s_1, s_2, \sigma_1, \sigma_2) = s_1 \times (K_2 - K_1) \times \mathcal{S}\left(\beta_0 + \beta_1 \frac{s_2}{s_1} + \beta_2 \left(\frac{s_2}{s_1}\right)^2 + \beta_3 \sigma_1 + \beta_4 \sigma_1 \frac{s_2}{s_1} + \beta_5 \sigma_1 \left(\frac{s_2}{s_1}\right)^2 + \beta_6 \sigma_2 + \beta_7 \sigma_2 \frac{s_2}{s_1} + \beta_8 \sigma_2 \left(\frac{s_2}{s_1}\right)^2\right)$$

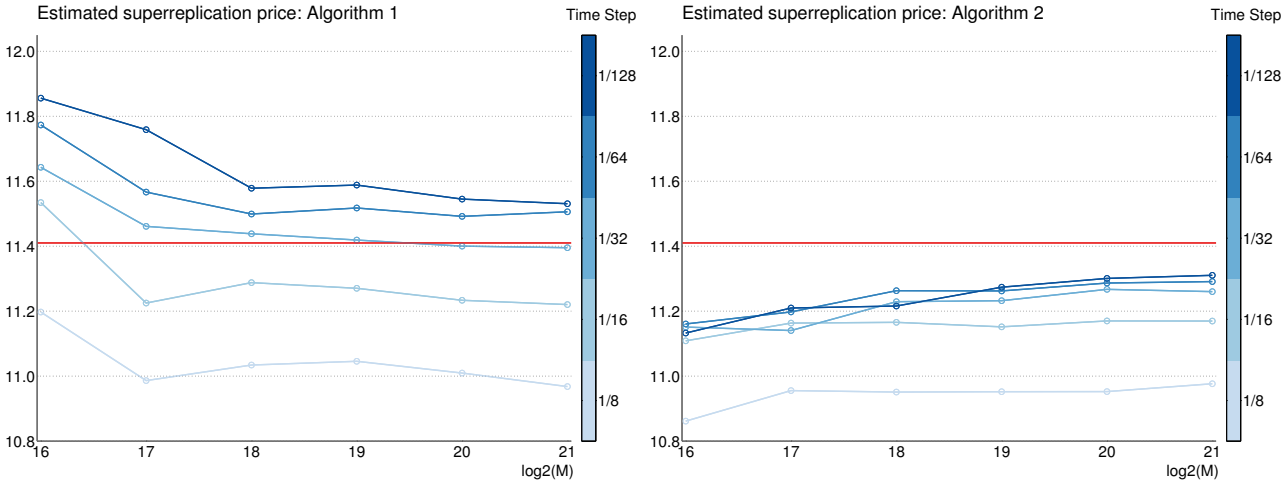

(a) First Algorithm

(b) Second Algorithm

Figure 5.4.8: Price of Outperformer Spread Option ($\rho = -0.5$)

In this example, one can see that the time discretization produces a large downward bias (as in the call spread and digital option examples), but both algorithms behave as expected (the

first algorithm produces high estimates, the second produces low estimates, and both are close to the true price (11.53 (+1%) and 11.31 (−0.9%))). Moreover, the mid-estimate 11.42 is very accurate.

In [62] is reported the estimate 10.83 for $(1/\Delta_t, \log_2(M)) = (8, 20)$, which is slightly worse than our estimates for the same choice of M and Δ_t (11.01 and 10.95), but the difference can be due to the different choice of basis. However, the three estimates are well below the true price, and our numerical results indicate that the reason is that $\Delta_t = 1/8$ is too large a time step.

This suggests that the estimates from [62] could be improved by considering smaller time steps. However, as acknowledged in their paper, the second-order BSDE method does not work properly when Δ_t is too small. Indeed, their BSDE scheme makes use of the first order component Z and the second order component Γ. The problem here is that, for fixed M, the variance of the estimators of Z and Γ tends to infinity when Δ_t tends to zero. However, as detailed in [4], this problem can be completely solved by amending the estimators using appropriate variance reduction terms. Therefore, in our opinion, a fair comparison of the jump-constrained BSDE approach and the second-order BSDE approach would require the use of the variance reduction method from [4] to allow for smaller time steps for the second-order BSDE approach.

As a final numerical example, we consider again the same outperformer spread option, with the exception that the correlation ρ is now considered uncertain, within $[-0.5, 0.5]$. The true (PDE) price is $\mathcal{C}_{\text{PDE}} = 12.83$, and the Black-Scholes price with mid-volatility is $\mathcal{C}_{BS} = 9.24$. We use the following basis functions:

$$\phi(t, s_1, s_2, \sigma_1, \sigma_2, \rho) = s_1 \times (K_2 - K_1) \times \mathcal{S}\left(\beta_0 + \beta_1 \frac{s_2}{s_1} + \beta_2 \left(\frac{s_2}{s_1}\right)^2 + \beta_3 \sigma_1 + \beta_4 \sigma_1 \frac{s_2}{s_1} + \beta_5 \sigma_1 \left(\frac{s_2}{s_1}\right)^2 + \beta_6 \sigma_2 + \beta_7 \sigma_2 \frac{s_2}{s_1} + \beta_8 \sigma_2 \left(\frac{s_2}{s_1}\right)^2 + \beta_6 \rho + \beta_7 \rho \frac{s_2}{s_1} + \beta_8 \rho \left(\frac{s_2}{s_1}\right)^2\right)$$

Remark that at each time step we perform here a five-dimensional regression.

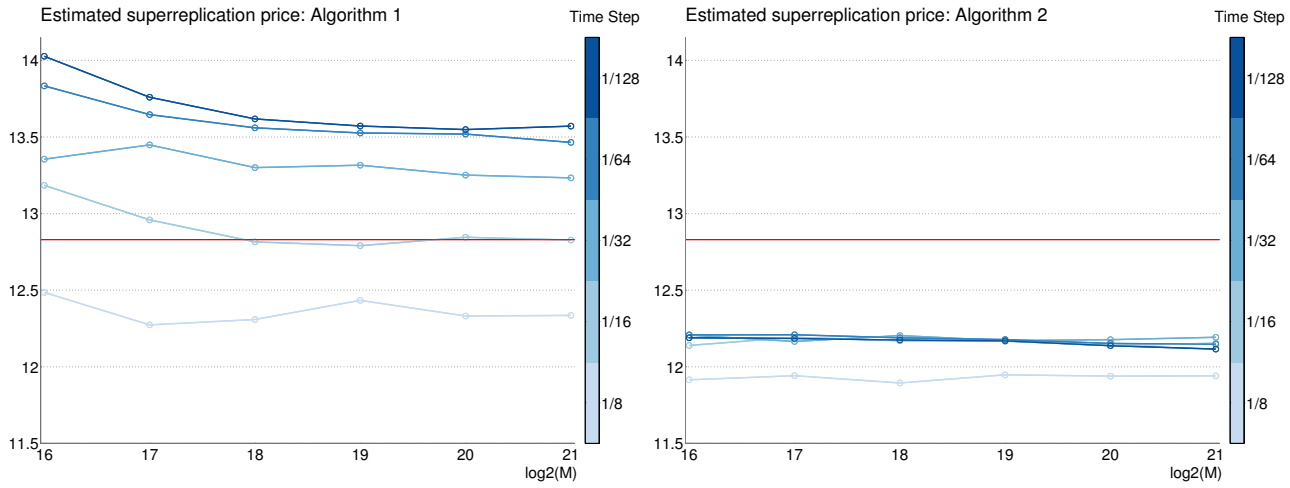

(a) First Algorithm

(b) Second Algorithm

Figure 5.4.9: Price of Outperformer Spread Option (ρ uncertain $\in [-0.5, 0.5]$)

On this example, we observe a wide gap between the two estimates 13.57 (+5.8%) and 12.12 (−5.6%) ($(1/\Delta_t, \log_2(M)) = (128, 21)$). As neither the number of Monte Carlo simulation nor

the discretization time step seem able to narrow the gap, it means that it is due to the chosen regression basis. Indeed, our basis is such that the optimal volatilities and correlation are of bang-bang type, as in the previous examples. However, unlike the previous examples, here both the volatilities and the correlation are uncertain, and in this case it is known (cf. [88] for instance) that the optimum is not of a bang-bang type. Therefore, one should look for a richer regression basis in order to narrow the estimation gap on this specific example. Remark however that the mid-estimated 12.84 remains very accurate. On the same example and with another regression basis, [62] manage to reach a price of 12.54 for $(1/\Delta_t, \log_2(M)) = (8, 20)$.

To conclude these subsection, here are the differences we could notice between the jump-constrained BSDE approach and the second-order BSDE approach applied to the problem of pricing by simulation under uncertain volatility model:

- Both are forward-backward schemes. Thus, the first step is to simulate the forward process. At this stage the jump-constrained BSDE approach is advantaged, because its forward process is a simple Markov process, therefore easy to simulate. Its randomization of the control is fully justified mathematically. On the contrary, the second-order BSDE requires to resort to heuristics in order to simulate the forward process despite the fact that the control is involved in its dynamics. [62] propose to use an arbitrary constant volatility (the mid-volatility) to simulate the forward process, and they notice that the specific choice of prior-volatility does impact substantially the resulting estimates.

- Then comes the estimation of the backward process. If both schemes require to perform regressions, this step is more difficult in the jump-constrained BSDE approach, because the dimensionality of the regressions is higher as the state process contains the randomized controls. In particular the choice of regression basis is more difficult.

- On the set of options considered here and within the same range of numerical parameters M and Δ_t we could not detect any significant and systematic difference between the two algorithms. Nevertheless, we strongly suggest the following two points:

- First, the second-order BSDE approach would strongly benefit from the use of the variance reduction method from [4]. It would allow for smaller time steps to be considered, and therefore allow for a sounder and more precise numerical comparison between the two approaches. Indeed, the accurate estimates recorded in [62] for very large time steps may be, as in Figure 5.4.9a for $\Delta_t = 1/16$, an incidental cancellation of biases of opposite signs. The significant quantity is the level where the estimates converge for small Δ_t.

- Second, to complement the downward biased, "Longstaff-Schwartz like" estimator considered in [62], we suggest the computation of the upward biased, "Tsitsiklis-van Roy like" estimator, as we did in this chapter, as both estimators appear to be informative in a complementary fashion, and the mid-estimator proposed here (which requires both estimators) seems to perform staggeringly well.

5.5 Conclusion

We proposed in this chapter a general probabilistic numerical algorithm, combining Monte Carlo simulations and empirical regressions, which is able to solve numerically very general HJB equations in high dimension. That includes general stochastic control problems with controlled volatility, possibly degenerate, but more generally, it can solve any jump-constrained BSDE ([72]).

We thoroughly analyzed the time-discretization step, including the establishment of a convergence rate. We initiated a partial analysis of the theoretical error of the scheme, and we provided several numerical application of the scheme on the problem of pricing under uncertain volatility, the results of which are very promising.

In the future, we would like to extend this work in the following direction:

- First, we would like to manage to obtain a comprehensive analysis of the error of the scheme, including the empirical regression step.
- Then, we would like to perform a more systematic numerical comparison with the alternative scheme described in [62], taking into account our empirical findings.
- Finally, we would like to extend the general methodology of control randomization and subsequent constraint on resulting jumps to more general problems, like HJB-Isaacs equations or even mean-fields games, with possible advances on the numerical solution of such problems.

Bibliography

[1] R. Aïd, L. Campi, and N. Langrené. A structural risk-neutral model for pricing and hedging power derivatives. *Mathematical Finance*, 23(3):387–438, July 2013. (Cited on pages 8, 26, 87, 114, 115, 117, and 124)

[2] R. Aïd, L. Campi, N. Langrené, and H. Pham. A probabilistic numerical method for optimal multiple switching problem in high dimension. *SIAM Journal on Financial Mathematics*, 5(1):191–231, 2014. (Cited on pages 13, 30, 158, and 172)

[3] R. Aïd, L. Campi, A. Nguyen Huu, and N. Touzi. A structural risk-neutral model of electricity prices. *International Journal of Theoretical and Applied Finance*, 12(7):925–947, 2009. (Cited on pages 9, 10, 26, 27, 28, 41, 42, 43, 53, 75, 87, 114, 115, and 117)

[4] S. Alanko and M. Avellaneda. Reducing variance in the numerical solution of BSDEs. *Comptes Rendus Mathematique*, 351(3-4):135–138, 2013. (Cited on pages 177 and 178)

[5] M. Avellaneda, A. Levy, and A. Parás. Pricing and hedging derivative securities in markets with uncertain volatilities. *Journal of Applied Mathematical Finance*, 2(2):73–88, 1995. (Cited on pages 8 and 26)

[6] V. Bally, L. Caramellino, and A. Zanette. Pricing and hedging American options by Monte Carlo methods using a Malliavin calculus approach. *Monte Carlo Methods and Applications*, 11(2):97–133, 2005. (Cited on pages 14 and 32)

[7] V. Bally and G. Pagès. Error analysis of the optimal quantization algorithm for obstacle problems. *Stochastic processes and their applications*, 106(1):1–40, 2003. (Cited on pages 131, 147, and 154)

[8] V. Bally, G. Pagès, and J. Printems. A quantization method for pricing and hedging multi-dimensional American style options. *Mathematical Finance*, 15(1):119–168, 2005. (Cited on pages 14 and 32)

[9] G. Barles, R. Buckdahn, and E. Pardoux. Backward stochastic differential equations and integral-partial differential equations. *Stochastics and Stochastics Reports*, 60(1-2):57–83, 1997. (Cited on pages 137, 138, 139, 140, 142, 150, and 152)

[10] G. Barles and E. Jakobsen. Error bounds for monotone approximation schemes for parabolic Hamilton-Jacobi-Bellman equations. *Mathematics of Computation*, 76(260):1861–1893, 2007. (Cited on pages 131, 136, 146, 147, 148, 149, and 155)

[11] M. T. Barlow. A diffusion model for electricity prices. *Mathematical Finance*, 12(4):287–298, 2002. (Cited on pages 42, 43, and 47)

[12] D. Belomestny, A. Kolodko, and J. Schoenmakers. Regression methods for stochastic control problems and their convergence analysis. *SIAM Journal on Control and Optimization*, 48(5):3562–3588, 2010. (Cited on pages 14 and 32)

[13] G. Benedetti and L. Campi. Utility indifference valuation for non-smooth payoffs with an application to power derivatives. *Applied Mathematics & Optimization*, 73(2):349–389, 2016. (Cited on pages 10 and 28)

[14] G. Benmenzer, E. Gobet, and C. Jérusalem. Arbitrage free cointegrated models in gas and oil future markets. Technical report, GDF SUEZ and Laboratoire Jean Kuntzmann, Grenoble, 2007. (Cited on pages 42, 116, and 117)

[15] F. Benth, J. Benth, and S. Koekebakker. *Stochastic Modeling of Electricity and Related Markets*. Advanced Series on Statistical Science and Applied Probability. World Scientific Publishing Company, 2008. (Cited on pages 8 and 26)

[16] F. E. Benth, L. Ekeland, R. Hauge, and B. F. Nielsen. A note on arbitrage-free pricing of forward contracts in energy markets. *Applied Mathematical Finance*, 10(4):325–336, 2003. (Cited on page 41)

[17] F. E. Benth, J. Kallsen, and T. Meyer-Brandis. A non-gaussian Ornstein-Uhlenbeck process for electricity spot price modeling and derivatives pricing. *Applied Mathematical Finance*, 14(2):153–169, 2007. (Cited on page 41)

[18] F. E. Benth and S. Koekebakker. Stochastic modeling of financial electricity contracts. *Energy Economics*, 30(3):1116–1157, 2008. (Cited on page 41)

[19] F. E. Benth and L. Vos. Pricing of forwards and options in a multivariate non-Gaussian stochastic volatility model for energy markets. *Advances in Applied Probability*, 45(2):572–594, 2013. (Cited on page 41)

[20] V. I. Bogachev. *Measure Theory 1*. Springer, 2007. (Cited on page 124)

[21] F. Bonnans and H. Zidani. Consistency of generalized finite difference schemes for the stochastic HJB equation. *SIAM Journal of Numerical Analysis*, 41(3):1008–1021, 2003. (Cited on page 131)

[22] A. Boogert and D. Dupont. When supply meets demand: the case of hourly spot electricity prices. *IEEE Transactions on Power Systems*, 23(2):389–398, 2008. (Cited on page 42)

[23] A. Botterud, M. Ilic, and I. Wangensteen. Optimal investments in power generation under centralized and decentralized decision making. *IEEE Transactions on Power Systems*, 20(1):254–263, 2005. (Cited on page 87)

[24] B. Bouchard and J.-F. Chassagneux. Discrete-time approximation for continuously and discretely reflected BSDEs. *Stochastic Processes and their Applications*, 118(12):2269–2293, 2008. (Cited on page 131)

[25] B. Bouchard and R. Elie. Discrete-time approximation of decoupled forward-backward SDE with jumps. *Stochastic Processes and their Applications*, 118(1):53–75, 2008. (Cited on pages 130, 141, 143, and 150)

[26] B. Bouchard and A. Nguyen Huu. No marginal arbitrage of the second kind for high production regimes in discrete time production-investment models with proportional transaction costs. *Mathematical Finance*, 23(2):366–386, 2013. (Cited on page 48)

[27] B. Bouchard and N. Touzi. Discrete-time approximation and Monte-Carlo simulation of backward stochastic differential equations. *Stochastic Processes and their Applications*, 111(2):175–206, 2004. (Cited on pages 19, 36, 87, 95, 99, 105, 130, and 154)

[28] B. Bouchard and X. Warin. Monte-Carlo valorisation of American options: facts and new algorithms to improve existing methods. In R. Carmona, P. Del Moral, P. Hu, and N. Oudjane, editors, *Numerical Methods in Finance*, volume 12 of *Springer Proceedings in Mathematics*, 2012. (Cited on pages 15, 17, 32, 34, 87, 91, 92, 95, 96, 124, 158, and 172)

[29] M. Burger, B. Klar, A. Müller, and G. Schindlmayr. A spot market model for pricing derivatives in electricity markets. *Quantitative Finance*, 4(1):109–122, 2004. (Cited on page 41)

[30] R. Carmona and M. Coulon. A survey of commodity markets and structural models for electricity prices. In F. Benth, V. Kholodnyi, and P. Laurence, editors, *Quantitative Energy Finance: Modeling, Pricing and Hedging in Energy and Commodity Markets*, 2013. (Cited on pages 8, 11, 26, 29, and 115)

[31] R. Carmona, P. Del Moral, N. Oudjane, and P. Hu. *Numerical Methods in Finance*. Springer, 2012. (Cited on pages 7, 25, and 87)

[32] R. Carmona and M. Ludkovski. Pricing asset scheduling flexibility using optimal switching. *Applied Mathematical Finance*, 15(5):405–447, 2008. (Cited on pages 14, 32, 87, 91, 99, and 109)

[33] J. Carriere. Valuation of the early-exercise price for options using simulations and nonparametric regression. *Insurance: Mathematics and Economics*, 19(1):19–30, 1996. (Cited on page 92)

[34] Á. Cartea and M. Figueroa. Pricing in electricity markets: a mean reverting jump diffusion model with seasonality. *Applied Mathematical Finance*, 12(4):313–335, 2005. (Cited on page 41)

[35] Á. Cartea and P. Villaplana. Spot price modeling and the valuation of electricity forward contracts: the role of demand and capacity. *Journal of Banking and Finance*, 32(12):2502–2519, 2008. (Cited on pages 41, 42, 43, and 47)

[36] R. Chan, Y. Chen, and K.-M. Yeung. A memory reduction method in pricing American options. *Journal of Statistical Computation and Simulation*, 74(7):501–511, 2004. (Cited on pages 17, 34, 87, 110, and 113)

[37] R. Chan, C.-Y. Wong, and K.-M. Yeung. Pricing multi-asset American-style options by memory reduction Monte Carlo methods. *Applied Mathematics and Computation*, 179(2):535–544, 2006. (Cited on pages 87, 110, and 113)

[38] R. Chan and T. Wu. Memory-reduction method for pricing American-style options under exponential Lévy processes. *East Asian Journal on Applied Mathematics*, 1(1):20–34, 2011. (Cited on pages 87, 110, and 113)

[39] J. Chassagneux, R. Elie, and I. Kharroubi. Discrete-time approximation of multidimensional BSDEs with oblique reflections. *Annals of Applied Probability*, 22(3):971–1007, 2012. (Cited on page 99)

[40] M. A. Chaudhry and S. M. Zubair. *On a class of incomplete gamma functions with applications*. Chapman & Hall, 2002. (Cited on page 68)

[41] P. Cheridito, M. Soner, N. Touzi, and N. Victoir. Second order backward stochastic differential equations and fully non-linear parabolic PDEs. *Communications on Pure and Applied Mathematics*, 60(7):1081–1110, 2007. (Cited on page 169)

[42] L. Clewlow and C. Strickland. *Energy Derivatives*. Lacima Group, 2000. (Cited on page 41)

[43] M. Coulon and S. Howison. Stochastic behaviour of the electricity bid stack: from fundamental drivers to power prices. *Journal of Energy Markets*, 2(1):29–69, 2009. (Cited on pages 42, 43, and 47)

[44] M. Crandall, H. Ishii, and P.-L. Lions. User's guide to viscosity solutions of second order partial differential equations. *Bulletin of the American Methematical Society*, 27(1):1–67, 1992. (Cited on pages 136 and 144)

[45] A. Deaño and N. M. Temme. Analytical and numerical aspects of a generalization of the complementary error function. *Applied Mathematics and Computation*, 216(12):3680–3693, 2010. (Cited on pages 12, 29, and 67)

[46] F. Delbaen, P. Monat, W. Schachermayer, M. Schweizer, and C. Stricker. Weighted norm inequalities and hedging in incomplete markets. *Finance and Stochastics*, 1(3):181–227, 1997. (Cited on pages 51 and 52)

[47] S. Deng. Stochastic models of energy commodity prices and their applications: Mean-reversion with jumps and spikes. Technical Report PWP-073, University of California Energy Institute, 2000. (Cited on page 41)

[48] B. El Asri. Optimal multi-modes switching problem in infinite horizon. *Stochastics and Dynamics*, 10(2):231–261, 2010. (Cited on page 91)

[49] A. Fahim, N. Touzi, and X. Warin. A probabilistic numerical method for fully nonlinear parabolic PDEs. *The Annals of Applied Probability*, 21(4):1322–1364, 2011. (Cited on pages 21, 22, 23, 38, 40, 129, 131, 147, 155, and 169)

[50] M. Fischer and G. Nappo. On the moments of the modulus of continuity of Itô processes. *Stochastic Analysis and Applications*, 28(1):103–122, 2009. (Cited on page 100)

[51] W. Fleming and M. Soner. *Controlled Markov Processes and Viscosity Solutions*, volume 25 of *Stochastic Modelling and Applied Probability*. Springer, 2nd edition, 2006. (Cited on page 136)

[52] H. Föllmer and M. Schweizer. Hedging of contingent claims under incomplete information. *Applied stochastic analysis (London, 1989), Stochastics Monogr., 5, Gordon and Breach, New York*, pages 389–414, 1991. (Cited on pages 10, 27, 42, 48, 50, 51, and 55)

[53] A. Friedman. *Stochastic differential equations and applications*. Academic Press New York, 1975. (Cited on page 60)

[54] N. Frikha and V. Lemaire. Joint modelling of gas and electricity spot prices. *Applied Mathematical Finance*, 20(1):69–93, 2013. (Cited on page 42)

[55] P. Gassiat, I. Kharroubi, and H. Pham. Time discretisation and quantization methods for optimal multiple switching problem. *Stochastic Processes and their Applications*, 122(5):2019–2052, 2012. (Cited on pages 14, 15, 32, 87, 99, and 100)

[56] H. Geman and V. Nguyen. Soybean inventory and forward curve dynamics. *Management Science*, 51(7):1076–1091, 2005. (Cited on page 42)

[57] E. Gobet, J.-P. Lemor, and X. Warin. A regression-based Monte Carlo method to solve Backward Stochastic Differential Equations. *The Annals of Applied Probability*, 15(3):2172–2202, 2005. (Cited on page 95)

[58] E. Gobet and P. Turkedjiev. Approximation of discrete BSDE using least-squares regression. Preprint, 2011. (Cited on pages 19, 22, 36, 39, 158, 162, 164, and 169)

[59] C. Gouriéroux and P. Valéry. Estimation of a Jacobi process. Preprint, 2004. (Cited on page 116)

[60] S. Goutte, N. Oudjane, and F. Russo. Variance optimal hedging for continuous time additive processes and applications. *Stochastics*, 86(1):147–185, 2014. (Cited on pages 41 and 48)

[61] C. Grobbel. *Competition in Electricity Generation in Germany and Neighboring Countries from a System Dynamics Perspective: Outlook Until 2012*. Peter Lang Publishing, 1999. (Cited on pages 8 and 26)

[62] J. Guyon and P. Henry-Labordère. Uncertain volatility model: a Monte Carlo approach. *The Journal of Computational Finance*, 14(3):37–71, 2011. (Cited on pages 6, 23, 24, 40, 131, 169, 171, 172, 173, 174, 175, 177, 178, and 179)

[63] S. Hamadène. Optimal switching systems of reflected BSDEs and systems of variational inequalities with interconnected obstacles. In J. Blath, P. Imkeller, and S. Roelly, editors, *Surveys in Stochastic Processes*, Series of Congress Reports, 2011. (Cited on page 91)

[64] S. Hamadène, J.-P. Lepeltier, and Z. Wu. Infinite horizon reflected BSDEs and applications in mixed control and game problems. *Probability and Mathematical Statistics*, 19(2):211–234, 1999. (Cited on page 91)

[65] D. Heath, E. Platen, and M. Schweizer. A comparison of two quadratic approaches to hedging in incomplete markets. *Mathematical Finance*, 11(4):385–413, 2001. (Cited on page 60)

[66] D. Hobson. Bounds for the utility-indifference prices of non-traded assets in incomplete markets. *Decisons in Economics and Finance*, 28(1):33–52, 2005. (Cited on page 51)

[67] Y. Hu and S. Tang. Multi-dimensional BSDE with oblique reflection and optimal switching. *Probability Theory and Related Fields*, 147:89–121, 2010. (Cited on page 99)

[68] T. Kanamura and K. Ohashi. A structural model for electricity prices with spikes: Measurement of spike risk and optimal policies for hydropower plant operation. *Energy Economics*, 29(5):1010–1032, 2007. (Cited on pages 42, 43, and 47)

[69] I. Karatzas and S. E. Shreve. *Brownian Motion and Stochastic Calculus*. Springer, 2nd edition, 1991. (Cited on page 50)

[70] I. Kharroubi, N. Langrené, and H. Pham. A numerical algorithm for fully nonlinear HJB equations: an approach by control randomization. *Monte Carlo Methods and Applications*, 20(2):145–165, 2014. (Cited on pages 18 and 35)

[71] I. Kharroubi, N. Langrené, and H. Pham. Discrete-time approximation of fully nonlinear HJB equations via BSDEs with nonpositive jumps. *Annals of Applied Probability*, 25(4):2301–2338, 2015. (Cited on pages 18 and 35)

[72] I. Kharroubi and H. Pham. Feynman-Kac representation for Hamilton-Jacobi-Bellman IPDE. *Annals of Probability*, 43(4):1823–1865, 2015. (Cited on pages 7, 19, 20, 25, 36, 37, 129, 130, 136, 147, and 178)

[73] P. Kloeden and E. Platen. *Numerical Solution of Stochastic Differential Equations*, volume 23 of *Stochastic Modelling and Applied Probability*. Springer, 3rd edition, 1999. (Cited on pages 90 and 133)

[74] M. Kohler. A review on regression based Monte Carlo methods for pricing American options. In L. Devroye, B. Karasözen, M. Kohler, and R. Korn, editors, *Recent Developments in Applied Probability and Statistics*, Physica-Verlag, pages 39–61, 2010. (Cited on pages 14, 32, and 92)

[75] V. Kolodnyi. Valuation and hedging of european contingent claims on power with spikes: a non-markovian approach. *Journal of Engineering Mathematics*, 49(3):233–252, 2004. (Cited on page 41)

[76] D. Kroese, T. Taimre, and Z. Botev. *Handbook of Monte Carlo methods*, volume 706 of *Wiley series in probability and statistics*. Wiley, 2011. (Cited on pages 110 and 111)

[77] N. Krylov. Approximating value functions for controlled degenerate diffusion processes by using piece-wise constant policies. *Electronic Journal of Probability*, 4:1–19, 1999. (Cited on pages 21, 38, and 148)

[78] N. Krylov. On the rate of convergence of finite-difference approximations for bellman's equations with variable coefficients. *Probability theory and related fields*, 117(1):1–16, 2000. (Cited on pages 21, 38, 131, 136, 146, 147, 149, and 155)

[79] H. Kushner and P. Dupuis. *Numerical methods for stochastic control problems in continuous time*, volume 24 of *Stochastic Modelling and Applied Probability*. Springer, 1992. (Cited on page 131)

[80] N. Langrené, W. van Ackooij, and F. Bréant. Dynamic constraints for aggregated units: Formulation and application. *IEEE Transactions on Power Systems*, 26(3):1349–1356, August 2011. (Cited on pages 8, 26, and 114)

[81] D. Lautier and F. Raynaud. Statistical properties of derivatives: A journey in term structures. *Physica A: Statistical Mechanics and its Applications*, 390(11):2009–2019, June 2011. (Cited on page 55)

[82] J.-P. Lemor, E. Gobet, and X. Warin. Rate of convergence of an empirical regression method for solving generalized backward stochastic differential equations. *Bernoulli*, 12(5):889–916, 2006. (Cited on pages 16, 19, 22, 33, 36, 39, 96, 130, 154, 158, and 162)

[83] F. Longstaff and E. Schwartz. Valuing American options by simulation: a simple least-squares approach. *Review of Financial Studies*, 14(1):113–147, 2001. (Cited on pages 14, 32, 87, 92, and 172)

[84] M. R. Lyle and R. J. Elliott. A "simple" hybrid model for power derivatives. *Energy Economics*, 31(5):757–767, 2009. (Cited on pages 41, 42, 43, and 47)

[85] B. A. Mamedov. Evaluation of the generalized Goodwin-Staton integral using binomial expansion theorem. *Journal of Quantitative Spectroscopy and Radiative Transfer*, 105:8–11, 2007. (Cited on page 67)

[86] J. Marabel. Pricing digital outperformance options with uncertain correlation. *International Journal of Theoretical and Applied Finance*, 14(5):709–722, 2011. (Cited on pages 169 and 171)

[87] B. Mo, J. Hegge, and I. Wangensteen. Stochastic generation expansion planning by means of stochastic dynamic programming. *IEEE Transactions on Power Systems*, 6(2):662–668, 1991. (Cited on page 87)

[88] M. Mrad. *Méthodes numériques d'évaluation et de couverture des options exotiques multi-sous-jacents : modèles de marché et modèles à volatilité incertaine*. PhD thesis, University of Paris 1 Pantheon-Sorbonne, 2008. (Cited on pages 14, 32, 169, and 178)

[89] A. Nguyen Huu and N. Oudjane. Hedging expected losses on derivatives in electricity Futures markets. In *Commodities, Energy and Environmental Finance*, volume 74 of *Fields Institute Communications*, pages 149–181. Springer, 2015. (Cited on page 169)

[90] J. Obermayer. An analysis of the fundamental price drivers of EU ETS carbon credits. Master's thesis, KTH Royal Institute of Technology, Stockholm, 2009. (Cited on page 116)

[91] H. Pham. On quadratic hedging in continuous time. *Mathematical Methods of Operations Research*, 51(2):315–339, 2000. (Cited on page 50)

[92] H. Pham. *Continuous-time stochastic control and optimization with financial applications*, volume 61 of *Stochastic Modelling and Applied Probability*. Springer, 2009. (Cited on page 136)

[93] G. Pirrong and M. Jermakyan. The price of power - the valuation of power and weather derivatives. *Journal of Banking and Finance*, 32(12):2520–2529, 2008. (Cited on pages 41, 42, and 47)

[94] D. Revuz and M. Yor. *Continuous Martingales and Brownian Motion*. Springer, 1991. (Cited on page 50)

[95] D. Revuz and M. Yor. *Continuous Martingales and Brownian Motion*. Springer, 3rd edition, 1999. (Cited on pages 123 and 124)

[96] E. Schwartz. The stochastic behavior of commodity prices: Implications for valuation and hedging. *The Journal of Finance*, 52(3):923–973, 1997. (Cited on page 41)

[97] M. Schweizer. A guided tour through quadratic hedging approaches. In E. Jouini, J. Cvitanic, and M. Musiela, editors, *Option Pricing, Interest Rates and Risk Management*, pages 538–574, 2001. (Cited on page 50)

[98] R. Seydel. Existence and uniqueness of viscosity solutions for QVI associated with impulse control of jump-diffusions. *Stochastic Processes and their Applications*, 119(10):3719–3748, 2009. (Cited on page 91)

[99] D. Talay. Model risk in finance: some modelling and numerical analysis issues. In A. Bensoussan and Q. Zhang, editors, *Mathematical Modeling and Numerical Methods in Finance*, volume 15 of *Handbook of Numerical Analysis*, pages 3–28, 2008. (Cited on page 169)

[100] X. Tan. A splitting method for fully nonlinear degenerate parabolic PDEs. *Electronic Journal of Probability*, 18(15):1–24, 2013. (Cited on pages 16, 33, 87, 95, 105, 119, and 147)

[101] N. Todorović. *Bewertung Amerikanischer Optionen mit Hilfe von regressionbasierten Monte-Carlo-Verfahren*. PhD thesis, University of Saarland, 2007. (Cited on pages 14, 32, and 92)

[102] J. Tsitsiklis and B. Van Roy. Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks*, 12(4):694–703, 2001. (Cited on pages 14, 32, 87, 92, and 172)

[103] A. Veraart and L. Veraart. Stochastic volatility and stochastic leverage. *Annals of Finance*, 8(2-3):205–233, 2012. (Cited on page 116)

[104] V. Volpe. *The Electricity price modelling and derivatives pricing in the Nord Pool market*. PhD thesis, Università della Svizzera italiana, 2009. (Cited on page 110)

[105] A. Wagner. Residual demand modeling and application to electricity pricing. Technical Report 213, Fraunhofer ITWM, Department for Financial Mathematics, 2012. (Cited on page 116)

[106] H. Working. Price relations between may and new-crop wheat futures at Chicago since 1885. *Wheat Studies of the Food Research Institute*, 10(5):183–228, 1934. (Cited on page 42)

[107] J. Yong and X. Zhou. *Stochastic Controls: Hamiltonian Systems and HJB Equations*. Springer, 1999. (Cited on pages 23, 40, 166, and 168)

[108] D. Zanger. Quantitative error estimates for a least-squares Monte Carlo algorithm for American option pricing. *Finance and Stochastics*, 17(3):503–534, 2013. (Cited on page 158)

[109] J. Zhang. A numerical scheme for BSDEs. *Annals of Applied Probability*, 14(1):459–488, 2004. (Cited on pages 130 and 154)