



HAL
open science

Asymptotic optimal pricing with asymmetric risk and applications in finance

Isaque Santa Brigida Pimentel

► **To cite this version:**

Isaque Santa Brigida Pimentel. Asymptotic optimal pricing with asymmetric risk and applications in finance. Computational Finance [q-fin.CP]. Université Paris Saclay (COmUE), 2018. English. NNT : 2018SACLX059 . tel-01982408

HAL Id: tel-01982408

<https://pastel.hal.science/tel-01982408v1>

Submitted on 15 Jan 2019

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

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

THÈSE DE DOCTORAT

de

L'UNIVERSITÉ PARIS-SACLAY

École doctorale de mathématiques Hadamard (EDMH, ED 574)

Établissement d'inscription : Ecole polytechnique

Laboratoire d'accueil : Centre de mathématiques appliquées de Polytechnique, UMR 7641
CNRS

Spécialité de doctorat : Mathématiques appliquées

ISAQUE SANTA BRIGIDA PIMENTEL

VALORISATION OPTIMALE ASYMPTOTIQUE AVEC RISQUE ASYMÉTRIQUE ET APPLICATIONS EN FINANCE

Date de soutenance : 16 octobre 2018

Rapporteurs : CHRISTIAN BENDER (Saarland University)
JEAN-FRANÇOIS CHASSAGNEUX (Université Paris Diderot)

Jury de soutenance :

FRÉDÉRIC ABERGEL	(CentraleSupélec) Examineur
MIREILLE BOSSY	(Inria Sophia Antipolis) Président
JEAN-FRANÇOIS CHASSAGNEUX	(Université Paris Diderot) Rapporteur
EMMANUEL GOBET	(Ecole Polytechnique) Directeur de thèse
NIZAR TOUZI	(Ecole Polytechnique) Examineur
XAVIER WARIN	(EDF Lab) Co-directeur de thèse

à Edilaine et à mes parents

Remerciements

Suivant une tradition séculaire, je commencerai par honorer la dimension collective de cette expérience doctorale à l'aide d'un inventaire détaillé de remerciements.

Tout d'abord, je tiens à remercier mon directeur de thèse, Emmanuel Gobet, pour m'avoir fait découvrir le monde de la recherche en mathématiques financières et en particulier l'analyse stochastique et les techniques de simulation Monte Carlo. Je lui suis particulièrement reconnaissant pour sa grande patience et son soutien lors de ces trois années, ses nombreux conseils m'ont beaucoup appris.

Ensuite, j'adresse mes remerciements à mon encadrant industriel, Xavier Warin, pour sa disponibilité, sa confiance et son aide inestimable face à plusieurs problèmes numériques complexes. Je lui suis très reconnaissant pour m'avoir encouragé, conseillé et corrigé tout au long de mes travaux de recherche.

Je voudrais adresser tous mes remerciements à Christian Bender et Jean-François Chassagneux d'avoir accepté de rapporter cette thèse. Sachant qu'ils ont des emplois du temps très chargés, je leur suis extrêmement reconnaissant du temps qu'ils m'ont consacré. Je remercie également Mireille Bossy, Frédéric Abergel et Nizar Touzi d'avoir gentiment accepté de faire partie de mon jury.

Je remercie tous les membres du CMAP, du laboratoire FiME, et du département OSIRIS de la R&D d'EDF, plus particulièrement le groupe R33, que j'ai eu la chance de côtoyer. Je tiens à remercier les chercheurs avec qui j'ai eu la chance de collaborer lors de ces années de thèse : Clémence Alasseur, Mathieu Laurière, Huyen Pham, Jorge Zubelli.

Je souhaite remercier les assistantes du laboratoire et d'EDF, en particulier Nasséra et Anne-Laure, pour leur disponibilité, leur efficacité et leur bonne humeur. Je remercie aussi tous les doctorants que j'ai pu côtoyer pendant ces trois années : Alessandro, Benjamin, Côme et David (au CEMRACS à Marseille), Thomas, Emma et Michael (au département OSIRIS d'EDF).

Je souhaite aussi formellement remercier tous les chercheurs m'ayant précédé sur mes travaux de recherche. Également, je remercie les futurs lecteurs, et félicite ceux qui auraient le courage de décodifier les lemmes les plus techniques de ces travaux.

Pour finir, je remercie très affectueusement mon épouse, Edilaine, pour son soutien réconfortant pendant ces trois années. Je remercie mes parents dont la confiance lors de mes nombreuses années d'étude a été très précieuse et toutes les personnes non précédemment mentionnées ayant directement ou indirectement contribué à la réalisation effective de cette thèse.

Valorisation optimale asymptotique avec risque asymétrique et applications en finance

Résumé : Cette thèse est constituée de deux parties qui peuvent être lues indépendamment.

Dans la première partie de la thèse, nous étudions des problèmes de couverture et de valorisation d'options liés à une mesure de risque. Notre approche principale est l'utilisation d'une fonction de risque asymétrique et d'un cadre asymptotique dans lequel nous obtenons des solutions optimales à travers des équations aux dérivées partielles (EDP) non-linéaires.

Dans le premier chapitre, nous nous intéressons à la valorisation et la couverture des options européennes. Nous considérons le problème de l'optimisation du risque résiduel généré par une couverture à temps discret en présence d'un critère asymétrique de risque. Au lieu d'analyser le comportement asymptotique de la solution du problème discret associé, nous avons étudié la mesure asymétrique du risque résiduel intégré dans un cadre Markovian. Dans ce contexte, nous montrons l'existence de cette mesure de risque asymptotique. Ainsi, nous décrivons une stratégie de couverture asymptotiquement optimale via la solution d'une EDP totalement nonlinéaire.

Le deuxième chapitre est une application de cette méthode de couverture au problème de valorisation de la production d'une centrale. Puisque la centrale génère de coûts de maintenance qu'elle soit allumée ou non, nous nous sommes intéressés à la réduction du risque associé aux revenus incertains de cette centrale en se couvrant avec des contrats à terme. Nous avons étudié l'impact d'un coût de maintenance dépendant du prix d'électricité dans la stratégie couverture.

Dans la seconde partie de la thèse, nous considérons plusieurs problèmes de contrôle liés à l'économie et la finance.

Le troisième chapitre est dédié à l'étude d'une classe de problème du type McKean-Vlasov (MKV) avec bruit commun, appelée MKV polynomiale conditionnelle. Nous réduisons cette classe polynomiale par plongement de Markov à des problèmes de contrôle en dimension finie. Nous comparons trois techniques probabilistes différentes pour la résolution numérique du problème réduit : la quantification, la régression par randomisation du contrôle et la régression différée. Nous fournissons de nombreux exemples numériques, comme par exemple, la sélection de portefeuille avec incertitude sur une tendance du sous-jacent.

Dans le quatrième chapitre, nous résolvons des équations de programmation dynamique associées à des valorisations financières sur le marché de l'énergie. Nous considérons qu'un modèle calibré pour les sous-jacents n'est pas disponible et qu'un petit échantillon obtenu des données historiques est accessible. En plus, dans ce contexte, nous supposons que les contrats à terme sont souvent gouvernés par des facteurs cachés modélisés par des processus de Markov. Nous proposons une méthode nonintrusive pour résoudre ces équations à travers les techniques de régression empirique en utilisant seulement l'historique du log du prix des contrats à terme observables.

Mot-clés : Couverture discrète, Risque asymétrique, Optimalité asymptotique, Gestion des coûts, Equations aux Dérivées Partielles Non-linéaires Contrôle du type McKean-Vlasov, Classe polynomiale, Quantification, Régression différée, Randomisation du contrôle, Equations de programmation dynamique discrètes, Régression empirique, Méthodes de ré-échantillonnage, Modèles multi-facteurs, Marché d'électricité.

Asymptotic optimal pricing with asymmetric risk and applications in finance

Abstract: This thesis is constituted by two parts that can be read independently.

In the first part, we study several problems of hedging and pricing of options related to a risk measure. Our main approach is the use of an asymmetric risk function and an asymptotic framework in which we obtain optimal solutions through nonlinear partial differential equations (PDE).

In the first chapter, we focus on pricing and hedging European options. We consider the optimization problem of the residual risk generated by discrete-time hedging in the presence of an asymmetric risk criterion. Instead of analyzing the asymptotic behavior of the solution to the associated discrete problem, we study the integrated asymmetric measure of the residual risk in a Markovian framework. In this context, we show the existence of the asymptotic risk measure. Thus, we describe an asymptotically optimal hedging strategy via the solution to a fully nonlinear PDE.

The second chapter is an application of the hedging method to the valuation problem of the power plant. Since the power plant generates maintenance costs whether it is on or off, we are interested in reducing the risk associated with its uncertain incomes by hedging with forwards contracts. We study the impact of a maintenance cost depending on the electricity price into the hedging strategy.

In the second part, we consider several control problems associated with economy and finance.

The third chapter is dedicated to the study of a McKean-Vlasov (MKV) problem class with common noise, called polynomial conditional MKV. We reduce this polynomial class by a Markov embedding to finite-dimensional control problems. We compare three different probabilistic techniques for numerical resolution of the reduced problem: quantization, control randomization and regress later. We provide numerous numerical examples, such as the selection of a portfolio under drift uncertainty.

In the fourth chapter, we solve dynamic programming equations associated with financial valuations in the energy market. We consider that a calibrated underlying model is not available and that a limited sample of historical data is accessible. In this context, we suppose that hidden factors modeled by Markov processes govern the forward contracts. We propose a non-intrusive method to solve these equations through empirical regression techniques using only the log price history of observable futures contracts.

Keywords: Discrete hedging, Asymmetric risk, Asymptotic optimality, Cost management, Non-linear Partial Differential Equations, McKean-Vlasov control, Polynomial class, Quantization, Regress later, Control randomization, Discrete dynamic programming equations, Empirical regression, Resampling methods, Multi-factor models, Electricity market.

Contents

Introduction générale	1
General Introduction	15
I Asymmetric risk measure for pricing and hedging of options	27
1 Option pricing and hedging using asymmetric risk measure: Asymptotic optimality	29
1.1 Introduction	29
1.2 Notations and assumptions	32
1.3 Asymptotic risk: Existence of the limit $\mathbf{R}_{N,\gamma}$	34
1.3.1 Asymptotic optimality through the PDE nonlinearity	35
1.3.2 Optimal PDE nonlinearity: Study in dimension 1	37
1.4 Asymptotic risk: Proof of the main result	39
1.4.1 Rescaling and conditioning	40
1.4.2 Stochastic expansion and approximation of sensitivities	41
1.4.3 Aggregation and passage to the limit	42
1.4.4 Proof of the stochastic expansion in Subsection 1.4.2.	44
1.4.5 Proof of the almost sure convergence in Subsection 1.4.3	50
1.5 Numerical results	51
1.5.1 Optimal PDE valuation of European options	51
1.5.2 Asymptotic risk: Dependence on the PDE nonlinearity	54
1.5.3 Optimal PDE valuation/hedging: Comparison with the discrete-time solution	56
1.6 Appendix	59
1.6.1 Technical results about Section 1.4	59
1.6.2 Proof of the convergence of the sensitivities	61
2 Asymptotic asymmetric risk measure: Application to physical asset valuation	65
2.1 Introduction	65

2.2	Setting	69
2.2.1	An asymmetric risk valuation with cost management	69
2.2.2	Main assumptions	70
2.2.3	Main results	71
2.2.4	Asymptotic risk with cost management	71
2.2.5	Optimal PDE non-linearity	72
2.3	Existence of the asymptotic risk.	74
2.4	Convex-concave staggered cost models	80
2.4.1	Increasing costs	81
2.4.2	Decreasing costs	81
2.5	Numerical experiments	81
2.5.1	Optimal valuation/hedging rule for increasing costs	83
2.5.2	Optimal valuation/hedging rule for decreasing costs	84
2.6	Conclusions and extensions	86
II	Numerical methods in stochastic control	89
3	Polynomial conditional McKean-Vlasov control problems: Some probabilistic numerical methods	91
3.1	Introduction	91
3.2	Notations and Assumptions	93
3.2.1	Main assumptions	93
3.2.2	Markovian embedding	95
3.3	Probabilistic numerical methods	95
3.3.1	Value and Performance iteration	96
3.3.2	Approximation of conditional expectations	97
	Regression Monte Carlo	97
	Quantization	100
3.3.3	Training points design	103
	Exploitation only strategy	103
	Explore first, exploit later	103
3.3.4	Optimal control searching	104
	Low cardinality control set	104
	High cardinality/continuous control space	104
3.3.5	Upper and lower bounds	105
3.3.6	Pseudo-codes	105
	Pseudo-code for a Regress-Later-based algorithm	105

	Pseudo-code for a Control Randomization-based algorithm	105
	Pseudo-code for a Quantization-based algorithm	106
3.4	Applications	106
3.4.1	Portfolio Optimization under drift uncertainty	106
3.4.2	Interbank systemic risk with partial observation	110
3.5	Numerical results	111
3.5.1	Portfolio Optimization	111
3.5.2	Interbank systemic risk	116
3.6	Discussion	121
4	A non-intrusive stratified resampler for multi-factor models: Application in energy market	123
4.1	Introduction	123
4.2	Setting	126
4.2.1	Dynamic programming equation	126
4.2.2	Modeling of the observable and hidden processes	126
4.2.3	Non-Intrusive resampling	129
4.3	The NISR scheme	130
4.3.1	Extraction	130
4.3.2	Stratification	130
4.3.3	NISR-regression Monte Carlo algorithm	131
4.4	Applications to Energy Market	133
4.4.1	Bermudan options	133
4.4.2	Swing options	137
	Bibliography	139

Introduction générale

Le présent travail est consacré à la solution des problèmes de valorisation et de couverture des options liées au marché de l'électricité. Il est également dédié aux méthodes numériques pour une classe de problèmes de contrôle stochastique du type McKean-Vlasov liés à des questions financières, comme par exemple, la sélection de portefeuille sous incertitude de tendance du sous-jacent. Dans une dernière partie, n'ayant aucun accès à un modèle multifactoriel complètement calibré pour les prix à terme, nous proposons un algorithme Monte Carlo stratifié non-intrusif pour résoudre le problème de valorisation des options Bermudéennes et Swing.

Dans la Partie I, nous étudions des problèmes de couverture et de valorisation des options européennes en tenant compte d'une mesure de risque asymétrique. L'approche principale est l'utilisation d'un cadre asymptotique afin d'obtenir des politiques de valorisation/couverture optimales à travers des équations aux dérivées partielles nonlinéaires.

Dans le Chapitre 1, nous considérons l'erreur résiduelle produite par un trading en temps discret. Nous supposons qu'un trader veut trouver des stratégies pour le portefeuille de couverture de sorte que l'erreur de discrétisation soit faible pénalisant les scénarios de pertes plus ceux de gains : nous cherchons à obtenir une politique de valorisation/couverture avec un grand nombre de dates de rebalancement du portefeuille pour laquelle l'erreur de couverture induite dans le cadre d'une mesure de risque asymétrique est minimisée. Plus précisément, nous définissons une stratégie optimale en terme de la solution d'un problème d'optimisation lié à l'erreur asymptotique. Dans ce cas, la fonction valeur optimale (resp., couverture) est la solution (resp., le gradient de la solution) d'une EDP totalement non linéaire, dont le terme nonlinéaire dépend de la dérivée seconde de la solution. En fait, ce terme de nonlinéarité caractérise notre stratégie asymptotiquement optimale. En dimension 1, une formulation quasi-explicite pour ce terme est obtenue.

Dans le Chapitre 2, nous étudions le problème d'évaluation d'une centrale électrique en tant qu'une application de la méthode de couverture précédente avec de faibles coûts proportionnels par pas de temps. Considérant que certains coûts (d'entretien) doivent être payés indépendamment du fait qu'une centrale électrique produise ou non de l'électricité, nous sommes intéressés à valoriser un revenu associé à la vente de sa production dans le futur. En fait, le propriétaire de l'usine est prêt à réduire le risque d'un revenu incertain en effectuant des opérations de couverture sur le marché à terme. Par conséquent, la stratégie de couverture optimale est obtenue en minimisant le risque asymptotique associé à une erreur de discrétisation, y compris les coûts proportionnels par pas de temps. Nous avons étudié l'impact des coûts (de maintenance) dans la stratégie de couverture, lorsque ces coûts dépendent du prix de l'électricité.

Dans la Partie II, nous avons étudié plusieurs méthodes numériques pour résoudre une classe de problèmes de contrôle associés à des interactions à champ moyen avec informations partielles. Nous avons également étudié une méthode de ré-échantillonnage stratifiée non-intrusive pour résoudre des équations de programmation dynamique discrètes en présence de peu de données

historiques.

Dans le Chapitre 3, nous étudions une classe de contrôle du type McKean-Vlasov (MKV) en présence d'un bruit commun, appelé MKV polynomial conditionnel. Cette classe est une généralisation de la classe de contrôle MKV stochastique linéaire quadratique. Ces problèmes découlent de la formulation asymptotique d'un équilibre coopératif dans une large population de joueurs avec une interaction à champ moyen ou d'un problème de contrôle avec fonction de coût dépendant d'une fonctionnelle nonlinéaire de la distribution du processus d'état. Après avoir réduit cette classe polynomiale par un plongement de Markov à des problèmes de contrôle en dimension finie, nous proposons une discussion et une comparaison de trois méthodes probabilistes différentes pour la résolution numérique du problème réduit : quantification, régression par contrôle aléatoire et régression différée. Nous avons résolu numériquement plusieurs problèmes liés à ce contrôle MKV polynomial : la sélection et la liquidation du portefeuille sous incertitude de la tendance ainsi qu'un modèle de risque interbancaire systémique avec observation partielle.

Dans le Chapitre 4, nous nous intéressons à la résolution d'équations de programmation dynamique (DPE) associées à une valorisation financière sur le marché de l'énergie. En particulier, nous nous concentrons sur le pricing des options Bermudéennes et Swing, où les actifs sous-jacents sont des contrats à terme. Nous recherchons une méthode stratifiée non-intrusive pour résoudre ces équations en utilisant uniquement des données historiques, sans un modèle sous-jacent calibré. Sur le marché de l'énergie, les contrats à terme sont souvent gouvernés par des facteurs de Markov cachés (modélisation multifactorielle). Par conséquent, le problème de valorisation de ces options produit des DPE dépendants de ces facteurs non observables. Puisque nous avons accès aux prix à terme mais pas aux facteurs cachés, nous proposons une méthode de rééchantillonnage uniquement sur les contrats observables.

I. Mesure de risque asymétrique pour la valorisation et couverture d'options

1. Valorisation et couverture d'options à l'aide d'une mesure de risque asymétrique : Optimalité asymptotique

En finance, la valorisation et la couverture des flux financiers futurs sont des préoccupations majeures. D'un point de vue théorique, ces problèmes sont bien établis (voir par exemple [KS98]). En pratique, les traders couvrent leurs contrats en échangeant uniquement à des dates discrètes, disons $t_0 = 0 < t_1 < \dots < t_N = T$, ce qui génère un risque résiduel. Ici, ils ont l'intention de couvrir le flux H_T à l'instant T en utilisant d produits de couverture dont les prix sont indiqués par $X = (X^{(1)}, \dots, X^{(d)})$. Pour cela, nous définissons le *risque local* \mathcal{E}_n associé aux dates de trading t_n et t_{n+1} donnés par

$$\mathcal{E}_n = V_{t_{n+1}} - V_{t_n} - \vartheta_{t_n} \cdot X_{t_{n+1}} - X_{t_n},$$

où V désigne le processus d'évaluation, $\vartheta = (\vartheta^{(1)}, \dots, \vartheta^{(d)})$ indique le processus de couverture avec $\vartheta^{(i)}$ étant le nombre d'actions investi dans le i -ème instrument de couverture.

Nous cherchons à trouver les règles d'évaluation/couverture (V, ϑ) minimisant ce risque résiduel en utilisant une fonction de risque ℓ . En comparaison des résultats existants (voir par exemple, [Sch99], pour la minimisation du risque local quadratique), nous choisissons une fonction ℓ pénalisant les profits ($\mathcal{E}_n < 0$) et les pertes ($\mathcal{E}_n > 0$) de manière asymétrique. Dans ce

contexte, nous étudions le *risque local intégré* \mathbf{R}_N sous la forme

$$\mathbf{R}_N(V, \vartheta) = \sum_{n=0}^{N-1} \mathbb{E}[\ell(\mathcal{E}_n)],$$

où la **fonction de risque** ℓ est choisie comme

$$\ell_\gamma(y) = (1 + \gamma \operatorname{sgn}(y))^2 y^2 / 2, \quad \gamma \in (0, 1).$$

Le paramètre γ , dans la fonction de risque ℓ_γ , représente l'aversion au risque des investisseurs pénalisant davantage les pertes que les profits (voir la Figure 1.1).

Approche asymptotique standard. Dans ce contexte, nous visons à étudier l'asymptotique du minimum

$$\min_{(V, \vartheta) \in \mathcal{A}} \mathbf{R}_N(V, \vartheta), \quad (1)$$

avec un grand nombre N de dates de couverture dans le cas équidistant $t_n = n\Delta t$ avec pas de temps $\Delta t = T/N$. Le minimum (1) est calculé sur l'ensemble \mathcal{A} de toutes les paires de processus (V, ϑ) intégrable, adaptées à la filtration $(\mathcal{F}_t)_{t \geq 0}$, sous la contrainte de réplcation $V_T = H_T$.

Il existe dans la littérature quelques résultats dans cette direction :

- Dans [Pha00], l'auteur étudie une fonction \mathbb{L}_p de risque des pertes et avec un nombre fixe des dates de couverture ;
- Dans [AM11], les auteurs analysent des stratégies pseudo-optimales et obtiennent des résultats asymptotiques dans le cas d'une fonction de risque de classe C^3 .

Nous remarquons que la **discontinuité de la dérivée seconde** ℓ''_γ rend l'analyse difficile et change complètement la nature des résultats ultérieurs. En bref, les références précédentes prennent en compte des contextes et difficultés différents des nôtres.

Approche asymptotique avec évaluation par EDP. Au lieu de résoudre le problème (1) dans le régime asymptotique $N \rightarrow +\infty$, nous abordons un problème modifié en utilisant une évaluation par EDP. Nous supposons que les instruments de couverture X sont modélisés par

$$dX_t = \mu(t, X_t) + \sigma(t, X_t) dW_t,$$

où W est un mouvement brownien et les flux financiers futurs sont de la forme $H_T = g(X_T)$. Ensuite, nous supposons que le flux aléatoire est **évalué de manière exogène** par un processus d'évaluation $V_t = v(t, X_t)$, pour une fonction v donnée. La fonction v joue le rôle d'un prix de référence. Compte tenu de cette valorisation exogène, le trader déterminera comment se couvrir sur chaque intervalle $[t_n, t_{n+1}]$ en choisissant une règle d'évaluation/couverture adaptée $(\tilde{V}_{t_n}, \tilde{\vartheta}_{t_n})$ et en considérant le *risque local conditionnel* $\mathbf{R}_{n,\gamma}$

$$\mathbf{R}_{n,\gamma} = \mathbb{E}[\ell_\gamma(V_{t_{n+1}} - \tilde{V}_{t_n} - \tilde{\vartheta}_{t_n} \cdot X_{t_{n+1}} - X_{t_n}) \mid \mathcal{F}_{t_n}].$$

Inspiré par la relation entre les valorisations dynamiques des risques et les équations différentielles stochastiques rétrogrades nonlinéaires (voir par exemple [Cre13]), le trader paramétrise sa règle de valorisation/couverture par une fonction f , éventuellement non linéaire, telle que

$$\tilde{V}_{t_n} = u^{t_{n+1}}(t_n, X_{t_n}), \quad \tilde{\vartheta}_{t_n} = D_x u^{t_{n+1}}(t_n, X_{t_n}),$$

où $u^{t_{n+1}}$ est la solution définie sur l'intervalle $[t_n, t_{n+1}]$ de la f -EDP suivante :

- La fonction $u^{t_{n+1}} : [t_n, t_{n+1}] \times \mathbb{R}^d \rightarrow \mathbb{R}$ satisfait

$$u_t^{t_{n+1}}(t, x) + \frac{1}{2} \text{Tr}[\sigma \sigma^\top D_x^2 u^{t_{n+1}}](t, x) + f(t, x, u^r(t, x), D_x u^{t_{n+1}}(t, x), D_x^2 u^{t_{n+1}}(t, x)) = 0, \quad (2)$$

pour tout $(t, x) \in [t_n, t_{n+1}] \times \mathbb{R}^d$, avec condition terminale $u^{t_{n+1}}(t_{n+1}, x) = v(t_{n+1}, x)$ à l'instant t_{n+1} ;

- Pour des fonctions continues $\sigma : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$, $f : [0, T] \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$; et un prix de référence $v(t_{n+1}, \cdot)$ à l'instant t_{n+1} ;

Dans ce contexte, il est maintenant naturel d'analyser le comportement asymptotique du *risque local conditionnel intégré* après une renormalisation appropriée

$$\mathbf{R}_{N,\gamma}(v, f) = \frac{1}{\Delta t} \sum_{n=0}^{N-1} \mathbb{E}[\mathbf{R}_{n,\gamma}].$$

Notre contribution principale au Chapitre 1 est le résultat suivant (voir Théorème 1.6, page 34).

Résultat Principal 1 (Theorème 1.6). *Soit $B : [0, 1] \times \Omega \rightarrow \mathbb{R}^d$ un autre mouvement brownien standard indépendant de W . Sous les hypothèses 1.1-1.5, la limite*

$$\mathbf{R}_\gamma(v, f) := \lim_{N \rightarrow \infty} \mathbf{R}_{N,\gamma}(v, f)$$

existe et est donnée par

$$\mathbf{R}_\gamma(v, f) = \mathbb{E} \left[\int_0^T \int_0^1 \ell''_\gamma \left(\int_0^\theta B_\theta^\top G_t dB_{\theta'} - F_t \theta \right) \left(F_t^2 \theta - F_t \int_0^\theta B_\theta^\top G_t dB_{\theta'} + \|G_t B_\theta\|^2 / 2 \right) d\theta dt \right],$$

où

$$F_t = f(t, X_t, v(t, X_t), D_x v(t, X_t), D_x^2 v(t, X_t)) \in \mathbb{R},$$

$$G_t = (\sigma^\top (D_x^2 v) \sigma)(t, X_t) \in \mathbb{R}^{d \times d}.$$

Par conséquent, nous discutons de l'existence d'un terme de nonlinéarité optimal pour la f^* -EDP telle que l'évaluation par f^* -EDP minimise le risque asymptotique dans le sens : $\mathbf{R}_\gamma(v, f^*) \leq \mathbf{R}_\gamma(v, f)$, pour toute nonlinéarité admissible f . De plus, ce terme de nonlinéarité optimal f^* est **quasi-explicit en dimension 1** (voir (1.22), page 39) et dépend du paramètre de risque γ , de la dérivée seconde valorisation de référence et de la volatilité du processus de prix. Par conséquent, un choix naturel et consistant pour la règle de valorisation/couverture peut être la solution à la f^* -EDP (2) avec le payoff $g : \mathbb{R}^d \rightarrow \mathbb{R}$ étant la condition terminale à l'horizon T .

Ce chapitre est tiré d'un article rédigé en collaboration avec Emmanuel Gobet et Xavier Warin [GPW18].

2. Mesure de risque asymétrique asymptotique : Application à valorisation des actifs physiques

Dans le Chapitre 2, nous fournissons des politiques d'évaluation/couverture pour les revenus futurs liés à la production de centrales électriques à l'aide d'une évaluation asymétrique du risque . Depuis la déréglementation des marchés de l'énergie, plusieurs marchés spot et futurs

ont été créés pour échanger de l'électricité. Après cela, les propriétaires de centrales électriques commencent à faire face au problème de **l'évaluation la production de leur centrales**, qui dépend du prix spot de l'électricité dans le futur.

En pratique, les propriétaires des centrales effectuent une couverture discrète à des dates, $t_0 = 0 < t_1 < \dots < t_N = T$ en utilisant un actif de couverture X afin de réduire le risque associé à un aléa. revenu H_T à une date future T . Cette couverture discrète génère un risque résiduel :

$$\mathcal{E} = H_T - \sum_{n=0}^{N-1} \vartheta_{t_n} (X_{t_{n+1}} - X_{t_n}),$$

où

- ϑ indique le nombre d'actions investies dans le actif de couverture ;
- X est un contrat à terme $F(\cdot, T)$ avec livraison T .

Ensuite, ils cherchent à trouver la règle d'évaluation/couverture en tenant compte du bilan local \mathcal{E}_n pénalisant ($\mathcal{E}_n < 0$) via une fonction de risque asymétrique ℓ :

$$\mathcal{E}_n = V_{t_{n+1}} - V_{t_n} - \vartheta_{t_n} (X_{t_{n+1}} - X_{t_n}),$$

où V représente le processus d'évaluation sous la contrainte de réplication $V_T = H_T$. Notez que les propriétaires des centrales :

- veulent obtenir un risque résiduel \mathcal{E} avec un petit écart-type ;
- préfèrent les scénarios profitables où $\mathcal{E}_n > 0$.

En réalité, la centrale génère des coûts fixes, qu'elle produise de l'électricité ou non. Ces coûts sont appelés ici **coûts fixes** :

- ils comprennent les coûts d'investissement et d'amortissement ;
- ils excluent les coûts de carburants (charbon, gaz, uranium).

Les producteurs d'électricité sont contraints de trouver aujourd'hui un prix équivalent au revenu futur généré par leur centrale afin de l'enregistrer dans leur bilan comptable.

En fait, ils **ne font face à aucun risque financier** en raison de ce revenu. D'autre part, parce qu'ils ne peuvent pas augmenter leur revenu en influençant le prix spot de l'électricité S , ils préfèrent obtenir une valeur certaine en vendant leur production future au lieu de recevoir le montant aléatoire positif $H_T = g(S_T)$.

Sur le plan économique, lorsque le prix spot de l'électricité est élevé, nous exigeons que la centrale fasse face à une demande plus élevée. En effet, un prix élevé est produit par une demande croissante. Par conséquent, nous augmentons la production en démarrant plus souvent la centrale électrique. C'est la raison pour laquelle nous supposons une dépendance des coûts c sur le prix au comptant de l'électricité S .

Valorisation asymétrique de risque. Considérant les temps de couverture équidistants $t_n = n\Delta t$ avec le pas de temps $\Delta t = T/N$, nous prenons en compte les coûts fixes proportionnels au pas de temps Δt en soustrayant $c(S_{t_n})\Delta t$ du bilan locale \mathcal{E}_n . Dans un cadre analogue au Chapitre 1, nous étudions l'asymptotique du *risque intégré* \mathbf{R}_N

$$\mathbf{R}_N = \frac{1}{\Delta t} \sum_{n=0}^{N-1} \mathbb{E} \left[\ell \left(V_{t_{n+1}} - V_{t_n} - \vartheta_{t_n} (X_{t_{n+1}} - X_{t_n}) - c(S_{t_n})\Delta t \right) \right],$$

Ici, nous définissons

$$\ell_\gamma(y) = (1 + \gamma \operatorname{sgn}(y))^2 y^2 / 2, \quad \gamma \in (-1, 0),$$

où le paramètre γ représente le désir du producteur de réduire l'aléa des revenus futurs. Contrairement au paramètre de risque du Chapitre 1, ce paramètre γ prend des valeurs dans $(-1, 0)$ (voir Figure 2.1).

En vue d'étudier le risque intégré $\mathbf{R}_{N,\gamma}$ dans le régime asymptotique $N \rightarrow +\infty$ avec une fonction de risque asymétrique ℓ_γ , nous suivons une approche d'évaluation de EDP décrite au Chapitre 1. Nous supposons l'existence d'un processus d'évaluation exogène $V_t = v(t, X_t)$ défini par une fonction v . Par exemple, v est un prix de référence obtenu par le producteur d'électricité lorsque aucun coût n'est pris en compte. Maintenant, le propriétaire de la centrale détermine sa position de couverture sur chaque intervalle $[t_n, t_{n+1}]$ en choisissant une règle d'évaluation/couverture adaptée $(\tilde{V}_{t_n}, \tilde{\vartheta}_{t_n})$. Pour cela, nous considérons le *risque conditionnel* $\mathbf{R}_{n,\gamma}$ en utilisant la fonction de risque ℓ_γ

$$\mathbf{R}_{n,\gamma} = \mathbb{E} \left[\ell_\gamma \left(V_{t_{n+1}} - \tilde{V}_{t_n} - \tilde{\vartheta}_{t_n} (X_{t_{n+1}} - X_{t_n}) - c(S_{t_n}) \Delta t \right) \middle| \mathcal{F}_{t_n} \right].$$

Cette fois, la règle d'évaluation/couverture $(\tilde{V}_{t_n}, \tilde{\vartheta}_{t_n})$ du producteur d'électricité est donnée par

$$\tilde{V}_{t_n} = u^{t_{n+1}}(t_n, X_{t_n}), \quad \tilde{\vartheta}_{t_n} = u_x^{t_{n+1}}(t_n, X_{t_n}),$$

et est paramétrée par une fonction f , éventuellement non linéaire, à travers la solution $u^{t_{n+1}} : [t_n, t_{n+1}] \times \mathbb{R}$ de la f -EDP avec condition terminale $u^{t_{n+1}}(t_{n+1}, \cdot) = v(t_{n+1}, \cdot)$ à l'instant t_{n+1} (de façon analogue à la version unidimensionnelle de (1.5) mais avec terme de diffusion égal à $\bar{\sigma}(t, T)x$, avec une fonction de volatilité $\bar{\sigma}(t, T) = \bar{\sigma}_0 e^{-a_0(T-t)}$).

Ainsi comme dans le Chapitre 1, nous analysons le régime asymptotique du risque conditionnel intégré

$$\mathbf{R}_{N,\gamma}(v, f) = \frac{1}{\Delta t} \sum_{n=0}^{N-1} \mathbb{E}[\mathbf{R}_{n,\gamma}],$$

lorsque le nombre N des dates de couverture tend vers l'infini.

Notre contribution principale au Chapitre 2 est le résultat suivant (voir Théorème 2.9, page 71). C'est une version du Théorème 1.6 en présence de coûts fixes proportionnels au pas de temps.

Résultat Principal 2 (Théorème 2.9). *Soit $B : [0, 1] \times \Omega \rightarrow \mathbb{R}$ un autre mouvement Brownien standard indépendant de W . Sous les hypothèses 2.3-2.8, la limite de $\mathbf{R}_{N,\gamma}(v, f)$ quand $N \rightarrow \infty$ existe et est donnée par*

$$\mathbf{R}_\gamma(v, f) = \mathbb{E} \left[\int_0^T \int_0^1 \ell_\gamma'' \left(J_1(v, (t, X_t)) \frac{B_\theta^2 - \theta}{2} - J_2(v, (t, X_t)) \theta \right) \left(|J_2(v, (t, X_t))|^2 \theta - J_2(v, (t, X_t)) J_1(v, (t, X_t)) \frac{B_\theta^2 - \theta}{2} + |J_1(v, (t, X_t))|^2 \frac{B_\theta^2}{2} \right) d\theta dt \right],$$

où

$$\begin{aligned} J_2(v, (t, X_t)) &= f(t, X_t, v(t, X_t), v_x(t, X_t), v_{xx}(t, X_t)) + c(X_t) \in \mathbb{R}, \\ J_1(v, (t, X_t)) &= \bar{\sigma}^2(t, T) X_t^2 v_{xx}(t, X_t) \in \mathbb{R}. \end{aligned}$$

Comme dans le Chapitre 1, l'idée est d'obtenir un terme de nonlinéarité optimal f^* minimisant

le risque asymptotique :

$$\mathbf{R}_\gamma(v, f^*) \leq \mathbf{R}_\gamma(v, f),$$

pour toute nonlinéarité admissible f . Une fois de plus, ce terme de nonlinéarité optimal f^* est explicite et dépend de γ , de la dérivée seconde de v et de la fonction de coût $c(x)$. Enfin, le producteur d'électricité choisira la règle suivante

$$V_{t_n} = v^*(t_n, X_{t_n}), \quad \vartheta_{t_n} = v_x^*(t_n, X_{t_n}),$$

où v^* est la solution $v^* : [0, T] \times \mathbb{R}^d$ de la f^* -EDP avec la condition terminale $v^*(T, \cdot) = g(\cdot)$ à l'instant T .

Pour les résultats numériques, nous prenons un coût fixe dépendant du niveau du prix spot de l'électricité comme une fonction convexe (ou concave) $c(x)$. Nous calculons la solution numérique de la f^* -EDP pour une fonction payoff du type call et différents modèles pour $c(x)$.

Ce chapitre est tiré d'un article en cours d'écriture en collaboration avec Clémence Alasseur, Emmanuel Gobet et Xavier Warin.

II. Méthodes numériques en contrôle stochastique

3. Problème de contrôle du type McKean-Vlasov avec bruit commun : Quelques méthodes probabilistes numériques

Dans le Chapitre 3, nous nous intéressons au problème de contrôle du type McKean-Vlasov (MKV) sous observation partielle et bruit commun. La formulation est décrite de la manière suivante. Sur un espace de probabilité $(\Omega, \mathcal{F}, \mathbb{P})$ équipé de deux mouvements browniens indépendants B et W^0 , nous considérons la dynamique stochastique contrôlée du type McKean-Vlasov dans \mathbb{R}^n :

$$dX_s = b\left(X_s, \mathbb{P}_{X_s}^{W^0}, \alpha_s\right)ds + \sigma\left(X_s, \mathbb{P}_{X_s}^{W^0}, \alpha_s\right)dB_s + \sigma_0\left(X_s, \mathbb{P}_{X_s}^{W^0}, \alpha_s\right)dW_s^0,$$

où

- $\mathbb{P}_{X_s}^{W^0}$ indique la distribution de X_s conditionnellement à W^0 ;
- $\mathbb{F}^0 = (\mathcal{F}_t^0)_{t \geq 0}$ est la filtration naturelle générée par W^0 ;
- le contrôle α est \mathbb{F}^0 -progressive à valeurs dans un espace polonais A .

La fonction de coût associée à l'équation stochastique du type McKean-Vlasov (9) pour un contrôle α est

$$J(\alpha) = \mathbb{E} \left[\int_0^T f(X_t, \mathbb{P}_{X_t}^{W^0}, \alpha_t) dt + g(X_T, \mathbb{P}_{X_T}^{W^0}) \right],$$

et l'objectif est de minimiser sur un ensemble admissible \mathcal{A} de contrôles la fonctionnel de coût :

$$V_0 = \inf_{\alpha \in \mathcal{A}} J(\alpha).$$

Notre but est **d'étudier certaines classes de problèmes de contrôle MKV**, qui peuvent être réduits à des problèmes de dimension finie en vue d'une résolution numérique.

Une classe de modèles polynomiaux et un plongement markovien. Premièrement, nous considérons une classe de modèles où les coefficients de l'équation MKV sont linéaires par rapport

la variable d'état X (voir (3.9)), alors que les fonctions de coût courant et de coût terminal sont polynomiales dans la variable d'état dans le sens suivant

$$f(x, \mu, a) = f_0(\mu, a) + \sum_{k=1}^p f_k(\mu, a)x^k, \quad g(x, \mu) = g_0(\mu) + \sum_{k=1}^p g_k(\mu)x^k,$$

pour un entier $p \geq 1$. Deuxièmement, nous supposons que tous les coefficients dépendent de μ à travers ses premiers moments p (voir (3.10)).

Étant donné le processus contrôlé $X = X^\alpha$ solution de la dynamique stochastique du type McKean-Vlasov (9), nous notons

$$Y_t^k = \mathbb{E}[X_t^k | W^0], \quad k = 1, \dots, p.$$

A partir des hypothèses linéaires et polynomiales (par la formule d'Itô et les espérances conditionnelles), nous dérivons la dynamique de (Y^1, Y^2, \dots, Y^p) comme

$$\left\{ dY_t^k = B_k(Y_t^1, Y_t^2, \dots, Y_t^p, \alpha_t)dt + \Sigma_k(Y_t^1, Y_t^2, \dots, Y_t^p, \alpha_t)dW_t^0, \quad k = 1, \dots, p, \right. \quad (3)$$

pour certaines fonctions B_k, Σ_k (voir page 95) dépendant des moments conditionnels (Y^1, Y^2, \dots, Y^p) , alors que la fonctionnelle de coût est écrite sous une forme réduite

$$J(\alpha) = \mathbb{E} \left[\int_0^T \bar{f}(Y_t^1, Y_t^2, \dots, Y_t^p, \alpha_t)dt + \bar{g}(Y_T^1, Y_T^2, \dots, Y_T^p) \right]. \quad (4)$$

Le problème de contrôle du type McKean-Vlasov est **réduit** dans ce cadre polynomial à des problèmes de contrôle **fini-dimensionnels** avec des variables contrôlées (Y^1, Y^2, \dots, Y^p) qui sont \mathbb{F}^0 -adaptées.

Méthodes numériques probabilistes. En vue de résoudre le problème réduit (3)-(4), nous choisissons la formulation suivante. Soit $Z = (Y^1, Y^2, \dots, Y^p)$ un processus contrôlé par un contrôle adapté α prenant des valeurs dans A , solution de l'équation

$$dZ_t^\alpha = b(Z_t^\alpha, \alpha_t)dt + \sigma_0(Z_t^\alpha, \alpha_t)dW_t^0$$

et

$$J(t, z, \alpha) = \mathbb{E} \left[\int_t^T f(Z_s^\alpha, \alpha_s)ds + g(Z_T^\alpha) \middle| Z_t^\alpha = z \right],$$

une version dynamique de la fonctionnelle de coût $J(\alpha)$.

En ayant une discrétisation en temps $0 = t_0, t_1, \dots, t_N = T$, nous écrivons l'approximation d'Euler de Z_t^α :

$$Z_{t_{n+1}}^\alpha = Z_{t_n}^\alpha + b(Z_{t_n}^\alpha, \alpha_{t_n})\Delta t + \sigma_0(Z_{t_n}^\alpha, \alpha_{t_n})\Delta W_{t_n}^0$$

et l'équivalent discret de $J(t, z, \alpha)$:

$$J(t_n, z, \alpha) = \mathbb{E} \left[\sum_{i=n}^N f(Z_{t_i}^\alpha, \alpha_{t_i})\Delta t + g(Z_{t_N}^\alpha) \middle| Z_{t_n}^\alpha = z \right].$$

Maintenant, la fonction valeur $V(t_n, z) = \sup_{(\alpha_{t_i})_{i=n}^N \in \mathcal{A}} J(t_n, z, \alpha)$ est représentée à travers l'équation de programmation dynamique suivante, étant donnée la condition terminale $g(z)$

connue,

$$\begin{aligned} V(T_N, z) &= g(z), \\ V(t_n, z) &= \sup_{\alpha} \left\{ f(Z_{t_n}^{\alpha}, \alpha) \Delta t + \mathbb{E}_{\alpha} [V(t_{n+1}, Z_{t_{n+1}}) | Z_{t_n} = z] \right\}. \end{aligned} \quad (5)$$

L'équation de programmation dynamique (5) conduit à des méthodes numériques qui approchent itérativement la fonction valeur de manière rétrograde, à partir de la condition terminale. La principale difficulté dans la mise en œuvre d'une telle approche réside dans **l'estimation des espérances conditionnelles** $\mathbb{E}_{\alpha} [V(t_{n+1}, Z_{t_{n+1}}) | Z_{t_n} = z]$.

Notre première contribution dans le Chapitre 3 est la description (voir Section 3.3) et l'application (voir Section 3.4) suivantes.

Résultat Principal 3. *Nous décrivons trois méthodes numériques résolvant des problèmes du type MKV conditionnels (dans leur version réduite) :*

- *les techniques de régression de Monte Carlo (par randomisation de contrôle et par régression différée) sont une famille d'algorithmes dont l'efficacité repose sur le choix des fonctions de base utilisées pour projeter les futures fonctions valeur (voir Algorithmes 1 et 2) ;*
- *les techniques de quantification approchent le processus contrôlé $Z_{t_n}^{\alpha}$ avec une chaîne de Markov particulière à états finis pour laquelle les espérances peuvent être rapidement estimées (voir Algorithme 3).*

Nous appliquons chacune de ces méthodes à trois applications issues de problèmes de contrôle du type polynomial MKV sous hypothèse d'observation partielle et de bruit commun :

- l'Optimisation de portefeuille sous l'incertitude de tendance (voir Sous-section 3.4.1), où la tendance de l'actif sous-jacent est inconnue et inobservable. D'abord, nous considérons le problème d'un trader prêt à liquider un grand nombre de parts d'actif dans un temps fini T en faisant face à un coût d'exécution et à un impact sur les prix du marché. Ensuite, nous regardons le problème de sélection d'une stratégie de portefeuille maximisant l'utilité de la richesse finale.
- le Risque intersystémique avec observation partielle (voir Sous-section 3.4.2). Nous supposons que les réserves monétaires des N banques, prêtant et empruntant les unes aux autres, satisfont un système de diffusion avec une interaction à champ moyen et un bruit commun. Ensuite, un planificateur social (la banque centrale, par exemple) en observant seulement le bruit commun agit sur l'intensité des interactions interbancaires afin de minimiser l'écart entre la réserve de chaque banque et la moyenne globale.

Résultats numériques. Notre deuxième contribution au Chapitre 3 est les expériences numériques suivantes (voir Section 3.5). Nous avons présenté les résultats de trois exemples différents d'applications résumées ci-dessous

- Nous avons constaté que les algorithmes de régression de Monte Carlo fonctionnent correctement dans les problèmes de contrôle de la tendance. Dans ces problèmes, ils sont beaucoup plus rapides que la quantification pour une précision similaire. En particulier, nous avons remarqué que la Régression différée (regress later) est plus fiable que la randomisation de contrôle : le choix d'une distribution uniforme des points d'apprentissage sur un intervalle approprié suffit pour obtenir des estimations de haute qualité.

- D'autre part, la Randomisation du contrôle est sensible au choix de la distribution du contrôle randomisé et peu de répétitions sont nécessaires avant de trouver une bonne distribution de contrôle. Nous avons également essayé d'utiliser les méthodes d'itération de performance ou de re-calcul de trajectoires. Mais, sur les exemples considérés, ces méthodes prenaient beaucoup de temps et ne contribuaient pas beaucoup à la précision.
- Les techniques de Quantification ont fourni les résultats les plus stables et les plus précis pour les trois différents cas de problèmes de contrôle. De plus, nous pouvons choisir la grille pour quantifier le processus contrôlé. Il est possible d'exploiter cette fonctionnalité lorsque nous avons une idée approximative de l'endroit où le processus contrôlé doit être guidé par la stratégie optimale (voir le problème de Liquidation de portefeuille). Par conséquent, nous devrions construire une grille avec de nombreux points situés là où le processus est censé aller.

Ce chapitre est tiré d'un article rédigé en collaboration avec Alessandro Balata, Côme Huré, Mathieu Laurière et Huyen Pham [BHL⁺18].

4. Un schéma de ré-échantillonnage stratifié non-intrusif pour les modèles multi-facteurs : Application en marché de l'énergie

Les équations stochastiques de programmation dynamique sont liées à la résolution de problèmes non-linéaires (contrôles stochastiques ou EDP non-linéaires) qui se posent dans presque tous les domaines de la science, de la gestion des réservoirs d'eau à la finance.

Dans le Chapitre 4, nous visons à résoudre des équations de programmation dynamique (DPE) liées à des évaluations financières sur le marché de l'énergie. Nous nous intéressons ici, par exemple, à la valorisation des options Bermudéennes ou Swing, où les actifs sous-jacents sont des contrats à terme. Ensuite, nous prévoyons de développer un algorithme nonintrusif pour résoudre ces DPEs en utilisant uniquement les données observables sans la calibration complète d'un modèle.

Dans un premier temps, nous traitons une DPE discrète prenant la forme suivante :

$$Y_N = \tilde{g}_N(X_N),$$

$$Y_i = \mathbb{E}[\tilde{g}_i(Y_{i+1}, \dots, Y_N, X_i, \dots, X_N) \mid X_i], \quad i = N - 1, \dots, 0,$$

où X est une chaîne de Markov dans \mathbb{R}^n ; et \tilde{g}_N, \tilde{g}_i sont des fonctions réelles dépendant du problème considéré.

Dans notre contexte, le nombre M de données historiques est généralement **petit** comparé à N_{MC} . Ainsi, l'étape de calibration donne une erreur plus grande que celle de régression empirique. C'est la raison pour laquelle nous effectuons une approche directe consistant à **rééchantillonner les données observées** et obtenir directement les fonctions de régression (voir Figure 4.1 décrivant le passage de l'approche statistique à l'approche avec rééchantillonnage) .

Modèle multi-factoriel pour les contrats à terme. Dans le marché de l'énergie, nous modélisons habituellement le prix à l'instant t d'un contrat à terme $F(t, T)$ comme étant dirigé par des facteurs de Markov cachés X . Nous obtenons du marché des données observables O à différents moments t ; typiquement O_t est l'ensemble des log prix à terme $\log F(t, T)$. Lorsque nous récrivons le DPE en termes de $F(\cdot, T)$ sous-jacent (aussi une fonction de X), nous obtenons des fonctions de contrôle ou des fonctions valeur dépendant des facteurs cachés X . Puisque ces processus de Markov X ne sont pas observables (nous n'avons aucun accès direct), nous ne pouvons pas appliquer une méthode de rééchantillonnage (proposée dans [GLZ18], sous le nom

de **Non-Intrusive Stratified Resampler**) régénérant chacun de ces facteurs individuellement. Pour traiter ce problème, nous **modifions** le schéma stratifié non-intrusif mentionné ci-dessus et **proposons une méthode de rééchantillonnage** uniquement sur les données observables O . Maintenant, l'équation de programmation dynamique discrète (DDPE) prend une autre forme en préparation à la méthode de rééchantillonnage sur des données observables O :

$$\begin{aligned} Y_N &= g_N(O_N), \\ Y_i &= \mathbb{E}[g_i(Y_{i+1}, \dots, Y_N, O_i, \dots, O_N) \mid O_i], \quad i = N-1, \dots, 0, \end{aligned} \tag{6}$$

où le processus O est à valeurs dans \mathbb{R}^d .

Approche par rééchantillonnage des données historiques. Pour garantir un schéma précis, les méthodes de régression Monte Carlo nécessitent généralement que le nombre de simulations N_{MC} soit beaucoup plus grand que la dimension de l'espace vectoriel \mathcal{L} (nombre de coefficients). Notez que les entrées $O^{1:N_{MC}} := O^1, \dots, O^{N_{MC}}$ sont échantillonnées à partir d'un modèle estimé à partir de données de taille M seulement. Dans ce contexte, puisque M est petit, l'erreur du modèle peut être une préoccupation importante.

Dans [GLZ18], les auteurs conçoivent un schéma où les N_{MC} simulations sont remplacées par les données observées de taille M . Pour surmonter le problème des nombreux coefficients à calculer (malgré le petit nombre de données), ils combinent l'approche de rééchantillonnage avec une stratégie **de stratification** et **d'approximation locale**, appelée schéma NISR (Non-Intrusive Stratified Resampler). Le point clé dans la partie non-intrusive du schéma est de supposer que seule la structure du modèle pour les données observables O est connue, mais pas leurs valeurs de paramètre/coefficient. Par ailleurs, la force de la partie stratifiée est le fait que nous avons des problèmes locaux indépendants sur un espace d'approximation de faible dimension. Ensuite, nous avons une double amélioration : (1) un petit nombre de coefficients à calculer sur chaque problème local et (2) la qualité de l'approximation agrégée n'est pas détériorée lorsque le nombre K de problèmes locaux augmente.

Brièvement, si nous avons observé les facteurs X , nous les aurions rééchantillonné et utilisé directement dans le schéma NISR. Dans notre cas, nous n'avons accès qu'à O , une transformation affine et partiellement inconnue de X (voir Hypothèse 4.1). Nous avons ensuite proposé une extension du schéma NISR natif prenant seulement O comme entrées. De plus, nous considérons que **le modèle multi-factoriel** est décrit comme un processus d'Ornstein-Uhlenbeck (OU) multidimensionnel (voir Hypothèse 4.2), dont les paramètres de retour à la moyenne sont supposés **connus** (mais pas ceux de volatilité/corrélation).

Notre contribution au Chapitre 4 est le schéma NISR suivant (voir Section 4.3). Nous présentons l'algorithme de régression Monte Carlo NISR en supposant l'accès complet aux données observables.

Résultat Principal 4 (Schéma NISR sur le processus observable O). *Nous avons les étapes suivantes pour la résolution du DDPE (6)*

- *A partir de la relation linéaire entre O et X (voir (4.3)), la dynamique de O est aussi un processus OU (voir Proposition 4.4) :*

$$O_j = e^{-\tilde{\alpha}(t_j - t_i)} O_i + \int_{t_i}^{t_j} e^{-\alpha(t-s)} (\tilde{\mu} ds + \tilde{\sigma} dW_s), \quad \text{for all } 0 \leq i \leq j \leq N,$$

pour certains coefficients $\tilde{\alpha} \in \mathbb{R}^d$, $\tilde{\mu} \in \mathbb{R}^d$ and $\tilde{\sigma} \in \mathbb{R}^{d \times n}$ (où $\tilde{\alpha}$ est supposé connu).

- *En suivant la solution de O (voir (4.7)), nous extrayons la source aléatoire U d'une réalisation*

de l'échantillon de base $\{O_i^m : 0 \leq i \leq N\}$ de la manière suivante

$$U_{i,j}^m := O_j^m - e^{-\bar{\alpha}(t_j-t_i)} O_i^m, \quad \text{for all } 0 \leq i \leq j \leq N.$$

- Inversement, à partir de la seule réalisation de la source aléatoire U^m , nous obtenons une trajectoire correspondante partant à l'instant t_i du point $z \in \mathbb{R}^d$

$$\Theta_{i,j}(z, U^m) := e^{-\bar{\alpha}(t_j-t_i)} z + U_{i,j}^m, \quad \text{for all } 0 \leq i \leq j \leq N.$$

Maintenant nous pouvons rééchantillonner plusieurs trajectoires partant de différentes conditions initiales (t_i, z) grâce à la fonction flux $\Theta_{i,j}$ (see (4.12)).

- Dans K ensembles différents \mathcal{H}_k (appelés strates) d'une partition de l'espace \mathbb{R}^d , nous effectuons une approximation par régression Monte Carlo. Pour cela, nous considérons une mesure de probabilité ν telle que sa restriction ν_k sur \mathcal{H}_k

$$\nu_k(dz) = \frac{\mathbf{1}_{\mathcal{H}_k}(z)}{\nu(\mathcal{H}_k)} \nu(dz).$$

est utilisée pour rééchantillonner des trajectoires à partir de différents points initiaux sur \mathcal{H}_k .

- Selon la Proposition 4.7, la distribution tensorielle du type Pareto ν donnée dans l'Equation (4.15) satisfait à une propriété de stabilité de la norme : il existe $K \geq 1$ telle que, pour toute fonction $\varphi : \mathbb{R}^d \rightarrow \mathbb{R} \in \mathbb{L}^2(\nu)$, on a

$$\int_{\mathbb{R}^d} \mathbb{E} |\phi(O_{i+1}^{i,z})|^2 \nu(dz) \leq K \int_{\mathbb{R}^d} |\phi(z)|^2 \nu(dz), \quad i \in \{0, \dots, N-1\},$$

nécessaire pour obtenir une correcte propagation des erreurs.

- En utilisant la distribution ν_k , nous obtenons un M -échantillon sur chaque strate \mathcal{H}_k à partir du temps t_i jusqu'au temps t_n :

$$O_j^{i,k,m} = \Theta_{i,j}(O_i^{i,k,m}, U^m), \quad \text{for all } 1 \leq m \leq M,$$

où $O_i^{i,k,m}$ est échantillonné selon la distribution ν_k indépendamment de i, k, m ;

- A partir du M -échantillon $O_j^{i,k,m}$, nous calculons l'approximation de la fonction valeur y_i dans l'espace des fonctions de base linéaires par morceaux $\mathcal{L}_k = \text{span}\{\mathbf{1}_{\mathcal{H}_k}, z_1 \mathbf{1}_{\mathcal{H}_k}, \dots, z_d \mathbf{1}_{\mathcal{H}_k}\}$. Ici, par simplicité, $\dim(\mathcal{L}_k)$ est indépendant de k , c'est-à-dire, $\dim(\mathcal{L}_k) = d + 1$.
- Pour l'approximation locale de y_i , nous introduisons l'opérateur Moindres Carrés Ordinaires (OLS) approximant la fonction $C : \mathbb{R}^{d \times 2} \rightarrow \mathbb{R}$ sur l'espace \mathcal{L}_k en utilisant l'échantillon $O_{i:i+1}^{k,1:M}$:

$$\mathbf{OLS} \left(C, \mathcal{L}_k, O_{i:i+1}^{k,1:M} \right) = \arg \min_{\phi \in \mathcal{L}_k} \sum_{m=1}^M \left| C(O_{i:i+1}^{i,k,m}) - \phi(O_i^{i,k,m}) \right|^2.$$

- De manière rétrograde en temps à partir de $\hat{y}_N^M = g_N$, nous définissons $\hat{y}_i^M = \sum_{k=1}^K \hat{y}_i^{M,k} \mathbf{1}_{\mathcal{H}_k}$, où

$$\begin{aligned} \hat{y}_i^{M,k} &= -|y_i|_\infty \vee \phi^{M,k} \wedge |y_i|_\infty, \\ \phi^{M,k} &= \mathbf{OLS}(C^M, \mathcal{L}_k, O_{i:i+1}^{i,k,1:M}), \\ C^M(z_{i:i+1}) &= g_i(\hat{y}_{i+1}^M(z_{i+1}), z_{i:i+1}). \end{aligned}$$

On obtient alors une approximation \hat{y}_i^M de la fonction de régression y_i pour tout i dans $\{0, \dots, N\}$.

Dans l’Algorithme 4, nous observons que l’échantillon source $O^{1:M}$ est une donnée d’entrée pour le schéma de rééchantillonnage stratifié nonintrusif (voir Définition 4.8) à travers la source aléatoire $U^{1:M}$. Par conséquent, l’existence de l’approximation de l’algorithme est conditionné au fait que, pour tout $i \in \{0, \dots, N\}$, les données observables O_i sont en effet **observées**. En outre, nous proposons une version modifiée de l’Algorithme 4 dans le cas où certains log prix à terme $\log F(t, T)$ sont manquants (voir Sous-section 4.4.1).

Ce chapitre est tiré d’un article en cours d’écriture en collaboration avec Emmanuel Gobet et Jorge Zubieli [GPZ18].

General Introduction

The present work is devoted to the solution of problems of pricing and hedging of options related to the electricity market. It is also dedicated to the numerical methods for a class of McKean-Vlasov stochastic control problems related to financial questions, e.g., the portfolio selection under drift uncertainty. In the last part, having no access to a fully calibrated multi-factor model for forward prices, we propose a non-intrusive stratified Monte Carlo algorithm to solve the pricing problem of Bermudan and Swing options.

In Part I, we study the pricing and hedging of European options taking into account an asymmetric risk measure. The main approach is the use of an asymptotic framework in order to obtain optimal pricing/hedging policies through nonlinear partial differential equations.

In Chapter 1, we consider the residual error produced by discrete-time trading. We suppose a trader wants to find strategies for the hedging portfolio such that the discretization error is small and penalizing the losses scenarios more the profits ones: we search to get a valuation/hedging policy at a large number of portfolio rebalancing dates for which the induced hedging error under an asymmetric risk measure is minimized. More precise, we define an optimal strategy in terms of the solution to the optimization problem related to the asymptotic expected-error. In this setting, the optimal value (resp., hedging) function is the solution (resp., the gradient of the solution) to a fully nonlinear PDE, whose nonlinear term depends on the second-order derivative of the solution. In fact, this nonlinearity term characterizes our asymptotically optimal strategy. In dimension 1, a quasi-explicit formulation for this term is derived.

In Chapter 2, we study the valuation problem of a power plant as an application of the previous hedging method with small costs proportional to the time step. Considering that some (maintenance) costs need to be paid regardless of the fact that a power plant is producing electricity or not, we are interested to evaluate an income associated to the selling of its production in the future. In fact, the plant's owner willing to reduce risk of a uncertain income by hedging in the forward market. Therefore, the optimal hedging strategy is obtained by minimizing the asymptotic risk associated to a discretization error including time-step-proportional costs. We studied impact of the (maintenance) costs into the hedging strategy, when those costs depends on the electricity price.

In Part II, we investigated several numerical methods to solve a class of control problems associated to mean-field interactions with partial information. We also studied a non-intrusive stratified resampler to solve discrete dynamic programming equations in the presence of a small historical data.

In Chapter 3, we study of a McKean-Vlasov (MKV) control class in the presence of a common noise, called polynomial conditional MKV. This class is a generalization of the linear quadratic stochastic MKV control class. Those problems raise from the asymptotic formulation of a cooperative equilibrium in a large population of players with mean-field interaction or from control problem with cost function depending on a nonlinear functional of the distribution

of the state process. After reducing this polynomial class by a Markov embedding to finite-dimensional control problems, we provide a discussion and a comparison of three different probabilistic methods for numerical resolution of the reduced problem: quantization, regression by control randomization and regression by regress later. We solved numerically several problems related to this polynomial MKV control: the selection and the liquidation of portfolio under drift uncertainty; and systemic interbank risk model with partial observation.

In Chapter 4, we are interested in solving dynamic programming equations (DPE) associated with a financial valuation in the energy market. In particular, we focus on the pricing of Bermudan and Swing options, where the underlying assets are forward contracts. We search for a non-intrusive stratified method to solve these equations using only historical data, without a full calibrated underlying model. In energy market, forward contracts are often driven by hidden Markov factors (multi-factor modeling). Consequently, the problem of pricing of those options produces DPEs depending on those unobservable factors. Since we have access to the forward prices but not to the hidden factors, we propose a resampling method only on observable contracts.

I. Asymmetric risk measure for pricing and hedging of options

1. Option pricing and hedging using asymmetric risk measure: Asymptotic optimality

In finance, the pricing and hedging of contingent claims are major concerns. In theoretical point of view, those problems is well established (see [KS98], for instance). In practice, traders hedge their contracts by trading only at discrete times, say $t_0 = 0 < t_1 < \dots < t_N = T$, yielding a residual risk. Here, they intend to hedge the claim H_T at time T using d hedging securities whose price are denoted by $X = (X^{(1)}, \dots, X^{(d)})$. For that, we define the *local risk* \mathcal{E}_n associated with the trading times t_n and t_{n+1} given by

$$\mathcal{E}_n = V_{t_{n+1}} - V_{t_n} - \vartheta_{t_n} \cdot X_{t_{n+1}} - X_{t_n},$$

where V denotes the valuation process; $\vartheta = (\vartheta^{(1)}, \dots, \vartheta^{(d)})$ denotes the hedging process with $\vartheta^{(i)}$ being the number of shares invested in the i -th hedging instrument.

We aim to find the valuation/hedging rules (V, ϑ) minimizing this residual risk using a risk function ℓ . From the existing results (for instance, [Sch99], for the quadratic local risk minimization), we choose a function ℓ **penalizing profits** ($\mathcal{E}_n < 0$) **and losses** ($\mathcal{E}_n > 0$) **asymmetrically**. In this context, we study the *integrated local risk* \mathbf{R}_N under the form

$$\mathbf{R}_N(V, \vartheta) = \sum_{n=0}^{N-1} \mathbb{E}[\ell(\mathcal{E}_n)],$$

where the **risk function** ℓ is chosen as

$$\ell_\gamma(y) = (1 + \gamma \operatorname{sgn}(y))^2 y^2 / 2, \quad \gamma \in (0, 1).$$

The parameter γ , in the risk function ℓ_γ , represents the investor risk aversion penalizing losses further than profits (see Figure 1.1).

Standard asymptotic approach. In this setting, we aim to study the asymptotics of the minimum

$$\min_{(V, \vartheta) \in \mathcal{A}} \mathbf{R}_N(V, \vartheta), \tag{7}$$

with a large number N of trading dates in the case of equidistant trading times $t_n = n\Delta t$ with time step $\Delta t = T/N$. The minimum (7) is computed over the set \mathcal{A} of all adapted to the underlying filtration $(\mathcal{F}_t)_{t \geq 0}$ and appropriately integrable pair (V, ϑ) , under the replication constraint $V_T = H_T$.

There exists in the literature a few results in that direction:

- In [Pha00], the author deals with a \mathbb{L}_p risk function of the losses and a fixed number trading dates;
- In [AM11], the authors study pseudo-optimal strategies and get asymptotic results in the case of a risk function in class C^3 .

We remark that the **discontinuity of the second derivative** ℓ''_γ complicates the analysis and fully changes the nature of subsequent results. In short, the previous references consider different settings and difficulties from ours.

Asymptotic approach with a PDE valuation. Instead of solving the problem (7) in the asymptotic regime $N \rightarrow +\infty$, we tackle a modified problem using a PDE valuation. We suppose the hedging instruments X are modeled by

$$dX_t = \mu(t, X_t) + \sigma(t, X_t) dW_t,$$

where W is a Brownian motion and the contingent claims is of the form $H_T = g(X_T)$. Then, we assume that the contingent claim is **evaluated exogenously** by a valuation process $V_t = v(t, X_t)$ for some function v . The function v acts as a reference price. Given this exogenous valuation, the trader will determine how to hedge on each interval $[t_n, t_{n+1}]$ by choosing an adapted valuation/hedging rule $(\tilde{V}_{t_n}, \tilde{\vartheta}_{t_n})$ and considering the related *conditional local risk* $\mathbf{R}_{n,\gamma}$

$$\mathbf{R}_{n,\gamma} = \mathbb{E}[\ell_\gamma(V_{t_{n+1}} - \tilde{V}_{t_n} - \tilde{\vartheta}_{t_n} \cdot X_{t_{n+1}} - X_{t_n}) \mid \mathcal{F}_{t_n}].$$

Inspired by the connection between dynamic risk valuations and nonlinear Backward Stochastic Differential Equations (BSDE) (see for instance [Cre13]), the trader parametrize its valuation/hedging rule through a function f , possibly nonlinear, such that

$$\tilde{V}_{t_n} = u^{t_{n+1}}(t_n, X_{t_n}), \quad \tilde{\vartheta}_{t_n} = D_x u^{t_{n+1}}(t_n, X_{t_n}),$$

where $u^{t_{n+1}}$ is the solution defined in the interval $[t_n, t_{n+1}]$ to the f -PDE below.

- The function $u^{t_{n+1}} : [t_n, t_{n+1}] \times \mathbb{R}^d \rightarrow \mathbb{R}$ satisfies

$$u_t^{t_{n+1}}(t, x) + \frac{1}{2} \text{Tr}[\sigma \sigma^\top D_x^2 u^{t_{n+1}}](t, x) + f(t, x, u^{t_{n+1}}(t, x), D_x u^{t_{n+1}}(t, x), D_x^2 u^{t_{n+1}}(t, x)) = 0, \quad (8)$$

for all $(t, x) \in [t_n, t_{n+1}] \times \mathbb{R}^d$ with the terminal condition $u^{t_{n+1}}(t_{n+1}, x) = v(t_{n+1}, x)$ at the time t_{n+1} ;

- For some continuous functions $\sigma : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$, $f : [0, T] \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$; and a reference price $v(t_{n+1}, \cdot)$ at t_{n+1} ;

In this context, it is now natural to analyze the asymptotic behavior of the *integrated conditional local risk* after appropriate renormalization

$$\mathbf{R}_{N,\gamma}(v, f) = \frac{1}{\Delta t} \sum_{n=0}^{N-1} \mathbb{E}[\mathbf{R}_{n,\gamma}].$$

Our main contribution in Chapter 1 is the following result (see Theorem 1.6, page 34).

Main Result 1 (Theorem 1.6). *Let $B : [0, 1] \times \Omega \rightarrow \mathbb{R}^d$ be another standard Brownian motion independent from W . Under the Assumptions 1.1-1.5,*

$$\mathbf{R}_\gamma(v, f) := \lim_{N \rightarrow \infty} \mathbf{R}_{N, \gamma}(v, f)$$

exists and is given by,

$$\mathbf{R}_\gamma(v, f) = \mathbb{E} \left[\int_0^T \int_0^1 \ell''_\gamma \left(\int_0^\theta B_{\theta'}^\top G_t dB_{\theta'} - F_t \theta \right) \left(F_t^2 \theta - F_t \int_0^\theta B_{\theta'}^\top G_t dB_{\theta'} + \|G_t B_\theta\|^2 / 2 \right) d\theta dt \right],$$

where

$$\begin{aligned} F_t &= f(t, X_t, v(t, X_t), D_x v(t, X_t), D_x^2 v(t, X_t)) \in \mathbb{R}, \\ G_t &= (\sigma^\top (D_x^2 v) \sigma)(t, X_t) \in \mathbb{R}^{d \times d}. \end{aligned}$$

Consequently, we discuss the existence of an optimal PDE nonlinearity f^* such that the f^* -PDE valuation minimizes the asymptotic risk in the sense: $\mathbf{R}_\gamma(v, f^*) \leq \mathbf{R}_\gamma(v, f)$, for any admissible nonlinearity f . Additionally, this optimal PDE nonlinearity f^* is **quasi-explicit in dimension 1** (see (1.22), page 39) and it depends on the risk parameter γ , on the reference valuation second derivative and on the price process volatility. Consequently, a natural and consistent choice for the valuation/hedging rule may be the solution to the f^* -PDE (8) with the payoff $g : \mathbb{R}^d \rightarrow \mathbb{R}$ being the terminal condition at the horizon time T .

This chapter is based on a paper written in collaboration with Emmanuel Gobet and Xavier Warin [GPW18].

2. Asymptotic asymmetric risk measure: Application to physical asset valuation

In Chapter 2, we provide valuation and hedging policies for future incomes due to the production of power plants using an asymmetric risk valuation. Since the deregulation of energy markets, several spot and future markets were created to exchange electricity. After that, power plant owners start to face the problem of **evaluating their plant production**, which depends on the electricity spot price in the future.

In practice, power plant owners perform a discrete hedging at trading times, $t_0 = 0 < t_1 < \dots < t_N = T$ using a hedging security X in order to reduce the risk associated to a random income H_T at a future time T . This discrete hedging produces a residual risk:

$$\mathcal{E} = H_T - \sum_{n=0}^{N-1} \vartheta_{t_n} (X_{t_{n+1}} - X_{t_n}),$$

where

- ϑ denotes the number of shares invested in the hedging security;
- X is a forward contract $F(\cdot, T)$ with delivery time T .

Then, they search to find the valuation/hedging rule taking into account the local balance \mathcal{E}_n penalizing $\mathcal{E}_n < 0$) through an asymmetric risk function ℓ :

$$\mathcal{E}_n = V_{t_{n+1}} - V_{t_n} - \vartheta_{t_n} (X_{t_{n+1}} - X_{t_n}),$$

where V stands for the valuation process under the replication constraint $V_T = H_T$. Notice that plant owners:

- want to obtain a residual risk \mathcal{E} with a small standard-deviation;
- prefer profitable scenarios where $\mathcal{E}_n > 0$.

In reality, power plant generates some fixed costs whether it is producing electricity or not. Those costs are called here **fixed costs**:

- they include the investment and depreciation costs;
- they exclude the fuel costs (coal, gas, uranium).

To record in their balance book, power producers are constrained to find today an equivalent price to the future income generated by their power plant.

In fact, they **do not face any financial risk** due to this income. On the other hand, because they can not increase their income by influencing the electricity spot price S , they prefer to obtain a certain value due to the selling of their future production instead to receive the random positive amount $H_T = g(S_T)$.

Economically speaking, when the electricity spot price S is high, we require the power plant to face a higher demand. In indeed, a high price is produced by an increasing demand. Therefore, we increase the production by starting more often the power plant. That is the reason why we assume a dependence of the costs c on the electricity spot price S .

Asymmetric risk valuation. Considering equidistant trading times $t_n = n\Delta t$ with time step $\Delta t = T/N$, we take into account the fixed costs proportional to the time step Δt by subtracting $c(S_{t_n})\Delta t$ from the local balance \mathcal{E}_n . In a setting analogous to Chapter 1, we study the asymptotic of the *integrated risk* \mathbf{R}_N

$$\mathbf{R}_N = \frac{1}{\Delta t} \sum_{n=0}^{N-1} \mathbb{E} \left[\ell \left(V_{t_{n+1}} - V_{t_n} - \vartheta_{t_n} (X_{t_{n+1}} - X_{t_n}) - c(S_{t_n})\Delta t \right) \right],$$

Here, we set

$$\ell_\gamma(y) = (1 + \gamma \operatorname{sgn}(y))^2 y^2 / 2, \quad \gamma \in (-1, 0),$$

where the parameter γ represents the producer's desire of reducing the randomness of future income. Unlike the risk parameter in Chapter 1, this parameter γ takes values in $(-1, 0)$ (see Figure 2.1).

In view of studying the integrated risk $\mathbf{R}_{N,\gamma}$ in the asymptotic regime $N \rightarrow +\infty$ with an asymmetric risk function ℓ_γ , we follow an approach by PDE valuation described in Chapter 1. We assume the existence of exogenous valuation process $V_t = v(t, X_t)$ defined by some function v . For example, v is a reference price obtained by the power producer when no cost is taken into account. Now, power plant owner determines its hedge position on each interval $[t_n, t_{n+1}]$ by choosing an adapted valuation/hedging rule $(\tilde{V}_{t_n}, \tilde{\vartheta}_{t_n})$. For that, we consider the related *conditional risk* $\mathbf{R}_{n,\gamma}$ using the risk function ℓ_γ

$$\mathbf{R}_{n,\gamma} = \mathbb{E} \left[\ell_\gamma \left(V_{t_{n+1}} - \tilde{V}_{t_n} - \tilde{\vartheta}_{t_n} (X_{t_{n+1}} - X_{t_n}) - c(S_{t_n})\Delta t \right) \middle| \mathcal{F}_{t_n} \right].$$

This time, the valuation/hedging rule $(\tilde{V}_{t_n}, \tilde{\vartheta}_{t_n})$ of the power producer is given by

$$\tilde{V}_{t_n} = u^{t_{n+1}}(t_n, X_{t_n}), \quad \tilde{\vartheta}_{t_n} = u_x^{t_{n+1}}(t_n, X_{t_n}),$$

and is parametrized by a function f , possibly nonlinear, through the solution $u^{t_{n+1}} : [t_n, t_{n+1}] \times \mathbb{R}$ to the f -PDE with terminal condition $u^{t_{n+1}}(t_{n+1}, \cdot) = v(t_{n+1}, \cdot)$ at time t_{n+1} (analogously to the one-dimensional version of (1.5) but with diffusion term equal to $\bar{\sigma}(t, T)x$, with a volatility function $\bar{\sigma}(t, T) = \bar{\sigma}_0 e^{-a_0(T-t)}$).

As in Chapter 1, we analyze the asymptotic regime of the integrated conditional risk

$$\mathbf{R}_{N,\gamma}(v, f) = \frac{1}{\Delta t} \sum_{n=0}^{N-1} \mathbb{E}[\mathbf{R}_{n,\gamma}],$$

as the number N of trading dates goes to infinity.

Our main contribution in Chapter 2 is the following result (see Theorem 2.9, page 71). This is a version of the Theorem 1.6 in the presence of fixed costs proportional to the time step.

Main Result 2 (Theorem 2.9). *Let $B : [0, 1] \times \Omega \rightarrow \mathbb{R}$ be another standard Brownian motion independent from W . Under the Assumptions 2.3-2.8, the limit of $\mathbf{R}_{N,\gamma}(v, f)$ as $N \rightarrow \infty$ exists and is given by*

$$\mathbf{R}_\gamma(v, f) = \mathbb{E} \left[\int_0^T \int_0^1 \ell''_\gamma \left(J_1(v, (t, X_t)) \frac{B_\theta^2 - \theta}{2} - J_2(v, (t, X_t)) \theta \right) \left(|J_2(v, (t, X_t))|^2 \theta - J_2(v, (t, X_t)) J_1(v, (t, X_t)) \frac{B_\theta^2 - \theta}{2} + |J_1(v, (t, X_t))|^2 \frac{B_\theta^2}{2} \right) d\theta dt \right],$$

where

$$\begin{aligned} J_2(v, (t, X_t)) &= f(t, X_t, v(t, X_t), v_x(t, X_t), v_{xx}(t, X_t)) + c(X_t) \in \mathbb{R}, \\ J_1(v, (t, X_t)) &= \bar{\sigma}^2(t, T) X_t^2 v_{xx}(t, X_t) \in \mathbb{R}. \end{aligned}$$

As in Chapter 1, the idea is to obtain an optimal PDE nonlinearity f^* minimizing the asymptotic risk:

$$\mathbf{R}_\gamma(v, f^*) \leq \mathbf{R}_\gamma(v, f),$$

for any admissible f . Once more, this optimal PDE nonlinearity f^* is explicit and depends on γ , on the second derivative of v and on cost function $c(x)$. Finally, the power producer will choose the following rule

$$V_{t_n} = v^*(t_n, X_{t_n}), \quad \vartheta_{t_n} = v_x^*(t_n, X_{t_n}),$$

where v^* is the solution to $v^* : [0, T] \times \mathbb{R}^d$ to the f^* -PDE with terminal condition $v^*(T, \cdot) = g(\cdot)$ at time T .

For the numerical experiments, we take a fixed cost depending on the level of the electricity spot price through a convex (or concave) function $c(x)$. We compute the numerical solution to the f^* -PDE for a call payoff and different models for $c(x)$.

This chapter is based on a work in progress in collaboration with Clémence Alasseur, Emmanuel Gobet and Xavier Warin.

II. Numerical methods in stochastic control

3. Polynomial conditional McKean-Vlasov control problems: Some probabilistic numerical methods

In Chapter 3, we are interested to the McKean-Vlasov (MKV) control problem under partial observation and common noise. The formulation is described as follows. On a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with two independent Brownian motions B and W^0 , we consider the controlled stochastic McKean-Vlasov dynamics in \mathbb{R}^n :

$$dX_s = b\left(X_s, \mathbb{P}_{X_s}^{W^0}, \alpha_s\right)ds + \sigma\left(X_s, \mathbb{P}_{X_s}^{W^0}, \alpha_s\right)dB_s + \sigma_0\left(X_s, \mathbb{P}_{X_s}^{W^0}, \alpha_s\right)dW_s^0, \quad (9)$$

where

- $\mathbb{P}_{X_s}^{W^0}$ denotes the conditional distribution of X_s given W^0 ;
- $\mathbb{F}^0 = (\mathcal{F}_t^0)_{t \geq 0}$ is the natural filtration generated by W^0 ;
- the control α is \mathbb{F}^0 -progressive valued in some Polish space A .

The cost functional associated to the stochastic McKean-Vlasov equation (9) for a control process α is

$$J(\alpha) = \mathbb{E} \left[\int_0^T f(X_t, \mathbb{P}_{X_t}^{W^0}, \alpha_t) dt + g(X_T, \mathbb{P}_{X_T}^{W^0}) \right],$$

and the objective is to minimize over an admissible set \mathcal{A} of control processes the cost functional:

$$V_0 = \inf_{\alpha \in \mathcal{A}} J(\alpha).$$

Our purpose is to **investigate some classes of MKV control problems**, which can be reduced to finite dimensional problems in view of numerical resolution.

A Class of polynomial models and Markovian embedding. First, we consider a class of models where the coefficients of the MKV equation are linear w.r.t. the state variable X (see (3.9)), while the running and terminal cost functions are polynomial in the state variable in the following sense

$$f(x, \mu, a) = f_0(\mu, a) + \sum_{k=1}^p f_k(\mu, a)x^k, \quad g(x, \mu) = g_0(\mu) + \sum_{k=1}^p g_k(\mu)x^k,$$

for some integer $p \geq 1$. Second, we assume all the coefficients depend on μ through its first p moments (see (3.10)).

Given the controlled process $X = X^\alpha$ solution to the stochastic McKean-Vlasov dynamics (9), we denote

$$Y_t^k = \mathbb{E}[X_t^k | W^0], \quad k = 1, \dots, p.$$

From the linear/polynomial assumptions (by Itô's formula and conditional expectations), we derive the dynamics of (Y^1, Y^2, \dots, Y^p) as

$$\{ dY_t^k = B_k(Y_t^1, Y_t^2, \dots, Y_t^p, \alpha_t)dt + \Sigma_k(Y_t^1, Y_t^2, \dots, Y_t^p, \alpha_t)dW_t^0, \quad k = 1, \dots, p, \quad (10)$$

for some functions B_k, Σ_k (see page 95) depending on the conditional moments (Y^1, Y^2, \dots, Y^p) ,

while the cost functional is written in a reduced form

$$J(\alpha) = \mathbb{E} \left[\int_0^T \bar{f}(Y_t^1, Y_t^2, \dots, Y_t^p, \alpha_t) dt + \bar{g}(Y_T^1, Y_T^2, \dots, Y_T^p) \right]. \quad (11)$$

The McKean-Vlasov control problem is **reduced** in this polynomial framework into a **finite-dimensional** control problem with \mathbb{F}^0 -adapted controlled variables (Y^2, Y^2, \dots, Y^p) .

Probabilistic numerical methods. In view of solving the reduced problem (10)-(11), we choose the following setting. Let $Z = (Y^1, Y^2, \dots, Y^p)$ be a controlled process by an adapted control α taking values in A , solution to

$$dZ_t^\alpha = b(Z_t^\alpha, \alpha_t) dt + \sigma_0(Z_t^\alpha, \alpha_t) dW_t^0$$

and

$$J(t, z, \alpha) = \mathbb{E} \left[\int_t^T f(Z_s^\alpha, \alpha_s) ds + g(Z_T^\alpha) \middle| Z_t^\alpha = z \right],$$

be a dynamic version of the cost functional $J(\alpha)$.

On a time discretization $0 = t_0, t_1, \dots, t_N = T$, we write the Euler approximation of Z_t^α :

$$Z_{t_{n+1}}^\alpha = Z_{t_n}^\alpha + b(Z_{t_n}^\alpha, \alpha_{t_n}) \Delta t + \sigma_0(Z_{t_n}^\alpha, \alpha_{t_n}) \Delta W_{t_n}^0$$

and the discrete equivalent of $J(t, z, \alpha)$:

$$J(t_n, z, \alpha) = \mathbb{E} \left[\sum_{i=n}^N f(Z_{t_i}^\alpha, \alpha_{t_i}) \Delta t + g(Z_{t_N}^\alpha) \middle| Z_{t_n}^\alpha = z \right].$$

Now, the value function $V(t_n, z) = \sup_{(\alpha_{t_i})_{i=n}^N \in \mathcal{A}} J(t_n, z, \alpha)$ is represented alternatively through the following dynamic programming equation, given the known terminal condition $g(z)$,

$$\begin{aligned} V(T_N, z) &= g(z), \\ V(t_n, z) &= \sup_{\alpha} \left\{ f(Z_{t_n}^\alpha, \alpha) \Delta t + \mathbb{E}_{\alpha} [V(t_{n+1}, Z_{t_{n+1}}^\alpha) | Z_{t_n}^\alpha = z] \right\}. \end{aligned} \quad (12)$$

The dynamic programming equation (12) inspires numerical methods that approximate the value function iteratively backward in time, starting from the terminal condition. The main difficulty in implementing such approach lies in the **estimation of conditional expectations** $\mathbb{E}_{\alpha} [V(t_{n+1}, Z_{t_{n+1}}^\alpha) | Z_{t_n}^\alpha = z]$.

Our first contribution in Chapter 3 is the following description (see Section 3.3) and application (see Section 3.4).

Main Result 3. *We describe three numerical methods solving conditional MKV problems (in their reduced version):*

- the Regression Monte Carlo techniques (by control randomization and by regress later) are a family of algorithms whose effectiveness relies on the choice of the basis functions used to project future value functions (see Algorithms 1 and 2);
- the Quantization techniques approximates the controlled process $Z_{t_n}^\alpha$ with a particular finite state Markov chain for which expectations can be approximated quickly (see Algorithm 3).

We apply each of these methods to three applications arising from polynomial MKV control problems under partial observation and common noise:

- the Portfolio optimization under drift uncertainty (see Subsection 3.4.1), where the drift of the underlying asset is unknown and unobservable. First, we consider the problem of a trader willing to liquidate a large number of shares within a finite time T and facing execution cost and market price impact. Then, we see the selection problem of a portfolio strategy maximizing the utility of the terminal wealth.
- the Intersystemic risk with partial observation (see Subsection 3.4.2). We suppose the monetary reserves of N banks lending to and borrowing from each other satisfy diffusion system with mean-field interaction and a common noise. Then a social planner (the central bank, for example) only observing the common noise acts on strength of the interbank interaction in order to minimize the spread between each bank's reserve and the average.

Numerical results. Our second contribution in Chapter 3 is the following numerical experiments (see Section 3.5). We presented the results of three different examples of applications summarized below

- We found that the Regression Monte Carlo algorithms perform correctly in problems of drift control. In those problems, they are much faster than Quantization for similar precision. In particular, we noticed that Regress Later is more reliable than Control Randomization: the choice of a uniform distribution of the training points on a suitable interval suffices to obtain high-quality estimations.
- On the other hand, Control Randomization is sensitive to the choice of the distribution of the randomized control and few repetitions are necessary before finding a proper control distribution. We have also tried to use the performance iteration or path re-computation methods. However, on the considered examples, those methods were time-consuming and did not help much in terms of accuracy.
- Quantization techniques provided the most stable and accurate results for the three different cases of control problems. Moreover, we are able to choose the grid to quantize the controlled process. It is possible to exploit this feature when we have a rough idea of where the controlled process should be driven by the optimal strategy (see the Portfolio liquidation problem). Therefore, we should build a grid with many points located where the process is supposed to go.

This chapter is based on a paper written in collaboration with Alessandro Balata, Côme Huré, Mathieu Laurière et Huyen Pham [BHL⁺18].

4. A non-intrusive stratified resampler for multi-factor models: Application in energy market

Stochastic dynamic programming equations are related to the resolution of non-linear problems (stochastic controls or nonlinear PDEs) arising in almost all areas of science, from water reservoir management to finance.

In Chapter 4, we aim to solve dynamic programming equations (DPE) related to a financial valuation in energy market. Here, we are concerned, for example, in the pricing of Bermudan or Swing options, where the underlying assets are forward contracts. Then, we plan to develop a non-intrusive algorithm to solve those DPE using only the observable data without a full model calibration.

In a first moment, we deal with a discrete DPE taking the following form:

$$Y_N = \tilde{g}_N(X_N), \tag{13}$$

$$Y_i = \mathbb{E}[\tilde{g}_i(Y_{i+1}, \dots, Y_N, X_i, \dots, X_N) \mid X_i], \quad i = N - 1, \dots, 0,$$

where X is a Markov chain in \mathbb{R}^n ; and \tilde{g}_N, \tilde{g}_i are some real functions depending on the problem under consideration.

To solve the discrete DPE (13), we first observe a historical data O of size M in view of calibrating the parameter values for some stochastic model describing X . Then, we perform regression Monte Carlo (MC) algorithms with N_{MC} sampled points X to obtain the regression functions \tilde{y}_i such that $Y_i = \tilde{y}_i(X_i)$.

In our context, the number M of historical data is usually **small** compared to N_{MC} . Thus, the calibration step gives a more substantial error than the empirical regression one. This is the reason why we perform a direct approach consisting in **resampling the observed data** and obtain directly the regression functions (see Figure 4.1 describing the change from the statistical approach to the resampling one).

Multi-factor model for forward contracts. In energy market, we usually model the price at time t of a T -forward contract $F(t, T)$ as driven in terms of hidden Markov factors X . We get from the market some observable data O at different times t ; typically O_t is the set of log forward prices $\log F(t, T)$. When we rewrite the DPE in terms of the underlying $F(\cdot, T)$ (also a function of X), we get control or value functions depending on the hidden factors X . Since these Markov process X are not observable (we have no direct access to them), we can not apply a resampling method (proposed in [GLZ18] under the name of Non-Intrusive Stratified Resampler) regenerating each one of those factors individually. To deal with this problem, we **modify** the non-intrusive stratified scheme mentioned above and **propose a resampling method** on the observable data O .

Now, the discrete dynamic programming equation (DDPE) takes another form in preparation for the resampling method on observable data O :

$$\begin{aligned} Y_N &= g_N(O_N), \\ Y_i &= \mathbb{E}[g_i(Y_{i+1}, \dots, Y_N, O_i, \dots, O_N) \mid O_i], \quad i = N - 1, \dots, 0, \end{aligned} \tag{14}$$

where the process O takes values in \mathbb{R}^d .

Approach by historical-data Resampling. To ensure accurate scheme, the regression Monte Carlo methods usually need that the number of simulations N_{MC} has to be much larger than the dimension of the vector space \mathcal{L} (number of coefficients). Note that the inputs $O^{1:N_{\text{MC}}} := O^1, \dots, O^{N_{\text{MC}}}$ are sampled from a model which is estimated from M -size data only. In our setting, M is small, then model error may be a significant concern.

In [GLZ18], the authors design a scheme where the N_{MC} simulations are replaced by the M -size observed data. To overcome the problem of numerous coefficients to compute (despite the small number of data), they combine resampling approach with a **stratification and local-approximation** strategy, called the Non-Intrusive Stratified Resampler (NISR) scheme. The critical point in the non-intrusive part of the scheme is to assume that only the model structure for observable data O is known, but not their parameter/coefficient values. While the strength of the stratified part is the fact that we have independent local problems on a low-dimensional approximation space. Then, we have a double-fold improvement: (1) a small number of coefficients to compute on each local problem and (2) the quality of the aggregated approximation is not deteriorated when the number K of local problems increases.

Briefly, had we observed the factors X , we would have resampled it and used directly in the NISR scheme. In our case, we have only access to O , a affine and partially unknown

transformation of X (see Assumption 4.1). Then we have proposed an extension of the native NISR scheme taking O as inputs. Additionally, we consider that the **factors model** is described as multidimensional Ornstein-Uhlenbeck (OU) (see Assumption 4.2), whose the mean-reverting parameters is supposed **known** (but not the volatilities/correlation).

Our contribution in Chapter 4 is the following NISR scheme (see Section 4.3). We present the NISR-regression Monte Carlo algorithm assuming the complete access to the observable data.

Main Result 4 (NISR scheme on the observable process O). *We have the following steps to the resolution of the DDPE (14)*

- From the linear relation between O and X (see (4.3)), the dynamics of O is also a OU (see Proposition 4.4):

$$O_j = e^{-\tilde{\alpha}(t_j-t_i)} O_i + \int_{t_i}^{t_j} e^{-\alpha(t-s)} (\tilde{\mu} ds + \tilde{\sigma} dW_s), \quad \text{for all } 0 \leq i \leq j \leq N,$$

for some coefficients $\tilde{\alpha} \in \mathbb{R}^d$, $\tilde{\mu} \in \mathbb{R}^d$ and $\tilde{\sigma} \in \mathbb{R}^{d \times n}$ (where $\tilde{\alpha}$ is supposed known).

- Following the solution of O (see (4.7)), we extract the random source U from one realization of the **root sample** $\{O_i^m : 0 \leq i \leq N\}$ as follows

$$U_{i,j}^m := O_j^m - e^{-\tilde{\alpha}(t_j-t_i)} O_i^m, \quad \text{for all } 0 \leq i \leq j \leq N.$$

- Inversely, from the one realization of the random source U^m , we obtain a corresponding path starting at time t_i from point $z \in \mathbb{R}^d$

$$\Theta_{i,j}(z, U^m) := e^{-\tilde{\alpha}(t_j-t_i)} z + U_{i,j}^m, \quad \text{for all } 0 \leq i \leq j \leq N.$$

Now we can resample multiple paths from different initial conditions (t_i, z) thanks to flow function $\Theta_{i,j}$ (see (4.12)).

- In K different sets \mathcal{H}_k (called strata) from a partition of the space \mathbb{R}^d , we perform a regression Monte Carlo approximation. For that, we consider a probability measure ν such that its restriction ν_k on \mathcal{H}_k

$$\nu_k(dz) = \frac{\mathbf{1}_{\mathcal{H}_k}(z)}{\nu(\mathcal{H}_k)} \nu(dz).$$

is used to resample paths starting from different initial points on \mathcal{H}_k .

- According to Proposition 4.7, the tensor product Pareto-type distribution ν given in (4.15) satisfies a norm-stability property: there exists a constant $K \geq 1$ such that, for any $\varphi : \mathbb{R}^d \rightarrow \mathbb{R} \in \mathbb{L}^2(\nu)$, it holds

$$\int_{\mathbb{R}^d} \mathbb{E} |\phi(O_{i+1}^{i,z})|^2 \nu(dz) \leq K \int_{\mathbb{R}^d} |\phi(z)|^2 \nu(dz), \quad i \in \{0, \dots, N-1\},$$

required to obtain the well-behaved propagation of errors.

- Using the distribution ν_k , we obtain a M -sample on each stratum \mathcal{H}_k from time t_i until time t_n :

$$O_j^{i,k,m} = \Theta_{i,j}(O_i^{i,k,m}, U^m), \quad \text{for all } 1 \leq m \leq M,$$

where $O_i^{i,k,m}$ is sampled according to the distribution ν_k independently from i, k, m ;

- From the M -sample $O_j^{i,k,m}$, we compute the approximation of the value function y_i in the space of piecewise linear basis functions $\mathcal{L}_k = \text{span}\{\mathbf{1}_{\mathcal{H}_k}, z_1 \mathbf{1}_{\mathcal{H}_k}, \dots, z_d \mathbf{1}_{\mathcal{H}_k}\}$. Here, for simplicity, $\dim(\mathcal{L}_k)$ is independent from k , i.e., $\dim(\mathcal{L}_k) = d + 1$.

- To the local approximation of y_i , we introduce the Ordinary Least Square (OLS) operator approximating the function $C : \mathbb{R}^{d \times 2} \rightarrow \mathbb{R}$ on the space \mathcal{L}_k using the sample $O_{i:i+1}^{k,1:M}$:

$$\mathbf{OLS} \left(C, \mathcal{L}_k, O_{i:i+1}^{k,1:M} \right) = \arg \min_{\phi \in \mathcal{L}_k} \sum_{m=1}^M \left| C(O_{i:i+1}^{i,k,m}) - \phi(O_i^{i,k,m}) \right|^2.$$

- Backwardly in time starting from $\hat{y}_N^M = g_N$, we set $\hat{y}_i^M = \sum_{k=1}^K \hat{y}_i^{M,k} \mathbf{1}_{\mathcal{H}_k}$, where

$$\begin{aligned} \hat{y}_i^{M,k} &= -|y_i|_\infty \vee \phi^{M,k} \wedge |y_i|_\infty, \\ \phi^{M,k} &= \mathbf{OLS}(C^M, \mathcal{L}_k, O_{i:i+1}^{i,k,1:M}), \\ C^M(z_{i:i+1}) &= g_i(\hat{y}_{i+1}^M(z_{i+1}), z_{i:i+1}). \end{aligned}$$

Then we obtain an approximation \hat{y}_i^M of regression function y_i for every i in $\{0, \dots, N\}$.

In the Algorithm 4, we observe that the root sample $O^{1:M}$ is an input to the non-intrusive stratified resampler (see Definition 4.8) through the random source $U^{1:M}$. Therefore, the **existence** of algorithm's approximation is conditioned to the fact that for any $i \in \{0, \dots, N\}$ the observable data O_i is indeed **observed**. In addition, we propose a modified version of Algorithm 4 in the case where some log forward prices $\log F(t, T)$ are missing (see Subsection 4.4.1).

This chapter is based on a work in progress in collaboration with Emmanuel Gobet and Jorge Zubelli [GPZ18].

Part I

Asymmetric risk measure for pricing and hedging of options

Option pricing and hedging using asymmetric risk measure: Asymptotic optimality

This chapter is based on the paper [GPW18].

Abstract

Discrete time hedging produces a residual risk, namely, the tracking error. The major problem is to get valuation/hedging policies minimizing this error. We evaluate the risk between trading dates through a function penalizing asymmetrically profits and losses. After deriving the asymptotics within a discrete time risk measurement for a large number of trading dates, we derive the optimal strategies minimizing the asymptotic risk in the continuous time setting. We characterize the optimality through a class of fully nonlinear Partial Differential Equations (PDE). Numerical experiments show that the optimal strategies associated with discrete and asymptotic approach coincides asymptotically.

1.1 Introduction

The valuation and hedging of contingent claims are major concerns in finance, both from a theoretical and a practical point of view. The continuous-time theory is well established (see [KS98], for instance). But, in practice, hedging can be performed only at discrete trading times, say $t_0 = 0 < t_1 < \dots < t_N = T$, yielding a residual risk. Here, we intend to hedge the claim H_T at time T using d hedging instruments with price processes $X = (X^{(1)}, \dots, X^{(d)})$. So the *local risk* \mathcal{E}_n associated with the trading times t_n and t_{n+1} writes

$$\mathcal{E}_n = V_{t_{n+1}} - V_{t_n} - \vartheta_{t_n} \cdot X_{t_{n+1}} - X_{t_n}. \quad (1.1)$$

Here, V stands for the valuation process and $\vartheta = (\vartheta^{(1)}, \dots, \vartheta^{(d)})$, for the hedging process. Also, $\vartheta^{(i)}$ denotes the number of shares invested in the i -th hedging instrument. Up to considering discounted prices, we suppose the non-risky asset has zero drift.

In high-frequency hedging, the impact of discrete-time hedging compared to continuous-time one is small (see, for instance, [GT01] for results about convergence rate). In low-frequency hedging such as in energy markets [CDM17], the local residual risk is slightly bigger and may become an issue. Our aim is to find the valuation/hedging rules (V, ϑ) minimizing this risk. We differ

from the existing results (for instance, those related to the quadratic local risk minimization [FS88, Sch99]) by dealing with a risk function ℓ penalizing asymmetrically profits ($\mathcal{E}_n < 0$) and losses ($\mathcal{E}_n > 0$). So the *integrated local risk* \mathbf{R}_N under study takes the form

$$\mathbf{R}_N(V, \vartheta) = \sum_{n=0}^{N-1} \mathbb{E}[\ell(\mathcal{E}_n)].$$

The simplest case of such a risk function ℓ is

$$\ell_\gamma(y) = (1 + \gamma \operatorname{sgn}(y))^2 y^2 / 2 \quad (1.2)$$

where $\gamma \in (0, 1)$ to penalize losses further than profits (see Figure 1.1). We define the above sign function as $\operatorname{sgn}(y) := \mathbf{1}_{y>0} - \mathbf{1}_{y<0}$.

In this setting, our aim is to study the asymptotics of the minimum

$$\min_{(V, \vartheta) \in \mathcal{A}} \mathbf{R}_N(V, \vartheta) \quad (1.3)$$

as the number N of trading dates becomes larger. To simplify we take equidistant trading times $t_n = n\Delta t$ with time step $\Delta t = T/N$. The minimum (1.3) is computed over the set \mathcal{A} of all adapted to the underlying filtration $(\mathcal{F}_t)_{t \geq 0}$ and appropriately integrable pair (V, ϑ) , under the replication constraint $V_T = H_T$.

There are a few results in that direction. In [Pha00], the author deals with a \mathbb{L}_p risk function of the losses and a fixed number trading dates. In [PB04], the authors consider expected shortfall risk function. Their research concentrates on numerics for a fixed number of dates and does not handle any asymptotic analysis. In [AM11], the authors study pseudo-optimal strategies and get asymptotic results under the condition that the risk function is of class C^3 . So their analysis discards the prototype risk function (1.2). Indeed, the discontinuity of the second derivative ℓ''_γ complicates the analysis and fully changes the nature of subsequent results. In short, the existing references consider different settings and difficulties from ours.

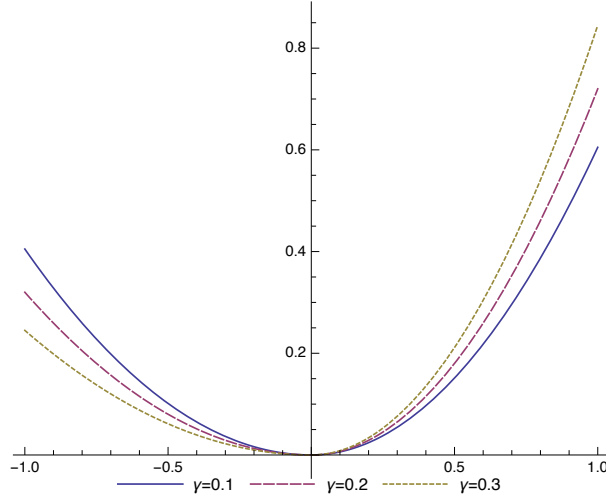
PDE valuation. The minimization problem (1.3) appears attractive, but its study in the asymptotic regime $N \rightarrow +\infty$ is tough in the case of asymmetric risk function (1.2). To tackle this problem, we slightly modify the approach.

First, we suppose the hedging instruments are modeled by a Stochastic Differential Equation (SDE) with drift μ and diffusion σ . We also consider contingent claims of the form $H_T = g(X_T)$. Second, we suppose that the contingent claim is evaluated exogenously by a valuation process $V_t = v(t, X_t)$ for some function v . For instance, v is given by a mark-to-model value promoted by the regulator or the Central Counterparty (CCP). The latter imposes its minimum margin requirement to which the hedging entity has to comply with. Given this exogenous reference valuation, the trader will determine how to hedge on each interval $[t_n, t_{n+1}]$ by choosing an adapted valuation/hedging rule $(\tilde{V}_{t_n}, \tilde{\vartheta}_{t_n})$ and considering the related *conditional local risk* $\mathbf{R}_{n,\gamma}$

$$\mathbf{R}_{n,\gamma} = \mathbb{E}[\ell_\gamma(V_{t_{n+1}} - \tilde{V}_{t_n} - \tilde{\vartheta}_{t_n} \cdot X_{t_{n+1}} - X_{t_n}) \mid \mathcal{F}_{t_n}]. \quad (1.4)$$

To clarify, the valuation/hedging rule of the trader will be parametrized by a function f , possibly nonlinear. Inspired by the connection between dynamic risk valuations, nonlinear Partial Differential Equations (PDE) and nonlinear Backward Stochastic Differential Equations (BSDE) [EPQ97, Pen04, Cre13], we introduce the concept of the *f-PDE valuation*. Let

$$\sigma : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}, \quad f : [0, T] \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{R}^{d \times d} \rightarrow \mathbb{R}$$

Figure 1.1 – Risk function ℓ_γ for different risk parameter γ .

be continuous functions. Let $\tau \in (0, T]$ be a time horizon and let $v(\tau, \cdot)$ be a reference valuation at the time τ . Given τ and $v(\tau, \cdot)$, the function $u^\tau : [0, \tau] \times \mathbb{R}^d \rightarrow \mathbb{R}$ is a solution to the f -PDE, if it satisfies

$$u_t^\tau(t, x) + \frac{1}{2} \text{Tr}[\sigma \sigma^\top D_x^2 u^\tau](t, x) + f(t, x, u^\tau(t, x), D_x u^\tau(t, x), D_x^2 u^\tau(t, x)) = 0, \quad (1.5)$$

for all $(t, x) \in [0, T] \times \mathbb{R}^d$ with the terminal condition $u^\tau(\tau, x) = v(\tau, x)$ at the time τ . The f -PDE valuation is the mapping from $(\tau, v(\tau, \cdot))$ to the f -PDE (1.5) solution u^τ ; this is typically the nonlinear valuation/hedging rule of the trader. We refer to f as the PDE nonlinearity and $f \equiv 0$ corresponds to the usual risk-neutral valuation [KS98], other nonlinearities appear in [EPQ97] for instance. Then, in the conditional local risk expression given by (1.4), we naturally set

$$\tilde{V}_{t_n} = u^{(n+1)}(t_n, X_{t_n}), \quad \tilde{\vartheta}_{t_n} = D_x u^{(n+1)}(t_n, X_{t_n}),$$

where we denote $u^{(n+1)} := u^{t_{n+1}}$.

Our contributions. Our first main result is to prove the existence (Theorem 1.6) of the following limit, called the *asymptotic risk*,

$$\mathbf{R}_\gamma(v, f) = \lim_{N \rightarrow +\infty} \frac{1}{\Delta t} \sum_{n=0}^{N-1} \mathbb{E}[\mathbf{R}_{n, \gamma}]. \quad (1.6)$$

Moreover, we give an explicit expression for $\mathbf{R}_\gamma(v, f)$ depending on γ , v , f , σ , X and T . Then, we discuss the existence of an optimal PDE nonlinearity f^* such that the f^* -PDE valuation minimizes the asymptotic risk in the sense

$$\mathbf{R}_\gamma(v, f^*) \leq \mathbf{R}_\gamma(v, f), \quad (1.7)$$

for any admissible f . In dimension 1, this optimal PDE nonlinearity f^* is explicit (see (1.22)) and it depends on the risk parameter γ , on the reference valuation second derivative and on the price process volatility.

Now, a natural choice for the reference valuation may be the solution to the f^* -PDE (1.5). Here, the payoff $g : \mathbb{R} \rightarrow \mathbb{R}$ is the f^* -PDE terminal condition at the time T . We denote by v^* the

resulting valuation. In dimension one, this PDE takes the form

$$v_t^*(t, x) + \frac{1}{2}\sigma^2(t, x)v_{xx}^*(t, x) + c_1\sigma^2(t, x)(v_{xx}^*(t, x))_+ - c_2\sigma^2(t, x)(v_{xx}^*(t, x))_- = 0$$

for some constants $c_1 \geq 0$ and $c_2 \leq 0$ depending on the risk parameter γ . In higher dimension, v^* solves a fully nonlinear PDE with a nonlinear term depending on the Hessian $D_x^2 v^*$ (see the nonlinear PDE (1.19)).

It gives somehow a consistent way to value the claim g by accounting for local hedging errors measured with the asymmetric risk function ℓ_γ . To the best of our knowledge, this work is an original contribution, where local hedging errors are analyzed with asymmetric risk function. We perform the asymptotics of a large number of trading dates and we derive an optimal valuation/hedging policy.

Summing up, instead of minimizing (1.3) and then taking the limit in N after rescaling by Δt , we take first the limit in N of the cumulated integrated local risk for a wide class of f -PDE valuation and then minimize over all nonlinearities f . We do not prove that inverting minimization and limit holds true in this setting. In other words, we do not state the limit of the minimum (1.3) rescaled by Δt corresponds to $\mathbf{R}_\gamma(v^*, f^*)$. However, our numerical tests in dimension one seem to corroborate this fact. Proving this result rigorously is, so far, an open problem, that we expect to handle in the next future.

This chapter is structured as follows. In Section 1.2, we present the notations, the stochastic setting and the assumptions. The Section 1.3 is reserved to the main result: the existence of the asymptotic risk. The proofs are gathered in Section 1.4. Section 1.5 contains our numerical experiments. Some technical results are collected in Appendix 1.6.

1.2 Notations and assumptions

Usual notations. Let $d \in \mathbb{N}^*$ and let a, b in \mathbb{R}^d . We denote by $a \cdot b = \sum_{i=1}^d a_i b_i$ the scalar product on \mathbb{R}^d , adopted for both row or column vectors a and b . We set $\|a\| = \sqrt{a \cdot a}$. We denote by \mathcal{M}^d the set of all $d \times d$ matrices with real entries. By \mathbb{S}^d we denote all symmetric matrices in \mathcal{M}^d . Let $A \in \mathcal{M}^d$, we denote $\text{Tr}[A]$ and A^\top respectively the trace and the transpose of a matrix A . For A in \mathcal{M}^d , we set $\|A\| = \sqrt{\text{Tr}[AA^\top]}$.

Let E, E' be two generic Euclidean space and let $\phi : [0, T] \times E$ be a E' -valued function. In this work, we say ϕ satisfies a *local regularity condition in time and space* if for some real $q > 0$ the coefficient

$$\|\phi\|_{C_{\text{loc}, \text{pol}}^{1/2, 1}} := \sup_{t \neq t' \in [0, T]} \sup_{x \neq x' \in E} \frac{\|\phi(t, x) - \phi(t', x')\|}{(|t - t'|^{1/2} + \|x - x'\|)(1 + \|x\|^q + \|x'\|^q)}$$

is finite, then ϕ is said to be in $C_{\text{loc}, \text{pol}}^{1/2, 1}$. We are aware that $\|\phi\|_{C_{\text{loc}, \text{pol}}^{1/2, 1}}$ depends on q but in the following, the precise value of q is unimportant and we prefer to avoid the reference to q in the notation $\|\phi\|_{C_{\text{loc}, \text{pol}}^{1/2, 1}}$ for the sake of simplicity.

Observe that $\phi \in C_{\text{loc}, \text{pol}}^{1/2, 1}$ means that ϕ is locally 1/2-Hölder continuous in time and Lipschitz continuous in space; and it has polynomial growth in space uniformly in time. Furthermore, we assert that for any ϕ_1 and ϕ_2 in $C_{\text{loc}, \text{pol}}^{1/2, 1}$, the product $\phi_1 \phi_2$, the pair (ϕ_1, ϕ_2) and the composition w.r.t. the space variable $\phi_1(t, \phi_2(t, \cdot))$ are also in $C_{\text{loc}, \text{pol}}^{1/2, 1}$.

The set $C^{1,2}([0, T] \times E, E')$ denotes the set of functions $\phi : [0, T] \times E \rightarrow E'$ such that the partial derivatives $\partial_t \phi, \partial_{x_i} \phi, \partial_{x_i} \partial_{x_j} \phi$ exist and are continuous, for any $1 \leq i, j \leq d$. When $E = \mathbb{R}^d$ and

the domain E' is unambiguous, we simply write $C^{1,2}([0, T] \times \mathbb{R}^d)$.

For every function $\phi \in C^{1,2}([0, T] \times \mathbb{R}^d, \mathbb{R})$, we denote its gradient in space by a row vector $D_x \phi = (\phi_{x_i})_{1 \leq i \leq d}$ and its Hessian by $D_x^2 \phi = (\phi_{x_i x_j})_{1 \leq i, j \leq d}$. Also, let $\mathcal{L}_t \phi : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ be given by

$$\mathcal{L}_t \phi(t, x) = \phi_t(t, x) + \frac{1}{2} \text{Tr}[\sigma \sigma^\top D_x^2 \phi](t, x).$$

Notice that $\phi, \phi_t, D_x \phi, D_x^2 \phi \in C_{\text{loc, pol}}^{1/2, 1}$ is a sufficient condition to have $\phi \in C^{1,2}$ and be able to apply Ito's formula.

Probabilistic model. We fix a finite time horizon $T > 0$. Let $W = (W^{(1)}, \dots, W^{(d)}) : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ be a standard Brownian motion on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$. Let $\mathbb{F} = \{\mathcal{F}_t, t \in [0, T]\}$ be the augmented and completed filtration generated by W . We consider the \mathbb{F} -adapted process $X = (X^{(1)}, \dots, X^{(d)}) : [0, T] \times \Omega \rightarrow \mathbb{R}^d$ satisfying the following stochastic differential equation (SDE)

$$dX_t = \mu(t, X_t) dt + \sigma(t, X_t) dW_t, \quad (1.8)$$

with initial value $X_0 = x_0 \in \mathbb{R}^d$. The coefficients $\mu : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\sigma : [0, T] \times \mathbb{R}^d \rightarrow \mathcal{M}^d$ are Lipschitz in space uniformly in time (see Assumption 1.1 later).

Given $N \in \mathbb{N}^*$ equidistant hedging times $\{t_0 = 0 < t_1 < \dots < t_N = T\}$ on the interval $[0, T]$, with $t_n = n\Delta t$ and $\Delta t = T/N$, we write

$$\varphi_t^N := \sup \{t_n \mid t_n \leq t\}, \quad \bar{\varphi}_t^N := \inf \{t_n \mid t_n > t\}$$

and the increment of X from t_n to t_{n+1} as $\Delta X_n = X_{t_{n+1}} - X_{t_n}$.

In the following we systematically consider the risk function ℓ_γ as defined in (1.2). It is a convex and continuously differentiable function satisfying $\ell_\gamma(0) = \ell'_\gamma(0) = 0$ and $\ell_\gamma(y) = \ell_{-\gamma}(-y)$. In addition, it is symmetric if and only if $\gamma = 0$. Further, ℓ'_γ is a piecewise continuously differentiable function with ℓ''_γ being discontinuous as soon as $\gamma \neq 0$:

$$\ell'_\gamma(y) = (1 + \gamma \text{sgn}(y))^2 y, \quad \ell''_\gamma(y) = (1 + \gamma \text{sgn}(y))^2, \quad (1.9)$$

where ℓ''_γ is extended to zero as $\ell''_\gamma(0) = 1$, owing to $\text{sgn}(0) = 0$ (see Figure 1.2). In all the sequel, we assume $\gamma \in [0, 1]$.

For a payoff function $g : \mathbb{R}^d \rightarrow \mathbb{R}$, the input of our approach are a reference valuation $v :$

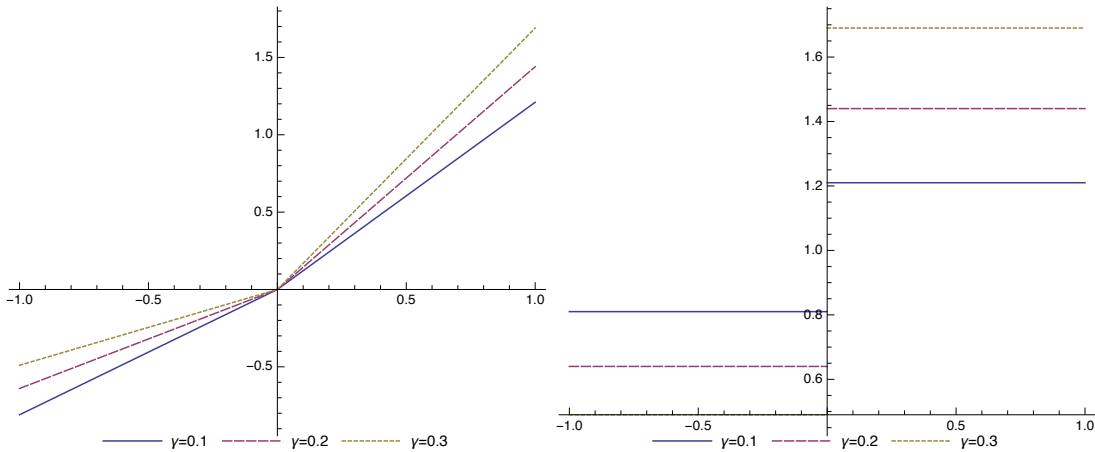


Figure 1.2 – Risk function ℓ_γ derivatives for different risk parameter γ : ℓ'_γ (left) and ℓ''_γ (right).

$[0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ such that $v(T, \cdot) = g(\cdot)$ and a PDE nonlinearity $f : [0, T] \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{S}^d \rightarrow \mathbb{R}$. Both are assumed to be smooth functions (see Assumptions 1.2 and 1.3). So we associate the f -PDE valuation giving rise to the family of functions $u^{t_{n+1}} : [0, t_{n+1}] \times \mathbb{R}^d \rightarrow \mathbb{R}$ indexed by hedging times t_{n+1} . These functions are the solutions to the PDE (1.5) with Cauchy boundary condition $u^{t_{n+1}}(t_{n+1}, \cdot) = v(t_{n+1}, \cdot)$ at time horizon t_{n+1} . Also, they are assumed to be smooth in the sense of Assumption 1.4. In this context, we set $u^{(n+1)} = u^{t_{n+1}}$ and define the local residual risk $\mathcal{E}_n : \Omega \rightarrow \mathbb{R}$ (see (1.1)) by

$$\mathcal{E}_n = u^{(n+1)}(t_{n+1}, X_{t_{n+1}}) - u^{(n+1)}(t_n, X_{t_n}) - D_x u^{(n+1)}(t_n, X_{t_n}) \Delta X_n \quad (1.10)$$

and the conditional local risk

$$\mathbf{R}_{n,\gamma} = \mathbb{E}[\ell_\gamma(\mathcal{E}_n) \mid \mathcal{F}_{t_n}]. \quad (1.11)$$

As explained in introduction, our aim is to analyze the asymptotic behavior of the integrated conditional local risk, after appropriate renormalization,

$$\mathbf{R}_{N,\gamma}(v, f) = \frac{1}{\Delta t} \sum_{n=0}^{N-1} \mathbb{E}[\mathbf{R}_{n,\gamma}]. \quad (1.12)$$

Main assumptions. Here, we study the asymptotic risk \mathbf{R}_γ (as defined in (1.6)) where a reference valuation v and a PDE nonlinearity f are given. We state the following assumptions.

Assumption 1.1. The coefficients $\mu : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}^d$ and $\sigma : [0, T] \times \mathbb{R}^d \rightarrow \mathcal{M}^d$ are in $C_{loc,pol}^{1/2,1}$.

Assumption 1.2. The reference valuation $v : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ is in $C_{loc,pol}^{1/2,1}$. Further, $D_x v$ and $D_x^2 v$ exist and are in $C_{loc,pol}^{1/2,1}$.

Assumption 1.3. The PDE nonlinearity $f : [0, T] \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{S}^d \rightarrow \mathbb{R}$ is in $C_{loc,pol}^{1/2,1}$.

Assumption 1.4. For all $\tau \in (0, T]$, there is a unique classical solution u^τ to the PDE (1.5) with the terminal condition $u^\tau(\tau, \cdot) = v(\tau, \cdot)$ at the time τ . In addition,

$$u_t^\tau, u_{x_i}^\tau, u_{x_i x_j}^\tau, u_{t x_i}^\tau, u_{x_i x_j x_k}^\tau,$$

exist and are in $C_{loc,pol}^{1/2,1}$.

Assumption 1.5. A non-degeneracy condition: the symmetric matrix $(\sigma^\top(D_x^2 v)\sigma)(t, X_t)$ is not 0 $dt \otimes d\mathbb{P}$ -a.e.

1.3 Asymptotic risk: Existence of the limit $\mathbf{R}_{N,\gamma}$

For stating the asymptotic result below, we need to introduce an extra Brownian motion B , independent of W , with the same dimension as W . All these are defined on an extended probability space with obvious definitions. Whenever necessary, the expectation w.r.t. the distribution of B , or W , or both, is denoted by \mathbb{E}^B , or \mathbb{E}^W , or $\mathbb{E}^{W \otimes B}$.

Theorem 1.6 (Existence of $\lim_{N \rightarrow +\infty} \mathbf{R}_{N,\gamma}$). Let $B = (B^{(1)}, \dots, B^{(d)}) : [0, 1] \times \Omega \rightarrow \mathbb{R}^d$ be another standard Brownian motion independent from W . Consider $\mathbf{R}_{N,\gamma}(v, f)$ given by (1.12) in the following form

$$\mathbf{R}_{N,\gamma}(v, f) = \frac{1}{\Delta t} \sum_{n=0}^{N-1} \mathbb{E}[\ell_\gamma(\mathcal{E}_n)],$$

where \mathcal{E}_n is given by (1.10). Under the Assumptions 1.1-1.5, the limit of $\mathbf{R}_{N,\gamma}(v, f)$ as $N \rightarrow \infty$ exists and is given by,

$$\begin{aligned} \mathbf{R}_\gamma(v, f) = \mathbb{E} \left[\int_0^T \int_0^1 \ell''_\gamma \left(\int_0^\theta B_{\theta'}^\top G_t dB_{\theta'} - F_t \theta \right) \right. \\ \left. \times \left(F_t^2 \theta - F_t \int_0^\theta B_{\theta'}^\top G_t dB_{\theta'} + \|G_t B_\theta\|^2 / 2 \right) d\theta dt \right], \end{aligned} \quad (1.13)$$

where

$$\begin{aligned} F_t &= f(t, X_t, v(t, X_t), D_x v(t, X_t), D_x^2 v(t, X_t)) \in \mathbb{R}, \\ G_t &= (\sigma^\top (D_x^2 v) \sigma)(t, X_t) \in \mathbb{S}^d \end{aligned}$$

The long and delicate proof is postponed to Section 1.4.

1.3.1 Asymptotic optimality through the PDE nonlinearity

Here, we study the optimization problem over the PDE nonlinearity f described in (1.7). To precise the definition of optimal PDE nonlinearity f^* , we rewrite the asymptotic risk in (1.13) as a functional $\mathbf{R}_\gamma : \Omega_v \times \Omega_f \rightarrow \mathbb{R}$ given by

$$\mathbf{R}_\gamma(v, f) = \mathbb{E} \left[\int_0^T \mathbf{L}_\gamma(G_t, F_t) dt \right],$$

where $\mathbf{L}_\gamma : \mathbb{S}^d \times \mathbb{R} \rightarrow \mathbb{R}$ is

$$\begin{aligned} \mathbf{L}_\gamma(S, a) = \mathbb{E} \left[\int_0^1 \ell''_\gamma \left((B_\theta^\top S B_\theta - \text{Tr}[S]\theta) / 2 - a\theta \right) \right. \\ \left. \times \left(a^2 \theta - a (B_\theta^\top S B_\theta - \text{Tr}[S]\theta) / 2 + (B_\theta^\top S^\top S B_\theta) / 2 \right) d\theta \right], \end{aligned} \quad (1.14)$$

with

$$\Omega_v = \left\{ v \in C_{\text{loc,pol}}^{1/2,1} \mid D_x v, D_x^2 v \in C_{\text{loc,pol}}^{1/2,1} \right\}, \quad \Omega_f = \left\{ f \in C_{\text{loc,pol}}^{1/2,1} \right\}.$$

We aim at proving the existence of minimizers to the variational problem

$$\min_{f \in \Omega_f} \mathbf{R}_\gamma(v, f), \quad (1.15)$$

for all $v \in \Omega_v$. Observe that the minimizer $f^\dagger(t, x, y, z, A)$ defined by (for any fixed (t, x, y, z, A))

$$f^\dagger(t, x, y, z, A) = \arg \min_{a \in \mathbb{R}} \mathbf{L}_\gamma(\sigma^\top(t, x) A \sigma(t, x), a),$$

is also a minimizer to (1.15) under the condition to be in Ω_f . Indeed, we just need to integrate and to take the expectation in both sides of

$$\mathbf{L}_\gamma(G(t, X_t), f^\dagger(t, X_t, v(t, X_t), D_x v(t, X_t), D_x^2 v(t, X_t))) \leq \mathbf{L}_\gamma(G(t, X_t), F(t, X_t)).$$

This is why we seek a minimizer to $a \mapsto \mathbf{L}_\gamma(S, a)$ for a given symmetric matrix S .

We now prove the existence of a minimizer.

Proposition 1.7. *Let $\gamma \in [0, 1)$ and $S \in \mathbb{S}^d$. Consider the minimization problem*

$$\min_{a \in \mathbb{R}} \mathbf{L}_\gamma(S, a). \quad (1.16)$$

Under the hypothesis of Theorem 1.6, there exists a global minimizer $a^ \in \mathbb{R}$ such that $\mathbf{L}_\gamma(S, a^*) \leq \mathbf{L}_\gamma(S, a)$ for all $a \in \mathbb{R}$.*

If a^* is unique for each S , we define the mapping $S \mapsto \chi^*(S) = a^*$. Then a natural candidate for f^* is given by

$$f^*(t, x, y, z, A) = \chi^*(\sigma^\top(t, x) A \sigma(t, x)), \quad (1.17)$$

for any $t, x, y, z, A \in [0, T] \times \mathbb{R}^d \times \mathbb{R} \times \mathbb{R}^d \times \mathbb{S}^d$.

Proof. First, we show that the function $\mathbf{L}_\gamma(S, a)$ is coercive and continuous in a . For any $\theta \in (0, 1]$, we consider $Z_\theta^{S,a} = (B_\theta^\top S B_\theta - \text{Tr}[S]\theta)/2 - a\theta$. Through simple computations, we check that $Z_\theta^{S,a}$ is continuous in a and integrable w.r.t. $d\mathbb{P}^B \otimes d\theta$.

1. *Coercivity.* We exhibit a coercive function which bounds $\mathbf{L}_\gamma(S, a)$ from below. Owing to the boundedness of ℓ''_γ , we estimate

$$\begin{aligned} \ell''_\gamma(Z_\theta^{S,a}) \left(-aZ_\theta^{S,a} + \|SB_\theta\|^2/2 \right) &\geq (1-\gamma)^2 \left(a^2\theta + \|SB_\theta\|^2/2 \right) \\ &\quad - (1+\gamma)^2 |a| \left(|B_\theta^\top S B_\theta| + |\text{Tr}[S]|\theta \right) / 2, \end{aligned}$$

$d\mathbb{P}^B \otimes d\theta$ -almost surely. By integrating in θ and taking the expectation of the previous estimate, we get

$$\mathbf{L}_\gamma(S, a) \geq (1-\gamma)^2 \left(a^2/2 + \text{Tr}[S^\top S]/4 \right) - (1+\gamma)^2 |a| \left(\mathbb{E}|G^\top S G| + |\text{Tr}[S]| \right) / 4,$$

where G is a standard normal random vector. Then we conclude that $a \mapsto \mathbf{L}_\gamma(S, a)$ is coercive.

2. *Continuity.* First we take $S = \mathbf{0}$ and we get

$$\mathbf{L}_\gamma(\mathbf{0}, a) = (1 + \gamma \text{sgn}(-a))a^2/2.$$

Therefore, $a \mapsto \mathbf{L}_\gamma(\mathbf{0}, a)$ is a continuous and strictly convex function. Then we conclude that there is a unique global minimizer given by $a^* = 0$. Now we take $S \neq \mathbf{0}$ and decompose $\mathbf{L}_\gamma(S, a)$ as follows

$$\mathbf{L}_\gamma(S, a) = -\mathbb{E} \left[\int_0^1 a Z_\theta^{S,a} \ell''_\gamma(Z_\theta^{S,a}) d\theta \right] + \frac{1}{2} \mathbb{E} \left[\int_0^1 \ell''_\gamma(Z_\theta^{S,a}) \|SB_\theta\|^2 d\theta \right]. \quad (1.18)$$

By replacing the expression of ℓ''_γ (see Equation (1.9)), we get

$$a \mapsto a Z_\theta^{S,a} \ell''_\gamma(Z_\theta^{S,a}) = (1 + \gamma^2) a Z_\theta^{S,a} + 2\gamma a |Z_\theta^{S,a}|,$$

which is continuous $d\mathbb{P}^B \otimes d\theta$ -almost surely and bounded by $(1 + \gamma^2)|a||Z_\theta^{S,a}|$ (integrable w.r.t. $d\mathbb{P}^B \otimes d\theta$ locally uniformly in a). By the dominated convergence theorem, we conclude the first term of the decomposition in (1.18) is continuous in a . Also, we estimate $|\ell''_\gamma(Z_\theta^{S,a}) \|SB_\theta\|^2| \leq (1 + \gamma^2) \|SB_\theta\|^2$, which is integrable uniformly in a . Following that $B_\theta^\top S B_\theta$ has a density w.r.t. the Lebesgue measure (see the proof of Proposition 1.19 in Appendix), we get $Z_\theta^{S,a} \neq 0$, $d\mathbb{P}^B \otimes d\theta$ -almost surely. It holds that

$$a \mapsto \ell''_\gamma(Z_\theta^{S,a}) \|SB_\theta\|^2$$

is continuous $d\mathbb{P}^B \otimes d\theta$ -almost surely, due to the continuity of ℓ''_γ on \mathbb{R}^* . Now, we conclude the second term of the decomposition in (1.18) is also continuous in a , by applying again the dominated convergence theorem. Therefore, we have proved that $\mathbf{L}_\gamma(S, a)$ is continuous in a .

Let $\alpha \in \mathbb{R}$ large enough such that $K = \{a : \mathbf{L}_\gamma(S, a) \leq \alpha\}$ is non-empty. Due to the continuity and coercivity of $\mathbf{L}_\gamma(S, a)$, K is compact. Then, by Weierstrass's Theorem, we conclude the announced result. \square

Here, we have just shown the existence of a minimizer a^* to the Problem (1.16) for a given symmetric matrix S . The regularity of $\chi^*(S)$ has not been analyzed, because the uniqueness has not been proved. In fact, the uniqueness and smoothness of f^* of the problem (1.15) is challenging in the general case. Certainly, if $\chi^*(S)$ is unique, then we could define f^* as in (1.17). Then, a natural candidate for the self-consistent valuation/hedging rule is given by the solution to nonlinear PDE

$$\begin{cases} v_t^*(t, x) + \frac{1}{2} \text{Tr}[\sigma \sigma^\top D_x^2 v^*](t, x) + \chi^*(\sigma^\top(t, x) D_x^2 v^*(t, x) \sigma(t, x)) = 0, \\ v^*(T, x) = g(x). \end{cases} \quad (1.19)$$

This PDE is fully nonlinear with a nonlinear term depending on the Hessian. Unfortunately, in full generality, we are not able to prove the existence/uniqueness of a solution v^* satisfying Assumption 1.2. Also proving that the new PDE nonlinearity f^* fulfills Assumption 1.3 is not straightforward. Fortunately, the one-dimensional case provides us a quasi-explicit formulation for χ^* , which hopefully is a first step in the analysis of the PDE (1.19). Further investigation is left to future research.

1.3.2 Optimal PDE nonlinearity: Study in dimension 1

Here, we present a quasi-explicit formulation of the optimal PDE nonlinearity f^* in the one-dimensional case. Here, $(B_\theta^\top S B_\theta - \text{Tr}[S]\theta)/2$ becomes $(B_\theta^2 - \theta)y/2$ for $y = S \in \mathbb{R}$. So, we rewrite the function $\mathbf{L}_\gamma(S, a)$ given by (1.14) as

$$\mathbf{L}_\gamma(y, a) = \mathbb{E} \left[\int_0^1 \ell''_\gamma \left(y (B_\theta^2 - \theta) / 2 - a\theta \right) \left(a^2\theta - ay (B_\theta^2 - \theta) / 2 + y^2 B_\theta^2 / 2 \right) d\theta \right].$$

Let $\chi^* \in \mathbb{R}$ a global minimizer of $\min_{a \in \mathbb{R}} \mathbf{L}_\gamma(y, a)$. In the following proposition, we sum up some interesting properties of χ^* . We denote by $\Phi_{\mathcal{N}}$ the cumulative distribution function (CDF) of the standard normal distribution and $\phi_{\mathcal{N}} = \Phi'_{\mathcal{N}}$ its density.

Proposition 1.8. *Let $\gamma \in [0, 1)$.*

(a) *Let $c_1^* \in \mathbb{R}$ and $c_2^* \in \mathbb{R}$ be global minimizers of*

$$\min_{c \in \mathbb{R}} \mathbf{L}_\gamma(1, c) \text{ and } \min_{c \in \mathbb{R}} \mathbf{L}_\gamma(-1, c),$$

respectively. Then $\chi^(y) = c_1^* y \mathbf{1}_{y>0} + c_2^* y \mathbf{1}_{y<0}$ is a global minimizer of*

$$\min_{a \in \mathbb{R}} \mathbf{L}_\gamma(y, a).$$

(b) *The mapping*

$$c \mapsto \mathbf{L}_\gamma(1, c) \text{ and } c \mapsto \mathbf{L}_\gamma(-1, c)$$

are strictly convex. Thus, c_1^* and c_2^* are unique characterized by

$$(1 + \gamma^2) c_1^* + \gamma T(c_1^*) = 0 \text{ and } (1 + \gamma^2) c_2^* - \gamma T(c_2^*) = 0,$$

respectively, where

$$T(c) = 2c \mathbf{1}_{2c+1 \leq 0} + \left(8c \Phi_{\mathcal{N}}(-\sqrt{2c+1}) - 4\phi_{\mathcal{N}}(\sqrt{2c+1}) \sqrt{2c+1} - 2c \right) \mathbf{1}_{2c+1 > 0}.$$

Therefore, the minimizer $\chi^*(y)$ (defined through c_1^* and c_2^*) is unique.

Proof. (a) We start by the special case $y = 0$, we get $\mathbf{L}_{\gamma}(0, a) = (1 + \gamma \operatorname{sgn}(-a))^2 a^2 / 2$. So, $\chi^*(0) = 0$. Now we consider the more interesting case $y \neq 0$. By setting $c = a/y$, we rewrite $\mathbf{L}_{\gamma}(y, a)$

$$\begin{aligned} \mathbf{L}_{\gamma}(y, cy) &= \mathbb{E} \left[\int_0^1 \ell''_{\gamma} \left((B_{\theta}^2 - \theta)/2 - c\theta \right) \left(c^2\theta - c(B_{\theta}^2 - \theta)/2 + B_{\theta}^2/2 \right) d\theta \right] y^2 \mathbf{1}_{y>0} \\ &+ \mathbb{E} \left[\int_0^1 \ell''_{\gamma} \left(-(B_{\theta}^2 - \theta)/2 + c\theta \right) \left(c^2\theta - c(B_{\theta}^2 - \theta)/2 + B_{\theta}^2/2 \right) d\theta \right] y^2 \mathbf{1}_{y<0} \end{aligned} \quad (1.20)$$

because $\ell''_{\gamma}(y\zeta) = \ell''_{\gamma}(\zeta)$ if $y > 0$ and $\ell''_{\gamma}(y\zeta) = \ell''_{\gamma}(-\zeta)$ if $y < 0$, for any $\zeta \in \mathbb{R}$.

Consider a global minimizer $c^*(y)$ of $\min_{c \in \mathbb{R}} \mathbf{L}_{\gamma}(y, cy)$, then $\chi^*(y) = c^*(y)y$ is also a global minimizer of $\min_{a \in \mathbb{R}} \mathbf{L}_{\gamma}(y, a)$. Because $(y, c) \mapsto \mathbf{L}_{\gamma}(y, cy)$ is multiplicatively separable on $y > 0$ and on $y < 0$, we write $c^*(y) = c_1^* \mathbf{1}_{y>0} + c_2^* \mathbf{1}_{y<0}$, where c_1^* and c_2^* are respectively global minimizers of $\min_{c \in \mathbb{R}} \mathbf{L}_{\gamma}(1, c)$ and $\min_{c \in \mathbb{R}} \mathbf{L}_{\gamma}(-1, c)$.

(b) Let G be a standard normal random variable. It will be useful later to know $\mathbb{E}[G^2 \mathbf{1}_{G < \alpha}]$ for any real α : we have

$$\begin{aligned} \mathbb{E}[G^2 \mathbf{1}_{G < \alpha}] &= -\alpha \phi_{\mathcal{N}}(\alpha) + \Phi_{\mathcal{N}}(\alpha), \\ \mathbb{E}[G^2 \mathbf{1}_{G > \alpha}] &= \alpha \phi_{\mathcal{N}}(-\alpha) + \Phi_{\mathcal{N}}(-\alpha), \\ \mathbb{E}[G^2 \mathbf{1}_{-\alpha < G < \alpha}] &= -2\alpha \phi_{\mathcal{N}}(-\alpha) + (\Phi_{\mathcal{N}}(\alpha) - \Phi_{\mathcal{N}}(-\alpha)). \end{aligned} \quad (1.21)$$

It holds that $B_{\theta} \sim \sqrt{\theta}G$ for all θ in $[0, 1]$. From (1.20), we get

$$\mathbf{L}_{\gamma}(1, c) = \frac{1 + \gamma^2}{2} T_1(c) + \gamma T_2(c), \quad \mathbf{L}_{\gamma}(-1, c) = \frac{1 + \gamma^2}{2} T_1(c) - \gamma T_2(c),$$

where

$$\begin{aligned} T_1(c) &= \mathbb{E} \left[c^2 - c(G^2 - 1)/2 + G^2/2 \right] = c^2 + 1/2, \\ T_2(c) &= \mathbb{E} \left[\operatorname{sgn} \left((G^2 - 1)/2 - c \right) \left(c^2 - c(G^2 - 1)/2 + G^2/2 \right) \right]. \end{aligned}$$

Considering $\alpha(c) = \sqrt{2c+1}$, it holds

$$\operatorname{sgn} \left((G^2 - 1)/2 - c \right) = \mathbf{1}_{2c+1 < 0} + \mathbf{1}_{2c+1 > 0} \left(\mathbf{1}_{G < -\alpha(c)} + \mathbf{1}_{G > \alpha(c)} - \mathbf{1}_{-\alpha(c) < G < \alpha(c)} \right).$$

From the expectations in (1.21), we deduce

$$\begin{aligned} T_2(c) &= \mathbf{1}_{2c+1 < 0} (c^2 + 1/2) \\ &+ \mathbf{1}_{2c+1 > 0} (c^2 + c/2) \mathbb{E} \left[\mathbf{1}_{G < -\alpha(c)} + \mathbf{1}_{G > \alpha(c)} - \mathbf{1}_{-\alpha(c) < G < \alpha(c)} \right] \\ &+ \mathbf{1}_{2c+1 > 0} (1/2 - c/2) \mathbb{E} \left[G^2 \mathbf{1}_{G < -\alpha(c)} + G^2 \mathbf{1}_{G > \alpha(c)} - G^2 \mathbf{1}_{-\alpha(c) < G < \alpha(c)} \right] \\ &= \mathbf{1}_{2c+1 < 0} (c^2 + 1/2) + \mathbf{1}_{2c+1 > 0} \beta(c), \end{aligned}$$

where

$$\beta(c) = (c^2 + 1/2) (3 - 4\Phi_{\mathcal{N}}(\alpha(c))) + 2(1 - c)\alpha(c)\phi_{\mathcal{N}}(\alpha(c)).$$

We easily check that $\mathbf{L}_{\gamma}(1, c)$ and $\mathbf{L}_{\gamma}(-1, c)$ are C^0 and piecewise C^2 . Let us compute their first derivatives for $c < -1/2$ and $c > -1/2$

$$\begin{aligned} \partial_c \mathbf{L}_{\gamma}(1, c) &= (1 + \gamma^2) c + \gamma (\mathbf{1}_{2c+1 < 0} 2c + \mathbf{1}_{2c+1 > 0} \beta'(c)) \\ &= \mathbf{1}_{2c+1 < 0} (1 + \gamma^2)^2 c + ((1 + \gamma^2) c + \gamma \beta'(c)) \mathbf{1}_{2c+1 > 0}, \\ \partial_c \mathbf{L}_{\gamma}(-1, c) &= (1 + \gamma^2) c - \gamma (\mathbf{1}_{2c+1 < 0} 2c + \mathbf{1}_{2c+1 > 0} \beta'(c)) \\ &= \mathbf{1}_{2c+1 < 0} (1 - \gamma^2)^2 c + ((1 + \gamma^2) c - \gamma \beta'(c)) \mathbf{1}_{2c+1 > 0}, \end{aligned}$$

where

$$\beta'(c) = 8c \Phi_{\mathcal{N}}(-\sqrt{2c+1}) - 4\phi_{\mathcal{N}}(\sqrt{2c+1})\sqrt{2c+1} - 2c.$$

Standard computations show that $\partial_c \mathbf{L}_{\gamma}(-1, c)$ and $\partial_c \mathbf{L}_{\gamma}(1, c)$ are continuous at $c = -1/2$. Moreover, we see that $\partial_c \mathbf{L}_{\gamma}(1, c)$ and $\partial_c \mathbf{L}_{\gamma}(-1, c)$ are strictly increasing on c under the condition that $|\beta''(c)| \leq 2$ on $2c + 1 > 0$. Indeed, we have

$$\beta''(c) = 6 - 8 \Phi_{\mathcal{N}}(\sqrt{2c+1}) \in [-2, 2].$$

due to $\Phi_{\mathcal{N}}(\sqrt{2c+1}) \in [1/2, 1]$ for all $2c + 1 > 0$. Because $\mathbf{L}_{\gamma}(1, c)$, $\mathbf{L}_{\gamma}(-1, c)$ are strictly convex, the optimal values c_1^* and c_2^* are unique and characterized respectively by $\partial_c \mathbf{L}_{\gamma}(1, c_1^*) = 0$ and $\partial_c \mathbf{L}_{\gamma}(-1, c_2^*) = 0$. \square

We depict the global minimizer χ^* in Figure 1.3. We show the approximate values of c_1^* and c_2^* calculated by a root finding algorithm in Table 1.1.

Therefore, in the spirit of Equation (1.17), we set

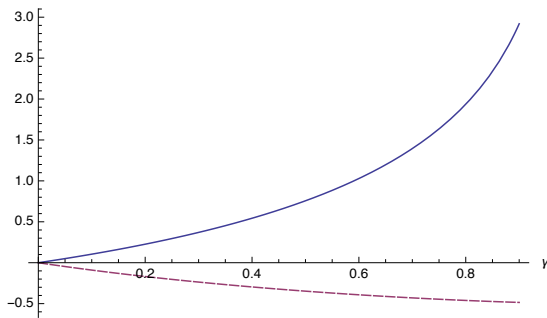
$$f^*(t, x, y, z, A) = f_{\gamma}^*(\sigma^2(t, x) A)$$

with f_{γ}^* denoting the optimal PDE nonlinearity in dimension 1:

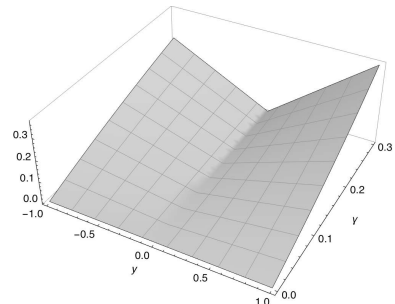
$$f_{\gamma}^*(y) := \chi^*(y) = c_1^* y \mathbf{1}_{y > 0} + c_2^* y \mathbf{1}_{y < 0}. \quad (1.22)$$

1.4 Asymptotic risk: Proof of the main result

The proof is long and technical. For this reason, we split it into different stages.



(a) Plot of $c_1^*(\gamma)$ (blue line) and $c_2^*(\gamma)$ (purple dashed) on $[0.0, 0.9]$.



(b) Plot of $\chi^*(\gamma, y)$ on $[0, 0.3] \times [-1, 1]$.

Figure 1.3 – Global minimizer $\chi^*(y)$.

Table 1.1 – Optimal slopes c_1^* and c_2^* .

γ	c_1^*	c_2^*
0.1	0.1043	-0.09013
0.2	0.2262	-0.1684
0.3	0.3702	-0.2366

- First, we study the conditional local risk $\mathbf{R}_{n,\gamma}$ on the interval $[t_n, t_{n+1}]$, by using a time-space rescaling argument (see Subsection 1.4.1). This rescaling turns out to be essential to pass to the limit later.
- Second, we derive an explicit approximation of the conditional local risk $\mathbf{R}_{n,\gamma}$ (see Subsection 1.4.2).
- Finally, we prove that the remainder terms converge almost surely towards 0. For this, we show that the Greeks of $u^\tau(t, \cdot)$ converge to those of $v(\tau, \cdot)$ as $t \uparrow \tau$ (see Subsection 1.4.2). Also, we show that the set of discontinuity points of ℓ_γ'' has measure zero under the Assumption 1.5.

In the proof, we use several constants $K_{n,N}(\xi)$ depending polynomially on the space variable ξ (uniformly in the interval $[t_n, t_{n+1}]$ and in the number of time steps). To simplify, we note $K_{n,N}(\xi) \in K_{\text{pol}}$ if for some real $q > 0$,

$$\sup_{N \in \mathbb{N}^*} \sup_{0 \leq n \leq N-1} \sup_{\xi \in \mathbb{R}^d} \frac{|K_{n,N}(\xi)|}{1 + \|\xi\|^q} < +\infty.$$

This upper bound depends on the polynomial bounds on the functions μ , σ , f , v and u .

1.4.1 Rescaling and conditioning

First, we start by a few observations.

- Thanks to the Markov property of the SDE and in view of our smoothness assumptions, $\mathbf{R}_{n,\gamma}$ is a continuous function of t_n and X_{t_n} only (see (1.11));
- $\mathbf{R}_{n,\gamma}$ goes to zero at rate Δt^2 , because we prove that the remainder of a second-order stochastic Taylor expansion will be inside ℓ_γ . Rescaling it by Δt , we expect to get a non-zero limit for the aggregated value of $\mathbf{R}_{n,\gamma}$ (see (1.12));
- Note that ℓ_γ'' has a jump discontinuity at zero (see (1.9)). Then to decompose the conditional local risk, we will need to apply a stronger version of Ito's formula, known as, the Ito-Tanaka formula.

In view of a Taylor-Ito expansion, we consider the process $X^{\Delta t} = \{X_\theta^{\Delta t}, \theta \in [0, 1]\}$ satisfying

$$dX_\theta^{\Delta t} = \Delta t \mu(t_n + \theta \Delta t, X_\theta^{\Delta t}) d\theta + \Delta t^{1/2} \sigma(t_n + \theta \Delta t, X_\theta^{\Delta t}) dB_\theta, \quad X_0^{\Delta t} = \xi \in \mathbb{R}^d, \quad (1.23)$$

where B is an extra Brownian motion independent from W . It is a time-space rescaling of the original process starting from ξ at t_n .

By denoting $X^{t,\xi}$ as the SDE solution starting from ξ at t , we notice that the processes $\{X_{t_n+\theta\Delta t}^{t_n,\xi}, \theta \in [0, 1]\}$ and $\{X_\theta^{\Delta t}, \theta \in [0, 1]\}$ have the same distribution. This is due to the fact both processes satisfy the same SDE generated by Brownian motions both independent from \mathcal{F}_{t_n} . Then we can

rewrite $\mathbf{R}_{n,\gamma}$ (see (1.11)) as a continuous function in terms of X_{t_n} and $X_\theta^{\Delta t}$. Setting

$$P^{\Delta t}(t_n, \xi) = \Delta t^{-2} \mathbb{E}^B \left[\ell_\gamma \left(u^{(n+1)}(t_{n+1}, X_1^{\Delta t}) - u^{(n+1)}(t_n, \xi) - D_x u^{(n+1)}(t_n, \xi) (X_1^{\Delta t} - \xi) \right) \right],$$

leads to

$$\mathbf{R}_{n,\gamma} = \Delta t^2 P^{\Delta t}(t_n, X_{t_n}). \quad (1.24)$$

1.4.2 Stochastic expansion and approximation of sensitivities

Proposition 1.9 (Stochastic expansion of $P^{\Delta t}(t_n, X_{t_n})$). *Assume notations and assumptions of Theorem 1.6. Denote $F^{(n+1)} : [0, t_{n+1}] \times \mathbb{R}^d \rightarrow \mathbb{R}$ and $G^{(n+1)} : [0, t_{n+1}] \times \mathbb{R}^d \rightarrow \mathbb{S}^d$ as*

$$\begin{aligned} F^{(n+1)}(t, \cdot) &= f(t, \cdot, u^{(n+1)}(t, \cdot), D_x u^{(n+1)}(t, \cdot), D_x^2 u^{(n+1)}(t, \cdot)); \\ G^{(n+1)}(t, \cdot) &= (\sigma^\top (D_x^2 u^{(n+1)}) \sigma)(t, \cdot). \end{aligned} \quad (1.25)$$

For any t_n and $\xi \in \mathbb{R}^d$, let $X_\theta^{\Delta t} : [0, 1] \times \Omega \rightarrow \mathbb{R}^d$ be the strong solution to the SDE (1.23) such that $X_0^{\Delta t} = \xi$ and let $\mathcal{E}_\theta^{\Delta t} : [0, 1] \times \Omega \rightarrow \mathbb{R}$ be the stochastic process defined by

$$\mathcal{E}_\theta^{\Delta t} = u^{(n+1)}(t_n + \theta \Delta t, X_\theta^{\Delta t}) - u^{(n+1)}(t_n, \xi) - D_x u^{(n+1)}(t_n, \xi) \cdot X_\theta^{\Delta t} - \xi \quad (1.26)$$

so that

$$P^{\Delta t}(t_n, \xi) = \Delta t^{-2} \mathbb{E}^B [\ell_\gamma(\mathcal{E}_1^{\Delta t})]. \quad (1.27)$$

The following local risk decomposition holds

$$\begin{aligned} P^{\Delta t}(t_n, \xi) &= \mathbb{E}^B \left[\int_0^1 \ell_\gamma'' \left(\mathcal{E}_\theta(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi)) + R_\theta^{\Delta t}(t_n, \xi) \right) \right. \\ &\quad \left. \times \mathcal{Q}_\theta \left(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi) \right) d\theta \right] + K_{n,N}(\xi) \Delta t^{1/2}, \end{aligned}$$

where

$$\mathcal{E}_\theta(S, y) = \int_0^\theta B_{\theta'}^\top S dB_{\theta'} - y\theta, \quad (1.28)$$

$$\mathcal{Q}_\theta(S, y) = y^2 \theta - y \int_0^\theta B_{\theta'}^\top S dB_{\theta'} + \|SB_\theta\|^2/2, \quad (1.29)$$

$$R_\theta^{\Delta t}(t_n, \xi) = \mathcal{E}_\theta^{\Delta t}/\Delta t - \mathcal{E}_\theta \left(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi) \right), \quad (1.30)$$

for some constant $K_{n,N}(\xi) \in K_{pol}$.

The proof of Proposition 1.9 is delicate. We postpone it to Subsection 1.4.4. In order to perform a second - order stochastic expansion, we need that $u^{(n+1)}$ and $D_x u^{(n+1)}$ be in $C^{1,2}$ to apply Ito's formula. Additionally, we require σ , $D_x u^{(n+1)}$, $D_x^2 u^{(n+1)}$, $D_x u_t^{(n+1)}$ and $D_x^2 D_x u^{(n+1)}$ to have polynomial growth to obtain proper integrability along the computations. Finally, we ask for σ and $D_x^2 u^{(n+1)}$ to be in $C_{loc,pol}^{1/2,1}$, which is useful at the stochastic expansion of the gradient $D_x u^{(n+1)}$. All the above conditions are satisfied thanks to our assumptions.

Notice that the above expansion of $P^{\Delta t}(t_n, \xi)$ depends on $u^{(n+1)}$, solution of the PDE (1.5) on the subinterval $[t_n, t_{n+1}]$, whose size goes to 0. Therefore, by invoking a small - time approximation argument, we replace $u^{(n+1)}$ and its first - second derivatives by its terminal value $v(t_{n+1}, \cdot)$ and its first - second derivatives. Notice that the reference valuation v is independent of Δt . This is

the matter of following statement, proved in Appendix 1.6.2.

Proposition 1.10 (Approximation of sensitivities). *Assume notations and assumptions of Theorem 1.6. Then, there exists some constant $K_{n,N}(\xi) \in K_{\text{pol}}$ such that*

$$\left| u^{(n+1)}(t_n, \xi) - v(t_{n+1}, \xi) \right| \leq K_{n,N}(\xi) \Delta t^{1/2}, \quad (1.31)$$

$$\left\| D_x u^{(n+1)}(t_n, \xi) - D_x v(t_{n+1}, \xi) \right\| \leq K_{n,N}(\xi) \Delta t^{1/2}, \quad (1.32)$$

$$\left\| D_x^2 u^{(n+1)}(t_n, \xi) - D_x^2 v(t_{n+1}, \xi) \right\| \leq K_{n,N}(\xi) \Delta t^{1/2}. \quad (1.33)$$

1.4.3 Aggregation and passage to the limit

We set

$$\begin{aligned} F(t, \xi) &= f(t, \xi, v(t, \xi), D_x v(t, \xi), (D_x^2 v)(t, \xi)) \in \mathbb{R}, \\ G(t, \xi) &= (\sigma^\top (D_x^2 v) \sigma)(t, \xi) \in \mathbb{S}^d. \end{aligned} \quad (1.34)$$

Replacing ξ by X_{t_n} in the expansion of $P^{\Delta t}(t_n, \xi)$ in Proposition 1.9 leads to

$$\begin{aligned} P^{\Delta t}(t_n, X_{t_n}) &= \mathbb{E}^B \left[\int_0^1 \ell''_\gamma \left(\mathcal{E}_\theta \left(G^{(n+1)}(t_n, X_{t_n}), F^{(n+1)}(t_n, X_{t_n}) \right) + R_\theta^{\Delta t}(t_n, X_{t_n}) \right) \right. \\ &\quad \left. \times \mathcal{Q}_\theta \left(G^{(n+1)}(t_n, X_{t_n}), F^{(n+1)}(t_n, X_{t_n}) \right) d\theta \right] + K_{n,N}(X_{t_n}) \Delta t^{1/2}, \end{aligned}$$

where $K_{n,N}(X_{t_n}) \in K_{\text{pol}}$. By substituting $u^{(n+1)}(t_n, \cdot)$ by its terminal value $v(t_{n+1}, \cdot)$ in $F^{(n+1)}(t_n, \cdot)$ and $G^{(n+1)}(t_n, \cdot)$ (see (1.25)), we get $F(t_{n+1}, \cdot)$ and $G(t_{n+1}, \cdot)$ (see (1.34)). Hence,

$$\begin{aligned} P^{\Delta t}(t_n, X_{t_n}) &= \mathbb{E}^B \left[\int_0^1 \ell''_\gamma \left(\mathcal{E}_\theta \left(G(t_{n+1}, X_{t_n}), F(t_{n+1}, X_{t_n}) \right) + \bar{R}_\theta^{\Delta t}(t_n, X_{t_n}) \right) \right. \\ &\quad \left. \times \mathcal{Q}_\theta \left(G(t_{n+1}, X_{t_n}), F(t_{n+1}, X_{t_n}) \right) d\theta \right] + \bar{C}^{\Delta t}(t_n, X_{t_n}) + K_{n,N}(X_{t_n}) \Delta t^{1/2}, \end{aligned}$$

where

$$\bar{R}_\theta^{\Delta t}(t_n, \xi) := \mathcal{E}_\theta(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi)) - \mathcal{E}_\theta(G(t_{n+1}, \xi), F(t_{n+1}, \xi)) + R_\theta^{\Delta t}(t_n, \xi) \quad (1.35)$$

$$\begin{aligned} \bar{C}^{\Delta t}(t_n, \xi) &:= \mathbb{E}^B \left[\int_0^1 \ell''_\gamma \left(\mathcal{E}_\theta(G(t_{n+1}, \xi), F(t_{n+1}, \xi)) + \bar{R}_\theta^{\Delta t}(t_n, \xi) \right) \right. \\ &\quad \left. \times \left(\mathcal{Q}_\theta(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi)) - \mathcal{Q}_\theta(G(t_{n+1}, \xi), F(t_{n+1}, \xi)) \right) d\theta \right]. \end{aligned} \quad (1.36)$$

In the sequel, we require estimates of $\bar{R}_\theta^{\Delta t}(t_n, X_{t_n})$ and $\bar{C}^{\Delta t}(t_n, X_{t_n})$, summarized in the following Proposition, proved later in Subsection 1.4.5.

Proposition 1.11 (Almost sure convergence of the remainder). *Under the assumptions of Theorem 1.6, for any $p \geq 1$, there exists a constant K_p such that*

$$(a) \quad \mathbb{E} \left[\sup_{0 \leq n \leq N-1} \sup_{\theta \in [0,1]} \right]$$

- $$\left| \mathcal{E}_\theta(G^{(n+1)}(t_n, X_{t_n}), F^{(n+1)}(t_n, X_{t_n})) - \mathcal{E}_\theta(G(t_{n+1}, X_{t_{n+1}}), F(t_{n+1}, X_{t_{n+1}})) \right|^p \leq K_p \Delta t^{p/2};$$
- (b) $\sup_{0 \leq n \leq N-1} \mathbb{E} |\bar{C}^{\Delta t}(t_n, X_{t_n})| \leq K_1 \Delta t^{1/2};$
- (c) $\sup_{0 \leq n \leq N-1} \sup_{\theta \in [0,1]} |\bar{R}_\theta^{\Delta t}(t_n, X_{t_n})| \xrightarrow[N \rightarrow \infty]{} 0, \text{ d}\mathbb{P}^W \otimes \text{d}\mathbb{P}^B\text{-a.s.}$

From the definition of $P^{\Delta t}$ in (1.24), we have $\Delta t^{-1} \mathbf{R}_{n,\gamma} = P^{\Delta t}(t_n, X_{t_n}) \Delta t$. By summing it for $0 \leq n \leq N-1$, we obtain

$$\begin{aligned} \Delta t^{-1} \mathbb{E} \left[\sum_{n=0}^{N-1} \mathbf{R}_{n,\gamma} \right] &= \mathbb{E} \left[\sum_{n=0}^{N-1} P^{\Delta t}(t_n, X_{t_n}) \Delta t \right] = \mathbb{E} \left[\int_0^T P^{\Delta t}(\varphi_t^N, X_{\varphi_t^N}) dt \right] \\ &= \mathbb{E}^{W \otimes B} \left[\int_0^T \int_0^1 \ell''_\gamma \left(\mathcal{E}_\theta(G(\bar{\varphi}_t^N, X_{\varphi_t^N}), F(\bar{\varphi}_t^N, X_{\varphi_t^N})) + \bar{R}_\theta^{\Delta t}(\varphi_t^N, X_{\varphi_t^N}) \right) \right. \\ &\quad \left. \times \mathcal{Q}_\theta(G(\bar{\varphi}_t^N, X_{\varphi_t^N}), F(\bar{\varphi}_t^N, X_{\varphi_t^N})) d\theta dt \right] + \sum_{n=0}^{N-1} \mathbb{E} \left[\bar{C}^{\Delta t}(t_n, X_{t_n}) \Delta t + K_{n,N}(X_{t_n}) \Delta t^{3/2} \right]. \end{aligned} \quad (1.37)$$

The last sum goes to 0 as $N \rightarrow +\infty$, owing to Proposition 1.11 and Proposition 1.9. It remains to determine the limit of the first term in (1.37). We achieve this result by applying the dominated convergence theorem.

Step 1. Because of $\sigma, v, D_x v, D_x^2 v, f \in C_{\text{loc, pol}}^{1/2,1}$ (therefore, they are continuous in time and space) and the path-continuity of X , we get $\text{d}\mathbb{P}^W$ -a.s. for any t

$$\begin{aligned} &(\sigma^\top(D_x^2 v)\sigma)(\bar{\varphi}_t^N, X_{\varphi_t^N}) \xrightarrow[N \rightarrow \infty]{} (\sigma^\top(D_x^2 v)\sigma)(t, X_t), \\ &f(\bar{\varphi}_t^N, X_{\varphi_t^N}, v(\bar{\varphi}_t^N, X_{\varphi_t^N}), D_x v(\bar{\varphi}_t^N, X_{\varphi_t^N}), D_x^2 v(\bar{\varphi}_t^N, X_{\varphi_t^N})) \\ &\quad \xrightarrow[N \rightarrow \infty]{} f(t, X_t, v(t, X_t), D_x v(t, X_t), D_x^2 v(t, X_t)). \end{aligned}$$

Hence, it holds $\text{d}\mathbb{P}^W \otimes \text{d}\mathbb{P}^B$ -a.s. for any θ, t

$$\begin{aligned} \mathcal{E}_\theta(G(\bar{\varphi}_t^N, X_{\varphi_t^N}), F(\bar{\varphi}_t^N, X_{\varphi_t^N})) &\xrightarrow[N \rightarrow \infty]{} \mathcal{E}_\theta(\Gamma(t, X_t), F(t, X_t)), \\ \mathcal{Q}_\theta(G(\bar{\varphi}_t^N, X_{\varphi_t^N}), F(\bar{\varphi}_t^N, X_{\varphi_t^N})) &\xrightarrow[N \rightarrow \infty]{} \mathcal{Q}_\theta(G(t, X_t), F(t, X_t)), \end{aligned}$$

because \mathcal{E}_θ and \mathcal{Q}_θ (see (1.28)-(1.29)) are continuous in S, y , $\text{d}\mathbb{P}^B \otimes \text{d}\theta$ -a.s. Also, from the item (c) of Proposition 1.11, we have

$$\sup_{0 \leq n \leq N-1} \sup_{\theta \in [0,1]} |\bar{R}_\theta^{\Delta t}(t_n, X_{t_n})| \xrightarrow[N \rightarrow \infty]{} 0,$$

$\text{d}\mathbb{P}^W \otimes \text{d}\mathbb{P}^B$ -almost surely.

Step 2. Seeing that the second derivative ℓ''_γ is discontinuous at 0 and the set

$$\mathcal{A} := \{(\omega, t, \theta) \in \Omega \times [0, T] \times [0, 1] : \mathcal{E}_\theta(G(t, X_t(\omega)), F(t, X_t(\omega))) (\omega) = 0\} \quad (1.38)$$

has measure zero (see Proposition 1.19 in Appendix), it holds

$$\ell''_\gamma \left(\mathcal{E}_\theta(G(\bar{\varphi}_t^N, X_{\varphi_t^N}), F(\bar{\varphi}_t^N, X_{\varphi_t^N})) + \bar{R}_\theta^{\Delta t}(\bar{\varphi}_t^N, X_{\varphi_t^N}) \right)$$

$$\xrightarrow{N \rightarrow \infty} \ell''_{\gamma} (\mathcal{E}_{\theta} (G(t, X_t), F(t, X_t))),$$

$d\mathbb{P}^W \otimes d\mathbb{P}^B \otimes dt \otimes d\theta$ -almost surely.

Step 3. Because of the boundedness of ℓ''_{γ} and the polynomial growth of $\sigma, v, D_x v, D_x^2 v$, we have

$$\begin{aligned} & \left| \ell''_{\gamma} \left(\mathcal{E}_{\theta} \left(\Gamma(\bar{\varphi}_t^N, X_{\varphi_t^N}), F(\bar{\varphi}_t^N, X_{\varphi_t^N}) \right) + \bar{R}_{\theta}^{\Delta t}(\bar{\varphi}_t^N, X_{\varphi_t^N}) \right) \mathcal{Q}_{\theta} \left(G(\bar{\varphi}_t^N, X_{\varphi_t^N}), F(\bar{\varphi}_t^N, X_{\varphi_t^N}) \right) \right| \\ & \leq C \left(1 + \sup_{t \in [0, T]} |X_t| + \|B_{\theta}\| + \left\| \int_0^{\theta} B_{\theta} dB_{\theta}^T \right\| \right)^q \end{aligned}$$

for some positive constants C and q .

By the dominated convergence theorem, we conclude

$$\begin{aligned} & \mathbb{E}^{W \otimes B} \left[\int_0^T \int_0^1 \ell''_{\gamma} \left(\mathcal{E}_{\theta} \left(\Gamma(\bar{\varphi}_t^N, X_{\varphi_t^N}), F(\bar{\varphi}_t^N, X_{\varphi_t^N}) \right) + \bar{R}_{\theta}^{\Delta t}(\bar{\varphi}_t^N, X_{\varphi_t^N}) \right) \right. \\ & \quad \left. \times \mathcal{Q}_{\theta} \left(\Gamma(\bar{\varphi}_t^N, X_{\varphi_t^N}), F(\bar{\varphi}_t^N, X_{\varphi_t^N}) \right) d\theta dt \right] \\ & \xrightarrow{N \rightarrow \infty} \mathbb{E}^{W \otimes B} \left[\int_0^T \int_0^1 \ell''_{\gamma} (\mathcal{E}_{\theta} (\Gamma(t, X_t), F(t, X_t))) \mathcal{Q}_{\theta} (\Gamma(t, X_t), F(t, X_t)) d\theta dt \right]. \end{aligned}$$

This completes the proof of Theorem 1.6.

1.4.4 Proof of the stochastic expansion in Subsection 1.4.2.

Now we present the proof of the results related to the stochastic expansion of $P^{\Delta t}(t_n, X_{t_n})$ in Subsection 1.4.2.

Proof of Proposition 1.9. For the sake of conciseness, we set $u = u^{(n+1)}$. By substituting $X_{\theta}^{\Delta t}$ in (1.23) into $\mathcal{E}_{\theta}^{\Delta t}$ in (1.26), we get

$$\begin{aligned} \mathcal{E}_{\theta}^{\Delta t} &= u(t_n + \theta \Delta t, X_{\theta}^{\Delta t}) - u(t_n, \xi) - \Delta t \int_0^{\theta} D_x u(t_n, \xi) \mu(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) d\theta' \\ & \quad - \Delta t^{1/2} \int_0^{\theta} D_x u(t_n, \xi) \sigma(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) dB_{\theta'}, \end{aligned} \quad (1.39)$$

where $D_x u(\cdot, \cdot)$ is a row vector.

In the proof, we use the Ito-Tanaka formula to $\ell_{\gamma}(\mathcal{E}_{\theta}^{\Delta t})$ between $\theta = 0$ and $\theta = 1$ and we perform some Taylor-Ito expansions in terms of Δt . Because $u, u_t, D_x u, D_x^2 u \in C_{\text{loc, pol}}^{1/2, 1}$, then $u \in C^{1, 2}([t_n, t_{n+1}] \times \mathbb{R}^d, \mathbb{R})$. Applying Ito's formula to $u(t_n + \theta \Delta t, X_{\theta}^{\Delta t})$ yields

$$\begin{aligned} u(t_n + \theta \Delta t, X_{\theta}^{\Delta t}) - u(t_n, \xi) &= \Delta t^{1/2} \int_0^{\theta} (D_x u \sigma)(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) dB_{\theta'} \\ & \quad + \Delta t \int_0^{\theta} (\mathcal{L}_{t_n + \theta' \Delta t} u)(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) d\theta' \\ & \quad + \Delta t \int_0^{\theta} ((D_x u) \mu)(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) d\theta'. \end{aligned} \quad (1.40)$$

Here, we denote $(\Delta D_x u)(t, \zeta) = D_x u(t, \zeta) - D_x u(t_n, \xi)$ for any $t \in [t_n, t_{n+1}]$ and $\zeta \in \mathbb{R}^d$. Replacing (1.40) in (1.39) leads to

$$\begin{aligned} \mathcal{E}_\theta^{\Delta t} &= \Delta t^{1/2} \int_0^\theta ((\Delta D_x u) \sigma)(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) dB_{\theta'} \\ &+ \Delta t \int_0^\theta (\mathcal{L}_{t_n + \theta' \Delta t} u)(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) d\theta' \\ &+ \Delta t \int_0^\theta ((\Delta D_x u) \mu)(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) d\theta'. \end{aligned} \quad (1.41)$$

Use that u solves the PDE (1.5) to simplify the second term above. Then, we apply the Ito-Tanaka formula to the convex function ℓ_γ (see Theorem 1.5 and Corollary 1.6 in [RY99, Chapter VI]) composed with the process $\mathcal{E}_\theta^{\Delta t}$ between $\theta = 0$ and $\theta = 1$. Because $\ell'_\gamma(y) = \ell''_\gamma(y)y$ for all $y \in \mathbb{R}$, we get

$$\begin{aligned} \ell_\gamma(\mathcal{E}_1^{\Delta t}) &= -\Delta t \int_0^1 \ell''_\gamma(\mathcal{E}_\theta^{\Delta t}) \mathcal{E}_\theta^{\Delta t} F^{(n+1)}(t_n + \theta \Delta t, X_\theta^{\Delta t}) d\theta \\ &+ \Delta t^{1/2} \int_0^1 \ell''_\gamma(\mathcal{E}_\theta^{\Delta t}) \mathcal{E}_\theta^{\Delta t} ((\Delta D_x u) \sigma)(t_n + \theta \Delta t, X_\theta^{\Delta t}) dB_\theta \\ &+ \Delta t \int_0^1 \ell''_\gamma(\mathcal{E}_\theta^{\Delta t}) \mathcal{E}_\theta^{\Delta t} ((\Delta D_x u) \mu)(t_n + \theta \Delta t, X_\theta^{\Delta t}) d\theta \\ &+ \frac{1}{2} \Delta t \int_0^1 \ell''_\gamma(\mathcal{E}_\theta^{\Delta t}) \|((\Delta D_x u) \sigma)(t_n + \theta \Delta t, X_\theta^{\Delta t})\|^2 d\theta. \end{aligned}$$

Considering $P^{\Delta t}(t_n, \xi)$ in (1.27), taking the expectation of the above expression and dividing by Δt^2 gives

$$P^{\Delta t}(t_n, \xi) =: P_1^{\Delta t}(t_n, \xi) + P_2^{\Delta t}(t_n, \xi) + P_3^{\Delta t}(t_n, \xi), \quad (1.42)$$

where

$$P_1^{\Delta t}(t_n, \xi) := -\Delta t^{-1} \mathbb{E}^B \left[\int_0^1 \ell''_\gamma(\mathcal{E}_\theta^{\Delta t}) \mathcal{E}_\theta^{\Delta t} F^{(n+1)}(t_n + \theta \Delta t, X_\theta^{\Delta t}) d\theta \right], \quad (1.43)$$

$$P_2^{\Delta t}(t_n, \xi) := \frac{1}{2} \Delta t^{-1} \mathbb{E}^B \left[\int_0^1 \ell''_\gamma(\mathcal{E}_\theta^{\Delta t}) \|((\Delta D_x u) \sigma)(t_n + \theta \Delta t, X_\theta^{\Delta t})\|^2 d\theta \right], \quad (1.44)$$

$$P_3^{\Delta t}(t_n, \xi) := \Delta t^{-1} \mathbb{E}^B \left[\int_0^1 \ell''_\gamma(\mathcal{E}_\theta^{\Delta t}) \mathcal{E}_\theta^{\Delta t} ((\Delta D_x u) \mu)(t_n + \theta \Delta t, X_\theta^{\Delta t}) d\theta \right]. \quad (1.45)$$

Here we have used that the stochastic integral in $\ell_\gamma(\mathcal{E}_1^{\Delta t})$ has expectation zero, following directly from $\mathbb{E} \int_0^1 |\mathcal{E}_\theta^{\Delta t}|^4 d\theta < +\infty$ and from the polynomial growth of σ and $D_x u$ (because $\sigma, D_x u \in C_{\text{loc}, \text{pol}}^{1/2, 1}$). Now we analyze the expansion of $\mathcal{E}_\theta^{\Delta t}$ and then apply it to $T_i^{\Delta t}(t_n, \xi)$ for $i = 1, 2, 3$.

Step 1: Expansion of $((\Delta D_x u) \sigma)(t_n + \theta \Delta t, X_\theta^{\Delta t})$ and $\mathcal{E}_\theta^{\Delta t}$. We approximate $\Delta D_x u$ up to order $\Delta t^{1/2}$, by setting

$$(\Delta D_x u)(t_n + \theta \Delta t, X_\theta^{\Delta t}) = \Delta t^{1/2} B_\theta^\top (\sigma^\top(D_x^2 u))(t_n, \xi) + r_\theta^{\Delta t}, \quad (1.46)$$

where $\Delta D_x u$ and $r_\theta^{\Delta t}$ are row vectors.

Lemma 1.12. *Let $p \geq 2$. Under the assumptions of Theorem 1.6, it holds*

- (a) $\sup_{0 \leq n \leq N-1} \sup_{\theta \in [0,1]} \mathbb{E}^B \|r_\theta^{\Delta t}\|^p \leq K_{n,N}(\xi) \Delta t^p$,
- (b) $\sup_{0 \leq n \leq N-1} \sup_{\theta \in [0,1]} \mathbb{E}^B \|(\Delta D_x u)(t_n + \theta \Delta t, X_\theta^{\Delta t})\|^p \leq K_{n,N}(\xi) \Delta t^{p/2}$,

for some constant $K_{n,N}(\xi) \in K_{pol}$.

Proof of Lemma 1.12. (a) Because $D_x u, D_x u_t, D_x^2 u, D_x^2 D_x u \in C_{loc,pol}^{1/2,1}$, then $D_x u \in C^{1,2}([t_n, t_{n+1}] \times \mathbb{R}^d)$. By applying Ito's formula to $D_x u(t_n + \theta \Delta t, X_\theta^{\Delta t})$, we get

$$\begin{aligned} r_\theta^{\Delta t} &= \Delta t \int_0^\theta (\mathcal{L}_{t_n + \theta' \Delta t} D_x u + \mu^\top(D_x^2 u))(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) d\theta' \\ &\quad + \Delta t^{1/2} \int_0^\theta dB_{\theta'}^\top ((\sigma^\top(D_x^2 u))(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) - (\sigma^\top(D_x^2 u))(t_n, \xi)). \end{aligned}$$

Owing to the Hölder inequality, the BDG inequality and the polynomial growth conditions on the functions (because $\sigma, D_x u, D_x u_t, D_x^2 u, D_x^2 D_x u \in C_{loc,pol}^{1/2,1}$), we estimate

$$\begin{aligned} \mathbb{E}^B \|r_\theta^{\Delta t}\|^p &\leq 2^{p-1} \Delta t^p \int_0^\theta \mathbb{E}^B \|(\mathcal{L}_{t_n + \theta' \Delta t} D_x u + \mu^\top(D_x^2 u))(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t})\|^p d\theta' \\ &\quad + 2^{p-1} C_{BDG} \Delta t^{p/2} \int_0^\theta \mathbb{E}^B \|(\sigma^\top(D_x^2 u))(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) - (\sigma^\top(D_x^2 u))(t_n, \xi)\|^p d\theta'. \end{aligned}$$

Using the growth conditions from the assumptions and applying bounds (1.58) in Lemma 1.17 to $\sigma^\top(D_x^2 u) \in C_{loc,pol}^{1/2,1}$ (because σ and $D_x^2 u$ are in $C_{loc,pol}^{1/2,1}$), we obtain the announced estimate.

(b) This item follows directly from Lemma 1.17 and from item (a), using standard computations. \square

By replacing the decomposition (1.46) into the expression of $\mathcal{E}_\theta^{\Delta t}$ given in (1.41), we obtain

$$\mathcal{E}_\theta^{\Delta t} = \Delta t \mathcal{E}_\theta \left(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi) \right) + \Delta t R_\theta^{\Delta t}(t_n, \xi), \quad (1.47)$$

where

$$\mathcal{E}_\theta \left(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi) \right) = \int_0^\theta B_{\theta'}^\top G^{(n+1)}(t_n, \xi) dB_{\theta'} - F^{(n+1)}(t_n, \xi) \theta, \quad (1.48)$$

and

$$\begin{aligned} R_\theta^{\Delta t}(t_n, \xi) &= - \int_0^\theta \left(F^{(n+1)}(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) - F^{(n+1)}(t_n, \xi) \right) d\theta' \\ &\quad + \int_0^\theta B_{\theta'}^\top (\sigma^\top(D_x^2 u))(t_n, \xi) (\sigma(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) - \sigma(t_n, \xi)) dB_{\theta'} \\ &\quad + \Delta t^{-1/2} \int_0^\theta r_{\theta'}^{\Delta t} \sigma(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) dB_{\theta'} \\ &\quad + \Delta t^{1/2} \int_0^\theta B_{\theta'}^\top (\sigma^\top(D_x^2 u))(t_n, \xi) \mu(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) d\theta' \\ &\quad + \int_0^\theta r_{\theta'}^{\Delta t} \mu(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) d\theta'. \end{aligned} \quad (1.49)$$

Lemma 1.13. *Under the assumptions of Theorem 1.6, it holds*

- (a) $\sup_{\theta \in [0,1]} \mathbb{E}^B |\mathcal{E}_\theta(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi))|^2 \leq K_{n,N}(\xi)$;

$$(b) \sup_{\theta \in [0,1]} \mathbb{E}^B |R_\theta^{\Delta t}(t_n, \xi)|^2 \leq K_{n,N}(\xi) \Delta t,$$

for some constant $K_{n,N}(\xi) \in K_{pol}$.

Proof of Lemma 1.13. (a) From (1.48) we get

$$\begin{aligned} & \mathbb{E}^B \left| \mathcal{E}_\theta \left((\sigma^\top (D_x^2 u) \sigma)(t_n, \xi), F^{(n+1)}(t_n, \xi) \right) \right|^2 \\ & \leq 2 |F^{(n+1)}(t_n, \xi)|^2 \theta^2 + 2 \mathbb{E}^B \left| \int_0^\theta B_{\theta'}^\top \left((\sigma^\top (D_x^2 u) \sigma)(t_n, \xi) \right) dB_{\theta'} \right|^2, \end{aligned}$$

and we conclude owing to the Ito isometry and the growth conditions on the coefficients (because $\sigma, u, D_x u, D_x^2 u, f \in C_{loc,pol}^{1/2,1}$).

(b) From (1.49) we estimate

$$\begin{aligned} & \mathbb{E}^B |R_\theta^{\Delta t}(t_n, \xi)|^2 \leq 5 \int_0^1 \mathbb{E}^B \left| F^{(n+1)}(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) - F^{(n+1)}(t_n, \xi) \right|^2 d\theta' \\ & + 5 \mathbb{E}^B \int_0^1 \left\| B_{\theta'}^\top \left(\sigma^\top (D_x^2 u) \right)(t_n, \xi) \left(\sigma(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) - \sigma(t_n, \xi) \right) \right\|_2^2 d\theta' \\ & + 5 \Delta t^{-1} \mathbb{E}^B \int_0^1 \left\| r_{\theta'}^{\Delta t} \sigma(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) \right\|_2^2 d\theta' \\ & + 5 \Delta t \int_0^1 \mathbb{E}^B \left| B_{\theta'}^\top \left(\sigma^\top (D_x^2 u) \right)(t_n, \xi) \mu(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) \right|^2 d\theta' \\ & + 5 \int_0^1 \mathbb{E}^B \left| r_{\theta'}^{\Delta t} \mu(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) \right|^2 d\theta', \end{aligned}$$

for all $\xi \in \mathbb{R}^d$, $n \in \{0, \dots, N-1\}$ and $\theta \in [0, 1]$. Now we conclude to the inequality (b) by using that $f, u, D_x u, D_x^2 u, \sigma$ are in $C_{loc,pol}^{1/2,1}$, and by applying Lemmas 1.17 and 1.12. \square

Step 2: Expansion of $P_1^{\Delta t}(t_n, \xi)$ and $P_3^{\Delta t}(t_n, \xi)$. From $\ell''_\gamma(y) = \ell''_\gamma(y/\Delta t)$, for all $y \in \mathbb{R}$, and the expansion of $\mathcal{E}_\theta^{\Delta t}$ in (1.47), we get

$$\ell''_\gamma(\mathcal{E}_\theta^{\Delta t}) = \ell''_\gamma \left(\mathcal{E}_\theta \left(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi) \right) + R_\theta^{\Delta t}(t_n, \xi) \right). \quad (1.50)$$

By combining this with (1.43) and (1.47), we obtain

$$\begin{aligned} P_1^{\Delta t}(t_n, \xi) &= -\mathbb{E}^B \left[\int_0^1 \ell''_\gamma \left(\mathcal{E}_\theta \left(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi) \right) + R_\theta^{\Delta t}(t_n, \xi) \right) \right. \\ & \quad \left. \times F(t_n, \xi) \mathcal{E}_\theta \left(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi) \right) d\theta \right] + C_1^{\Delta t}(t_n, \xi), \end{aligned}$$

where

$$\begin{aligned} C_1^{\Delta t}(t_n, \xi) &:= -\mathbb{E}^B \left[\int_0^1 \ell''_\gamma \left(\mathcal{E}_\theta \left(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi) \right) + R_\theta^{\Delta t}(t_n, \xi) \right) \right. \\ & \quad \left. \times R_\theta^{\Delta t}(t_n, \xi) F^{(n+1)}(t_n + \theta \Delta t, X_\theta^{\Delta t}) d\theta \right] \\ & \quad - \mathbb{E}^B \left[\int_0^1 \ell''_\gamma(\mathcal{E}_\theta^{\Delta t}) \left(F^{(n+1)}(t_n + \theta \Delta t, X_\theta^{\Delta t}) - F^{(n+1)}(t_n, \xi) \right) \right] \end{aligned}$$

$$\times \mathcal{E}_\theta \left(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi) \right) d\theta \Big]. \quad (1.51)$$

The estimates of $C_1^{\Delta t}(t_n, \xi)$ and $P_3^{\Delta t}(t_n, \xi)$ are summarized in the following lemma.

Lemma 1.14. *Under the assumptions of Theorem 1.6, it holds*

- (a) $\mathbb{E}^B |C_1^{\Delta t}(t_n, \xi)| \leq K_{n,N}(\xi) \Delta t^{1/2}$,
- (b) $\mathbb{E}^B |P_3^{\Delta t}(t_n, \xi)| \leq K_{n,N}(\xi) \Delta t^{1/2}$,

for some constant $K_{n,N}(\xi) \in K_{pol}$.

Proof of Lemma 1.14. (a) From (1.51), it readily follows that

$$\begin{aligned} |C_1^{\Delta t}(t_n, \xi)| &\leq K \mathbb{E}^B \left[\int_0^1 |F^{(n+1)}(t_n + \theta \Delta t, X_\theta^{\Delta t})| |R_\theta^{\Delta t}(t_n, \xi)| d\theta \right] \\ &+ K \mathbb{E}^B \left[\int_0^1 |F^{(n+1)}(t_n + \theta \Delta t, X_\theta^{\Delta t}) - F^{(n+1)}(t_n, \xi)| \left| E_\theta \left(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi) \right) \right| d\theta \right] \end{aligned}$$

where K is an upper bound of ℓ''_γ . For the first term above, use that $F^{(n+1)}$ has polynomial growth in its arguments (because $u^{(n+1)}, D_x u^{(n+1)}, D_x^2 u^{(n+1)}, f \in C_{loc, pol}^{1/2, 1}$), and Lemma 1.13(b). For the second term, apply the Cauchy-Schwarz with Lemmas 1.17 and 1.13(a). It yields

$$|C_1^{\Delta t}(t_n, \xi)| \leq K_{n,N}(\xi) \Delta t^{1/2}$$

as announced.

(b) Similarly, from (1.45) we write

$$\begin{aligned} |P_3^{\Delta t}(t_n, \xi)| &\leq K \int_0^1 \mathbb{E}^B \left[\left| \mathcal{E}_\theta \left(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi) \right) + R_\theta^{\Delta t}(t_n, \xi) \right| \left| ((\Delta D_x u) \mu)(t_n + \theta \Delta t, X_\theta^{\Delta t}) \right| \right] d\theta \\ &\leq K \int_0^1 \sqrt{\mathbb{E}^B \left[\left| \mathcal{E}_\theta \left(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi) \right) + R_\theta^{\Delta t}(t_n, \xi) \right|^2 \right]} \\ &\quad \times \sqrt{\mathbb{E}^B \left[\left| ((\Delta D_x u) \mu)(t_n + \theta \Delta t, X_\theta^{\Delta t}) \right|^2 \right]} d\theta. \end{aligned}$$

It is now straightforward to conclude that the above is bounded by $K_{n,N}(\xi) \Delta t^{1/2}$, using Lemmas 1.17, 1.12 and 1.13. \square

Step 3: Expansion of $C_2^{\Delta t}(t_n, \xi)$. Using the expansion of $\Delta D_x u$ in (1.46), we obtain

$$\begin{aligned} &\|((\Delta D_x u) \sigma)(t_n + \theta \Delta t, X_\theta^{\Delta t})\|^2 \\ &= \Delta t \|B_\theta^\top (\sigma^\top(D_x^2 u))(t_n, \xi) \sigma(t_n + \theta \Delta t, X_\theta^{\Delta t})\|^2 + \|r_\theta^{\Delta t} \sigma(t_n + \theta \Delta t, X_\theta^{\Delta t})\|^2 \\ &\quad + 2\Delta t^{1/2} (B_\theta^\top (\sigma^\top(D_x^2 u))(t_n, \xi) \sigma(t_n + \theta \Delta t, X_\theta^{\Delta t})) \cdot (r_\theta^{\Delta t} \sigma(t_n + \theta \Delta t, X_\theta^{\Delta t})). \end{aligned}$$

Replacing the identity $\sigma(t, \zeta) = \Delta \sigma(t, \zeta) + \sigma(t_n, \xi)$ into the first term of the previous equation, we get

$$\|((\Delta D_x u) \sigma)(t_n + \theta \Delta t, X_\theta^{\Delta t})\|^2 = \Delta t \|B_\theta^\top (\sigma^\top(D_x^2 u) \sigma)(t_n, \xi)\|^2 + c_\theta^{\Delta t}(t_n, \xi), \quad (1.52)$$

where

$$\begin{aligned}
 c_\theta^{\Delta t}(t_n, \xi) &= \Delta t \|B_\theta^\top (\sigma^\top(D_x^2 u))(t_n, \xi) \Delta \sigma(t_n + \theta \Delta t, X_\theta^{\Delta t})\|^2 + \|r_\theta^{\Delta t} \sigma(t_n + \theta \Delta t, X_\theta^{\Delta t})\|^2 \\
 &+ 2\Delta t (B_\theta^\top (\sigma^\top(D_x^2 u))(t_n, \xi) \Delta \sigma(t_n + \theta \Delta t, X_\theta^{\Delta t})) \cdot (B_\theta^\top (\sigma^\top(D_x^2 u) \sigma)(t_n, \xi)) \\
 &+ 2\Delta t^{1/2} (B_\theta^\top (\sigma^\top(D_x^2 u))(t_n, \xi) \sigma(t_n + \theta \Delta t, X_\theta^{\Delta t})) \cdot (r_\theta^{\Delta t} \sigma(t_n + \theta \Delta t, X_\theta^{\Delta t})). \quad (1.53)
 \end{aligned}$$

From (1.52) and (1.50), the expression of $P_2^{\Delta t}$ in (1.44) becomes

$$\begin{aligned}
 P_2^{\Delta t}(t_n, \xi) &= \frac{1}{2} \mathbb{E}^B \left[\int_0^1 \ell''_\gamma \left(\mathcal{E}_\theta \left(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi) \right) + R_\theta^{\Delta t}(t_n, \xi) \right) \|B_\theta^\top (\sigma^\top(D_x^2 u) \sigma)(t_n, \xi)\|^2 d\theta \right] \\
 &+ C_2^{\Delta t}(t_n, \xi),
 \end{aligned}$$

where

$$C_2^{\Delta t}(t_n, \xi) := \frac{1}{2} \mathbb{E}^B \left[\int_0^1 \ell''_\gamma \left(\mathcal{E}_\theta \left(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi) \right) + R_\theta^{\Delta t}(t_n, \xi) \right) \Delta t^{-1} c_\theta^{\Delta t}(t_n, \xi) d\theta \right]. \quad (1.54)$$

The estimate of $C_2^{\Delta t}(t_n, \xi)$ is summarized in the following lemma.

Lemma 1.15. *Under the assumptions of Theorem 1.6, it holds*

$$|C_2^{\Delta t}(t_n, \xi)| \leq K_{n,N}(\xi) \Delta t^{1/2},$$

for some constant $K_{n,N}(\xi) \in K_{pol}$.

Proof of Lemma 1.15. From the expression $c_\theta^{\Delta t}(t_n, \xi)$ in (1.53), we write

$$\begin{aligned}
 \mathbb{E}^B |c_\theta^{\Delta t}(t_n, \xi)| &\leq \Delta t \mathbb{E}^B \|B_\theta^\top (\sigma^\top(D_x^2 u))(t_n, \xi) \Delta \sigma(t_n + \theta \Delta t, X_\theta^{\Delta t})\|^2 \\
 &+ \mathbb{E}^B \|r_\theta^{\Delta t} \sigma(t_n + \theta \Delta t, X_\theta^{\Delta t})\|^2 \\
 &+ 2\Delta t \sqrt{\mathbb{E}^B \|B_\theta^\top (\sigma^\top(D_x^2 u))(t_n, \xi) \Delta \sigma(t_n + \theta \Delta t, X_\theta^{\Delta t})\|^2} \\
 &\quad \times \sqrt{\mathbb{E}^B \|B_\theta^\top (\sigma^\top(D_x^2 u) \sigma)(t_n, \xi)\|^2} \\
 &+ 2\Delta t^{1/2} \sqrt{\mathbb{E}^B \|B_\theta^\top (\sigma^\top(D_x^2 u))(t_n, \xi) \sigma(t_n + \theta \Delta t, X_\theta^{\Delta t})\|^2} \\
 &\quad \times \sqrt{\mathbb{E}^B \|r_\theta^{\Delta t} \sigma(t_n + \theta \Delta t, X_\theta^{\Delta t})\|^2} \\
 &\leq K_{n,N}(\xi) \Delta t^{3/2}.
 \end{aligned}$$

Again we have used the polynomial growth condition on $\sigma, D_x^2 u$ and the local regularity condition on $\sigma \in C_{loc, pol}^{1/2,1}$ with Lemma 1.17, and Lemma 1.12 - (a). Consequently and in view of the definition (1.54) of $C_2^{\Delta t}(t_n, \xi)$, we obtain the estimate

$$|C_2^{\Delta t}(t_n, \xi)| \leq \frac{1}{2} |\ell''_\gamma|_\infty \Delta t^{-1} \sup_{\theta \in [0,1]} \mathbb{E}^B |c_\theta^{\Delta t}(t_n, \xi)|,$$

which leads to the announced result. \square

Step 4: Expansion of $P^{\Delta t}(t_n, \xi)$. From (1.42) and the previous expansions of $T_i^{\Delta t}(t_n, \xi)$ for $i = 1, 2, 3$, we deduce

$$\begin{aligned} P^{\Delta t}(t_n, \xi) &= \mathbb{E}^B \left[\int_0^1 \ell''_\gamma \left(\mathcal{E}_\theta \left(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi) \right) + R_\theta^{\Delta t}(t_n, \xi) \right) \right. \\ &\quad \left. \times \mathcal{Q}_\theta \left(G^{(n+1)}(t_n, \xi), F^{(n+1)}(t_n, \xi) \right) d\theta \right] \\ &\quad + C_1^{\Delta t}(t_n, \xi) + C_2^{\Delta t}(t_n, \xi) + P_3^{\Delta t}(t_n, \xi), \end{aligned}$$

where \mathcal{Q}_θ is defined in (1.29). Since $C_1^{\Delta t}(t_n, \xi)$, $C_2^{\Delta t}(t_n, \xi)$ and $P_3^{\Delta t}(t_n, \xi)$ satisfy $\Delta t^{1/2}$ -bounds, we get the result of Proposition 1.9. \square

1.4.5 Proof of the almost sure convergence in Subsection 1.4.3

Now we present the proof of the results related to the almost sure convergence of the remainder $\bar{R}_\theta^{\Delta t}(t_n, X_{t_n})$ in Subsection 1.4.3.

Proof of Proposition 1.11. (a) From the definition of \mathcal{E}_θ in (1.28), it follows that

$$\begin{aligned} &\left| \mathcal{E}_\theta \left(G^{(n+1)}(t_n, X_{t_n}), F^{(n+1)}(t_n, X_{t_n}) \right) - \mathcal{E}_\theta \left(\Gamma(t_{n+1}, X_{t_n}), F(t_{n+1}, X_{t_n}) \right) \right| \\ &\leq \left| f(t_n, X_{t_n}, u^{(n+1)}(t_n, X_{t_n}), D_x u^{(n+1)}(t_n, X_{t_n}), D_x^2 u^{(n+1)}(t_n, X_{t_n})) \right. \\ &\quad \left. - f(t_{n+1}, X_{t_n}, v(t_{n+1}, X_{t_n}), D_x v(t_{n+1}, X_{t_n}), D_x^2 v(t_{n+1}, X_{t_n})) \right| \\ &\quad + \left\| (\sigma^\top (D_x^2 u^{(n+1)}) \sigma)(t_n, X_{t_n}) - (\sigma^\top (D_x^2 v) \sigma)(t_{n+1}, X_{t_n}) \right\| \left\| \int_0^\theta B_{\theta'} dB_{\theta'}^\top \right\| \\ &\leq K_{n,N}(X_{t_n}) \Delta t^{1/2} \left(1 + \left\| \int_0^\theta B_{\theta'} dB_{\theta'}^\top \right\| \right), \end{aligned}$$

for some constant $K_{n,N}(X_{t_n}) \in K_{\text{pol}}$, where we have used Proposition 1.10 and the assumptions on coefficients, prices and greeks. Owing to the Burkholder-Davis-Gundy (BDG) inequalities, we conclude the proof of (a).

(b) From $\bar{C}^{\Delta t}(t_n, \xi)$ in (1.36) we get

$$\begin{aligned} &\mathbb{E} \left| \bar{C}^{\Delta t}(t_n, X_{t_n}) \right| \leq \\ &\left| \ell''_\gamma \right|_\infty \int_0^1 \mathbb{E} \left| \mathcal{Q}_\theta \left(G^{(n+1)}(t_n, X_{t_n}), F^{(n+1)}(t_n, X_{t_n}) \right) - \mathcal{Q}_\theta \left(\Gamma(t_{n+1}, X_{t_n}), F(t_{n+1}, X_{t_n}) \right) \right| d\theta. \end{aligned}$$

Considering the expression of \mathcal{Q}_θ in (1.29), we are able to apply the same arguments as for (a). Further details are left to the reader. We are done with the estimate (b).

(c) Let $p \geq 1$ and set $Z_N := \sup_{0 \leq n \leq N-1} \sup_{\theta \in [0,1]} \left| \bar{R}_\theta^{\Delta t}(t_n, X_{t_n}) \right|^p$. From the definition (1.35) of $\bar{R}_\theta^{\Delta t}(t_n, \xi)$ we write

$$\begin{aligned} \mathbb{E}[Z_N] &\leq 2^{p-1} \mathbb{E} \left[\sup_{0 \leq n \leq N-1} \sup_{\theta \in [0,1]} \left| \mathcal{E}_\theta \left(G^{(n+1)}(t_n, X_{t_n}), F^{(n+1)}(t_n, X_{t_n}) \right) \right. \right. \\ &\quad \left. \left. - \mathcal{E}_\theta \left(\Gamma(t_{n+1}, X_{t_n}), F(t_{n+1}, X_{t_n}) \right) \right|^p \right] \end{aligned}$$

$$+ 2^{p-1} \mathbb{E} \left[\sup_{0 \leq n \leq N-1} \sup_{\theta \in [0,1]} |R_\theta^{\Delta t}(t_n, X_{t_n})|^p \right] \leq K_p N \Delta t^{p/2}$$

owing to (a) and Lemma 1.16 stated below. Last, apply Lemma 1.18 to the above Z_N with $p > 4$: we are done. \square

In the proof, we have used the following result, useful to justify the almost sure. convergence to 0 of remainder terms. We postpone its proof to Appendix 1.6.1.

Lemma 1.16. *Let $R_\theta^{\Delta t}(t_n, \xi)$ be given by (1.30) and $p \geq 1$. Under the assumptions of Theorem 1.6, there exists a finite positive constant K_p depending on the coefficients $\mu, \sigma, f, u^{(n+1)}$ and its derivatives such that*

$$\mathbb{E} \left[\sup_{0 \leq n \leq N-1} \sup_{\theta \in [0,1]} |R_\theta^{\Delta t}(t_n, X_{t_n})|^p \right] \leq K_p N \Delta t^{p/2}.$$

Proof. We claim the following upper bound holds

$$\mathbb{E}^B \left[\sup_{\theta \in [0,1]} |R_\theta^{\Delta t}(t_n, \xi)|_2^p \right] \leq K_{n,N}(\xi) \Delta t^p \quad (1.55)$$

for some constant $K_{n,N}(\xi) \in K_{\text{pol}}$.

With this control at hand, we complete the proof by using the rough inequality

$$\mathbb{E} \left[\sup_{0 \leq n \leq N-1} \sup_{\theta \in [0,1]} |R_\theta^{\Delta t}(t_n, X_{t_n})|^p \right] \leq \sum_{n=0}^{N-1} \mathbb{E} \left[\sup_{\theta \in [0,1]} |R_\theta^{\Delta t}(t_n, X_{t_n})|^p \right].$$

So, it is enough to show (1.55). Regarding the control with $R_\theta^{\Delta t}(t_n, \xi)$, we follow the proof of Lemma 1.13 item (b). The adaptation is obvious since instead of taking $p = 2$, we take $p \geq 1$. Then we handle the supremum over θ inside the expectation using BDG inequalities. Other arguments are unchanged, leading to the announced estimate. We leave to the reader the details. \square

1.5 Numerical results

In this section, we compute the numerical solution to the f -PDE (1.5) in dimension 1. In Subsection 1.3.2, we have obtained a formulation for the optimal PDE nonlinearity f_γ^* (see (1.22)). First, in Section 1.5.1, we present the numerical solution for different European options. Then, in Section 1.5.2, we compute the asymptotic risk $\mathbf{R}_\gamma(v^*, f)$ for different $f \in \{f_0^*, f_{0.1}^*, f_{0.2}^*, f_{0.3}^*\}$ confirming the optimality of f_γ^* .

Finally, in Section 1.5.3, we compare the optimal f^* -PDE solution with the solution to the minimization problem (1.3). We aim to check the conjecture whether one can interchange the limit in N and the minimization over strategies in our setting. In other words, we verify the solution to the minimization problem in discrete time (see (1.3)) corresponds, as N large, to the solution to the nonlinear f^* -PDE (1.5).

1.5.1 Optimal PDE valuation of European options

Here, we assume that the hedging instrument X satisfies the SDE (1.8) with

$$\mu(t, x) = 0, \quad \sigma(t, x) = \sigma x, \quad 0 \leq t \leq T, \quad x \in \mathbb{R}.$$

Then, we show the numerical solution to the PDE (1.5) for different European options.

The forward f -PDE valuation in dimension 1. Consider the value function (t, x) as the solution to the f -PDE valuation

$$\frac{\partial \mathbf{u}}{\partial t}(t, x) = \alpha(x) \mathbf{u}_{xx}(t, x) + f(2\alpha(x) \mathbf{u}_{xx}(t, x)), \quad 0 < t \leq T, \quad x \in \mathbb{R}. \quad (1.56)$$

in a forward form, where $\alpha(x) = \sigma^2 x^2 / 2$, $x \in \mathbb{R}$, and $f : \mathbb{R} \rightarrow \mathbb{R}$ is a real-valued function to be chosen.

Seeing that (1.56) has a second-order partial differential in space and first-order in time, we require for a numerical resolution one initial and two boundary conditions. Also, European option with maturity T , defined by a payoff function $g(x)$, are used as initial time condition.

European option payoffs. Denote the mapping $x \mapsto \max(x, 0)$ by x_+ . We choose the following options:

- Call and put option with strike price K_0 ;
- Asset-or-nothing call and put option with strike price K_0 ;
- Bull-spread and bear-spread option with strike prices $K_2 > K_1$;

which is summarized below

Type	Call	Put
Vanilla	$(x - K_0)_+$	$(K_0 - x)_+$
Digital	$x \mathbf{1}_{K_0 - x < 0}$	$x \mathbf{1}_{x - K_0 < 0}$
Spread	$(x - K_1)_+ - (x - K_2)_+$	$(K_2 - x)_+ - (K_1 - x)_+$

In the following, we consider:

Set	Type	(K_0, K_1, K_2)	σ	T
A	Vanilla and digital	$(100, -, -)$	0.3	1.0
B	Spread	$(-, 90, 110)$	0.3	1.0

We examine the digital and the spread options because of the discontinuity and the change of convexity, respectively, of their payoffs. Aware these payoffs do not satisfy the assumptions of Theorem 1.6, we believe that these hypotheses are only sufficient and the previous asymptotic analysis can also be applied to those payoffs.

Numerical scheme: space discretization. Let us detail our numerical scheme. We look for a second-order approximation to the PDE (1.56) solution on a finite domain $L = [0, L_{\max}]$. Let $I \in \mathbb{N}$. Therefore, we equally discretize L in $I + 1$ points $\{x_0, x_1, \dots, x_{I-1}, x_I\}$ such that $\Delta x = L_{\max}/I$ and $x_i = i\Delta x$ for each $0 \leq i \leq I$.

Assuming that \mathbf{u} is smooth enough, we get the second-order approximation of the second derivative of \mathbf{u}

$$\frac{\mathbf{u}_{i+1}(t) - 2\mathbf{u}_i(t) + \mathbf{u}_{i-1}(t)}{\Delta x^2} = \mathbf{u}(t, x)_{xx}|_{x=x_i} + \mathbf{O}(\Delta x^2),$$

for every $1 \leq i \leq I - 1$, with $\mathbf{u}_i(t)$ denoting $\mathbf{u}(t, x_i)$. Owing to the second-order approximation, we obtain a semi-discretization from (1.56):

$$\mathbf{u}_i(t)_t = \alpha_i \frac{\mathbf{u}_{i+1}(t) - 2\mathbf{u}_i(t) + \mathbf{u}_{i-1}(t)}{\Delta x^2} + f \left(2\alpha_i \frac{\mathbf{u}_{i+1}(t) - 2\mathbf{u}_i(t) + \mathbf{u}_{i-1}(t)}{\Delta x^2} \right), \quad (1.57)$$

for every $1 \leq i \leq I - 1$, where the factor α_i is $\alpha(x)$ evaluated in each x_i .

Assuming that f in (1.56) is Lipschitz continuous, the system (1.57) is a second-order approximation of the PDE (1.56) and can be viewed in matrix form as

$$\frac{dU(t)}{dt} = AU(t) + f(2AU(t)),$$

where A is the coefficient matrix and $U(t) = (\mathbf{u}_0(t), \dots, \mathbf{u}_I(t))^T$ the discrete solution.

Besides the system (1.57), $U(t)$ also satisfies $\mathbf{u}_0(t) = b_{\min}(t)$ and $\mathbf{u}_I(t) = b_{\max}(t)$, where b_{\min} and b_{\max} represent a Dirichlet-type boundary condition imposed to the numerical solution. Therefore, the matrix A is of form

$$A_{00} = 0, \quad A_{ii-1} = \frac{\alpha_i}{\Delta x^2}, \quad A_{ii} = -\frac{2\alpha_i}{\Delta x^2}, \quad A_{ii+1} = \frac{\alpha_i}{\Delta x^2}, \quad A_{II} = 0.$$

After the space discretization, it remains a system of ordinary differential equations

$$\frac{dU(t)}{dt} = AU(t) + F(t), \quad U(0) = g, \quad F(t) := f(2AU(t)).$$

Numerical scheme: Time discretization. Now we apply a second-order method in time. Let $J \in \mathbb{N}$. Divide the time interval $[0, T]$ in J intervals with a constant time step $\Delta t = T/J$.

Denote U^j (resp. F^j) as the vector $U(t)$ (resp. $F(t)$) evaluated at $t = j\Delta t$. Due to the nonlinearity of $F(t)$ regarding $U(t)$, we use Adams-Moulton (AM) methods together with Adams-Bashforth (AB) methods to construct a Predictor-Corrector algorithm with AM and AB of the same order. Here we apply the second-order Adams-Bashforth (AB2) method to predict F^{j+1} giving $F^{j+1/2}$ and we use $F^{j+1/2}$ within the second-order Adams-Moulton (AM2) method to correct F^{j+1} :

1. With AB2, we predict (U^{j+1}, F^{j+1}) giving us:

$$(F^{j+1/2}, U^{j+1/2}) = \left(f(2AU^{j+1/2}), U^j + \Delta t \left(\frac{3}{2}(AU^j + F^j) - \frac{1}{2}(AU^{j-1} + F^{j-1}) \right) \right).$$

2. With AM2, we correct (U^{j+1}, F^{j+1}) giving us:

$$(F^{j+1}, U^{j+1}) = \left(f(2AU^{j+1}), U^j + \Delta t \left(\frac{1}{2}(AU^{j+1} + F^{j+1/2}) + \frac{1}{2}(AU^j + F^j) \right) \right).$$

Since the algorithm looks two steps back, we will need some initialization steps. Therefore, we use the AB1 (Forward Euler) and the AM1 (Backward Euler) method for the prediction and correction part, respectively,

$$U^{1/2} = U^0 + \Delta t(AU^0 + f(2U^0)), \quad U^1 = U^0 + \Delta t(AU^1 + f(2U^{1/2})),$$

Notice that the chosen f is the optimal PDE nonlinearity f^*

$$f^*(y) = c_1^* y_+ - c_2^* y_-, \quad \forall y \in \mathbb{R},$$

as given in (1.22). Moreover, in Table 1.1, we find the computed constants c_1^* and c_2^* depending on the risk parameter γ .

Parameters. Regarding the boundary conditions, we stipulate a space domain $L = [0, L_{\max}]$, where L_{\max} is supposed to be large enough. Then, we use Dirichlet boundary conditions

$$\mathbf{u}(t, 0) = b_{\min}(t) \text{ and } \mathbf{u}(t, L_{\max}) = b_{\max}(t), \quad \forall t \in [0, T].$$

We set the left and right boundary

$$b_{\min}(t) = g(0) \text{ and } b_{\max}(t) = g(L_{\max}), \quad \forall t \in [0, T].$$

Regarding the numerical solution, we fix

L_{\max}	I	J
400	200	200

In Figure 1.4, we show the vanilla option value plotted for different risk parameters γ . We depict analogous plot for digital and spread options in Figure 1.5 and 1.6, respectively. We remark the numerical solutions are increasing in function of γ . Intuitively, whenever the seller's risk aversion increases, it will be more reasonable to asking him for a higher option price. According to Proposition 1.8, we have $y \mapsto f_\gamma^*(y)$ is nonnegative for all $\gamma \in (0, 1)$. Therefore, the nonlinear source of PDE (1.56) is nonnegative whatever the sign of second derivatives. Our risk-aversion valuation adds a risk premium to the risk-neutral one whenever the underlying price varies too quickly, i.e., proportionally to the Greek Gamma.

1.5.2 Asymptotic risk: Dependence on the PDE nonlinearity

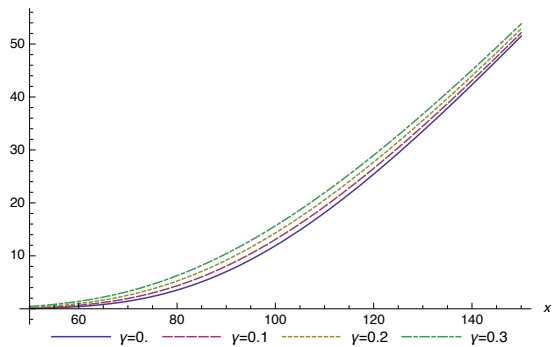
Here, we suppose the reference valuation v is given. Then, we test the asymptotic risk $\mathbf{R}_\gamma(v, f)$ (see (1.13)) for different PDE nonlinearity f .

Using the optimal f_γ^* -PDE valuation in dimension 1. Let us consider as reference valuation the function $v_\gamma^*(t, \cdot) = \mathbf{u}(T - t, \cdot)$, where \mathbf{u} is the PDE (1.56) solution using the optimal PDE nonlinearity f_γ^* (see Proposition 1.8).

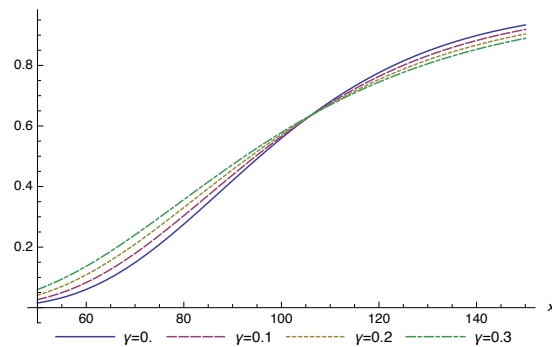
We confirm numerically the optimality of f_γ^* for the reference valuation v_γ^* by computing $\mathbf{R}_\gamma(v_\gamma^*, f_{\bar{\gamma}}^*)$ for a different risk parameter $\bar{\gamma}$. To achieve that, we approximate $\mathbf{R}_\gamma(v_\gamma^*, f_{\bar{\gamma}}^*)$ by forward Monte Carlo simulations of X . In addition, we use the numerical PDE solution to compute the partial derivatives of v_γ^* present in the expression of $\mathbf{R}_\gamma(v_\gamma^*, f_{\bar{\gamma}}^*)$ (see (1.13)). Denote its estimate by $\hat{\mathbf{R}}_{N,M}(\gamma, \bar{\gamma})$, where N is the number of time steps and M is the number of paths $\{X_{t_n}\}_{n=0}^N$.

Parameters. Set

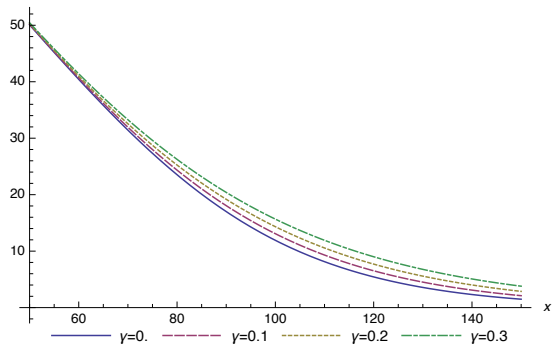
$$\sigma = 0.3, \quad N = 20, \quad M = 5 \times 10^5 \quad \text{and} \quad X_0 \in \{90, 110\}.$$



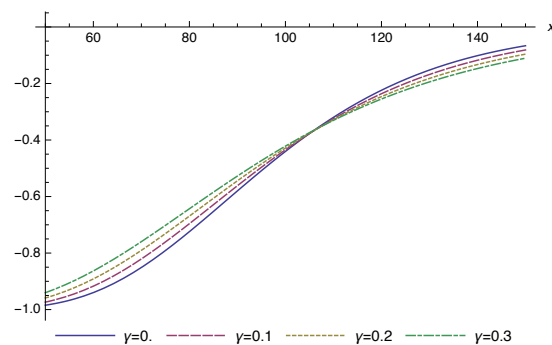
(a) Call option: value function u



(b) Call option: hedge function u_x

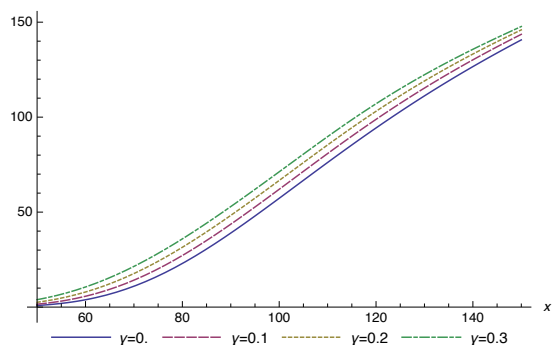


(c) Put option: value function u

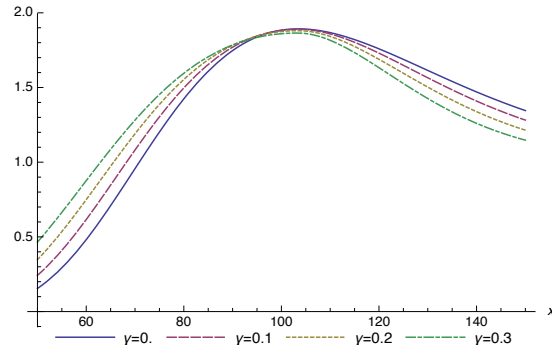


(d) Put option: hedge function u_x

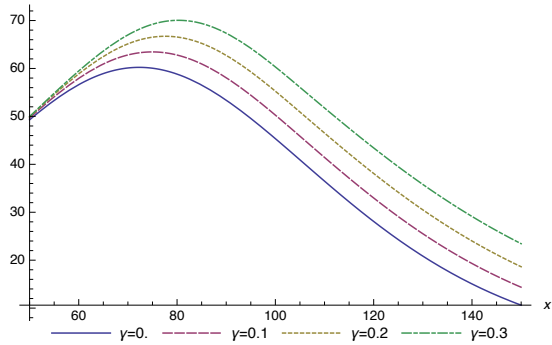
Figure 1.4 – Vanilla options: f^* -PDE solution u for different risk parameter γ .



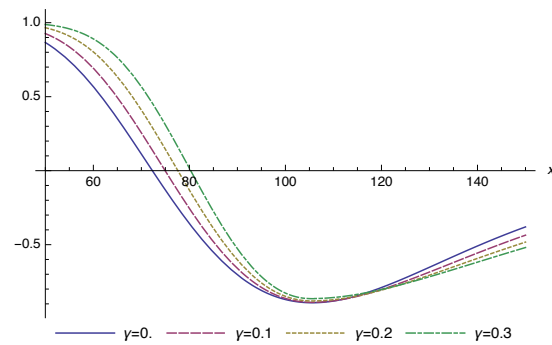
(a) Asset-or-nothing call: value function u



(b) Asset-or-nothing call: hedge function u_x



(c) Asset-or-nothing put: value function u



(d) Asset-or-nothing put: hedge function u_x

Figure 1.5 – Digital options: f^* -PDE solution u for different risk parameter γ .

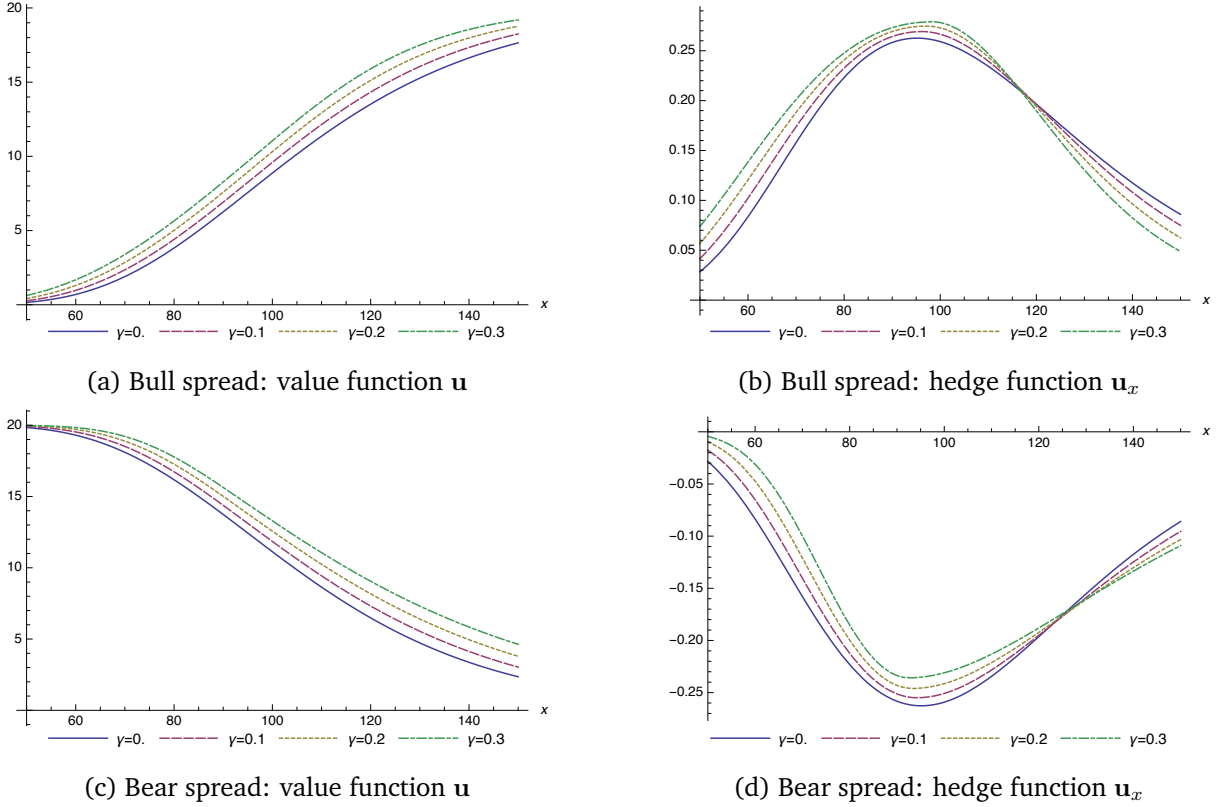


Figure 1.6 – Spread options: f^* -PDE solution u for different risk parameter γ .

The number of time steps used in the PDE resolution between each time step of the MC algorithm is 50. We study the following options:

Option	(K_0, K_1, K_2)	T
Call	$(100, -, -)$	1.0
Bear	$(-, 80, 120)$	1.0

Let $\gamma \in [0, 1)$ and $X_0 \in \mathbb{R}_+$ fixed. Thanks to Theorem 1.6 and Proposition 1.7, we expect the minimum of

$$\bar{\gamma} \mapsto \widehat{\mathbf{R}}_{N,M}(\gamma, \bar{\gamma})$$

is attained at $\bar{\gamma} = \gamma$. In Table 1.2, we compute the numerical approximation $\widehat{\mathbf{R}}_{N,M}(\gamma, \bar{\gamma})$ for all $\gamma, \bar{\gamma} \in \{0.0, 0.1, 0.2, 0.3\}$ to verify this claim.

1.5.3 Optimal PDE valuation/hedging: Comparison with the discrete-time solution

Here, we compare the optimal f_γ^* -PDE valuation/hedging rule

$$\varphi_\gamma^*(t, \cdot) = (v_\gamma^*(t, \cdot), \partial_x v_\gamma^*(t, \cdot)), \quad \forall t \in [0, T]$$

and the discrete time problem solution

$$\varphi_\gamma^N(t_n, \cdot) = (V_\gamma^N(t_n, \cdot), \vartheta_\gamma^N(t_n, \cdot)), \quad \forall n \in \{0, \dots, N\},$$

Table 1.2 – Asymptotic risk estimate $\widehat{\mathbf{R}}_{N,M}$ for $N = 20$ and $M = 5 \times 10^5$.

γ	$\bar{\gamma}$	$X_0 = 90$	$X_0 = 110$	γ	$\bar{\gamma}$	$X_0 = 90$	$X_0 = 110$
0	0	41.20±0.08	47.92±0.08	0	0	41.20±0.08	47.92±0.08
0	0.1	42.09±0.08	48.96±0.08	0	0.1	42.09±0.08	48.96±0.08
0	0.2	45.41±0.09	52.82±0.09	0	0.2	45.41±0.09	52.82±0.09
0	0.3	52.49±0.10	61.05±0.10	0	0.3	52.49±0.10	61.05±0.10
0.1	0	41.26±0.08	48.03±0.08	0.1	0	41.26±0.08	48.03±0.08
0.1	0.1	40.53±0.08	47.17±0.08	0.1	0.1	40.53±0.08	47.17±0.08
0.1	0.2	41.52±0.08	48.32±0.08	0.1	0.2	41.52±0.08	48.32±0.08
0.1	0.3	45.20±0.09	52.61±0.08	0.1	0.3	45.20±0.09	52.61±0.08
0.2	0	41.46±0.08	48.29±0.07	0.2	0	41.46±0.08	48.29±0.07
0.2	0.1	39.46±0.07	45.97±0.07	0.2	0.1	39.46±0.07	45.97±0.07
0.2	0.2	38.66±0.07	45.03±0.07	0.2	0.2	38.66±0.07	45.03±0.07
0.2	0.3	39.75±0.07	46.30±0.07	0.2	0.3	39.75±0.07	46.30±0.07
0.3	0	41.61±0.07	48.51±0.07	0.3	0	41.61±0.07	48.51±0.07
0.3	0.1	38.66±0.07	45.08±0.07	0.3	0.1	38.66±0.07	45.08±0.07
0.3	0.2	36.51±0.07	42.57±0.06	0.3	0.2	36.51±0.07	42.57±0.06
0.3	0.3	35.65±0.06	41.56±0.06	0.3	0.3	35.65±0.06	41.56±0.06

(a) $\widehat{\mathbf{R}}_{N,M}$ for a call option.

(b) $\widehat{\mathbf{R}}_{N,M}$ for a bear option.

where N is the number of hedging times. We approximate φ_γ^N by $\widehat{\varphi}_\gamma^{N,M}$ by using a Regression Monte Carlo (RMC) algorithm, where M is the number of Monte Carlo paths.

Principles of RMC methods. We present our RMC algorithm, which is a variation of the Hedged Monte Carlo algorithm (proposed in [PBS01]) with a fixed-point stage. We determine the option value by working step by step from $T = N\Delta t$ to the present $t = 0$, where Δt is the time interval. We denote the underlying asset price X at time $t_n = n\Delta t$ by X_n and the option value $V_n(X_n)$ at time t_n only depends on the current asset price X_n . We introduce the hedge $\vartheta_n(X_n)$, which is the amount of the underlying asset in the portfolio at time t_n , when the asset price is X_n .

The average risk, over all paths of the underlying process, is given by

$$\mathcal{R}_n = \langle \ell_\gamma (V_{n+1}(X_{n+1}) - V_n(X_n) - \vartheta_n(X_n)(X_{n+1} - X_n)) \rangle_M,$$

where the angled brackets $\langle \dots \rangle_M$ denote the average over the sampled asset values. The functional minimization of \mathcal{R}_n with respect to $V_n(X_n)$ and $\vartheta_n(X_n)$ gives us equations which allow us to determine the option value and hedge provided that V_{n+1} is known. We generate a set of M paths X_n^m , where n is the time index and m the path index. We decompose V_n and ϑ_n over a set of K basis functions L_k^n and C_k^n . The use of local basis function in RMC is presented in [GLW05]. Therefore, we set L_k^n and C_k^n , respectively, as a piecewise linear and constant function on each partition of real line. In addition, we use adaptive breakpoints as proposed in [BW12]:

$$V_n^K(x) := \sum_{k=1}^K a_k^n L_k^n(x), \quad \vartheta_n^K(x) := \sum_{k=1}^K b_k^n C_k^n(x).$$

In other words, we reduce the original functional optimization problem (find the functions V_n and ϑ_n) to a numerical optimization (find the coefficients a_k^n and b_k^n). We have a good approximation of the true functional solution conditionally to K be large enough. We then solve

N minimization problems backwardly in time from maturity T , where $V_N(x)$ is equal to the payoff function g . For each step n , we minimize

$$\frac{1}{M} \sum_{m=1}^M \ell_\gamma (\mathcal{E}_{n,m}^K (V_{n+1}, a^n, b^n))$$

where

$$\mathcal{E}_{n,m}^K (V, a, b) := V(X_{n+1}^m) - \sum_{k=1}^K a_k L_k^n (X_n^m) - \sum_{k=1}^K b_k C_k^n (X_n^m) (X_{n+1}^m - X_n^m).$$

Weighted Least Squares. Thanks to the choice of the risk function ℓ , we write $\ell_\gamma(y) = (y w_\gamma(y))^2$ with a weight function $w_\gamma(y) = 1 + \gamma \operatorname{sgn}(y)$. Then, for each $n \in \{N-1, \dots, 0\}$, we solve the following fixed-point problem starting from the quadratic optimal solution

$$(a^{n,0}, b^{n,0}) := \operatorname{argmin}_{(a,b)} \frac{1}{M} \sum_{m=1}^M (\mathcal{E}_{n,m}^K (V_{n+1}, a, b))^2,$$

$$(a^{n,p+1}, b^{n,p+1}) := \operatorname{argmin}_{(a,b)} \frac{1}{M} \sum_{m=1}^M (\mathcal{E}_{n,m}^K (V_{n+1}, a, b) w_\gamma (\mathcal{E}_{n,m}^K (V_{n+1}, a^{n,p}, b^{n,p})))^2,$$

for every $p \in \{0, \dots, P-1\}$, where $V_{n+1} = V_{n+1}^{K,P}$,

$$V_{n+1}^{K,P} := \sum_{k=1}^K a_k^{n+1,P} L_k^{n+1}, \quad \vartheta_{n+1}^{K,P} := \sum_{k=1}^K b_k^{n+1,P} C_k^{n+1}.$$

The least square problem with weights is solved using standard procedures. From a practical point of view, we have used a C++ library called StOpt (see the documentation in [GLW16]) to implement this previous RMC with local basis function and adaptive breakpoints. Even though we do not establish any theoretical convergence result, we know the previous algorithm is strongly related to a RMC method for computing generalized BSDEs proposed in [LGW06, GT16]. In following, we denote the optimal strategy $(V_n^{K,P}(\cdot), \vartheta_n^{K,P}(\cdot))$ as $\widehat{\varphi}_\gamma^{N,M}(t_n, \cdot)$.

Discussion. Regarding the RMC algorithm, we set

$$M = 8 \times 10^5, \quad N = 40, \quad K = 80 \quad \text{and} \quad P = 20.$$

For the underlying process, we set

$$\sigma = 0.3, \quad X_0 = 100 \quad \text{and} \quad T = 1.$$

Here, we compare the optimal valuation/hedging rule $\varphi_\gamma^*(t, \cdot)$ and the discrete-time solution $\varphi_\gamma^N(t_n, \cdot)$ for the call and bear option with

$$K_0 = 100, \quad K_1 = 80 \quad \text{and} \quad K_2 = 120.$$

Thanks to the algorithm described before, we compute the option value $\widehat{V}_\gamma^{N,M}(t_n, \cdot)$. Owing to the finite difference scheme in Subsection 1.5.1, we have the value function $v_\gamma^*(t_n, \cdot)$.

Here we consider the risk parameter γ varying in $\{0.0, 0.1, 0.2, 0.3\}$ and the following discrete

times $t_n \in \{0.1, 0.3\}$. In Figure 1.7, we present the relative error

$$x \mapsto \text{rerr}_{\gamma, t_n}^{N, M}(x) := \frac{\widehat{V}_{\gamma}^{N, M}(t_n, x)}{v_{\gamma}^*(t_n, x)} - 1$$

for a call option. We show analogous plot for a bear spread option in Figure 1.8. We observe that relative errors seem to confirm numerically the conjecture: the optimal price in discrete time for large number of hedging times coincides asymptotically with the f^* -PDE solution.

1.6 Appendix

In this appendix, we show the proof of some technical results used in previous sections.

1.6.1 Technical results about Section 1.4

Here, we gather the results related to the proof of Theorem 1.6 (see Section 1.4). First, we present the result about estimating the increment $\|X_{\theta}^{\Delta t} - x\|$.

Lemma 1.17. *Let $X_{\theta}^{\Delta t}$ be the solution of the SDE (1.23) starting from x at time t_n . It holds, for any $p \geq 1$, that*

$$\sup_{\theta \in [0, 1]} \mathbb{E}^B \|X_{\theta}^{\Delta t} - x\|^p \leq K_{n, N}(x) \Delta t^{p/2},$$

for some constant $K_{n, N}(x) \in K_{pol}$ depending on p .

Further, for any function $\phi : [0, T] \times \mathbb{R}^d \rightarrow \mathbb{R}$ in $C_{loc, pol}^{1/2, 1}$ it holds, for any $p \geq 1$, that

$$\sup_{\theta \in [0, 1]} \mathbb{E}^B |\phi(t_n + \theta \Delta t, X_{\theta}^{\Delta t}) - \phi(t_n, \xi)|^p \leq K_{n, N}(x) \Delta t^{p/2}, \quad (1.58)$$

for some constant $K_{n, N}(x) \in K_{pol}$ depending on p .

Because this is quite standard, the proof is left to the reader.

Next, the following lemma gives a sufficient condition to ensure the almost sure convergence of a sequence $(Y_N)_{N \geq 1}$. Indeed, the expectation of the general term Y_N decreasing with a speed strictly greater than 1 suffices to obtain the result.

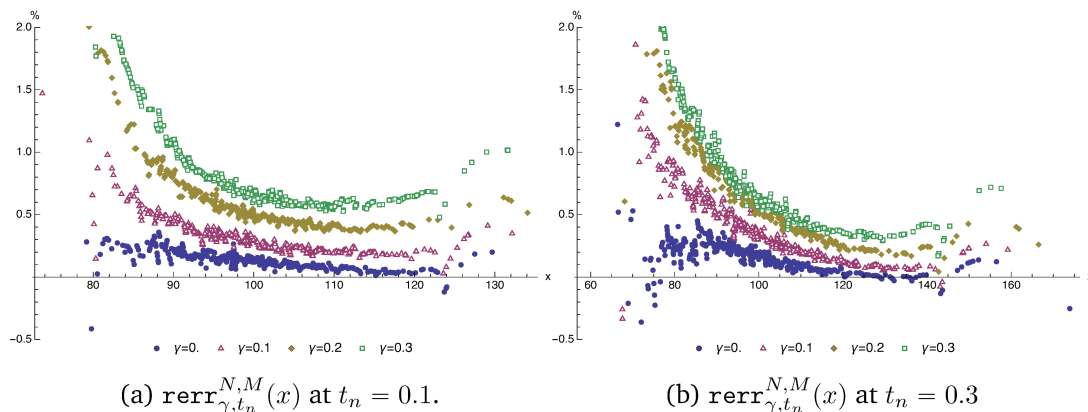


Figure 1.7 – Call option: Relative error function $\text{rerr}_{\gamma, t_n}^{N, M}$ for $M = 8 \times 10^5$, $N = 40$

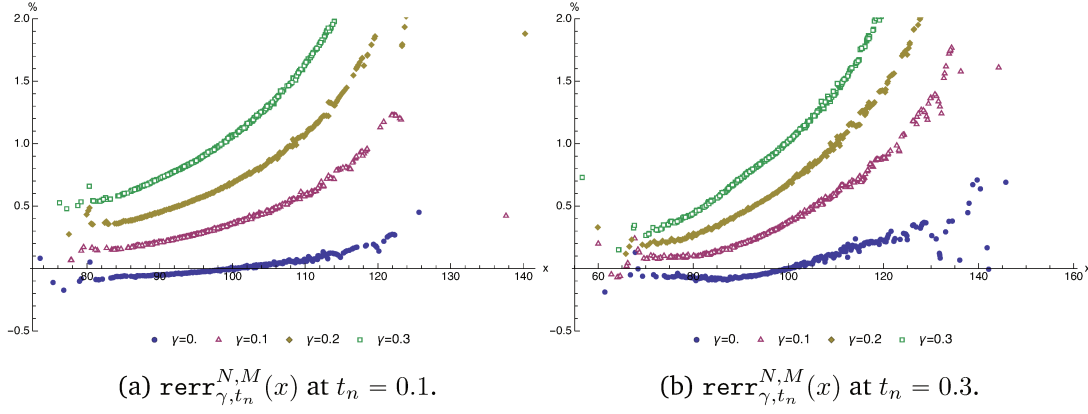


Figure 1.8 – Bear option: Relative error function $\text{rerr}_{\gamma, t_n}^{N, M}$ for $M = 8 \times 10^5$, $N = 40$.

Lemma 1.18. *Let $(Y_N)_{N \geq 1}$ be a sequence of positive real random variables. If Y_N satisfies $\mathbb{E}[Y_N] \leq c/N^{1+\delta}$ for some reals $c \geq 0$ and $\delta > 0$, then Y_N converges almost surely towards 0, i.e., $Y_N \xrightarrow[N \rightarrow \infty]{a.s.} 0$.*

Proof. The argument is quite standard, and for the sake of completeness, we give it. Summing up the expectation of Y_N , we get $\sum_{N \geq 1} \mathbb{E}[Y_N] \leq c \sum_{N \geq 1} 1/N^{1+\delta} < +\infty$ by hypothesis. Then the positive random variable $\sum_{N \geq 1} Y_N$ has a finite expectation. Owing to Fubini's theorem, this implies that $\sum_{N \geq 1} Y_N < +\infty$ with probability 1. Therefore, the general term Y_N converges almost surely towards 0. \square

This proposition is used to complete the proof of Theorem 1.6.

Proposition 1.19. *Let \mathcal{E}_θ be defined in (1.28). Let F and G be functions given by (1.34). Under the assumptions of Theorem 1.6 (especially, Assumption 1.5), the set \mathcal{A} given below*

$$\mathcal{A} = \left\{ (\omega, t, \theta) \in \Omega \times [0, T] \times [0, 1] \mid \mathcal{E}_\theta(G(t, X_t), F(t, X_t)) = 0 \right\}$$

and defined in (1.38) has measure zero w.r.t. $d\mathbb{P}^W \otimes d\mathbb{P}^B \otimes dt \otimes d\theta$.

Proof. From (1.28), we write \mathcal{E}_θ in the following form

$$\mathcal{E}_\theta(A, y) = (B_\theta^\top A B_\theta - \text{Tr}[A]\theta)/2 - y\theta, \quad \forall (\theta, A, y) \in [0, 1] \times \mathbb{S}^d \times \mathbb{R}.$$

From (1.34), we recall F and G

$$\begin{cases} F(t, x) = f(t, x, v(t, x), D_x v(t, x), D_x^2 v(t, x)), \\ G(t, x) = (\sigma^\top (D_x^2 v) \sigma)(t, x), \end{cases} \quad \forall (t, x) \in [0, 1] \times \mathbb{R}^d.$$

The announced result is equivalent to say

$$\mathbb{E}^{W \otimes B} \left[\int_0^T \int_0^1 \mathbf{1}_{\mathcal{E}_\theta(G(t, X_t), F(t, X_t))=0} d\theta dt \right] = 0.$$

Applying Fubini's theorem, the tower property of the conditional expectation and the Assumption 1.5, the previous condition is equivalent to

$$\int_0^T \int_0^1 \mathbb{E}^W \left[\mathbb{P}^B \left(\mathcal{E}_\theta(G(t, X_t), F(t, X_t)) = 0 \mid X_t \right) \mathbf{1}_{G(t, X_t) \neq 0} \right] d\theta dt = 0.$$

Claim 1.20. For $A \neq \mathbf{0} \in \mathbb{S}^d$, $y \in \mathbb{R}$ and $\theta \neq 0$, we claim that the random variable $\mathcal{E}_\theta(A, y)$ admits a density function w.r.t. the Lebesgue measure.

Therefore, for any triplet (A, y, θ) satisfying the claim, we obtain $\mathbb{P}^B(\mathcal{E}_\theta(A, y) = 0) = 0$. This proves the announced result. \square

Proof of the Claim 1.20. In view of the expression of $\mathcal{E}_\theta(A, y)$, we notice that showing $\mathcal{E}_\theta(A, y)$ admits a density function is equivalent to show that $B_\theta^\top A B_\theta / \theta$ has a density function. The latter has the same distribution as $Z^\top A Z$ where Z is a standard normal random vector. Consider the spectral decomposition of A : $A = \sum_{i=1}^I \lambda_i \mathbf{v}_i \mathbf{v}_i^\top$, where \mathbf{v}_i are orthonormal vectors and the eigenvalues λ_i are strictly positive. Since $A \neq \mathbf{0}$, $I \geq 1$. As a consequence, by setting $\bar{Z}_i = \mathbf{v}_i^\top Z$, we get $Z^\top A Z = \sum_{i=1}^I \lambda_i \bar{Z}_i^2$. The components \bar{Z}_i are independent and distributed as standard normal random variables. To sum up, we have decomposed $Z^\top A Z$ as a weighted sum of independent χ_1^2 random variables. Therefore, $Z^\top A Z$ has a probability density given by the convolution of χ_1^2 random variables. \square

1.6.2 Proof of the convergence of the sensibilities

Now we present the proof of the results related to the convergence of the sensibilities in Section 1.4.

Proof of Proposition 1.10. Because the estimates are made on $[t_n, t_{n+1}]$, we note $u^{(n+1)}$ by u . In the following, $K_{n,N}(x)$ denotes a generic constant with polynomial growth in x (defined in Section 1.4).

Initially, a strategy of proof could be to use the Feynman-Kac (FK) representations. Here, we would represent the PDE (1.5) in terms of the SDE with zero drift and diffusion coefficient σ . Although natural, this approach is cumbersome at some points, especially, when one has to deal with the derivatives of the SDE w.r.t. the initial condition. Instead, we take advantage of writing a FK formula using directly a Brownian motion. For any $\xi \in \mathbb{R}^d$, we set

$$W_t^{t_n, \xi} = \xi + W_t - W_{t_n}, \quad t_n \leq t \leq t_{n+1},$$

where W is a d -dimensional Brownian motion.

Now, the proof consists of applying Ito's formula to $u(\cdot, W_t^{t_n, \xi})$ and estimating $u(t_n, \xi) - v(t_{n+1}, \xi)$. Observe that $W_t^{t_n, \xi}$ is quite convenient for sensitivity computations because:

- the first-order derivative of $W_t^{t_n, \xi}$ w.r.t. ξ is the identity matrix \mathbf{I}_d in $\mathbb{R}^{d \times d}$,
- the second-order derivative of $W_t^{t_n, \xi}$ w.r.t. ξ is the null array $\mathbf{0}$ in $\mathbb{R}^{d \times d \times d}$.

Proof of (1.31). Applying Ito's formula to $u(\cdot, W_t^{t_n, \xi})$ and taking the expectation leads to

$$\mathbb{E}[u(t_{n+1}, W_{t_{n+1}}^{t_n, \xi})] = u(t_n, \xi) + \mathbb{E}\left[\int_{t_n}^{t_{n+1}} \left(u_t + \frac{1}{2} \text{Tr}[D_x^2 u]\right)(t, W_t^{t_n, \xi}) dt\right].$$

Since $D_x u(t, x)$ has a polynomial growth in space (because $D_x u(t, x) \in C_{\text{loc, pol}}^{1/2, 1}$), we have used that the stochastic integral is a martingale. By seeing that u is the solution of PDE (1.5) with terminal condition $u(t_{n+1}, \cdot) = v(t_{n+1}, \cdot)$, we get the FK representation

$$u(t_n, \xi) = \mathbb{E}\left[v(t_{n+1}, W_{t_{n+1}}^{t_n, \xi}) + \int_{t_n}^{t_{n+1}} \phi(t, W_t^{t_n, \xi}) dt\right], \quad (1.59)$$

where

$$\phi(t, x) := f(t, x, u(t, x), D_x u(t, x), D_x^2 u(t, x)) + \frac{1}{2} \text{Tr} [((\sigma \sigma^\top)(t, x) - \mathbf{I}_d)(D_x^2 u)(t, x)].$$

Subtracting $v(t_{n+1}, \xi)$ from (1.59) leads to

$$u(t_n, \xi) - v(t_{n+1}, \xi) = \mathbb{E} \left[v(t_{n+1}, W_{t_{n+1}}^{t_n, \xi}) - v(t_{n+1}, \xi) \right] + \int_{t_n}^{t_{n+1}} \mathbb{E} \left[\phi(t, W_t^{t_n, \xi}) \right] dt. \quad (1.60)$$

Under the assumptions, $\phi(t, x)$ and $v(t, x)$ are in $C_{\text{loc}, \text{pol}}^{1/2, 1}$. By using the estimates in Lemma 1.17 with $p = 1$, we get

$$|u(t_n, \xi) - v(t_{n+1}, \xi)| \leq \|v\|_{C_{\text{loc}, \text{pol}}^{1/2, 1}} K_{n, N}(\xi) \Delta t^{1/2} + \int_{t_n}^{t_{n+1}} K_{n, N}(\xi) dt \leq K_{n, N}(\xi) \Delta t^{1/2},$$

because the term related to the integral converges to zero at rate Δt . \square

To obtain the other announced results, we need to bound the derivatives of the following expectation $\mathbb{E}[\phi(t, W_t^{t_n, \xi})]$. The following lemma provides an estimate in the interval $[t_n, t_{n+1}]$.

Lemma 1.21. *Let $\alpha = (\alpha_1, \alpha_2, \dots, \alpha_d)$ be a d -dimensional multi-index with $|\alpha| = \sum_{i=1}^d \alpha_i$. For any function $\phi \in C_{\text{loc}, \text{pol}}^{1/2, 1}$ it holds*

$$|D_\xi^\alpha \mathbb{E}[\phi(t, W_t^{t_n, \xi})]| \leq K_{n, N}(\xi) (t - t_n)^{(1-|\alpha|)/2}, \quad t_n < t \leq t_{n+1}, \quad |\alpha| \geq 1,$$

for some constant $K_{n, N}(\xi) \in K_{\text{pol}}$.

Proof of (1.32). Next, we estimate the first derivative of $u(t_n, \xi) - v(t_{n+1}, \xi)$ w.r.t. the initial condition ξ . Differentiating (1.60) w.r.t. ξ yields

$$D_\xi u(t_n, \xi) - D_\xi v(t_{n+1}, \xi) = \mathbb{E} \left[D_x v(t_{n+1}, W_{t_{n+1}}^{t_n, \xi}) - D_\xi v(t_{n+1}, \xi) \right] + \int_{t_n}^{t_{n+1}} D_\xi \mathbb{E} \left[\phi(t, W_t^{t_n, \xi}) \right] dt. \quad (1.61)$$

From our assumptions, $\phi(t, x)$ and $D_x v(t, x)$ are in $C_{\text{loc}, \text{pol}}^{1/2, 1}$. By using the Lemma 1.21 with $|\alpha| = 1$ and the Lemma 1.17 with $p = 1$, we get

$$|D_\xi u(t_n, \xi) - D_\xi v(t_{n+1}, \xi)| \leq \|D_x v\|_{C_{\text{loc}, \text{pol}}^{1/2, 1}} K_{n, N}(\xi) \Delta t^{1/2} + \int_{t_n}^{t_{n+1}} K_{n, N}(\xi) dt \leq K_{n, N}(\xi) \Delta t^{1/2},$$

where the term related to the integral converges to zero at rate Δt . This implies the announced result. \square

Proof of (1.33). Analogously, we estimate the second derivative of $u(t_n, \xi) - v(t_{n+1}, \xi)$ by differentiating (1.61) w.r.t. ξ

$$D_\xi^2 u(t_n, \xi) - D_\xi^2 v(t_{n+1}, \xi) = \mathbb{E} \left[D_x^2 v(t_{n+1}, W_{t_{n+1}}^{t_n, \xi}) - D_\xi^2 v(t_{n+1}, \xi) \right] + \int_{t_n}^{t_{n+1}} D_\xi^2 \mathbb{E} \left[\phi(t, W_t^{t_n, \xi}) \right] dt.$$

From our assumptions, $\phi(t, x)$ and $D_x^2 v$ are in $C_{\text{loc}, \text{pol}}^{1/2, 1}$. By using the Lemma 1.21 with $|\alpha| = 2$

and the Lemma 1.17 with $p = 1$, we get

$$|D_\xi^2 u(t_n, \xi) - D_\xi^2 v(t_{n+1}, \xi)| \leq \|D_x^2 v\|_{C_{\text{loc, pol}}^{1/2, 1}} K_{n, N}(\xi) \Delta t^{1/2} + \int_{t_n}^{t_{n+1}} \frac{K_{n, N}(\xi)}{(t - t_n)^{1/2}} dt \leq K_{n, N}(\xi) \Delta t^{1/2},$$

where the integral term is of magnitude $\Delta t^{1/2}$. Therefore, we obtain the announced estimate. \square

Proof of Lemma 1.21. Let $|\alpha| \geq 1$. Differentiating $\mathbb{E}[\phi(t, W_t^{t_n, \xi})]$ w.r.t ξ yields

$$D_\xi^\alpha \mathbb{E}[\phi(t, W_t^{t_n, \xi})] = \int_{\mathbb{R}^d} \phi(t, x) D_\xi^\alpha p(t_n, \xi; t, x) dx \quad (1.62)$$

$$= \int_{\mathbb{R}^d} (\phi(t, x) - \phi(t_n, \xi)) D_\xi^\alpha p(t_n, \xi; t, x) dx, \quad (1.63)$$

where

$$p(t_n, \xi; t, x) := \frac{1}{(2\pi)^{d/2} (t - t_n)^{d/2}} \exp\left(-\frac{1}{2} \frac{\|x - \xi\|^2}{t - t_n}\right)$$

is the transition density function of $W_t^{t_n, \xi}$. To pass from (1.62) to (1.63), we used that

$$0 = D_\xi^\alpha \int_{\mathbb{R}^d} p(t_n, \xi; t, x) dx = \int_{\mathbb{R}^d} D_\xi^\alpha p(t_n, \xi; t, x) dx.$$

In the following, we use the result in [Fri08, Section 6, Chapter 9] related to the bounds for $p(t_n, \xi; t, x)$.

Proposition 1.22. Let $|\alpha| \geq 1$ and $p(t_n, \xi; t, x)$ be the transition density function of the Brownian motion $W_t^{t_n, \xi}$. There exist two positive constants c and K depending on α such that

$$|D_\xi^\alpha p(t_n, \xi; t, x)| \leq \frac{K}{(t - t_n)^{(|\alpha|+d)/2}} \exp\left(-c \frac{\|x - \xi\|^2}{t - t_n}\right), \quad \forall (\xi, t, x) \in \mathbb{R}^d \times]t_n, t_{n+1}] \times \mathbb{R}^d.$$

By using $\phi \in C_{\text{loc, pol}}^{1/2, 1}$ and the estimate in Proposition 1.22, we obtain from (1.63)

$$\begin{aligned} & \left| D_\xi^\alpha \mathbb{E}[\phi(t, W_t^{t_n, \xi})] \right| \\ & \leq K \|\phi\|_{C_{\text{loc, pol}}^{1/2, 1}} \int_{\mathbb{R}^d} (1 + \|x\|^q + \|\xi\|^q) \frac{|t - t_n|^{1/2} + \|x - \xi\|}{(t - t_n)^{d/2 + |\alpha|/2}} \exp\left(-c \frac{\|x - \xi\|^2}{t - t_n}\right) dx, \end{aligned} \quad (1.64)$$

for some real $q > 0$. In the following, we use that any exponential decrease crushes any polynomial growth: for any positive constants c and q , there exist two other positive constants $\bar{c} < c$ and \bar{K} such that, it holds

$$|(1 + |y|^q) e^{-y^2 c}| \leq \bar{K} e^{-y^2 \bar{c}}, \quad \forall y \in \mathbb{R}.$$

Therefore, we get from (1.64) that there exists a constant C such that

$$\left| D_\xi^\alpha \mathbb{E}[\phi(t, W_t^{t_n, \xi})] \right| \leq C(1 + \|\xi\|^p) (t - t_n)^{d/2} \frac{|t - t_n|^{1/2}}{(t - t_n)^{d/2 + |\alpha|/2}} = C(1 + \|\xi\|^p) (t - t_n)^{(1 - |\alpha|)/2},$$

for any $t_n < t \leq t_{n+1}$, which completes the proof.

Asymptotic asymmetric risk measure: Application to physical asset valuation

This chapter is based on the paper [AGPW18].

Abstract

Power producers are interested in valuing the future production of their power plant. They seek to reduce the risk associated with an uncertain income by trading into forward contracts. Also, their power plants generate fixed costs in the future despite the fact that they will be operating or not. After spreading those future fixed costs over the hedging period, we obtain staggered costs proportional to the time step, and we add it to the replication error. Thus, we obtain an optimal hedging strategy by minimizing the associated asymptotic risk with asymmetric risk function. We study the impact of staggered costs depending on the electricity spot price into the hedging strategy.

2.1 Introduction

This chapter provides valuation and hedging policies for future production of power plants using an *asymmetric risk* valuation. After the deregulation process of energy markets in early 1990, several spot and futures markets were created to exchange an almost non-storable commodity, the electricity. At the same time, power producers start to face the problem of evaluating their power plant production, which is dependent on the electricity spot price in the future.

To tackle this problem, there exist at least two different approaches. First, having a realistic model consistent with the electricity market data, we obtain a fair price by taking the expectation of the future income due to the production sale under the risk-neutral probability. Another approach consists of reducing the randomness of the future income by hedging with forward contracts (which are related to the electricity spot price).

Here, we consider the second approach as an application of asymmetric risk valuation and hedging to define an optimal pricing/hedging rule in an asymptotic framework (see [GPW18]). In practice, power plant owners perform discrete hedging at trading times, $t_0 = 0 < t_1 < \dots < t_N = T$ using a *hedging security* X to reduce the risk associated to a random *future income* G at

the time T . This discrete hedging produces a *residual risk*:

$$\mathcal{E} = G - \sum_{n=0}^{N-1} \vartheta_{t_n} (X_{t_{n+1}} - X_{t_n}),$$

where ϑ is the *hedging process* with ϑ_t denoting the number of shares invested in the hedging security at time t . Here, we consider the hedging security as a forward contract with delivery time T .

From a theoretical point of view, the residual risk related to the discrete hedging becomes narrower with the number N of trading times. However, in the electricity market context, the power producer wants to perform its hedging at low-frequency rate (once a week, for example). Then, we search to find the valuation/hedging rule taking into account the *local balance* \mathcal{E}_n depending on its sign through an asymmetric *risk function* ℓ :

$$\mathcal{E}_n = V_{t_{n+1}} - V_{t_n} - \vartheta_{t_n} (X_{t_{n+1}} - X_{t_n}), \quad 0 \leq n \leq N - 1,$$

where V stands for the valuation process under the replication constraint $V_T = G$. Notice those plant owners are willing to obtain a residual risk \mathcal{E} with a small standard deviation (i.e., they search for P&L histogram concentrated in a region). At the same time, they prefer the scenarios where the local balance $\mathcal{E}_n > 0$, because we are averse to negative residues.

Fixed cost accounting. In this context, we take into account the fact that a power plant generates some fixed costs, including *maintenance and investment costs, and excepting the fuel cost*, whether it is producing electricity or not. In an accounting point of view, the power producers are constrained to find today an equivalent price to the future income generated by their power plant to be written in the balance book. Actually, they do not face any financial risk due to a contingent claim. Since they can not influence the future electricity spot price S_T to obtain a higher income, they prefer to ensure a certain value for the positive amount $G = g(S_T)$ due to the sale of their future production.

For an illustration purpose, we give an estimation of those fixed costs for a nuclear plant. Following the assumptions¹ of the World Energy Outlook (WEO) 2016 in [Int16], we consider

- an *maintenance cost* of \$170 per kW (kilowatt) per yr (year);
- an *investment cost* of \$6600 per kW;
- and an *exchange rate* at 0.90 €/€ (for 2015 US dollar).

This means that a nuclear plant with 1 GW of installed capacity has cost \$6.6 billion (5.9 billion €). Also, we spend \$170 million per year (153 million €) in maintenance. After that, we spread those fixed costs over the plant lifespan (20 years). Thus, we obtain the following fixed costs of 51.40 € per MW of installed capacity and per life hour, which is detailed as:

- investment cost: 33.90 €/MWh
- maintenance cost: 17.50 €/MWh.

This means that for every 1 MW of installed capacity, one should pay each hour 33.90 € to repay the investment and 17.50 € to maintain it. Note that interest rates are not considered here and the investment cost is spread evenly over the life of the plant. Then, we look at the fixed costs generated by a power plant during a period of 24 hours: 1233.60 € for each 1 MW of installed power. Now, if we plan to hedge this production 1 year before, we must spread this cost over the 8760 previous hours and we get staggered costs of 0.14 € per MW (installed) per hour (hedging).

¹For details on the WEO assumptions used here, see the database in the *More Information on...* section and *Investment in power generation* subsection of the World Energy Model website: www.iea.org/weo/weomodel/.

In order to price the future income, they use the underlying forward market as a mark-to-market way of reducing the randomness of $g(S_T)$ by minimizing an asymmetric risk measure \mathbf{R}_N . To simplify the analysis, we consider equidistant trading times $t_n = n\Delta t$ with time step $\Delta t = T/N$. We include staggered costs proportional to the time step Δt by subtracting $c(S_{t_n})\Delta t$ from the local balance \mathcal{E}_n . Here, the staggered cost function $c : \mathbb{R} \rightarrow \mathbb{R}$, $c \in C(\mathbb{R})$ represents the dependency of the spread fixed costs on the electricity spot price S_{t_n} at the time t_n .

In this setting, we study the asymptotic of the *integrated risk* \mathbf{R}_N

$$\mathbf{R}_N = \frac{1}{\Delta t} \sum_{n=0}^{N-1} \mathbb{E}[\ell(\mathcal{E}_n^c)],$$

where

$$\mathcal{E}_n^c = V_{t_{n+1}} - V_{t_n} - \vartheta_{t_n}(X_{t_{n+1}} - X_{t_n}) - c(S_{t_n})\Delta t.$$

In the following, we give two economic points of view to justify the dependence of staggered costs c on the electricity spot price. First, when the electricity spot price is high, we use more the power plant to respond to higher demand. In indeed, a high price is generated by high demand, which corresponds to high production, meaning that we have to start more often the power plant. That is the reason why we assume a dependence of c on the spot price. For example, we can take $c = c(x)$ as an increasing convex or concave function of the spot price level x .

Second, the staggered costs related to maintenance are seasonal and higher during the refueling stop of a power plant. Also, the producer prefers to stop its power plant when the demand is low (i.e., when the spot price is low). In France, this occurs more often in summer than in winter. In this sense, the staggered costs and the spot prices are negatively correlated. Additionally, the maintenance team takes about 3 months to repair and refuel the power plant, and this refueling procedure for a nuclear plant occurs approximately every 12 months. Therefore, we consider a staggered costs c depending on the spot price level x such that $c(x)$ is decreasing, convex and takes a finite positive value at $x = 0$.

Asymptotic risk by PDE valuation. For the risk function $\ell : \mathbb{R} \rightarrow \mathbb{R}$, we take the same as in [GPW18]:

$$\ell_\gamma(y) = \frac{(1 + \gamma \operatorname{sgn}(y))^2}{2} y^2 \quad (2.1)$$

for $\gamma \in (-1, 0)$, where the sign function is defined as follows: $\operatorname{sgn}(y) := \mathbf{1}_{y>0} - \mathbf{1}_{y<0}$, $y \in \mathbb{R}$. Here, the parameter γ represents the producer's desire of reducing the future income more often than increasing it (see Figure 2.1).

In view of studying the integrated risk $\mathbf{R}_{N,\gamma}$ in the asymptotic regime $N \rightarrow +\infty$ using the *asymmetric risk* function ℓ_γ , we follow an approach by PDE (partial differential equation) valuation described below. We suppose the hedging security X is modeled by a stochastic differential equation (SDE) with diffusion coefficient σ . Then, we assume the existence of exogenous valuation process $V_t = v(t, X_t)$ defined by some function $v \in C([0, T] \times \mathbb{R})$. For example, v is a reference value function obtained by the power producer when no cost is taken into account. Now, with this exogenous reference valuation in hand, the power plant owner determines its hedge position on each interval $[t_n, t_{n+1}]$ by choosing an adapted valuation/hedging rule $(\tilde{V}_{t_n}, \tilde{\vartheta}_{t_n})$ and considering the related *conditional risk* $\mathbf{R}_{n,\gamma}$ using the risk function ℓ_γ :

$$\mathbf{R}_{n,\gamma} = \mathbb{E}\left[\ell_\gamma(\tilde{\mathcal{E}}_n^c) \middle| \mathcal{F}_{t_n}\right]. \quad (2.2)$$

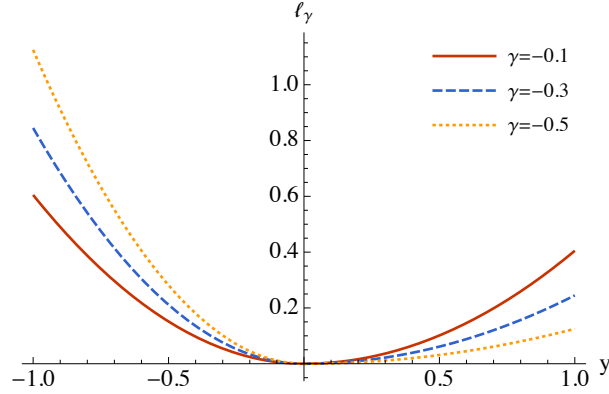


Figure 2.1 – Risk level l_γ depending on the local balance y with different parameter γ . For $\gamma \in (-1, 0)$, the risk function $l_\gamma(y)$ penalizes further the negative than positive local balances.

where

$$\tilde{\mathcal{E}}_n^c = V_{t_{n+1}} - \tilde{V}_{t_n} - \tilde{\vartheta}_{t_n}(X_{t_{n+1}} - X_{t_n}) - c(S_{t_n})\Delta t.$$

In fact, the valuation/hedging rule $(\tilde{V}_{t_n}, \tilde{\vartheta}_{t_n})$ of the power producer is given by

$$\tilde{V}_{t_n} = u^{t_{n+1}}(t_n, X_{t_n}), \quad \tilde{\vartheta}_{t_n} = u_x^{t_{n+1}}(t_n, X_{t_n}),$$

and is parametrized by a function f , possibly non-linear, through the solution $u^{t_{n+1}} : [t_n, t_{n+1}] \times \mathbb{R}$ to the f -PDE (2.3) with terminal condition $u^{t_{n+1}}(t_{n+1}, \cdot) = v(t_{n+1}, \cdot)$ at the time t_{n+1} .

Definition 2.1. Given a terminal time $\tau > 0$, a diffusion function $\sigma : [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$, $\sigma \in C([0, T] \times \mathbb{R})$, a PDE non-linearity $f : [0, T] \times \mathbb{R} \times \mathbb{R} \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$, $f \in C$, a terminal condition $v(\tau, \cdot) : \mathbb{R} \rightarrow \mathbb{R}$, $v(\tau, \cdot) \in C(\mathbb{R})$, the function $u^\tau : [0, \tau] \times \mathbb{R} \rightarrow \mathbb{R}$ is solution to the f -PDE if it satisfies

$$-u_t^\tau - f(t, x, u^\tau, u_x^\tau, u_{xx}^\tau) = \frac{1}{2}\sigma^2 u_{xx}^\tau \quad (2.3)$$

and the terminal condition

$$u^\tau(\tau, x) = v(\tau, x).$$

Following the same path as in [GPW18], we obtain the limit \mathbf{R}_γ of the *integrated conditional risk* $\mathbf{R}_{N,\gamma}$

$$\mathbf{R}_{N,\gamma} = \frac{1}{\Delta t} \sum_{n=0}^{N-1} \mathbb{E}[\mathbf{R}_{n,\gamma}],$$

as the number N of trading dates goes to infinity. After obtaining an explicit expression for \mathbf{R}_γ depending, among others, on γ , v , f , σ , we get an optimal PDE non-linearity f^* minimizing the asymptotic risk:

$$\mathbf{R}_\gamma[v, f^*] \leq \mathbf{R}_\gamma[v, f],$$

for any admissible f . Again, this optimal PDE non-linearity f^* is explicit in dimension 1, depending on γ , on the second-order derivative of v and on the staggered cost function $c(\cdot)$. In the end, the power producer will choose the following rule

$$V_{t_n} = v^*(t_n, X_{t_n}), \quad \vartheta_{t_n} = v_x^*(t_n, X_{t_n}),$$

where $v^* : [0, T] \times \mathbb{R}$ is the solution to the f^* -PDE with terminal condition satisfying $v^*(T, x) = g(x)$ at time T . It is a consistent way of reducing the randomness of the future income $G = g(S_T)$.

For the numerical experiments, we take staggered costs depending on the level of the electricity

spot price through convex-concave functions. As an example, we compute the numerical solution to the f^* -PDE for a call payoff given a function $c(x)$.

This chapter is structured as follows. In Section 2.2.1, we present the stochastic setting and define the problem. In Subsection 2.2.2, we gather the assumptions used here. Subsection 2.2.3 is reserved to the main result: the asymptotic risk with staggered fixed costs, and Section 2.3 summarizes important points of the proof. Finally, two different staggered cost models are provided in Section 2.4 and most numerical results are discussed in Section 2.5.

2.2 Setting

2.2.1 An asymmetric risk valuation with cost management

Here, we study the conditional risk $\mathbf{R}_{n,\gamma}$ using the risk function ℓ_γ in the case of staggered costs proportional to the time step. Analogous to [GPW18], we say ϕ satisfies a *local regularity condition in time and space* if, for some real $\alpha > 0$, the coefficient

$$\|\phi\|_{C_{\text{loc,pol}}^{1/2,1}} := \sup_{t \neq t' \in [0, T]} \sup_{x \neq x' \in \mathbb{R}} \frac{\|\phi(t, x) - \phi(t', x')\|}{(|t - t'|^{1/2} + |x - x'|)(1 + |x|^\alpha + |x'|^\alpha)}$$

is finite, then ϕ is said to be in $C_{\text{loc,pol}}^{1/2,1}$. Also, the set $C^{1,2}([0, T] \times \mathbb{R})$ denotes the set of functions $\phi : [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$ such that the partial derivatives $\phi_t, \phi_x, \phi_{xx}$ exist and are continuous.

Probabilistic model. We consider a fixed horizon time $T > 0$. Let $W = \{W_t : t \in [0, T]\}$ be a standard Brownian motion on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ endowed with the *natural filtration* $\mathbb{F} = \{\mathcal{F}_t : 0 \leq t \leq T\}$ generated by W up to time T . Moreover, we take a collection of equidistant trading times $\{0 = t_0 < t_1 < t_2 < \dots < t_N = T\}$ on the interval $[0, T]$ such that the time step is $\Delta t_n := t_{n+1} - t_n = T/N$, also denoted by Δt , for some $N \in \mathbb{N}$.

For simplicity, we assume that the interest rate is zero and that exist only one risk-neutral measure given by the *probability measure* \mathbb{P} . We consider the hedging security X as the forward contract $F(\cdot, T)$ with delivery time T . We model $F(\cdot, T) = \{F(t, T) : t \in [0, T]\}$ according to the dynamics given by the (one-factor model) stochastic differential equation:

$$dF(t, T) = \bar{\sigma}(t, T)F(t, T) dW_t, \quad 0 \leq t \leq T, \quad (2.4)$$

where the volatility function $t \in [0, T] \mapsto \bar{\sigma}(t, T)$ is continuous taking values in \mathbb{R}_+^* and satisfying $\bar{\sigma}(T, T) = \bar{\sigma}_0 > 0$.

For example, we could take $\bar{\sigma}(t, T) = \bar{\sigma}_0 e^{-a_0(T-t)}$ for some positive constant a_0 . The map $T \mapsto F(0, T)$ represents the initial forward curve. To be consistent with the definition of forward contracts, we take the electricity spot price S_t at the time t as being the forward price with immediate delivery, i.e., $S_t = F(t, t)$, for all $t \in [0, T]$.

For a *future income function* $g : \mathbb{R} \rightarrow \mathbb{R}$, $g \in C(\mathbb{R})$, we assume the existence of a *reference price* $v : [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$, $v \in C([0, T] \times \mathbb{R})$, satisfying a *compatibility condition* $v(T, \cdot) = g(\cdot)$. Parameterizing valuation/hedging rule of the power producer, we have a function $f : [0, T] \times \mathbb{R} \times \mathbb{R} \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ used as the source of non-linearity to the PDE (2.3). In fact, for each hedging time t_n , we associate a function $u^{t_{n+1}} : [t_n, t_{n+1}] \times \mathbb{R} \rightarrow \mathbb{R}$, the solution to

$$-u_t^{t_{n+1}} - f(t, x, u^{t_{n+1}}, u_x^{t_{n+1}}, u_{xx}^{t_{n+1}}) = \frac{1}{2} |\bar{\sigma}(t, T)x|^2 u_{xx}^{t_{n+1}}, \quad (2.5)$$

for all $(t, x) \in [0, t_{n+1}[\times \mathbb{R}$, with terminal condition $u^{t_{n+1}}(t_{n+1}, \cdot) = v(t_{n+1}, \cdot)$ at the next hedging time t_{n+1} . In this context, after replacing the valuation/hedging rule

$$\tilde{V}_{t_n} = u^{t_{n+1}}(t_n, X_{t_n}), \quad \tilde{\vartheta}_{t_n} = u_{xx}^{t_{n+1}}(t_n, X_{t_n}),$$

in (2.2), we obtain a conditional risk $\mathbf{R}_{n,\gamma}$ as follows:

$$\begin{aligned} \mathbf{R}_{n,\gamma} = \mathbb{E} \left[\ell_\gamma \left(u^{t_{n+1}}(t_{n+1}, X_{t_{n+1}}) - u^{t_{n+1}}(t_n, X_{t_n}) \right. \right. \\ \left. \left. - u_x^{t_{n+1}}(t_n, X_{t_n})(X_{t_{n+1}} - X_{t_n}) \right. \right. \\ \left. \left. - c(X_{t_n})\Delta t \right) \middle| \mathcal{F}_{t_n} \right], \end{aligned} \quad (2.6)$$

Thanks to the relation between the electricity spot price S and the T -forward contract X , for all $t \in [0, T]$, given an initial forward curve $t \mapsto F(0, t)$, for all $t \in [0, T]$, we have replace S_{t_n} by X_{t_n} in (2.2), for some continuous function $c : \mathbb{R} \rightarrow \mathbb{R}$.

Problem 2.2. Given a finite horizon time $T > 0$, a volatility function $\bar{\sigma}(\cdot, T) : [0, T] \rightarrow \mathbb{R}$, $\bar{\sigma}(\cdot, T) \in C([0, T])$, a reference price $v : [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$, $v \in C([0, T] \times \mathbb{R})$ to the future income function $g : \mathbb{R} \rightarrow \mathbb{R}$, a PDE non-linearity $f : [0, T] \times \mathbb{R} \times \mathbb{R} \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$, an equidistant partition $\{t_n = n\Delta t : n \in \{0, \dots, N-1\}\}$ of the trading interval $[0, T]$ with time step $\Delta t = T/N$, the solution $u^{t_{n+1}} : [t_n, t_{n+1}] \times \mathbb{R} \rightarrow \mathbb{R}$ to the f -PDE (2.5) with terminal condition $u^{t_{n+1}}(t_{n+1}, \cdot) = v(t_{n+1}, \cdot)$ for every $n \in \{0, \dots, N-1\}$ and a cost function $c : \mathbb{R} \rightarrow \mathbb{R}$, $c \in C(\mathbb{R})$, study the asymptotic regime of the integrated conditional risk $\mathbf{R}_{N,\gamma}$

$$\mathbf{R}_{N,\gamma}[v, f] = \frac{1}{\Delta t} \sum_{n=0}^{N-1} \mathbb{E}[\mathbf{R}_{n,\gamma}], \quad (2.7)$$

as the number N of trading dates goes to infinity, where the conditional risk $\mathbf{R}_{n,\gamma}$ is given in (2.6), the risk function ℓ_γ in (2.1) and the diffusion process $X = F(\cdot, T)$ in (2.4).

2.2.2 Main assumptions

In order to analyze the asymptotics of $\mathbf{R}_{N,\gamma}$ in (2.7) for a given reference valuation v and PDE non-linearity f , we state the following assumptions.

Assumption 2.3. *The volatility function $t \mapsto \bar{\sigma}(t, T)$ is continuously differentiable.*

Remark 2.4. Following the preceding assumption, we have that the map $t \mapsto \bar{\sigma}(t, T)$ is globally bounded, and that its derivative $t \mapsto \bar{\sigma}_t(t, T)$ is locally bounded. Therefore, the diffusion function $(t, x) \mapsto \sigma(t, x) = \bar{\sigma}(t, T)x$ is Lipschitz in space and time, and it has linear growth in space uniformly in time, which is in $C_{loc, pol}^{1/2, 1}$.

Assumption 2.5. *The reference price $v : [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$, its space derivatives $v_x, v_{xx} : [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$ and the PDE non-linearity $f : [0, T] \times \mathbb{R} \times \mathbb{R} \times \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ are in $C_{loc, pol}^{1/2, 1}$. Additionally, the cost function $c : \mathbb{R} \rightarrow \mathbb{R}$ is locally Lipschitz with polynomial growth.*

Remark 2.6. The previous assumption is equivalent to Assumptions 1.2-1.3. Moreover, it implies that the map

$$(t, x) \mapsto J_2(v, (t, x)) = f(t, x, v(t, x), v_x(t, x), v_{xx}(t, x)) + c(x)$$

is in $C_{loc, pol}^{1/2, 1}$. Now, any estimation or convergence result valid for

$$(t, x) \mapsto J_0(v, (t, x)) = f(t, x, v(t, x), v_x(t, x), v_{xx}(t, x)),$$

which is $C_{\text{loc,pol}}^{1/2,1}$, is also valid for $J_2(v, \cdot)$.

Assumption 2.7. Let $N \geq 1$. For every $n \in \{0, \dots, N-1\}$, there is a unique classical solution $u^{t_{n+1}} : [0, t_{n+1}] \times \mathbb{R} \rightarrow \mathbb{R}$ to the PDE (2.5) with the terminal condition $u^{t_{n+1}}(t_{n+1}, \cdot) = v(t_{n+1}, \cdot)$ at the time t_{n+1} . Moreover,

$$u_t^{t_{n+1}}, u_x^{t_{n+1}}, u_{xx}^{t_{n+1}}, u_{tx}^{t_{n+1}}, u_{xxx}^{t_{n+1}}$$

exist and are in $C_{\text{loc,pol}}^{1/2,1}$.

Assumption 2.8. The volatility function $\bar{\sigma}(\cdot, T) : [0, T] \rightarrow \mathbb{R}$, the diffusion process $X = F(\cdot, T)$ in (2.4), and the reference price $v : [0, T] \times \mathbb{R} \rightarrow \mathbb{R}$, $v(t, \cdot) \in C^2(\mathbb{R})$, satisfy a non-degeneracy condition:

$$|\bar{\sigma}(t, T)X_t|^2 v_{xx}(t, X_t) \neq 0,$$

$dt \otimes d\mathbb{P}$ -almost everywhere.

2.2.3 Main results

For stating the corresponding asymptotic result below, we need to introduce an extra Brownian motion B , independent of W . Also, the expectation w.r.t. the distribution of B , or W , or both, is denoted by \mathbb{E}^B , or \mathbb{E}^W , or $\mathbb{E}^{W \otimes B}$.

2.2.4 Asymptotic risk with cost management

Theorem 2.9. Consider the setting of Problem 2.2 and suppose that Assumptions 2.3-2.8 hold. Let $B = \{B_\theta : \theta \in [0, 1]\}$ be another standard Brownian motion independent from W . Then, the limit of $\mathbf{R}_{N,\gamma}[v, f]$ as $N \rightarrow \infty$ exists and is given by

$$\begin{aligned} \mathbf{R}_\gamma[v, f] = & \mathbb{E} \left[\int_0^T \int_0^1 \ell_\gamma'' \left(J_1(v, (t, X_t)) \frac{B_\theta^2 - \theta}{2} - J_2(v, (t, X_t)) \theta \right) \right. \\ & \times \left(|J_1(v, (t, X_t))|^2 \frac{B_\theta^2}{2} + |J_2(v, (t, X_t))|^2 \theta \right. \\ & \left. \left. - J_1(v, (t, X_t)) J_2(v, (t, X_t)) \frac{B_\theta^2 - \theta}{2} \right) d\theta dt \right], \end{aligned} \quad (2.8)$$

where the functionals $J_1(v, (t, x))$ and $J_2(v, (t, x))$ are given by

$$\begin{cases} J_1(v, (t, x)) = |\bar{\sigma}(t, T)x|^2 v_{xx}(t, x) \\ J_2(v, (t, x)) = f(t, x, v(t, x), v_x(t, x), v_{xx}(t, x)) + c(x). \end{cases} \quad (2.9)$$

Notice that Theorem 2.9 is a one-dimension version of Theorem 1.6 (Chapter 1) in the presence of staggered costs proportional to the time step. The proof relies on the same techniques as presented in Section 1.4 (Chapter 1) and is postponed to Section 2.3.

2.2.5 Optimal PDE non-linearity

Here, we obtain a quasi-explicit formulation for the optimal PDE non-linearity following the same approach as [GPW18]. First, we write the asymptotic risk in (2.8) as a functional

$$\mathbf{R}_\gamma[v, f] = \mathbb{E} \left[\int_0^T \mathbf{L}_\gamma(J_1(v, (t, X_t)), J_2(v, (t, X_t))) dt \right],$$

in terms of $\mathbf{L}_\gamma : \mathbb{R} \times \mathbb{R} \rightarrow \mathbb{R}$ given by

$$\begin{aligned} \mathbf{L}_\gamma(x_1, x_2) = \mathbb{E}^B \left[\int_0^1 \ell''_\gamma \left(x_1 \frac{B_\theta^2 - \theta}{2} - x_2 \theta \right) \right. \\ \left. \times \left(|x_1|^2 \frac{B_\theta^2}{2} + |x_2|^2 \theta - x_1 x_2 \frac{B_\theta^2 - \theta}{2} \right) d\theta \right], \end{aligned}$$

for all $x_1, x_2 \in \mathbb{R} \times \mathbb{R}$.

Second, we aimed to prove the existence of minimizers to the variational problem

$$\min_{f \in \Omega_f} \mathbf{R}_\gamma[v, f], \quad (2.10)$$

for all $v \in \Omega_v$. The sets Ω_f and Ω_v are the set of all functions satisfying the Assumption 2.5. Observe that $f^*(t, x, y_u, y_{u_x}, y_{u_{xx}})$ defined by

$$f^*(t, x, y_u, y_{u_x}, y_{u_{xx}}) := \arg \min_{x_2 \in \mathbb{R}} \mathbf{L}_\gamma(|\bar{\sigma}(t, T) x|^2 y_{u_{xx}}, x_2) - c(x),$$

for any fixed $(t, x, y_u, y_{u_x}, y_{u_{xx}}) \in [0, T] \times \mathbb{R} \times \mathbb{R} \times \mathbb{R} \times \mathbb{R}$, is also a minimizer to (2.10) under the condition to be in Ω_f . We integrate and take the expectation on both sides of

$$\mathbf{L}_\gamma(J_1(v, (t, X_t)), J_2^*(v, (t, X_t))) \leq \mathbf{L}_\gamma(J_1(v, (t, X_t)), J_2(v, (t, X_t))),$$

where

$$J_2^*(v, (t, x)) = f^*(t, x, v(t, x), v_x(t, x), v_{xx}(t, x)) + c(x).$$

Consequently, we seek a minimizer to the map $x_2 \mapsto \mathbf{L}_\gamma(x_1, x_2)$ for a given real x_1 . By denoting $\chi^* \in \mathbb{R}$ as a global minimizer of $\min_{x_2 \in \mathbb{R}} \mathbf{L}_\gamma(x_1, x_2)$, we finally showed (see Proposition 1.8) the existence and uniqueness χ^* , written in the following form

$$\chi^*(x_1) = a_+^*(x_1)^+ - a_-^*(x_1)^-, \quad (2.11)$$

for all $x_1 \in \mathbb{R}$, for some real a_+^* and a_-^* (depending on γ). Denoting by $\Phi_{\mathcal{N}}$ the cumulative distribution function of the standard normal distribution and $\phi_{\mathcal{N}} = \Phi'_{\mathcal{N}}$ its density, those constants a_+^* and a_-^* are respectively the unique solutions to

$$\begin{aligned} (1 + \gamma^2) a_+^* + \gamma T(a_+^*) &= 0 \\ (1 + \gamma^2) a_-^* - \gamma T(a_-^*) &= 0, \end{aligned} \quad (2.12)$$

with

$$\begin{aligned}
 T(a) &= 2a\mathbf{1}_{2a+1 \leq 0} + 2a\mathbf{1}_{2a+1 > 0} \\
 &\quad + 8a \Phi_{\mathcal{N}}(-\sqrt{2a+1}) \mathbf{1}_{2a+1 > 0} \\
 &\quad - 4\sqrt{2a+1} \phi_{\mathcal{N}}(\sqrt{2a+1}) \mathbf{1}_{2a+1 > 0}.
 \end{aligned}$$

Notice that Proposition 1.8 stands for γ in $[0, 1)$, but it is also for γ taking negatives values strictly greater than 1.

Because χ^* is unique, the optimal PDE non-linearity f^* is indeed given by

$$\begin{aligned}
 f^*(t, x, y_u, y_{u_x}, y_{u_{xx}}) &= \chi^*(|\bar{\sigma}(t, T) x|^2 y_{u_{xx}}) - c(x) \\
 &= a_+^* \left(|\bar{\sigma}(t, T) x|^2 y_{u_{xx}} \right)^+ - a_-^* \left(|\bar{\sigma}(t, T) x|^2 y_{u_{xx}} \right)^- - c(x),
 \end{aligned} \tag{2.13}$$

for any $t, x, y_u, y_{u_x}, y_{u_{xx}} \in [0, T] \times \mathbb{R} \times \mathbb{R} \times \mathbb{R} \times \mathbb{R}$.

For numerical purposes, we present, in Table 2.1, the approximate values of a_+^* and a_-^* defined in (2.12) calculated by a root finding algorithm. So we depict, in Figure 2.2, the global minimizer χ^* of $\min_{x_2 \in \mathbb{R}} \mathbf{L}_{\gamma}(x_1, x_2)$ in function of x_1 .

Table 2.1 – Optimal Positive/Negative-Part Constants

Parameter γ	PP constant a_+^*	NP constant a_-^*
-0.1	-0.0901	0.1044
-0.2	-0.1684	0.2262
-0.3	-0.2366	0.3702
-0.4	-0.2960	0.5434
-0.5	-0.3476	0.7567

Note. PP = Positive-Part. NP = Negative Part. Values obtained through the characteristic equation using the Mathematica FindRoot function.

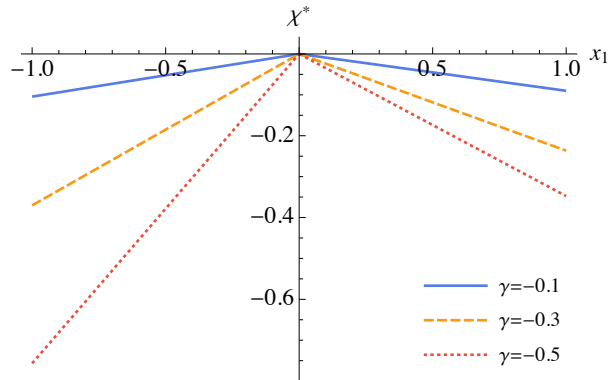


Figure 2.2 – Global minimizer χ^* depending on x_1 with different parameter γ . For $\gamma \in (-1, 0)$, the optimal function $\chi^*(x_1)$ is negative and piecewise linear, with a positive (resp. negative) slope for the strictly negative (resp. positive) values of x_1 .

2.3 Existence of the asymptotic risk.

Here, we review some important aspects of the computation of the asymptotic risk. First of all, we consider a time-space rescaling of the hedging security dynamics in order to study the conditional risk $\mathbf{R}_{n,\gamma}$ on each interval $[t_n, t_{n+1}]$. Then, we obtain an expansion of $\mathbf{R}_{n,\gamma}$ in terms of the time step Δt . With expansion in hands, we aggregate the expectation of those conditional $\mathbf{R}_{n,\gamma}$ and take the limit, after dividing by Δt .

As explained in Section 2.2.1, we deal with one hedging instrument X , which is the forward contract with delivery time T , satisfying the following SDE

$$dX_t = \bar{\sigma}(t, T)X_t dW_t, \quad 0 \leq t \leq T, \quad (2.14)$$

where W is a standard Brownian motion.

Rescaling and conditioning aspect. In view of the small-time approximations, we define a version $X_\theta^{\Delta t}$ of the solution X of the SDE (2.14)

$$X_\theta^{\Delta t} = x + \Delta t^{1/2} \int_0^\theta \bar{\sigma}(t_n + \theta' \Delta t, T) X_{\theta'}^{\Delta t} dB_{\theta'}, \quad 0 \leq \theta \leq 1, \quad (2.15)$$

where $B = \{B_\theta : \theta \in [0, 1]\}$ is an extra standard Brownian motion independent from W . Depending on the trading dates, $X^{\Delta t}$ is also dependent on the number N of time steps. Also, it was supposed the original probability space $(\Omega, \mathcal{F}, \mathbb{P})$ is large enough to contain an additional Brownian motion B .

Stochastic expansion aspect. In view of expanding the conditional risk $\mathbf{R}_{n,\gamma}$

$$\mathbb{E} \left[\ell_\gamma \left(u^{t_{n+1}}(t_{n+1}, X_{t_{n+1}}) - u^{t_{n+1}}(t_n, X_{t_n}) \right. \right. \\ \left. \left. - u_x^{t_{n+1}}(t_n, X_{t_n})(X_{t_{n+1}} - X_{t_n}) - c(X_{t_n})\Delta t \right) \middle| \mathcal{F}_{t_n} \right],$$

for any $n \in \{0, \dots, N-1\}$, we consider the stochastic process $\mathcal{E}_\theta^{\Delta t}$ defined by

$$\mathcal{E}_\theta^{\Delta t} = u^{t_{n+1}}(t_n + \theta \Delta t, X_\theta^{\Delta t}) - u^{t_{n+1}}(t_n, x) \\ - u_x^{t_{n+1}}(t_n, x)(X_\theta^{\Delta t} - x) - c(x)\theta \Delta t, \quad (2.16)$$

where $X_\theta^{\Delta t}$ is the solution to the SDE (2.15) starting from $x \in \mathbb{R}$. Using the fact that the processes $\{X_{t_n + \theta \Delta t}^{t_n, x} : \theta \in [0, 1]\}$ and $\{X_\theta^{\Delta t} : \theta \in [0, 1]\}$ have the same distribution, we rewrite $\mathbf{R}_{n,\gamma}$ as a continuous function in terms of $X^{t_n, x}$ and $X^{\Delta t}$ at $\theta = 0$, respectively. Indeed, by setting

$$P^{\Delta t}(t_n, x) = \frac{1}{\Delta t^2} \mathbb{E}[\ell_\gamma(\mathcal{E}_1^{\Delta t})] \\ = \frac{1}{\Delta t^2} \mathbb{E} \left[\ell_\gamma \left(u^{t_{n+1}}(t_{n+1}, X_1^{\Delta t}) - u^{t_{n+1}}(t_n, x) \right. \right. \\ \left. \left. - u_x^{t_{n+1}}(t_n, x)(X_1^{\Delta t} - x) - c(x)\Delta t \right) \right], \quad (2.17)$$

we obtain

$$\mathbf{R}_{n,\gamma} = \Delta t^2 P^{\Delta t}(t_n, X_{t_n}) = \mathbb{E}[\ell_\gamma(\mathcal{E}_1^{\Delta t})] \Big|_{x=X_{t_n}}. \quad (2.18)$$

In the following, we consider several constants $K_{n,N}(x)$ depending polynomially on x , uniformly in the interval $[t_n, t_{n+1}]$ and in the number N of time steps. Whenever we note $K_{n,N}(x) \in K_{\text{pol}}$, we mean that, for some real $\alpha > 0$, it holds

$$\sup_{N \geq 1} \sup_{0 \leq n \leq N-1} \sup_{x \in \mathbb{R}} \frac{|K_{n,N}(x)|}{1 + |x|^\alpha} < +\infty.$$

In the following proposition, we state the stochastic expansion of $P^{\Delta t}$ at point (t_n, x) in terms of the time step Δt in the presence of a cost function $c(x)$ proportional to Δt .

Proposition 2.10. *Consider the setting of Problem 2.2 and suppose that assumptions of Theorem 2.9 hold. Let $t_n = n\Delta t$, $x \in \mathbb{R}$, $X_\theta^{\Delta t}$ given in (2.15), $\mathcal{E}_\theta^{\Delta t}$ given in (2.16) and $P^{\Delta t}(t_n, x)$ be defined in (2.17). Then, it holds*

$$\begin{aligned} P^{\Delta t}(t_n, x) = & \mathbb{E}^B \left[\int_0^1 \ell_\gamma'' \left(J_1(u^{t_{n+1}}, (t_n, x)) \frac{B_\theta^2 - \theta}{2} - J_2(u^{t_{n+1}}, (t_n, x))\theta + R_\theta^{\Delta t}(t_n, x) \right) \right. \\ & \times \left(|J_1(u^{t_{n+1}}, (t_n, x))|^2 \frac{B_\theta^2}{2} + |J_2(u^{t_{n+1}}, (t_n, x))|^2 \theta \right. \\ & \left. \left. - J_1(u^{t_{n+1}}, (t_n, x)) J_2(u^{t_{n+1}}, (t_n, x)) \frac{B_\theta^2 - \theta}{2} \right) d\theta \right] + K_{n,N}(x) \Delta t^{1/2}, \end{aligned}$$

where

$$R_\theta^{\Delta t}(t_n, x) = \frac{\mathcal{E}_\theta^{\Delta t}}{\Delta t} - J_1(u^{t_{n+1}}, (t_n, x)) \frac{B_\theta^2 - \theta}{2} + J_2(u^{t_{n+1}}, (t_n, x))\theta,$$

for some constant $K_{n,N}(x) \in K_{\text{pol}}$.

Proof. Following the same argument as in Proposition 1.9, we use Ito's lemma in $u^{t_{n+1}}$ and its gradient $u_x^{t_{n+1}}$ to obtain the dynamics of $\mathcal{E}_\theta^{\Delta t}$. Then, we apply to the Ito-Tanaka formula to $\ell_\gamma(\mathcal{E}_\theta^{\Delta t})$ in order to obtain the announced result.

Indeed, replacing $X_\theta^{\Delta t}$ in (2.15) into $\mathcal{E}_\theta^{\Delta t}$ in (2.16) leads to

$$\begin{aligned} \mathcal{E}_\theta^{\Delta t} = & u^{t_{n+1}}(t_n + \theta\Delta t, X_\theta^{\Delta t}) - u^{t_{n+1}}(t_n, x) - c(x)\theta\Delta t \\ & - \Delta t^{1/2} \int_0^\theta u_x^{t_{n+1}}(t_n, x) \bar{\sigma}(t_n + \theta'\Delta t, T) X_{\theta'}^{\Delta t} dB_{\theta'}. \end{aligned}$$

Applying Ito's formula to $u^{t_{n+1}}(t_n + \theta\Delta t, X_{\theta}^{\Delta t})$, we obtain

$$\begin{aligned}
 \mathcal{E}_{\theta}^{\Delta t} &= -\Delta t \int_0^{\theta} c(x) d\theta' \\
 &+ \Delta t \int_0^{\theta} \left(u_t^{t_{n+1}} + \frac{1}{2} |\bar{\sigma}(\cdot, T) \cdot|^2 u_{xx}^{t_{n+1}} \right) (t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) d\theta' \\
 &+ \Delta t^{1/2} \int_0^{\theta} \left(u_x^{t_{n+1}}(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) - u_x^{t_{n+1}}(t_n, x) \right) \\
 &\times \bar{\sigma}(t_n + \theta' \Delta t, T) X_{\theta'}^{\Delta t} dB_{\theta'}.
 \end{aligned} \tag{2.19}$$

Now we replace the PDE (2.5) satisfied by the function $u^{t_{n+1}}$ in (2.19). Again applying Ito's formula to the gradient $u_x^{t_{n+1}}(t_n + \theta\Delta t, X_{\theta}^{\Delta t})$ in (2.19), we get

$$\begin{aligned}
 \mathcal{E}_{\theta}^{\Delta t} &= \Delta t |\bar{\sigma}(t_n, T)x|^2 u_{xx}^{t_{n+1}}(t_n, x) \frac{B_{\theta}^2 - \theta}{2} + \Delta t R_{\theta}^{\Delta t}(t_n, x) \\
 &- \Delta t \left(f(t_n, x, u^{t_{n+1}}(t_n, x), u_x^{t_{n+1}}(t_n, x), u_{xx}^{t_{n+1}}(t_n, x)) + c(x) \right) \theta,
 \end{aligned} \tag{2.20}$$

where the remainder process $R_{\theta}^{\Delta t}$ is given by

$$\begin{aligned}
 R_{\theta}^{\Delta t}(t_n, x) &= \int_0^{\theta} \left(u_t^{t_{n+1}}(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) - u_t^{t_{n+1}}(t_n, x) \right) d\theta' \\
 &+ \frac{1}{2} \int_0^{\theta} \left(|\bar{\sigma}(t_n + \theta' \Delta t, T) X_{\theta'}^{\Delta t}|^2 u_{xx}^{t_{n+1}}(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) \right. \\
 &\quad \left. - |\bar{\sigma}(t_n, T)x|^2 u_{xx}^{t_{n+1}}(t_n, x) \right) d\theta' \\
 &+ \int_0^{\theta} \left(\frac{u_x^{t_{n+1}}(t_n + \theta' \Delta t, X_{\theta'}^{\Delta t}) - u_x^{t_{n+1}}(t_n, x)}{\Delta t^{1/2}} |\bar{\sigma}(t_n, T) X_{\theta'}^{\Delta t}|^2 \right. \\
 &\quad \left. - u_{xx}^{t_{n+1}}(t_n, x) |\bar{\sigma}(t_n, T)x|^2 B_{\theta'} \right) dB_{\theta'}.
 \end{aligned}$$

Due to the result in Lemma 1.13 (see Chapter 1), $R_{\theta}^{\Delta t}$ is supposed to be small in the following sense:

$$\sup_{\theta \in [0,1]} \mathbb{E}^B |R_{\theta}^{\Delta t}(t_n, x)|^2 \leq K_{n,N}(x) \Delta t,$$

for some constant $K_{n,N}(x) \in K_{\text{pol}}$.

Thanks to the relation $\ell_{\gamma}''(y)y = \ell_{\gamma}'(y)$ for all $y \in \mathbb{R}$, we get by applying the Ito-Tanaka formula

$$\ell_{\gamma}(\mathcal{E}_{\theta}^{\Delta t}) = \int_0^{\theta} \ell_{\gamma}''(\mathcal{E}_{\theta'}^{\Delta t}) \left(\mathcal{E}_{\theta'}^{\Delta t} d\mathcal{E}_{\theta'}^{\Delta t} + \frac{1}{2} d\langle \mathcal{E}^{\Delta t} \rangle_{\theta'} \right).$$

After replacing the expression of $\mathcal{E}^{\Delta t}$ in (2.20), we take the expectation to obtain

$$\mathbb{E}^B [\ell_{\gamma}(\mathcal{E}_1^{\Delta t})] = \Delta t^2 P^{\Delta t}(t_n, x) \text{ (see relation in (2.18))}$$

$$\begin{aligned}
 &= \Delta t^2 \mathbb{E}^B \left[\int_0^1 \ell''_\gamma \left(J_1(u^{t_{n+1}}, (t_n, x)) \frac{B_\theta^2 - \theta}{2} \Delta t \right. \right. \\
 &\quad \left. \left. - J_2(u^{t_{n+1}}, (t_n, x)) \theta \Delta t + R_\theta^{\Delta t}(t_n, x) \Delta t \right) \right. \\
 &\quad \times \left(|J_1(u^{t_{n+1}}, (t_n, x))|^2 \frac{B_\theta^2}{2} + |J_2(u^{t_{n+1}}, (t_n, x))|^2 \theta \right. \\
 &\quad \left. \left. - J_1(u^{t_{n+1}}, (t_n, x)) J_2(u^{t_{n+1}}, (t_n, x)) \frac{B_\theta^2 - \theta}{2} \right) d\theta \right] \\
 &\quad + K_{n,N}(x) \Delta t^{5/2},
 \end{aligned}$$

where $J_2(u^{t_{n+1}}, \cdot)$ and $J_1(u^{t_{n+1}}, \cdot)$ are given in (2.9).

We omit the study of the terms in the remainder, because they were already treated in the proof of Proposition 1.9 in Chapter 1. Using the relation $\ell''_\gamma(\varepsilon y) = \ell''_\gamma(y)$ for any $\varepsilon > 0$, we obtain the announced result. \square

Aggregating aspect. From the relation between $P^{\Delta t}(t_n, X_{t_n})$ and $\mathbf{R}_{n,\gamma}$ in (2.18), we write $\mathbf{R}_{N,\gamma}[v, f]$ using the expansion of $P^{\Delta t}(t, x)$ obtained in Proposition 2.10:

$$\begin{aligned}
 \mathbf{R}_{N,\gamma}[v, f] &= \frac{1}{\Delta t} \sum_{n=0}^{N-1} \mathbb{E}[\mathbf{R}_{n,\gamma}] \\
 &= \sum_{n=0}^{N-1} \mathbb{E} \left[P^{\Delta t}(t_n, X_{t_n}) \right] \Delta t \text{ (see relation in (2.18))} \\
 &= \sum_{n=0}^{N-1} \mathbb{E} \left[\int_0^1 \ell''_\gamma \left(J_1(u^{t_{n+1}}, (t_n, X_{t_n})) \frac{B_\theta^2 - \theta}{2} \right. \right. \\
 &\quad \left. \left. - J_2(u^{t_{n+1}}, (t_n, X_{t_n})) \theta + R_\theta^{\Delta t}(t_n, X_{t_n}) \right) \right. \\
 &\quad \times \left(|J_1(u^{t_{n+1}}, (t_n, X_{t_n}))|^2 \frac{B_\theta^2}{2} + |J_2(u^{t_{n+1}}, (t_n, X_{t_n}))|^2 \theta \right. \\
 &\quad \left. \left. - J_1(u^{t_{n+1}}, (t_n, X_{t_n})) J_2(u^{t_{n+1}}, (t_n, X_{t_n})) \frac{B_\theta^2 - \theta}{2} \right) d\theta \right] \Delta t \\
 &\quad + \sum_{n=0}^{N-1} \mathbb{E} [K_{n,N}(X_{t_n}) \Delta t^{3/2}].
 \end{aligned} \tag{2.21}$$

In the previous expression, we substitute the solution $u^{t_{n+1}}(t_n, \cdot)$ by its terminal condition $v(t_{n+1}, \cdot)$. This is equivalent to replace $J_2(u^{t_{n+1}}, (t_n, \cdot))$ and $J_1(u^{t_{n+1}}, (t_n, \cdot))$ by $J_2(v, (t_{n+1}, \cdot))$

and $J_1(v, (t_{n+1}, \cdot))$. Then, we rewrite $\mathbf{R}_{N,\gamma}[v, f]$ in (2.21) as follows

$$\begin{aligned}
 & \mathbf{R}_{N,\gamma}[v, f] \\
 &= \sum_{n=0}^{N-1} \mathbb{E} \left[\int_0^1 \ell''_{\gamma} \left(J_1(v, (t_{n+1}, X_{t_n})) \frac{B_{\theta}^2 - \theta}{2} \right. \right. \\
 & \quad \left. \left. - J_2(v, (t_{n+1}, X_{t_n}))\theta + \bar{R}_{\theta}^{\Delta t}(t_n, X_{t_n}) \right) \right. \\
 & \quad \times \left(|J_1(v, (t_{n+1}, X_{t_n}))|^2 \frac{B_{\theta}^2}{2} + |J_2(v, (t_{n+1}, X_{t_n}))|^2 \theta \right. \\
 & \quad \left. - J_1(v, (t_{n+1}, X_{t_n})) J_2(v, (t_{n+1}, X_{t_n})) \frac{B_{\theta}^2 - \theta}{2} \right) d\theta \Big] \Delta t \\
 &+ \sum_{n=0}^{N-1} \mathbb{E} \left[(\bar{C}^{\Delta t}(t_n, X_{t_n}) + K_{n,N}(X_{t_n}) \Delta t^{3/2}) \right],
 \end{aligned} \tag{2.22}$$

where

$$\begin{aligned}
 \bar{R}_{\theta}^{\Delta t}(t_n, x) &:= R_{\theta}^{\Delta t}(t_n, x) \\
 &+ (J_1(u^{t_{n+1}}, (t_n, x)) - J_1(v, (t_{n+1}, x))) \frac{B_{\theta}^2 - \theta}{2} \\
 &- (J_2(u^{t_{n+1}}, (t_n, x)) - J_2(v, (t_{n+1}, x)))\theta
 \end{aligned}$$

and

$$\begin{aligned}
 \bar{C}^{\Delta t}(t_n, x) &:= \\
 &\mathbb{E}^B \left[\int_0^1 \ell''_{\gamma} \left(J_1(u^{t_{n+1}}, (t_n, x)) \frac{B_{\theta}^2 - \theta}{2} \right. \right. \\
 & \quad \left. \left. - J_2(u^{t_{n+1}}, (t_n, x))\theta + R_{\theta}^{\Delta t}(t_n, x) \right) \right. \\
 & \quad \times \left\{ \left(|J_1(u^{t_{n+1}}, (t_n, x))|^2 - |J_1(v, (t_{n+1}, x))|^2 \right) \frac{B_{\theta}^2}{2} \right. \\
 & \quad \left. + \left(|J_2(u^{t_{n+1}}, (t_n, x))|^2 - |J_2(v, (t_{n+1}, x))|^2 \right) \theta \right. \\
 & \quad \left. - \left(J_1(u^{t_{n+1}}, (t_n, x)) J_2(u^{t_{n+1}}, (t_n, x)) \right. \right. \\
 & \quad \left. \left. - J_1(v, (t_{n+1}, x)) J_2(v, (t_{n+1}, x)) \right) \frac{B_{\theta}^2 - \theta}{2} \right\} d\theta \Big].
 \end{aligned}$$

Noticed that the terms $\bar{R}_{\theta}^{\Delta t}(t_n, X_{t_n})$ and $\bar{C}^{\Delta t}(t_n, X_{t_n})$ need to be estimated in order to pass to limit $N \rightarrow \infty$ in the expression (2.22). From the result in Proposition 1.11 (in Chapter 1), we have the previous terms converge to zero in the following sense:

- $\sup_{0 \leq n \leq N-1} \mathbb{E} |\bar{C}^{\Delta t}(t_n, X_{t_n})| \leq K \Delta t^{1/2}$, for some constant $K > 0$.
- $\sup_{0 \leq n \leq N-1} \sup_{\theta \in [0,1]} |\bar{R}_{\theta}^{\Delta t}(t_n, X_{t_n})| \xrightarrow[N \rightarrow \infty]{} 0$, $d\mathbb{P}^W \otimes d\mathbb{P}^B$ -almost surely.

Now we write the integrated conditional risk $\mathbf{R}_{N,\gamma}$ in (2.22) as the expectation of a double integrable. By setting

$$\varphi_N(t) := \sup\{t_n : t \geq t_n\} \text{ and } \bar{\varphi}_N(t) := \inf\{t_n : t < t_n\},$$

we obtain from equation (2.21)

$$\begin{aligned}
 & \mathbf{R}_{N,\gamma}[v, f] \\
 &= \mathbb{E} \left[\int_0^T \int_0^1 \ell''_{\gamma} \left(J_1(v, (\bar{\varphi}_N(t), X_{\varphi_N(t)})) \frac{B_{\theta}^2 - \theta}{2} \right. \right. \\
 &\quad \left. \left. - J_2(v, (\bar{\varphi}_N(t), X_{\varphi_N(t)}))\theta + \bar{R}_{\theta}^{\Delta t}(\bar{\varphi}_N(t), X_{\varphi_N(t)}) \right) \right. \\
 &\quad \times \left\{ |J_1(v, (\bar{\varphi}_N(t), X_{\varphi_N(t)}))|^2 \frac{B_{\theta}^2}{2} + |J_2(v, (\bar{\varphi}_N(t), X_{\varphi_N(t)}))|^2 \theta \right. \\
 &\quad \left. - J_1(v, (\bar{\varphi}_N(t), X_{\varphi_N(t)})) J_2(v, (\bar{\varphi}_N(t), X_{\varphi_N(t)})) \frac{B_{\theta}^2 - \theta}{2} \right\} d\theta dt \Big] \\
 &\quad + \sum_{n=0}^{N-1} \mathbb{E}[\bar{C}^{\Delta t}(t_n, X_{t_n}) + K_{n,N}(X_{t_n})\Delta t^{3/2}].
 \end{aligned} \tag{2.23}$$

When the number of trading dates $N \rightarrow +\infty$, the last term of $\mathbf{R}_{N,\gamma}[v, f]$ in (2.23) goes to zero, due to the estimates in Lemma 1.11 and Proposition 1.9. On the other hand, the limit of the first term in (1.37) is achieved by applying the dominated convergence theorem.

Step 1. Because of the continuity of the volatility $\bar{\sigma}(t, T)$ in time t , the reference price v and its derivatives v_x, v_{xx} in time t and in space x , the PDE non-linearity f , the cost function c , and the path-continuity of X , we get

$$\begin{aligned}
 J_1(v, (\bar{\varphi}_N(t), X_{\varphi_N(t)})) &\xrightarrow{N \rightarrow \infty} J_1(v, (t, X_t)) \\
 J_2(v, (\bar{\varphi}_N(t), X_{\varphi_N(t)})) &\xrightarrow{N \rightarrow \infty} J_2(v, (t, X_t)),
 \end{aligned}$$

\mathbb{P}^W -almost surely, for all $t \in [0, T]$. Moreover, the result of item (c) of Lemma 1.11 states

$$\sup_{0 \leq n \leq N-1} \sup_{\theta \in [0,1]} |\bar{R}_{\theta}^{\Delta t}(t_n, X_{t_n})| \xrightarrow{N \rightarrow \infty} 0,$$

$\mathbb{P}^W \otimes \mathbb{P}^B$ -almost surely. Then, we get

$$\begin{aligned}
 & J_1(v, (\bar{\varphi}_N(t), X_{\varphi_N(t)})) \frac{B_{\theta}^2 - \theta}{2} - J_2(v, (\bar{\varphi}_N(t), X_{\varphi_N(t)}))\theta \\
 &+ \bar{R}_{\theta}^{\Delta t}(\bar{\varphi}_N(t), X_{\varphi_N(t)}) \xrightarrow{N \rightarrow \infty} \\
 & J_1(v, (t, X_t)) \frac{B_{\theta}^2 - \theta}{2} - J_2(v, (t, X_t))\theta,
 \end{aligned}$$

and

$$\begin{aligned}
 & |J_2(v, (\bar{\varphi}_N(t), X_{\varphi_N(t)}))|^2 \theta + |J_1(v, (\bar{\varphi}_N(t), X_{\varphi_N(t)}))|^2 \frac{B_{\theta}^2}{2} \\
 &- J_2(v, (\bar{\varphi}_N(t), X_{\varphi_N(t)})) J_1(v, (\bar{\varphi}_N(t), X_{\varphi_N(t)})) \frac{B_{\theta}^2 - \theta}{2} \xrightarrow{N \rightarrow \infty} \\
 & |J_2(v, (t, X_t))|^2 \theta + |J_1(v, (t, X_t))|^2 \frac{B_{\theta}^2}{2} \\
 &- J_2(v, (t, X_t)) J_1(v, (t, X_t)) \frac{B_{\theta}^2 - \theta}{2},
 \end{aligned}$$

$\mathbb{P}^W \otimes \mathbb{P}^B$ -almost surely, for all $(\theta, t) \in [0, 1] \times [0, T]$.

Step 2. Because of the second derivative ℓ''_γ is discontinuous at 0 but the set

$$\mathcal{A} := \left\{ (\omega, t, \theta) \in \Omega \times [0, T] \times [0, 1] : \right. \\ \left. J_1(v, (t, X_{t,\omega})) \frac{B_\theta(\omega)^2 - \theta}{2} - J_2(v, (t, X_{t,\omega}))\theta = 0 \right\}$$

has measure zero (due to Proposition 1.19), it holds

$$\ell''_\gamma \left(J_1(v, (\bar{\varphi}_N(t), X_{\varphi_N(t)})) \frac{B_\theta^2 - \theta}{2} - J_2(v, (\bar{\varphi}_N(t), X_{\varphi_N(t)}))\theta \right. \\ \left. + \bar{R}_\theta^{\Delta t}(\bar{\varphi}_N(t), X_{\varphi_N(t)}) \right) \xrightarrow{N \rightarrow \infty} \\ \ell''_\gamma \left(J_1(v, (t, X_t)) \frac{B_\theta^2 - \theta}{2} - J_2(v, (t, X_t))\theta \right),$$

$d\mathbb{P}^W \otimes d\mathbb{P}^B \otimes dt \otimes d\theta$ -almost surely.

Step 3. Because of the boundedness of ℓ''_γ , the polynomial growth of v , v_x , v_{xx} , and c in space x , we have that the integrand in (2.23) is bounded by

$$C \left(1 + \sup_{t \in [0, T]} |X_t| + |B_\theta| \right)^\alpha,$$

for some positive constants C and α . Finally, we take the limit of the first term of $\mathbf{R}_{N,\gamma}[v, f]$ in (2.23) to get

$$\mathbf{R}_\gamma[v, f] = \\ \mathbb{E} \left[\int_0^T \int_0^1 \ell''_\gamma \left(J_1(v, (t, X_t)) \frac{B_\theta^2 - \theta}{2} - J_2(v, (t, X_t))\theta \right) \right. \\ \times \left(|J_1(v, (t, X_t))|^2 \frac{B_\theta^2}{2} + |J_2(v, (t, X_t))|^2 \theta \right. \\ \left. - J_1(v, (t, X_t)) J_2(v, (t, X_t)) \frac{B_\theta^2 - \theta}{2} \right) d\theta dt \Big],$$

by dominated convergence theorem, which concludes the proof of Theorem 2.9.

2.4 Convex-concave staggered cost models

In this section, we consider convex-concave staggered cost models. First, we introduce the case of increasing costs with continuous differentiable dependence on the spot price. This corresponds to the economic point of view where the costs are split according to the power plant usability. Then, we examine decreasing costs being a piecewise convex function of the spot price. In this case, we consider the seasonality of the refueling and maintenance stops, which occur usually when the demand is in the lowest level, then when the spot price is also low.

2.4.1 Increasing costs

First, we take a convex-concave staggered cost function c depending on the spot price level

$$c_\alpha(x) = c_0 + c_1(x/c_2)^\alpha, \quad (2.24)$$

for $x \in \mathbb{R}$, where the constants $c_0 \geq 0$, $c_1 > 0$ and $c_2 > 0$. The parameter α controls the convexity ($\alpha > 1$) or concavity ($0 < \alpha < 1$) of the function c_α . Notice that c_0 represents the minimal maintenance cost and c_1 , the cost increment from 0 up to level price c_2 (see Figure 2.3).

2.4.2 Decreasing costs

Now, we consider staggered costs decreasing with the spot price level. Thus, we take a staggered cost function c as follows

$$\begin{aligned} c_{\beta,0}(x) &= b_0 + b_1 e^{-(x/b_2)^\beta}, \\ c_{\beta,1}(x) &= b_0 + b_1(1 - (x/b_2)^\beta)^+, \end{aligned}$$

for $x \in \mathbb{R}$, where the constants $b_0 \geq 0$, $b_1 > 0$ and $b_2 > 0$. Here, the parameter β characterizes the convexity ($0 < \beta \leq 1$) of the costs. Analogously to the previous model, notice that b_0 represents the minimal staggered costs (related to the investment), and b_1 , the staggered cost increment (related to the maintenance) (see Figure 2.4).

This part of the fixed costs are influenced mostly by the maintenance costs, which are greater during the refueling stops. From an economic point of view, power producers prefer to schedule the maintenance of their power plant when the electricity demand is low, increasing the stability grid in the periods of high need of power. Thus, by avoiding to stop when the electricity is also high, they seek for higher profits.

2.5 Numerical experiments

In this section, we present the producer valuation/hedging policies in the presence of fixed costs. Those costs are spread over the hedging period and are dependent on electricity spot price according to the models proposed in Section 2.4.

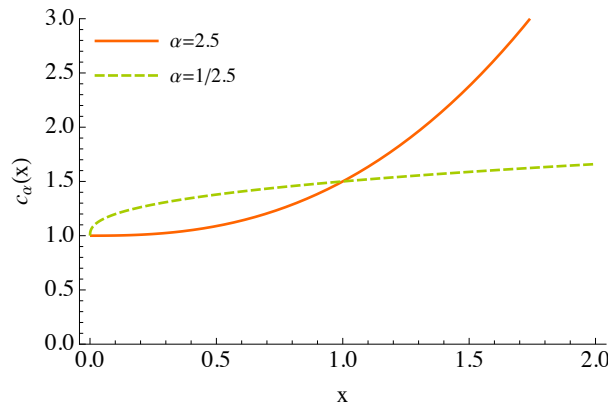


Figure 2.3 – Cost c_α depending on the price level x with $c_0 = 1$, $c_1 = 0.5$, and $c_2 = 1$. Around the price level c_2 , the convex function c_α presents a steep slope while the concave function $c_{1/\alpha}$ is flatter.

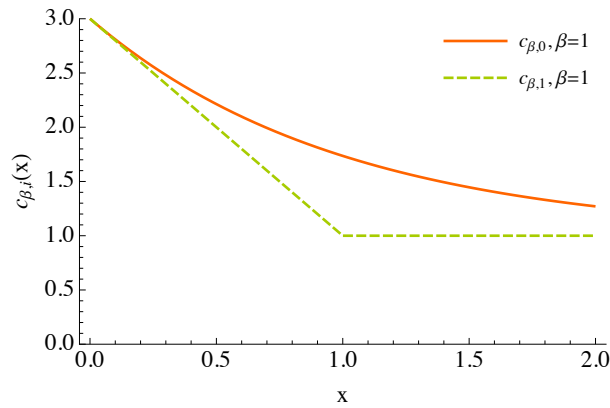


Figure 2.4 – Cost $c_{\beta,0}$ and $c_{\beta,1}$ depending on the price level x with $b_0 = 1$, $b_1 = 2$ and $b_2 = 1$.

Now we describe how those valuation/hedging policies are computed. Denote by g the future income payoff due the power plant production. With the optimal PDE non-linearity f^* (see (2.13)) in hands, the power producer selects the following rule

$$V_{t_n} = v^*(t_n, X_{t_n}), \quad \vartheta_{t_n} = v_x^*(t_n, X_{t_n}),$$

where v^* is the solution to the f^* -PDE

$$-v_t^* - f^*(t, x, v^*, v_x^*, v_{xx}^*) = \frac{1}{2} |\bar{\sigma}(t, T)x|^2 v_{xx}^*, \quad (2.25)$$

for all $(t, x) \in [0, t_{n+1}] \times \mathbb{R}$, with the terminal condition $v^*(T, \cdot) = g(\cdot)$ at the time T . Notice that v^* depend on the parameter γ and on the staggered cost function $c(x)$, through the PDE non-linearity f^* .

To solve PDE (2.25), we use standard finite difference methods as described in Subsection 1.5.1 (Chapter 1). Also, we consider the following volatility function

$$\bar{\sigma}(t, T) = \bar{\sigma}_0 e^{-a_0(T-t)}, \quad 0 \leq t \leq T,$$

for some constants $\bar{\sigma}_0 > 0$ and $a_0 \geq 0$ (see Figure 2.5).

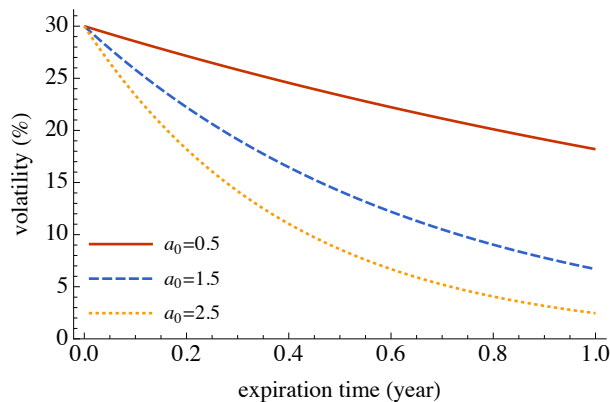


Figure 2.5 – Volatility $\bar{\sigma}(T - \tau, T)$ decreasing in expiration time τ with $\bar{\sigma}_0 = 30\%$. The parameter a_0 represents the decay rate. In this case, the one-year-to-expiration volatility is 60.65% of the initial volatility $\bar{\sigma}_0$ for $a_0 = 0.5$.

2.5.1 Optimal valuation/hedging rule for increasing costs

Here, we study the impact of an increasing cost function $c(x)$ on the producer valuation/hedging rule. At the horizon time T , power producers choose to turn their power plants on if the electricity spot price S_T is greater than K_0 , representing the starting costs. Here, we assume those costs are due to the fuel consumed to produce electricity. So, the future income function describing this option is a call payoff with strike price K_0 given by $g(S_T) = (S_T - K_0)^+$ at horizon time $T = 1$. Also, we consider the following set of parameters:

γ	α	c_0	c_1	c_2	$\bar{\sigma}_0$	a_0
-0.3	2.5	0.2	0.5	10	30%	0.5

Discussion about convex staggered costs. First, we consider the convex cost function c_α defined in (2.24) for $\alpha > 1$. Economically, those increasing fixed costs model a positive correlation between maintenance costs and the electricity spot price. This is a result from the assumption that maintenance costs increase with a longer and an intense use of the power plant.

In Figure 2.6, we show the valuation/hedging rule at the time $t = 0$ for future income function (described as call option) with starting costs $K_0 = 10$. In the plot, we observe three different rules:

- the red (full) line represents the Black-Scholes valuation corresponding to $\gamma = 0$ and $c(x) = 0$;
- the green (dashed) one represents the Black-Scholes valuation with staggered costs corresponding to $\gamma = 0$ and $c(x) = c_\alpha(x)$;
- the blue (dotted) one represents the *asymmetric risk* valuation with staggered costs corresponding to $\gamma = -0.3$ and $c(x) = c_\alpha(x)$.

Considering the value function, we remark that both asymmetric risk prices are inferior to the Black-Scholes one when the parameter $\gamma \in (-1, 0)$. By increasing the risk aversion parameter γ , we obtain an asymmetric risk price with a greater convexity with respect to the hedging instrument X_0 (a forward contract with delivery time T), corresponding to a smaller volatility coefficient in a Black-Scholes model. When the constant c_0 get larger, we have the price reduced due to costs incurred along the hedging.

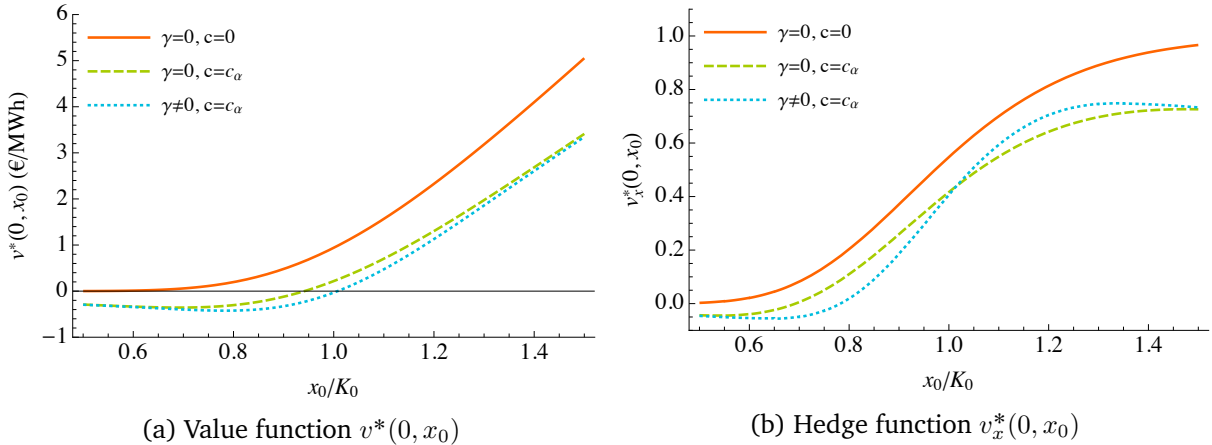


Figure 2.6 – Optimal valuation/hedging rule for a call payoff with $K_0 = 10$ using a convex staggered cost function $c_\alpha(x)$ with $\alpha > 1$.

Also, we notice that, when the staggered fixed costs depend on the electricity spot price level through a convex function, we obtain a hedging function inferior to hedge one in the case of zero cost. The first reason for that is the initial cost c_0 incurred independently of the price level. Secondly, due to the convexity of $c(x)$, the hedging function is reduced further for a large price level, then the parameter α controls the increment of the gap in the right-hand side of the graph in Figure 2.6 (b).

Discussion about concave staggered costs. Now, we consider the concave cost function defined in (2.24) by $0 < \alpha < 1$. In Figure 2.7, we show the valuation/hedging rule at the time $t = 0$ for call payoff with starting costs $K_0 = 10$. Similarly to the convex case, the hedging function is now reduced further for a small price level because it is the region where the costs are more important.

When power producers penalize further the negative local balance ($\mathcal{E}_n < 0$), as described in the introduction, through a more negative parameter γ , their optimal hedging for forward prices X_0 much larger than the strike price K_0 is to take a greater long position for call payoffs compared to Black-Scholes.

2.5.2 Optimal valuation/hedging rule for decreasing costs

Now, we examine the influence of staggered costs decreasing with the electricity spot price on the the producer valuation/hedging rule. Here, power producers will receive a future income $g(S_T)$ at time T due to the trading of electricity on the market. If the starting costs K_0 are covered by the sale of power, they will turn their plants on, obtaining a gain of $g(S_T) = (S_T - K_0)^+$.

Fixed cost accounting. However, their power plants generate fixed costs (for example, due to maintenance stops) despite of their future decision. Therefore, those fixed costs K_1 should be subtracted from the previous gain generating then the following balance: $(S_T - K_0)^+ - K_1$. In view of reduce the randomness of this future income, we trade on the forward contract $F(t, T)$ with delivery time T at equidistant dates. Then, the amount K_1 is spread over the time interval $[0, T]$.

In order to use realistic values for those costs K_0 and K_1 , it is interesting to rewrite the payoff in terms of the initial forward price. With this formulation, we could compare the results for different values of initial forward price $F(0, T)$ by changing the strike price in the rewritten payoff.

Rewriting in terms of the initial forward price We write the forward price $F(t, T)$ solution to the SDE (2.14) as follows

$$F(t, T) = F(0, T)Y(t, T),$$

where $F(0, T)$ is the initial forward price and $Y(t, T)$ is the stochastic deformation. Now, we consider the following transformation given by

$$\begin{aligned}\tilde{v}^*(t, y) &= v^*(t, F(0, T)y)/F(0, T) \\ \tilde{g}(y) &= g(F(0, T)y)/F(0, T), \\ \tilde{c}(y) &= c(F(0, T)y)/F(0, T),\end{aligned}$$

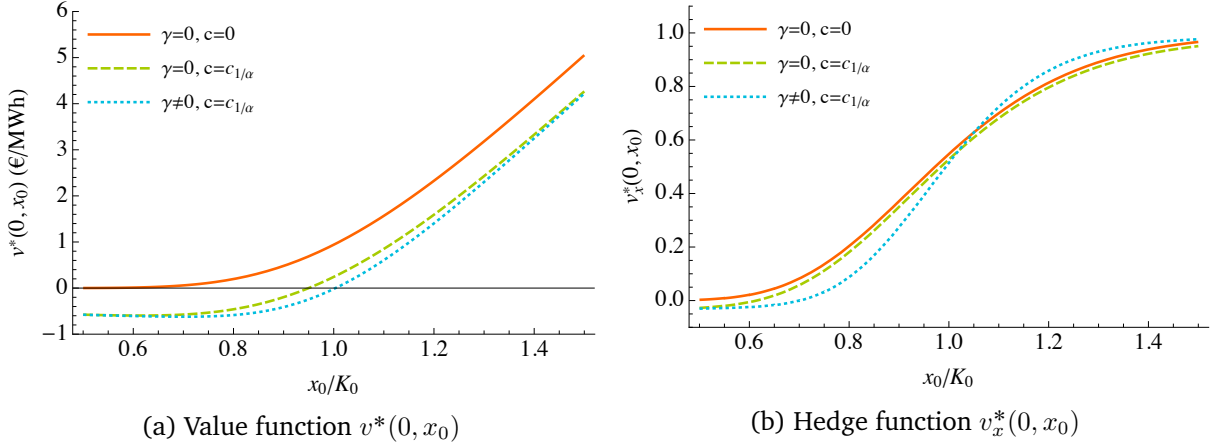


Figure 2.7 – Optimal valuation/hedging rule for a call payoff with $K_0 = 10$ using a concave staggered cost function $c_{1/\alpha}(x)$ with $0 < 1/\alpha < 1$

for an initial forward price $F(0, T) \neq 0$. Then, we get the following derivatives of \tilde{v}^* in terms of the derivatives of the value function v^*

$$\begin{aligned}\tilde{v}_y^*(t, y) &= v_x^*(t, F(0, T)y), \\ \tilde{v}_{yy}^*(t, y) &= v_{xx}^*(t, F(0, T)y)F(0, T).\end{aligned}$$

With this transformations in hands, we write the PDE (2.25) in v^* with terminal condition $v^*(T, x) = g(x)$ in terms of \tilde{v}^*

$$\tilde{v}_t^* + \tilde{f}^*(t, y, \tilde{v}^*, \tilde{v}_y^*, \tilde{v}_{yy}^*) = \frac{1}{2}|\bar{\sigma}(t, T)y|^2 \tilde{v}_{yy}^*,$$

with terminal condition $\tilde{v}^*(T, y) = \tilde{g}(y)$ at time T , where

$$\begin{aligned}\tilde{f}^*(t, y, \tilde{v}^*, \tilde{v}_y^*, \tilde{v}_{yy}^*) &= \chi^*(F(0, T)|\bar{\sigma}(t, T)y|^2 \tilde{v}_{yy}^*)/F(0, T) - \tilde{c}(y) \\ &= \chi^*(|\bar{\sigma}(t, T)y|^2 \tilde{v}_{yy}^*) - \tilde{c}(y),\end{aligned}$$

because χ^* is homogenous of degree 1 (see (2.11)).

Discussion about decreasing staggered costs Now, we consider a transformed payoff $\tilde{g}(y) = (y - \tilde{K}_0)^+$ and transformed costs $\tilde{c}_{\beta,0}(y) = \tilde{b}_0 + \tilde{b}_1 \exp(-(y/\tilde{b}_2)^\beta)$, $\tilde{c}_{\beta,1}(y) = \tilde{b}_0 + \tilde{b}_1(1 - (y/\tilde{b}_2)^\beta)^+$, where $\tilde{K}_0 = K_0/F(0, T)$, $\tilde{b}_i = b_i/F(0, T)$, for every $i \in \{0, 1, 2\}$. Also, we take the following sets of parameters corresponding to a growth of the initial forward price from $F(0, T) = 40$ to $F(0, T) = 60$:

	γ	\tilde{K}_0	\tilde{b}_0	\tilde{b}_1	\tilde{b}_2	$\bar{\sigma}_0$	a_0
A	-0.1	10/40	35/40	15/40	40/40	30%	0.5
B	-0.3	10/60	35/60	15/60	40/60	30%	0.5

In Figures 2.8 and 2.9, we show the valuation/hedging rule at the time $t = 0$ for the transformed call option with starting costs $\tilde{K}_0 = 1/4$ and $\tilde{K}_0 = 1/6$, respectively. In the plot, we observe three different rules:

- the red (full) line represents the *asymmetric risk* with staggered costs corresponding to $\gamma \in \{-0.1, -0.3\}$ and $c(x) = c_{\beta,0}(x)$;

- the green (dotted) one represents the *asymmetric risk* valuation with staggered costs corresponding to $\gamma \in \{-0.1, -0.3\}$ and $c(x) = c_{\beta,1}(x)$;
- the gray (dotted) one represents the Black-Scholes valuation corresponding to $\gamma = 0$ and $c(x) = 0$;

Regarding the value function, we notice a difference due to the transformed staggered costs \tilde{b}_0 and \tilde{b}_1 from the set A are lesser than the set B. At time t , the power producer observes the forward price $F(t, T)$ at level x and computes the ratio x/K_0 . Following the relation $x/K_0 = y/\tilde{K}_0$, the producer expects a positive option value at time t , if the value of $\tilde{v}^*(t, y)$ at point y/\tilde{K}_0 is positive (analogous to the function $\tilde{v}^*(0, y_0)$ in Figure 2.9). Considering the hedging function, we observe that the asymmetric risk hedging rule is greater than the Black-Scholes one, because of the decreasing property of staggered cost functions $c_{\beta,i}$. Observing a forward price x at time t , the power producer can compare the different hedging rules by looking the hedging function $\tilde{v}_y^*(t, y)$ at point x/K_0 . Finally, we notice that the parameter $\gamma < 0$ reduces the diffusion term in the PDE (2.25) on convex regions, and increases it on concave parts, influencing directly the value function.

2.6 Conclusions and extensions

Here, we described a valuation and hedging strategies for future production of power plants with fixed costs. Using an *asymmetric risk* valuation, we proved the existence of the asymptotic risk in the case of staggered fixed costs depending on the electricity spot price. Then, we provided methods to solve the associated PDE and to find the so-called optimal policies. Our model considered time-independent deterministic staggered costs which could be either a convex or concave function of the spot price level.

Several extensions are of interest. First, we can consider the case where the staggered costs c also depends on time. Since the maintenance costs increase in certain periods of the year (due to the scheduled maintenance stops), the problem is solved in a similar way. In addition, the optimal PDE in (2.25) becomes

$$-v_t^* - \chi^*(|\bar{\sigma}(t, T)x|^2 v_{xx}^*) + c(t, x) = \frac{1}{2}|\bar{\sigma}(t, T)x|^2 v_{xx}^*.$$

Finally, a natural extension is the case where c is a random function. This is important if the

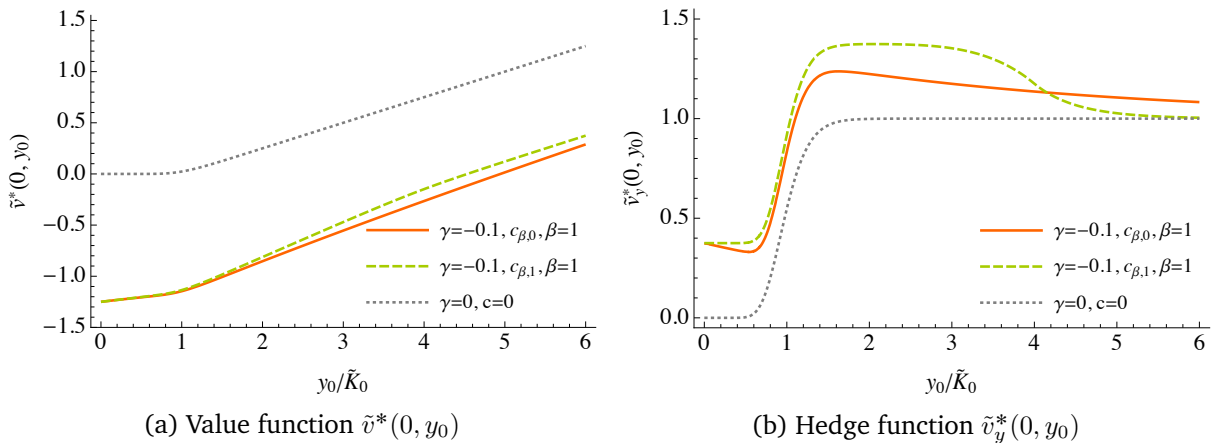


Figure 2.8 – Optimal valuation/hedging rule for a call payoff using a decreasing staggered cost function $c_{\beta,i}(x)$ with parameter set A.

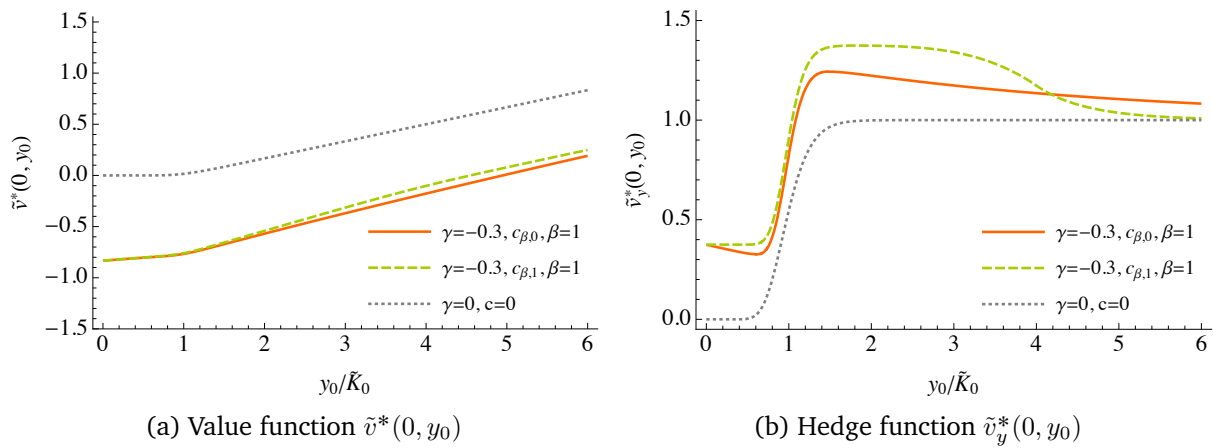


Figure 2.9 – Optimal valuation/hedging rule for a call payoff using a decreasing staggered cost function $c_{\beta,i}(x)$ with parameter set B.

maintenance costs are subject to unpredictable factors like when a nuclear plant has a breakdown or when the maintenance periods has to be prolonged.

Part II

Numerical methods in stochastic control

Polynomial conditional McKean-Vlasov control problems: Some probabilistic numerical methods

This chapter is based on the paper [BHL⁺18].

Abstract

We address a class of McKean-Vlasov (MKV) control problems with common noise, called polynomial conditional MKV, and extending the known class of linear quadratic stochastic MKV control problems. We show how this polynomial class can be reduced by suitable Markov embedding to finite-dimensional stochastic control problems. We provide a discussion and comparison of three probabilistic numerical methods for solving the reduced control problem: quantization, regression by control randomization, and regress later methods. Our numerical results are illustrated on various examples: portfolio selection and liquidation under drift uncertainty, and a interbank systemic risk model with partial observation.

3.1 Introduction

The optimal control of McKean-Vlasov (also called mean-field) dynamics is a rather new topic in the area of stochastic control and applied probability, which has been knowing a surge of interest with the emergence of the mean-field game theory. It is motivated on the one hand by the asymptotic formulation of cooperative equilibrium for a large population of particles (players) in mean-field interaction, and on the other hand from control problems with cost functional involving nonlinear functional of the law of the state process (e.g., the mean-variance portfolio selection problem or risk measure in finance).

In this chapter, we are interested in McKean-Vlasov (MKV) control problems under partial observation and common noise, whose formulation is described as follows. On a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ equipped with two independent Brownian motions B and W^0 , let us consider the controlled stochastic MKV dynamics in \mathbb{R}^n :

$$dX_s = b(X_s, \mathbb{P}_{X_s}^{W^0}, \alpha_s)ds + \sigma(X_s, \mathbb{P}_{X_s}^{W^0}, \alpha_s)dB_s + \sigma_0(X_s, \mathbb{P}_{X_s}^{W^0}, \alpha_s)dW_s^0, \quad X_0 = x_0 \in \mathbb{R}^n, \quad (3.1)$$

where $\mathbb{P}_{X_s}^{W^0}$ denotes the conditional distribution of X_s given W^0 (or equivalently given \mathcal{F}_s^0 where $\mathbb{F}^0 = (\mathcal{F}_t^0)_t$ is the natural filtration generated by W^0), and the control α is \mathbb{F}^0 -progressive valued in some Polish space A . This measurability condition on the control means that the controller has

a partial observation of the state, in the sense that she can only observe the common noise. We make the standard Lipschitz condition on the coefficients $b(x, \mu, a)$, $\sigma(x, \mu, a)$, $\sigma_0(x, \mu, a)$ with respect to (x, μ) in $\mathbb{R}^n \times \mathcal{P}_2(\mathbb{R}^n)$, uniformly in $a \in A$, where $\mathcal{P}_2(\mathbb{R}^n)$ is the set of all probability measures on $(\mathbb{R}^n, \mathcal{B}(\mathbb{R}^n))$ with a finite second-order moment, endowed with the 2-Wasserstein metric \mathcal{W}_2 . This ensures the well-posedness of the controlled MKV stochastic differential equation (SDE) (3.1). The cost functional over a finite horizon T associated to the stochastic MKV equation (3.1) (sometimes called conditional MKV equation) for a control process α , is

$$J(\alpha) = \mathbb{E} \left[\int_0^T f(X_t, \mathbb{P}_{X_t}^{W^0}, \alpha_t) dt + g(X_T, \mathbb{P}_{X_T}^{W^0}) \right],$$

and the objective is to maximize over an admissible set \mathcal{A} of control processes the cost functional:

$$V_0 = \sup_{\alpha \in \mathcal{A}} J(\alpha).$$

The set \mathcal{A} of admissible controls usually requires some integrability conditions depending on the growth conditions on f, g , in order to ensure that $J(\alpha)$ is well-defined for $\alpha \in \mathcal{A}$ (more details will be given in the examples, see Section 3.5). Notice that classical partial observation control problem (without MKV dependence on the coefficients) arises as a particular case of (3.1)-(3.2). We refer to the introduction in [PW17] for the details.

Let us recall from [PW17] the dynamic programming equation associated to the conditional MKV control problem (3.2). We start by defining a suitable dynamic version of this problem. Let us consider \mathcal{F}_0 a sub σ -algebra of \mathcal{F} independent of B, W^0 . It is assumed w.l.o.g. that \mathcal{F}_0 is rich enough in the sense that $\mathcal{P}_2(\mathbb{R}^n) = \{\mathcal{L}(\xi) : \xi \in L^2(\mathcal{F}_0; \mathbb{R}^n)\}$, where $\mathcal{L}(\xi)$ denotes the law of ξ . Given a control $\alpha \in \mathcal{A}$, we consider the dynamic version of (3.1) starting from $\xi \in L^2(\mathcal{F}_0; \mathbb{R}^n)$ at time $t \in [0, T]$, and written as:

$$\begin{aligned} X_s^{t, \xi, \alpha} &= \xi + \int_t^s b(X_u^{t, \xi, \alpha}, \mathbb{P}_{X_u^{t, \xi, \alpha}}^{W^0}, \alpha_u) du + \int_t^s \sigma(X_u^{t, \xi, \alpha}, \mathbb{P}_{X_u^{t, \xi, \alpha}}^{W^0}, \alpha_u) dB_u \\ &\quad + \int_t^s \sigma_0(X_u^{t, \xi, \alpha}, \mathbb{P}_{X_u^{t, \xi, \alpha}}^{W^0}, \alpha_u) dW_u^0, \quad t \leq s \leq T. \end{aligned}$$

Let us then define the dynamic cost functional:

$$J(t, \xi, \alpha) = \mathbb{E} \left[\int_t^T f(X_s^{t, \xi, \alpha}, \mathbb{P}_{X_s^{t, \xi, \alpha}}^{W^0}, \alpha_s) ds + g(X_T^{t, \xi, \alpha}, \mathbb{P}_{X_T^{t, \xi, \alpha}}^{W^0}) \right],$$

for $(t, \xi) \in [0, T] \times L^2(\mathcal{F}_0; \mathbb{R}^n)$, $\alpha \in \mathcal{A}$, and notice by the law of conditional expectations, and as α is \mathbb{F}^0 -progressive that

$$J(t, \xi, \alpha) = \mathbb{E} \left[\int_t^T \hat{f}(\mathbb{P}_{X_s^{t, \xi, \alpha}}^{W^0}, \alpha_s) ds + \hat{g}(\mathbb{P}_{X_s^{t, \xi, \alpha}}^{W^0}) \right],$$

where $\hat{f} : \mathcal{P}_2(\mathbb{R}^n) \times A \rightarrow \mathbb{R}$, $\hat{g} : \mathcal{P}_2(\mathbb{R}^n) \rightarrow \mathbb{R}$ are defined by

$$\begin{aligned} \hat{f}(\mu, a) &= \mu(f(\cdot, \mu, a)) = \int_{\mathbb{R}^n} f(x, \mu, a) \mu(dx), \\ \hat{g}(\mu) &= \mu(g(\cdot, \mu)) = \int_{\mathbb{R}^n} g(x, \mu) \mu(dx). \end{aligned}$$

Moreover, notice that the conditional law of $X_s^{t, \xi, \alpha}$ given W^0 depends on ξ only through its law

$\mathcal{L}(\xi)$, and we can then define for $\alpha \in \mathcal{A}$:

$$\rho_s^{t,\mu,\alpha} = \mathbb{P}_{x_s^{t,\xi,\alpha}}^{W^0}, \quad \text{for } t \leq s, \mu = \mathcal{L}(\xi) \in \mathcal{P}_2(\mathbb{R}^n).$$

Therefore, the dynamic cost functional $J(t, \xi, \alpha)$ depends on $\xi \in L^2(\mathcal{F}_0; \mathbb{R}^n)$ only through its law $\mathcal{L}(\xi)$, and by an abuse of notation, we write $J(t, \mu, \alpha) = J(t, \xi, \alpha)$ when $\mu = \mathcal{L}(\xi)$. We then consider the value function for the conditional MKV control problem (3.2), defined on $[0, T] \times \mathcal{P}_2(\mathbb{R}^n)$ by

$$v(t, \mu) = \sup_{\alpha \in \mathcal{A}} J(t, \mu, \alpha) = \sup_{\alpha \in \mathcal{A}} \mathbb{E} \left[\int_t^T \hat{f}(\rho_s^{t,\mu,\alpha}, \alpha_s) ds + \hat{g}(\rho_T^{t,\mu,\alpha}) \right],$$

and notice that at time $t = 0$, when $\xi = x_0$ is a constant, then $V_0 = v(0, \delta_{x_0})$.

It is shown in [PW17] that dynamic programming principle (DPP) for the conditional MKV control problem (3.5) holds: for $(t, \mu) \in [0, T] \times \mathcal{P}_2(\mathbb{R}^n)$,

$$v(t, \mu) = \sup_{\alpha \in \mathcal{A}} \mathbb{E} \left[\int_t^\theta \hat{f}(\rho_s^{t,\mu,\alpha}, \alpha_s) ds + v(\theta, \rho_\theta^{t,\mu,\alpha}) \right],$$

for any \mathbb{F}^0 -stopping time θ valued in $[t, T]$. Next, by relying on the notion of differentiability with respect to probability measures introduced by P. L. Lions [Lio12] (see also the lecture notes [Car10]) and the chain rule (Itô's formula) along flow of probability measures (see [BLPR17], [CCD14]), we derive the HJB equation for v :

$$\begin{cases} \partial_t v + \sup_{a \in \mathcal{A}} \left[\hat{f}(\mu, a) + \mu(\mathbb{L}^a v(t, \mu)) + \mu \otimes \mu(M^a v(t, \mu)) \right] = 0, & (t, \mu) \in [0, T) \times \mathcal{P}_2(\mathbb{R}^n), \\ v(T, \mu) = \hat{g}(\mu), & \mu \in \mathcal{P}_2(\mathbb{R}^n), \end{cases} \quad (3.6)$$

where for $\phi \in \mathcal{C}_b^2(\mathcal{P}_2(\mathbb{R}^n))$, $a \in \mathcal{A}$, and $\mu \in \mathcal{P}_2(\mathbb{R}^n)$, $\mathbb{L}^a \phi(\mu)$ is the function $\mathbb{R}^n \rightarrow \mathbb{R}$ defined by

$$\mathbb{L}^a \phi(\mu)(x) = \partial_\mu \phi(\mu)(x) \cdot b(x, \mu, a) + \frac{1}{2} \text{tr}(\partial_x \partial_\mu \phi(\mu)(x) (\sigma \sigma^\top + \sigma_0 \sigma_0^\top)(x, \mu, a)),$$

and $M^a \phi(\mu)$ is the function $\mathbb{R}^n \times \mathbb{R}^n \rightarrow \mathbb{R}$ defined by

$$M^a \phi(\mu)(x, x') = \frac{1}{2} \text{tr}(\partial_\mu^2 \phi(\mu)(x, x') \sigma_0(x, \mu, a) \sigma_0^\top(x', \mu, a)).$$

The HJB equation (3.6) is a fully nonlinear partial differential equation (PDE) in the infinite-dimensional Wasserstein space. In general, this PDE does not have an explicit solution except in the notable important class of linear-quadratic MKV control problem. Numerical resolution for MKV control problem or equivalently for the associated HJB equation is a challenging problem due to the nonlinearity of the optimization problem and the infinite-dimensional feature of the Wasserstein space. In this work, our purpose is to investigate a class of MKV control problems which can be reduced to finite-dimensional problems in view of numerical resolution.

3.2 Notations and Assumptions

3.2.1 Main assumptions

We make two kinds of assumptions on the coefficients of the model: one on the dependence on x and the other on the dependence on μ .

Assumptions: dependence on x : we consider a class of models where the coefficients of the MKV equation are linear w.r.t. the state variable X , i.e., they are in the form

$$\begin{cases} b(x, \mu, a) &= b_0(\mu, a) + b_1(\mu, a)x, \\ \vartheta(x, \mu, a) &= \vartheta_0(\mu, a) + \vartheta_1(\mu, a)x, \\ \sigma(x, \mu, a) &= \gamma_0(\mu, a) + \gamma_1(\mu, a)x, \end{cases} \quad (3.9)$$

while the running and terminal cost functions are polynomial in the state variable in the sense that they are in the form (for simplicity we present here the one-dimensional case $n = 1$)

$$\begin{aligned} f(x, \mu, a) &= f_0(\mu, a) + \sum_{k=1}^p f_k(\mu, a)x^k, \\ g(x, \mu) &= g_0(\mu) + \sum_{k=1}^p g_k(\mu)x^k, \end{aligned}$$

for some integer $p \geq 1$.

Assumptions: dependence on μ : we assume that all the coefficients depend on μ through its first p moments, i.e., they are in the form

$$\begin{cases} b_0(\mu, a) = \bar{b}_0(\bar{\mu}_1, \bar{\mu}_2, \dots, \bar{\mu}_p, a), & b_1(\mu, a) = \bar{b}_1(\bar{\mu}_1, \bar{\mu}_2, \dots, \bar{\mu}_p, a) \\ \vartheta_0(\mu, a) = \bar{\vartheta}_0(\bar{\mu}_1, \bar{\mu}_2, \dots, \bar{\mu}_p, a), & \vartheta_1(\mu, a) = \bar{\vartheta}_1(\bar{\mu}_1, \bar{\mu}_2, \dots, \bar{\mu}_p, a) \\ \gamma_0(\mu, a) = \bar{\gamma}_0(\bar{\mu}_1, \bar{\mu}_2, \dots, \bar{\mu}_p, a), & \gamma_1(\mu, a) = \bar{\gamma}_1(\bar{\mu}_1, \bar{\mu}_2, \dots, \bar{\mu}_p, a) \\ f_k(\mu, a) = \bar{f}_k(\bar{\mu}_1, \bar{\mu}_2, \dots, \bar{\mu}_p, a), & g_k(\mu) = \bar{g}_k(\bar{\mu}_1, \bar{\mu}_2, \dots, \bar{\mu}_p), \quad k = 0, \dots, p, \end{cases} \quad (3.10)$$

where, given $\mu \in \mathcal{P}_p(\mathbb{R})$, we denote by

$$\bar{\mu}_k = \int x^k \mu(dx), \quad k = 1, \dots, p.$$

We assume that the coefficients $\bar{b}_0, \bar{b}_1, \bar{\vartheta}_0, \bar{\vartheta}_1, \bar{\gamma}_0, \bar{\gamma}_1$ are Lipschitz w.r.t. the p first arguments uniformly w.r.t. the control argument $a \in A$. This condition will ensure existence and uniqueness of a solution to the finite-dimensional MKV SDE defined later in (3.11).

Notice that in this case, the functions \hat{f} and \hat{g} defined in (3.3)-(3.4) are given by

$$\begin{aligned} \hat{f}(\mu, a) &= \bar{f}_0(\bar{\mu}_1, \bar{\mu}_2, \dots, \bar{\mu}_p, a) + \sum_{k=1}^p \bar{f}_k(\bar{\mu}_1, \bar{\mu}_2, \dots, \bar{\mu}_p, a) \bar{\mu}_k \\ &=: \bar{f}(\bar{\mu}_1, \bar{\mu}_2, \dots, \bar{\mu}_p, a), \\ \hat{g}(\mu) &= \bar{g}_0(\bar{\mu}_1, \bar{\mu}_2, \dots, \bar{\mu}_p) + \sum_{k=1}^p \bar{g}_k(\bar{\mu}_1, \bar{\mu}_2, \dots, \bar{\mu}_p) \bar{\mu}_k \\ &=: \bar{g}(\bar{\mu}_1, \bar{\mu}_2, \dots, \bar{\mu}_p). \end{aligned}$$

Remark 3.1. In the multidimensional case, we should consider a class of multi-polynomial functions f and g of degree p in the form

$$f(x, \mu, a) = \sum_{|\mathbf{k}|=0}^p f_{\mathbf{k}} \left((\mu^{\mathbf{k}'})_{|\mathbf{k}'| \leq p}, a \right) x^{\mathbf{k}}, \quad g(x, \mu) = \sum_{|\mathbf{k}|=0}^p g_{\mathbf{k}} \left((\mu^{\mathbf{k}'})_{|\mathbf{k}'| \leq p} \right) x^{\mathbf{k}},$$

where we use multi-index notations $\mathbf{k} = (k_1, \dots, k_n) \in \mathbb{N}^n$, $|\mathbf{k}| = k_1 + \dots + k_n$, $x^{\mathbf{k}} = x_1^{k_1} \dots x_n^{k_n}$

for $x = (x_1, \dots, x_n) \in \mathbb{R}^n$ and

$$\mu^{\mathbf{k}} = \int_{\mathbb{R}^n} x^{\mathbf{k}} \mu(dx).$$

3.2.2 Markovian embedding

Given the controlled process $X = X^\alpha$ solution to the stochastic MKV dynamics (3.1), denote by

$$Y_t^{(k)} = \mathbb{E}[X_t^k | W^0], \quad k = 1, \dots, p.$$

To alleviate the notations, we assume that $n = 1$ (otherwise multi-indices should be used). From the linear/polynomial assumptions (3.9)-(3.10), by Itô's formula and taking conditional expectations, we can derive the dynamics of $(Y^{(1)}, Y^{(2)}, \dots, Y^{(p)})$ as

$$\begin{aligned} dY_t^{(k)} &= B_k(Y_t^{(1)}, Y_t^{(2)}, \dots, Y_t^{(p)}, \alpha_t) dt + \Sigma_k(Y_t^{(1)}, Y_t^{(2)}, \dots, Y_t^{(p)}, \alpha_t) dW_t^0, \\ Y_0^{(k)} &= x_0^k, \quad k = 1, \dots, p, \end{aligned} \quad (3.11)$$

where, by convention $y_0 = 1, y_{-1} = 0$,

$$\begin{aligned} B_k(y_1, y_2, \dots, y_p, a) &= k\bar{b}_0(y_1, \dots, y_p, a)y_{k-1} + k\bar{b}_1(y_1, \dots, y_p, a)y_k \\ &\quad + \frac{k(k-1)}{2}(\bar{\vartheta}_0(y_1, \dots, y_p, a))^2 y_{k-2} + \frac{k(k-1)}{2}(\bar{\vartheta}_1(y_1, \dots, y_p, a))^2 y_k \\ &\quad + k(k-1)\bar{\vartheta}_0(y_1, \dots, y_p, a)\bar{\vartheta}_1(y_1, \dots, y_p, a)y_{k-1} \\ &\quad + \frac{k(k-1)}{2}(\bar{\gamma}_0(y_1, \dots, y_p, a))^2 y_{k-2} + \frac{k(k-1)}{2}(\bar{\gamma}_1(y_1, \dots, y_p, a))^2 y_k \\ &\quad + k(k-1)\bar{\gamma}_0(y_1, \dots, y_p, a)\bar{\gamma}_1(y_1, \dots, y_p, a)y_{k-1}, \quad k = 1, \dots, p, \\ \Sigma_k(y_1, y_2, \dots, y_p, a) &= k(\bar{\gamma}_0(y_1, \dots, y_p, a)y_{k-1} + \bar{\gamma}_1(y_1, \dots, y_p, a)y_k), \quad k = 1, \dots, p, \end{aligned}$$

while the cost functional is written as

$$J(\alpha) = \mathbb{E}\left[\int_0^T \bar{f}(Y_t^{(1)}, Y_t^{(2)}, \dots, Y_t^{(p)}, \alpha_t) dt + \bar{g}(Y_T^{(1)}, Y_T^{(2)}, \dots, Y_T^{(p)})\right].$$

The MKV control problem is then reduced in this polynomial framework into a finite-dimensional control problem with \mathbb{F}^0 -adapted controlled variables $(Y^{(1)}, Y^{(2)}, \dots, Y^{(p)})$. In the next section, we describe three probabilistic numerical methods for solving finite-dimensional stochastic control problems and will apply in section 3.5 each of these methods to three examples arising from polynomial MKV control problems under partial observation and common noise.

3.3 Probabilistic numerical methods

In this section, we introduce our numerical methods for the resolution of the reduced problem (3.11)-(3.12).

Let us introduce the process Z^α , valued in \mathbb{R}^d , controlled by an adapted process α taking values in A , solution to

$$dZ_t^\alpha = b(Z_t^\alpha, \alpha_t) dt + \sigma_0(Z_t^\alpha, \alpha_t) dW_t^0, \quad Z_0^\alpha = z_0 \in \mathbb{R}^d,$$

and the performance measure

$$J(t, z, \alpha) = \mathbb{E} \left[\int_t^T f(Z_t^\alpha, \alpha_t) dt + g(Z_T^\alpha) \middle| Z_t^\alpha = z \right],$$

which assesses the average performance of the control.

Introduce now a time discretization $t_n = n\Delta t$, $n = 0, \dots, N$, $\Delta t = T/N$, and denote by $\mathcal{A}_{\Delta t}$ the space of discrete processes $(\alpha_{t_n})_{n=0}^{N-1}$ such that for all n , $n = 0, \dots, N-1$, α_{t_n} is $\mathcal{F}_{t_n}^0$ -measurable.

We can write the Euler approximation of the SDE governing the process $Z = Z^\alpha$, with $\alpha \in \mathcal{A}_{\Delta t}$ (to alleviate notations, we sometimes omit the dependence on α when there is no ambiguity, and keep the same notation Z for the discrete and continuous process)

$$Z_{t_{n+1}} = Z_{t_n} + b(Z_{t_n}, \alpha_{t_n})\Delta t + \sigma_0(Z_{t_n}, \alpha_{t_n})\Delta W_{t_n}^0, \quad (3.13)$$

where $\Delta W_{t_n}^0 \sim \mathcal{N}(0, \Delta t)$ is an increment of W^0 .

The discrete time approximation of $J(t_n, z, \alpha)$ is defined as:

$$J_{\Delta t}(t_n, z, \alpha) = \mathbb{E} \left[\sum_{k=n}^{N-1} f(Z_{t_k}, \alpha_{t_k})\Delta t + g(Z_{t_N}) \middle| Z_{t_n} = z \right],$$

where $\alpha \in \mathcal{A}_{\Delta t}$.

3.3.1 Value and Performance iteration

For $n = 0, \dots, N$, consider $V_{\Delta t}(t_n, z) = \sup_{\alpha \in \mathcal{A}_{\Delta t}} J_{\Delta t}(t_n, z, \alpha)$, the discrete time approximation of the value function at time t_n : $V(t_n, z) = \sup_{\alpha \in \mathcal{A}} J(t_n, z, \alpha)$. The dynamic programming principle states that $(V_{\Delta t}(t_n, \cdot))_{0 \leq n \leq N}$ is solution to the Bellman equation:

$$\begin{cases} V_{\Delta t}(t_N, z) = g(z) \\ V_{\Delta t}(t_n, z) = \sup_{a \in A} \left\{ f(z, a)\Delta t + \mathbb{E}_{n,z}^a [V_{\Delta t}(t_{n+1}, Z_{t_{n+1}})] \right\}, \quad n = N-1, \dots, 0, \end{cases} \quad (3.14)$$

where $\mathbb{E}_{n,z}^a[\cdot]$ denotes the expectation conditioned on the event $\{Z_{t_n} = z\}$ and when using the control $\alpha_{t_n} = a$ at time t_n . Observe that for $n = 0, \dots, N-1$, the equation (3.14) provides a backward procedure to recursively compute the $V_{\Delta t}(t_n, \cdot)$ if we know how to analytically compute the conditional expectations $\mathbb{E}_{n,z}^a[V_{\Delta t}(t_{n+1}, Z_{t_{n+1}})]$ for all $z \in \mathbb{R}^d$ and all control $a \in A$. We refer to the procedure in (3.14) as value iteration.

An alternative approach to compute $V_{\Delta t}(t_n, \cdot)$, for $n = 0, \dots, N-1$, is to notice that once again by the dynamic programming principle, it holds that $(V_{\Delta t}(t_n, \cdot))_{0 \leq n \leq N}$ is solution to the backward equation

$$\begin{cases} V_{\Delta t}(t_N, z) = g(z) \\ V_{\Delta t}(t_n, z) = \sup_{a \in A} \left\{ f(z, a)\Delta t + \mathbb{E}_{n,z}^a \left[\sum_{k=n+1}^{N-1} f(Z_{t_k}, \alpha_{t_k}^*(Z_{t_k}))\Delta t + g(Z_{t_N}) \right] \right\}, \quad n = N-1, \dots, 0, \end{cases} \quad (3.15)$$

where for $k = n+1, \dots, N-1$, the control $\alpha_{t_k}^*$ is the optimal control at time t_k defined as

follows:

$$\alpha_{t_k}^*(z) = \operatorname{argmax}_{a \in A} \left\{ f(z, a) \Delta t + \mathbb{E}_{k,z}^a \left[\sum_{\ell=k+1}^{N-1} f(Z_{t_\ell}^*, \alpha_{t_\ell}^*(Z_{t_\ell}^*)) \Delta t + g(Z_{t_N}^*) \right] \right\},$$

and where $(Z_{t_k}^*)_{n \leq k \leq N}$ is the process Z controlled by the following control α from time t_n to t_N :

$$\begin{cases} \alpha_{t_n} = a, \\ \alpha_{t_k} = \alpha_{t_k}^* \text{ for } n+1 \leq k \leq N-1. \end{cases}$$

For $n = 0, \dots, N-1$, the scheme (3.15) provides once again a backward procedure to compute $V_{\Delta t}(t_n, \cdot)$, assuming that we know how to analytically compute the conditional expectations $\mathbb{E}_{n,z}^a \left[\sum_{k=n+1}^{N-1} f(Z_{t_k}, \alpha_{t_k}^*(Z_{t_k})) \Delta t + g(Z_{t_N}) \right]$ for all $z \in \mathbb{R}^d$ and all control $a \in A$. We refer to the procedure in (3.15) as performance iteration.

Except for trivial cases, closed-form formulas for the conditional expectations appearing in the value and performance iteration procedures are not available, and they have to be approximated, which is the main difficulty when implementing both approaches to compute the value functions. In the next section, we discuss different ways to approximate conditional expectations and derive the corresponding estimations of the value functions $V_{\Delta t}(t_n, \cdot)$ for $n = 0, \dots, N-1$.

3.3.2 Approximation of conditional expectations

In this subsection, we present three numerical methods that we apply later to conditional MKV problems. Two of these methods belong to the class of Regression Monte Carlo techniques, a family of algorithms whose effectiveness highly relies on the choice of the basis functions used to approximate conditional expectations; the third algorithm, Quantization, approximate the controlled process $Z_{t_n}^\alpha$ with a particular finite state Markov chain for which expectations can be approximated quickly.

Regression Monte Carlo

In the simpler uncontrolled case, the family of Regression Monte Carlo algorithms is based on the idea of approximating the conditional expectation $\mathbb{E}[V_{\Delta t}(t_{n+1}, Z_{t_{n+1}}) | Z_{t_n}]$, for $n = 0, \dots, N-1$, by the orthogonal projection of $V_{\Delta t}(t_{n+1}, Z_{t_{n+1}})$ onto the space generated by a finite family of $\{\phi_k(Z_{t_n})\}_{k \geq 1}$ where $(\phi_k)_{k \geq 1}$ is a family of *basis functions*, i.e., a family of measurable real-valued functions defined on \mathbb{R}^d such that $(\phi_k(Z_{t_n}))_{k \geq 1}$ is total in $L^2(\sigma(Z_{t_n}))$ ¹ and such that for all scalars β_k and all $K \geq 1$, if $\sum_{k=1}^K \beta_k \phi_k(Z_{t_n}) = 0$ a.s. then $\beta_k = 0$, for $k = 1, \dots, K$.

The expectation $\mathbb{E}[V_{\Delta t}(t_{n+1}, Z_{t_{n+1}}) | Z_{t_n}]$ should then be approximated as follows:

$$\mathbb{E}[V_{\Delta t}(t_{n+1}, Z_{t_{n+1}}) | Z_{t_n}] \approx \sum_{k=1}^K \beta_k^n \phi_k(Z_{t_n}),$$

¹ $L^2(\sigma(Z_{t_n}))$ is the space of the square-integrable $\sigma(Z_{t_n})$ -measurable r.v.

where $K \geq 1$ is fixed and $\beta^n = (\beta_1^n, \dots, \beta_K^n)^T$ is defined as:

$$\beta^n = \operatorname{argmin}_{\beta \in \mathbb{R}^K} \left\{ \mathbb{E} \left[\left| V_{\Delta t}(t_{n+1}, Z_{t_{n+1}}) - \sum_{k=1}^K \beta_k \phi_k(Z_{t_n}) \right|^2 \right] \right\}. \quad (3.16)$$

Notice that β^n is defined in (3.16) as the minimizer of a quadratic function, and can then be rewritten by straightforward calculations as:

$$\beta^n = \mathbb{E} [\phi(Z_{t_n}) \phi(Z_{t_n})^T]^{-1} \mathbb{E} [V_{\Delta t}(t_{n+1}, Z_{t_{n+1}}) \phi(Z_{t_n})], \quad (3.17)$$

where we use the notation $\phi = (\phi_1, \dots, \phi_K)^T$, and where we assumed that $\mathbb{E} [\phi(Z_{t_n}) \phi(Z_{t_n})^T]$ is invertible.

In order to estimate a solution to (3.17) we rely on Monte Carlo simulations to approximate expectations with finite sums. Consider the training set $\{(Z_{t_n}^m, Z_{t_{n+1}}^m)\}_{m=1}^M$ at time t_n obtained by running $M \geq 1$ forward simulations of the process Z from time $t_0 = 0$ to t_{n+1} . β^n defined in (3.17) can then be estimated by

$$\hat{\beta}^n = \left(\hat{\mathcal{A}}_n^M \right)^{-1} \frac{1}{M} \sum_{m=1}^M V_{\Delta t}(t_{n+1}, Z_{t_{n+1}}^m) \phi(Z_{t_n}^m),$$

where we denote by $\hat{\mathcal{A}}_n^M$ the estimator $\frac{1}{M} \sum_{m=1}^M \phi(Z_{t_n}^m) \phi(Z_{t_n}^m)^T$ of the covariance matrix $\mathcal{A}_n = \mathbb{E} [\phi(Z_{t_n}) \phi(Z_{t_n})^T]$.

The procedure presented above offers a convenient mean to approximate conditional expectations when the dynamics of the process Z are uncontrolled. When controlled, however, one has to account for the effect of the control on the conditional expectations either explicitly, via Control Randomization, or implicitly, via Regress-Later.

Control Randomization

In order to explicitly account for the effect of the control, one could directly introduce dependence on the control in the basis function. This basic idea of Control Randomization consists in replacing in the dynamics of Z the endogenous control by an exogenous control $(I_{t_n})_{0 \leq n \leq N}$, as introduced in [KLP14]. Its trajectories can then be simulated from time t_0 to time t_N . Consider the training set $\{Z_{t_n}^m, I_{t_n}^m\}_{n=0, m=1}^{N, M}$, with $M \geq 1$, where $I_{t_n}^m$ are i.i.d. samples from a “training distribution” μ_n with support in A . The training set will be used to estimate the optimal β^n coefficients for $n = 0, \dots, N - 1$. In the case of value iteration, $\{V_{\Delta t}(t_{n+1}, Z_{t_{n+1}}^m)\}_{m=1}^M$ is regressed against basis functions (which are, in this context, functions of the state and the control) evaluated at the points $\{Z_{t_n}^m, I_{t_n}^m\}_{m=1}^M$, as follows:

$$\mathbb{E}_{n,z}^a [V_{\Delta t}(t_{n+1}, Z_{t_{n+1}})] \approx \sum_{k=1}^K \hat{\beta}_k^n \phi_k(z, a),$$

with

$$\begin{aligned} \hat{\beta}^n &= \operatorname{argmin}_{\beta \in \mathbb{R}^K} \left\{ \sum_{m=1}^M \left[V_{\Delta t}(t_{n+1}, Z_{t_{n+1}}^m) - \sum_{k=1}^K \beta_k \phi_k(Z_{t_n}^m, I_{t_n}^m) \right]^2 \right\} \\ &\approx \left(\hat{\mathcal{A}}_n^M \right)^{-1} \frac{1}{M} \sum_{m=1}^M \left[V_{\Delta t}(t_{n+1}, Z_{t_{n+1}}^m) \phi(Z_{t_n}^m, I_{t_n}^m) \right], \end{aligned}$$

and where $\phi = (\phi_1, \dots, \phi_K)^T$ and

$$\hat{\mathcal{A}}_n^M = \frac{1}{M} \sum_{m=1}^M \phi(Z_{t_n}^m, I_{t_n}^m) \phi(Z_{t_n}^m, I_{t_n}^m)^T \quad (3.18)$$

is an estimator of the covariance matrix $\mathcal{A}_n = \mathbb{E}[\phi(Z_{t_n}, I_{t_n}) \phi(Z_{t_n}, I_{t_n})^T]$.

Notice that the basis functions take state and action variables as input in the case of Control Randomization-based method, i.e., their domain is $\mathbb{R}^d \times A$. Also, observe that the estimated conditional expectation highly depends on the choice of the randomization for the control².

An optimal feedback control at time t_n given $Z_{t_n} = z$ is approximated by the expression (see Subsection 3.3.4 for more practical details on the computation of the argmax):

$$\hat{\alpha}_{t_n}(z) = \operatorname{argmax}_{a \in A} \left\{ f(z, a) \Delta t + \sum_{k=1}^K \hat{\beta}_k^n \phi_k(z, a) \right\}. \quad (3.19)$$

The value function at time t_n is then estimated using Control Randomization method and value iteration procedure as

$$\hat{V}_{\Delta t}^{\text{CR}}(t_n, z) = f(z, \hat{\alpha}_{t_n}(z)) \Delta t + \sum_{k=1}^K \hat{\beta}_k^n \phi_k(z, \hat{\alpha}_{t_n}(z)), \quad z \in \mathbb{R}^d.$$

Notice that Control Randomization can be easily employed in a performance iteration procedure by computing controls (3.19), keeping in mind that at each time t_n we need to re-simulate new trajectories $\{\tilde{Z}_{t_k}^m\}_{k=n, m=1}^{N, M}$ iteratively from the initial condition $\tilde{Z}_{t_n}^m = z$, using the estimated optimal strategies $(\hat{\alpha}_{t_k})_{k=n+1}^{N-1}$ to compute the quantities $\sum_{k=n}^{N-1} f(t_k, \tilde{Z}_{t_k}^m, \hat{\alpha}_{t_k}(\tilde{Z}_{t_k}^m)) + g(\tilde{Z}_{t_N}^m)$, for $1 \leq m \leq M$.

Regress-Later

We present now a regress-later idea in which conditional expectation with respect to Z_{t_n} is computed in two stages. First, a conditional expectation with respect to $Z_{t_{n+1}}$ is approximated in a regression step by a linear combination of basis functions of $Z_{t_{n+1}}$. Then, analytical formulas are applied to condition this linear combination of functions of future values on the present value Z_{t_n} . For further details see [GY04], [NMS17] or [BP17]. With this approach, the effect of the control is factored in implicitly, through its effect on the (conditional) distribution of $Z_{t_{n+1}}$ conditioned on Z_{t_n} .

Unlike the traditional Regress-Now method for approximating conditional expectations (which we discussed so far in the uncontrolled case and in Control Randomization), the Regress-Later approach, as studied in [BP17], imposes conditions on basis functions:

Assumption 3.2. For each basis function ϕ_k , $k = 1, \dots, K$, the conditional expectation

$$\hat{\phi}_k^n(z, a) = \mathbb{E}_{n, z}^a[\phi_k(Z_{t_{n+1}})]$$

can be computed analytically.

Using the Regress-Later approximation of the conditional expectation and recalling Assumption 3.2 we obtain the optimal control $\alpha_{t_n}^m$ corresponding to the point $Z_{t_n}^m$, sampled independently

²Basically, different randomized controls may drive the process Z to very different locations, and the estimations will suffer from inaccuracy on the states that have been rarely visited.

from a “training distribution” μ_n (see Subsection 3.3.3 for further details):

$$\alpha_{t_n}^m = \operatorname{argmax}_{a \in A} \left\{ f(Z_{t_n}^m, a) \Delta t + \sum_{k=1}^K \hat{\beta}_k^{n+1} \hat{\phi}_k^n(Z_{t_n}^m, a) \right\}.$$

Notice that we are able to exploit the linearity of conditional expectations because

$$\hat{\beta}^{n+1} = \operatorname{argmin}_{\beta \in \mathbb{R}^K} \left\{ \sum_{m=1}^M \left[V_{\Delta t}(t_{n+1}, Z_{t_{n+1}}^m) - \sum_{k=1}^K \beta_k \phi_k(Z_{t_{n+1}}^m) \right]^2 \right\} \quad (3.20)$$

is a constant once the training sets at times t_k , $k = n + 1, \dots, N$, are fixed.

The value function at time t_n , is then estimated using Regress-Later method and value iteration procedure as

$$\hat{V}_{\Delta t}^{\text{RL}}(t_n, Z_{t_n}^m) = f(Z_{t_n}^m, \alpha_{t_n}^m) \Delta t + \sum_{k=1}^K \hat{\beta}_k^{n+1} \hat{\phi}_k^n(Z_{t_n}^m, \alpha_{t_n}^m).$$

Notice that contrary to Control Randomization, Regress-Later does not require the training points to be distributed as $Z_{t_{n+1}}$ conditioned on Z_{t_n} because the projection (3.20) is an approximation to an expectation conditional to the measure μ_n which can be chosen freely to optimize the precision of the sample estimation. On the other hand Regress-Later, similarly to Control Randomization, can be easily employed in a performance iteration procedure by generating forward trajectories at each time step.

Remark 3.3. Recall that the Regress-Later method uses training points that are i.i.d at each time step and independent across time steps. Contrary to other Regression Monte Carlo approaches, Regress-Later does not require to use the information about the conditional distribution during the regression step as that is accounted for in the second step of the method, when conditional expectations are computed analytically.

Quantization

We propose in this section a quantization-based algorithm to numerically solve control problems. We may also refer to the latter as the Q-algorithm or Q in all the numerical examples considered in Section 3.5, where Q stands for Quantization. Let us first introduce some ingredients of Quantization, and then propose different ways of using them to approximate conditional expectations.

Let $(E, |\cdot|)$ be a Euclidean space. We denote by $\hat{\varepsilon}$ a L -quantizer of an E -valued random variable ε , that is a discrete random variable on a grid $\Gamma = \{e_1, \dots, e_L\} \subset E^L$ defined by

$$\hat{\varepsilon} = \operatorname{Proj}_{\Gamma}(\varepsilon) = \sum_{\ell=1}^L e_{\ell} 1_{\varepsilon \in C_{\ell}(\Gamma)},$$

where $C_1(\Gamma), \dots, C_L(\Gamma)$ are the Voronoi cells corresponding to Γ , i.e., they form a Borel partition of E satisfying

$$C_{\ell}(\Gamma) \subset \left\{ e \in E : |e - e_{\ell}| = \min_{j=1, \dots, L} |e - e_j| \right\},$$

and where $\operatorname{Proj}_{\Gamma}$ stands for the Euclidian projection on Γ .

The discrete law of $\hat{\varepsilon}$ is then characterized by

$$p_\ell = \mathbb{P}[\hat{\varepsilon} = e_\ell] = \mathbb{P}[\varepsilon \in C_\ell(\Gamma)], \quad \ell = 1, \dots, L.$$

The grid of points $(e_\ell)_{\ell=1}^L$ which minimizes the L^2 -quantization error $\|\varepsilon - \hat{\varepsilon}\|_2$ leads to the so-called optimal L -quantizer, and can be obtained by a stochastic gradient descent method, known as Kohonen algorithm or competitive learning vector quantization (CLVQ) algorithm, which also provides as a byproduct an estimation of the discrete law $(p_\ell)_{\ell=1}^L$. We refer to [PPP04] for a description of the algorithm, and mention that for the normal distribution, the optimal grids and the weights of the Voronoi tessellations are precomputed for dimension up to 10 and are available on the website <http://www.quantize.maths-fi.com>.

In practice, optimal grids of the Gaussian random variable $\mathcal{N}_1(0, 1)$ of dimension 1 with 25 to 50 points, have been used to solve the control problems considered in Section 3.5.

We now propose two ways to approximate conditional expectations. The first approximation belongs to the family of the constant piecewise approximation, and the other one is an improvement on the first one, where the continuity of the approximation w.r.t. the control variable is preserved.

In the sequel, assume that for $n = 0, \dots, N - 1$ we have a set Γ_n of points in \mathbb{R}^d that should be thought of as a training set used for estimating $V(t_n, \cdot)$. See Subsection 3.3.3 for more details on how to build Γ_n .

Piecewise constant interpolation

We assume here that we already have an estimate of $V_{\Delta t}(t_{n+1}, \cdot)$, the value function at time t_{n+1} , for $n \in \{0, \dots, N - 1\}$, and we denote by $\hat{V}_{\Delta t}^Q(t_{n+1}, \cdot)$ the estimate.

The conditional expectation is then approximated as

$$\mathbb{E}_{n,z}^a[\hat{V}_{\Delta t}^Q(t_{n+1}, Z_{t_{n+1}})] \approx \sum_{\ell=1}^L p_\ell \hat{V}_{\Delta t}^Q\left(t_{n+1}, \text{Proj}_{\Gamma_{n+1}}(G_{\Delta t}(z, a, e_\ell))\right), \quad \text{for } z \in \Gamma_n,$$

where:

- $G_{\Delta t}$ is defined, using the notations introduced in (3.13), as

$$G_{\Delta t}(z, a, \varepsilon) = z + b(z, a)\Delta t + \sigma_0(z, a)\sqrt{\Delta t} \varepsilon. \quad (3.21)$$

- $\text{Proj}_{\Gamma_n}(\cdot)$ stands for the Euclidean projection on Γ_n .
- $\Gamma = \{e_1, \dots, e_L\}$ and $\{p_\ell\}_{1 \leq \ell \leq L}$ are respectively the optimal L -quantizer and its associated sequence of weights of the exogenous noise ε . See above for more details.

An optimal feedback control at time t_n and point $z \in \Gamma_n$ is approximated by the expression (see Subsection 3.3.4 for more practical details on the computation of the argmax):

$$\hat{\alpha}_{t_n}^Q(z) = \operatorname{argmax}_{a \in A} \left\{ f(z, a)\Delta t + \sum_{\ell=1}^L p_\ell \hat{V}_{\Delta t}^Q\left(t_{n+1}, \text{Proj}_{\Gamma_{n+1}}(G_{\Delta t}(z, a, e_\ell))\right) \right\}.$$

The value function at time t_n , is then estimated using the piecewise constant approximation and value iteration procedure as

$$\hat{V}_{\Delta t}^Q(t_n, z) = f(Z_{t_n}^m, \hat{\alpha}_{t_n}^Q(z))\Delta t + \sum_{\ell=1}^L p_\ell \hat{V}_{\Delta t}^Q\left(t_{n+1}, \text{Proj}_{\Gamma_{n+1}}(G_{\Delta t}(z, \hat{\alpha}_{t_n}^Q(z), e_\ell))\right).$$

Remark 3.4. Clearly, the constant piecewise approximation can be easily extended to control problems of all dimensions $d \geq 1$. However the latter is, in most cases, not continuous w.r.t. the control variable since it remains constant on each Voronoi cells (see, e.g., Figure 3.1 p.114). As a result, the optimization process over the control space suffers from high instability and inaccuracy, which implies a poor estimation of the value function $V(t_n, \cdot)$.

Semi-linear interpolation Once again, we assume here that we already have $\widehat{V}_{\Delta t}^Q(t_{n+1}, \cdot)$, an estimate of the value function at time t_{n+1} , with $n \in \{0, \dots, N-1\}$, and wish to provide an estimation of the conditional expectation in the particular case where the controlled process lies in dimension $d=1$. Consider the following piecewise linear approximation of the conditional expectation, which is continuous w.r.t. the control variable a :

$$\mathbb{E}_{n,z}^a[\widehat{V}_{\Delta t}^Q(t_{n+1}, Z_{t_{n+1}})] \approx \sum_{\ell=1}^L p_{\ell} \left[\lambda_a^{e_{\ell},z} \widehat{V}_{\Delta t}^Q(t_{n+1}, z_+) + (1 - \lambda_a^{e_{\ell},z}) \widehat{V}_{\Delta t}^Q(t_{n+1}, z_-) \right], \quad (3.22)$$

for $z \in \Gamma_n$, where for all $\ell = 1, \dots, L$, z_- and z_+ are defined as follows:

- z_- and z_+ are the two closest states in Γ_{n+1} from $G_{\Delta t}(z, a, e_{\ell})$, such that $z_- < G_{\Delta t}(z, a, e_{\ell}) < z_+$, if such states exist; and, in this case, we define $\lambda_a^{e_{\ell},z} = \frac{G_{\Delta t}(z, a, e_{\ell}) - z_-}{z_+ - z_-}$.
- Otherwise, z_- and z_+ are equal and defined as the closest state in Γ_{n+1}^Z from $G_{\Delta t}(z, a, e_{\ell})$ and we define $\lambda_a^{e_{\ell},z} = 1$.

Remark 3.5. This second approximation is continuous w.r.t. the control variable, which brings stability and accuracy to the optimal control task (see Subsection 3.3.4), and also ensures an accurate estimate of the value function at time t_n . We will mainly use this approximation in the numerical tests (see Section 3.5).

Remark 3.6. Although the dimension $d = 1$ plays a central role to define clearly the states z_- and z_+ in (3.22), the semi-linear approximation can actually be generalized to a certain class of control problems of dimension greater than 1, using multi-dimensional Quantization (see, e.g., the comments on the Q-algorithm designed to solve the Portfolio Optimization example, in Subsection 3.5.1). However, it is not well-suited to solve numerically general control problems in dimension greater than 1. For these cases, other interpolating methods such as the use of Gaussian processes are more appropriated (see, e.g., [LM18] for an introduction on the use of Gaussian processes in Regression Monte Carlo).

The optimal feedback control at time t_n and point $z \in \Gamma_n$ is approximated as (see Subsection 3.3.4 for more practical details on the computation of the argmax):

$$\hat{\alpha}_{t_n}^Q(z) = \operatorname{argmax}_{a \in A} \left\{ f(z, a) \Delta t + \sum_{\ell=1}^L p_{\ell} \left[\lambda_a^{e_{\ell},z} \widehat{V}_{\Delta t}^Q(t_{n+1}, z_+) + (1 - \lambda_a^{e_{\ell},z}) \widehat{V}_{\Delta t}^Q(t_{n+1}, z_-) \right] \right\}.$$

The value function at time t_n is then estimated using the semi-linear approximation and value iteration procedure as

$$\widehat{V}_{\Delta t}^Q(t_n, z) = f(z, \hat{\alpha}_{t_n}^Q(z)) \Delta t + \sum_{\ell=1}^L p_{\ell} \left[\lambda_{\hat{\alpha}_{t_n}^Q(z)}^{e_{\ell},z} \widehat{V}_{\Delta t}^Q(t_{n+1}, z_+) + (1 - \lambda_{\hat{\alpha}_{t_n}^Q(z)}^{e_{\ell},z}) \widehat{V}_{\Delta t}^Q(t_{n+1}, z_-) \right],$$

where z_+ and z_- are defined using the control $\hat{\alpha}_{t_n}^Q(z)$. See (3.22) for their definitions.

3.3.3 Training points design

We discuss here the choice of the training measure μ and the sets $(\Gamma_n)_{n=0,\dots,N-1}$ used to compute the numerical approximations in Regression Monte Carlo and Quantization. Two cases are considered in this section. The first one is a knowledge-based selection, relevant when the controller knows with a certain degree of confidence where the process has to be driven in order to optimize her reward functional. The second case, on the other hand, is when the controller has no idea where or how to drive the process to optimize the reward functional.

Exploitation only strategy

In the knowledge-based setting there is no need for exhaustive and expensive (in time mainly) exploration of the state space, and the controller can directly choose training sets Γ constructed from distributions μ that assign more points to the parts of the state space where the optimal process is likely to be driven.

In practice, at time t_n , assuming we know that the optimal process is likely to stay in the ball centered around the point m_n and with radius r_n , we chose a training measure μ_n centered around m_n as, for example $\mathcal{N}(m_n, r_n^2)$, and build the training set as sample of the latter. In the Regress-Later setting this can be done straightforwardly, while Control Randomization requires one to select a measure for the random control such that the controlled process Z is driven in such area of the state space.

Taking samples according to μ to build grids makes them random. Another choice, which we used in the Quantization-based algorithm, is to use the (deterministic) optimal grid of $\mathcal{N}(m_n, \sigma_n^2)$ with reduced size (typically take 50 points for a problem in dimension 1, 250 for one of dimension 2 when $\sigma_n^2 = 1, \dots$), which can be found at www.quantize.maths-fi.com, to reduce the size of the training set and alleviate the complexity of the algorithms.

Remark 3.7. As the reader will see, we chose the training sets based on the “exploitation only strategy” procedure, i.e. by guessing where to drive optimally the process, when solving the Liquidation Problem introduced in Subsection 3.4.1.

Explore first, exploit later

Explore first: If the agent has no idea of where to drive the process to receive large rewards, she can always proceed to an exploration step to discover favorable subsets of the state space. To do so, the Γ_n , for $n = 0, \dots, N - 1$, can be built as uniform grids that cover a large part of the state space, or μ can be chosen uniform on such domain. It is essential to explore far enough to have a well understanding of where to drive and where not to drive the process.

Exploit later: The estimates for the optimal controls at time t_n , $n = 0, \dots, N - 1$, that come up from the *Explore first* step, are relatively good in the way that they manage to avoid the wrong areas of state space when driving the process. However, the training sets that have been used to compute the estimated optimal control are too sparse to ensure accuracy on the estimation. In order to improve the accuracy, the natural idea is to build new training sets by simulating M times the process using the estimates on the optimal strategy computed from the *Explore first* step, and then proceed to another estimation of the optimal strategies using the new training sets. This trick can be seen as a two steps algorithm that improves the estimate of the optimal control.

Remark 3.8. In Control Randomization, multiple runs of the method are often needed to obtain precise estimates, because the initial choice of the dummy control could drive the training

points far from where the optimal control would have driven them. In practice, after having computed an approximate policy backward in time, such policy is used to drive M simulations of the process forward in time, which in turn produce control paths that can be fed as a random controls in a new backward procedure, leading to more accurate results.

Remark 3.9. We applied the “explore first, exploit later” idea to solve the Portfolio Optimization problem introduced in Subsection 3.4.1.

3.3.4 Optimal control searching

Assume in this section that we already have the estimates $\widehat{V}_{\Delta t}(t_k, \cdot)$ for the value function at time t_k , for $k = n + 1, \dots, N$, and want to estimate $V(t_n, \cdot)$ the value function at time t_n .

The optimal control searching task consists in optimizing the function³

$$\widehat{Q}_n : (z, \cdot) \mapsto f(z, a)\Delta t + \widehat{\mathbb{E}}_{n,z}^a[\widehat{V}_{\Delta t}(t_{n+1}, Z_{t_{n+1}})]$$

over the control space A , for each $z \in \Gamma_n$, and where we denote by $\widehat{\mathbb{E}}_{n,z}^a[\widehat{V}_{\Delta t}(t_{n+1}, Z_{t_{n+1}})]$ an approximation of $\mathbb{E}_{n,z}^a[\widehat{V}_{\Delta t}(t_{n+1}, Z_{t_{n+1}})]$ using Regress-Later, or Control Randomization or Quantization-based methods (see Subsection 3.3.2). Once again, we remind that importance of this task is motivated by the dynamic programming principle stating that for all $n = 0, \dots, N - 1$, we can approximate the value function at time n as follows

$$\widehat{V}_{\Delta t}(t_n, z) = \sup_{a \in A} \widehat{Q}_n(z, a),$$

where $\widehat{V}_{\Delta t}(t_n, \cdot)$ is our desired estimate of the value function at time n .

Low cardinality control set

In the case where the control space A is discrete (with a relatively small cardinality), one can solve the optimization problem by an exhaustive search over all the available controls without compromising the computational speed.

Remark 3.10. Note that in the case where the control space is continuous, one can always discretize the latter in order to rely on the effectiveness of extensive search to solve the optimal control problem. However, the control space discretization brings an error. So the control might have to include a high number of points in the discretization in order to reduce the error thereby causing a considerable slow down of the computations.

High cardinality/continuous control space

If we assume differentiability almost everywhere, as follows from the semi-linear approximation in Quantization, and most choices of basis functions in Regression Monte Carlo, we can carry on the optimization step by using some gradient-based algorithm for optimization of differentiable functions. Actually, many optimizing algorithms (Brent, Golden-section Search, Newton gradient-descent, ...) are already implemented in standard libraries of most programming languages like Python (see, e.g., package `scipy.optimize`), Julia (see, e.g., package `Optim.jl`), C and C++ (see, e.g., package `NLOpt`).

³often referred to as the Q -function, or action-value function, in the reinforcement learning literature. Be aware that Q stands here for the “Quality” of an action taken in a given state, and in particular does not refer to Quantization.

Remark 3.11. When the control is of dimension 1, polynomials of order smaller than 5 are employed as basis functions in Regression Monte Carlo as well as for the running reward f . The optimal control can then be computed analytically as a function of the regression coefficients, since every polynomial equation of order smaller than 4 can be solved by radicals.

Concretely, in all the examples considered in Section 3.5, we used the Golden-section Search or the Brent methods when testing Quantization-based algorithm to find the optimal controls at each point of the grids. These algorithms were very accurate to find the optimal controls, and we made use of Remark 3.11 to find the optimal controls using the Regress-Later-based algorithm.

3.3.5 Upper and lower bounds

After completing the backward procedure, we can compute an unbiased estimation of the value of the control policy by using Monte Carlo simulations and sample average. Assume already computed (or simply available) the matrix of regression coefficients, in the case of Regression Monte Carlo, and discrete probability law \hat{p} for Quantization, we can use this information to implicitly compute the control and simulate forward many trajectories of the controlled process starting from a common initial condition. We can then evaluate the average performance measure by computing the sample average of the rewards collected on each trajectory. Denoting such approximation by $\hat{V}_{\Delta t}(0, z)$, and recalling that by definition $J_{\Delta t}(0, z, \alpha) \leq J_{\Delta t}(0, z, \alpha^*)$, for all $\alpha \in \mathcal{A}_{\Delta t}$ and where α^* represents the optimal control process; it holds $\hat{V}_{\Delta t}(0, z) \leq V_{\Delta t}(0, z)$, for $z \in \mathbb{R}^d$.

The argument above implies that, neglecting the time-discretization error, we obtain a lower bound for $V_{\Delta t}(0, \cdot)$ by evaluating the estimated policy. On the other hand, see [Lab17], based on [Rog02], to get an upper bound of the value function via duality.

3.3.6 Pseudo-codes

In this section, we present the pseudo-code for the three approaches presented in the previous sections. For simplicity, we will only show the algorithms designed using value iteration procedure. However, the performance iteration update rule can be substituted in the codes below provided that forward simulations are run to obtain a pathwise realization of the controlled process and associated rewards.

Pseudo-code for a Regress-Later-based algorithm

We present in Algorithm 1 a pseudo-code to estimate $V_{\Delta t}(t_n, \cdot)$, for $n = 0, \dots, N - 1$, using Value Iteration and based on Regress-Later method. For $n = 0, \dots, N - 1$, we denote by $\hat{V}_{\Delta t}^{\text{RL}}(t_n, \cdot)$ the derived estimation of $V_{\Delta t}(t_n, \cdot)$, and will refer to it as the RLMC algorithm in the numerical tests presented in Section 3.5.

Note that we use the same training measure μ at each time step so that there is only one covariance matrix to estimate (since \mathcal{A}_{t_n} is the same for all $n = 0, \dots, N - 1$). Denote by $\hat{\mathcal{A}}^M$ the estimator, as defined in (3.18).

Pseudo-code for a Control Randomization-based algorithm

We present in Algorithm 2 a pseudo-code to estimate $V_{\Delta t}(t_n, \cdot, \cdot)$, for $n = 0, \dots, N - 1$, using Value Iteration and based on Control Randomization method. For $n = 0, \dots, N - 1$, we denote

Algorithm 1 Regress-Later Monte Carlo algorithm (RLMC) - Value iteration

Inputs:

- M : number of training points,
- μ : distribution of training points,
- K : number of basis functions,
- $\{\phi_k\}_{k=1}^K$: family of basis functions,

- 1: Estimate the covariance matrix $\hat{\mathcal{A}}^M$
- 2: Generate i.i.d. training points $\{Z_{t_N}^m\}_{m=1}^M$ accordingly to the distribution μ .
- 3: Initialize the value function $\hat{V}_{\Delta t}^{\text{RL}}(t_N, Z_{t_N}^m) = g(Z_{t_N}^m), \quad \forall m = 1, \dots, M$
- 4: **for** $n = N - 1$ to 0 **do**
- 5: $\hat{\beta}^n = \hat{\mathcal{A}}_M^{-1} \frac{1}{M} \sum_{m=1}^M \left[\hat{V}_{\Delta t}^{\text{RL}}(t_{n+1}, Z_{t_{n+1}}^m) \phi(Z_{t_{n+1}}^m) \right]$
- 6: Generate a new layer of i.i.d. training points $\{Z_{t_n}^m\}_{m=1}^M$ accordingly to the distribution μ .
- 7: **For all** $m = 1, \dots, M$ **do**

$$\hat{V}_{\Delta t}^{\text{RL}}(t_n, Z_{t_n}^m) = \sup_{a \in A} \left\{ f(Z_{t_n}^m, a) \Delta t + \sum_{k=1}^K \hat{\beta}_k^n \hat{\phi}_k^n(Z_{t_n}^m, a) \right\}$$

- 8: **Evaluate the policy to obtain** $\hat{V}_{\Delta t}^{\text{RL}}$

Outputs: $\{\hat{\beta}_k^n\}_{n,k=1}^{N,K}, \hat{V}_{\Delta t}^{\text{RL}}(0, z)$ for $z \in \mathbb{R}^d$.

by $\hat{V}_{\Delta t}^{\text{CR}}(t_n, \cdot)$ the derived estimation of $V_{\Delta t}(t_n, \cdot)$, and will refer to it as the CR algorithm in the numerical tests presented in Section 3.5.

Pseudo-code for a Quantization-based algorithm

We present in Algorithm 3 a pseudo-code to estimate $V_{\Delta t}(t_n, \cdot, \cdot)$, for $n = 0, \dots, N - 1$, using value iteration procedure and based on Quantization method. For $n = 0, \dots, N - 1$, we denote by $\hat{V}_{\Delta t}^{\text{Q}}(t_n, \cdot)$ the derived estimation of $V_{\Delta t}(t_n, \cdot)$, and will refer to it as the Q-algorithm in the numerical tests presented in Section 3.5.

Note that we made use of a piecewise constant approximation of conditional expectations to approximate $\hat{V}_{\Delta t}^{\text{Q}}(t_n, \cdot)$ in order to keep the algorithm simple. Also, note that, as said previously, in most of the numerical tests run in Section 3.5, we will use optimal grids available at www.quantize.maths-fi.com and will take $L = 25$ to 50 points for the size of the optimal grid of the Gaussian noise ε .

3.4 Applications

3.4.1 Portfolio Optimization under drift uncertainty

We consider a financial market model with one risk-free asset, assumed to be equal to one, and d risky assets of price process $S = (S^1, \dots, S^d)$ governed

$$dS_t = \text{diag}(S_t)(\beta_t dt + \sigma dB_t^0), \quad S_0 = s_0 \in \mathbb{R}^d,$$

Algorithm 2 Control Randomization algorithm (CR) - Value iteration**Inputs:**

- M : number of training points,
- μ : initial distribution of dummy control,
- K : number of basis functions,
- $\{\phi_k\}_{k=1}^K$: family of basis functions,

- 1: Estimate the covariance matrix $\hat{\mathcal{A}}^M$
- 2: Generate m trajectories, $\{Z_{t_n}^m, I_{t_n}^m\}_{n=0, m=1}^{N, M}$, where $Z_{t_n}^m$ is driven by $I_{t_n}^m$, and the $I_{t_n}^m$ are i.i.d with distribution μ .
- 3: Initialize the value function $\hat{V}_{\Delta t}^{\text{CR}}(t_N, Z_{t_N}^m) = g(Z_{t_N}^m)$, $m = 1, \dots, M$
- 4: **for** $n = N - 1$ to 0 **do**
- 5: $\hat{\beta}^n = (\hat{\mathcal{A}}^M)^{-1} \frac{1}{M} \sum_{m=1}^M \left[\hat{V}_{\Delta t}^{\text{CR}}(t_{n+1}, Z_{t_{n+1}}^m) \phi(Z_{t_n}^m, I_{t_n}^m) \right]$
- 6: **For all** $m = 1, \dots, M$ **do**

$$\hat{V}_{\Delta t}^{\text{CR}}(t_n, Z_{t_n}^m) = \sup_{a \in A} \left\{ f(Z_{t_n}^m, a) \Delta t + \sum_{k=1}^K \hat{\beta}_k^n \phi_k(Z_{t_n}^m, a) \right\}$$

- 7: **Evaluate the policy to obtain** $\hat{V}_{\Delta t}^{\text{CR}}$

Outputs: $\{\hat{\beta}_k^n\}_{n=0, k=1}^{N, K}$, $\hat{V}_{\Delta t}^{\text{CR}}(0, z)$ for $z \in \mathbb{R}^d$.

Algorithm 3 Quantization algorithm (Q) - Value iteration**Inputs:**

- $\Gamma_k, k = 0, \dots, N$: grids of training points in \mathbb{R}^d ,
- $\Gamma = \{e_1, \dots, e_L\}$, $(p_\ell)_{1 \leq \ell \leq L}$: the L-optimal grid of the exogenous noise ε , and its associated weights.

- 1: Initialize the estimated value function at time N : $\hat{V}_{\Delta t}^{\text{Q}}(t_N, z) = g(z)$, $\forall z \in \Gamma_N$.
- 2: **for** $n = N - 1$ to 0 **do**
- 3: **Estimate the value function at time** t_n **as follows:**

$$\hat{V}_{\Delta t}^{\text{Q}}(t_n, z) = \max_{a \in A} \left[f(z, a) \Delta t + \sum_{\ell=1}^L p_\ell \hat{V}_{\Delta t}^{\text{Q}}(t_{n+1}, \text{Proj}_{\Gamma_{n+1}}(G_{\Delta t}(z, a, e_\ell))) \right], \quad \forall z \in \Gamma_n. \quad (3.23)$$

- 4: **Compute the optimal strategies** $\hat{\alpha}(t_n, z)$, $z \in \Gamma_n$, **as maximizer of (3.23):**

$$\hat{\alpha}(t_n, z) \in \operatorname{argmax}_{a \in A} \left[f(z, a) \Delta t + \sum_{\ell=1}^L p_\ell \hat{V}_{\Delta t}^{\text{Q}}(t_{n+1}, \text{Proj}_{\Gamma_{n+1}}(G_{\Delta t}(z, a, e_\ell))) \right], \quad \forall z \in \Gamma_n.$$

- 5: **Evaluate the policy to obtain** $\hat{V}_{\Delta t}^{\text{Q}}$

Outputs: $(\hat{\alpha}(t_n, z))_{z \in \Gamma_n, 0 \leq n \leq N-1}$, $(\hat{V}_{\Delta t}^{\text{Q}}(0, z))_{z \in \Gamma_0}$.

where B^0 is a d -dimensional Brownian motion on a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P}^0)$, σ is the $d \times d$ invertible matrix volatility coefficient, assumed to be known and constant. However, the drift (β_t) of the asset (which is typically a diffusion process governed by another independent Brownian motion B) is unknown and unobservable like the Brownian motion B^0 . The agent can actually only observe the stock prices S , and we denote by \mathbb{F}^S the filtration generated by the price process S , which should be view as the available information.

In this context, we shall consider two important classes of optimization problems in finance:

1. Portfolio Liquidation. We consider the problem of an agent (trader) who has to liquidate a large number y_0 of shares in some asset (we consider one stock, $d = 1$) within a finite time T , and faces execution costs and market price impact. In contrast with frictionless Merton problem, we do not consider mark-to-market value of the portfolio and instead consider separately the amount on the cash account and the inventory Y , i.e., the position or number of shares held at any time. The strategy of the agent is then described by a real-valued \mathbb{F}^S -adapted process α , representing the velocity at which she buys ($\alpha_t > 0$) or sells ($\alpha_t < 0$) the asset, and the inventory is thus given by

$$Y_t = y_0 + \int_0^t \alpha_u du, \quad 0 \leq t \leq T.$$

The objective of the trader is to minimize over α the total liquidation cost

$$J_1(\alpha) = \mathbb{E}^0 \left[\int_0^T \alpha_t (S_t + f(\alpha_t)) dt + \ell(Y_T) \right]$$

where $f(\cdot)$ is an increasing function with $f(0) = 0$, representing a temporary price impact, and $\ell(\cdot)$ is a loss function, i.e., a convex function with $\ell(0) = 0$, penalizing the trader when she does not succeed to liquidate all her shares.

2. Portfolio Selection. The set \mathcal{A} of portfolio strategies, representing the amount invested in the assets, consists in all \mathbb{F}^S -adapted processes α valued in some set A of \mathbb{R}^d , and satisfying $\int_0^T |\alpha_t|^2 dt < \infty$. The dynamics of wealth process $X = X^\alpha$ associated to a portfolio strategy α is then governed by

$$\begin{aligned} dX_t &= \alpha_t S_t^{-1} dS_t \\ &= \alpha_t \cdot \beta_t dt + \alpha_t^\top \sigma dB_t^0, \quad X_0 = x_0 \in \mathbb{R}, \end{aligned}$$

and as in Merton Portfolio Selection problem, the objective of the agent is to maximize over portfolio strategies the utility of terminal wealth

$$J_2(\alpha) = \mathbb{E}^0[U(X_T)],$$

where U is a utility function on \mathbb{R} , e.g., CARA function $U(x) = -\exp(-px)$, $p > 0$.

Let us show how one can reformulate the above problems into a McKean-Vlasov type problem under partial observation and common noise as described in Section 3.1. We first introduce the so-called probability reference \mathbb{P} , which makes the observation price process a martingale. Let us then define the process

$$Z_t = \exp \left(- \int_0^t \sigma^{-1} \beta_u dB_u^0 - \frac{1}{2} \int_0^t |\sigma^{-1} \beta_u|^2 du \right), \quad 0 \leq t \leq T,$$

which is a $(\mathbb{P}^0, \mathbb{F})$ -martingale (under suitable integrability conditions on β), and defines a

probability measure $\mathbb{P} \sim \mathbb{P}^0$ through its density: $\frac{d\mathbb{P}}{d\mathbb{P}^0}\Big|_{\mathcal{F}_t} = Z_t$, and under which the process

$$W_t^0 = B_t^0 + \int_0^t \sigma^{-1} \beta_u du, \quad 0 \leq t \leq T,$$

is a (\mathbb{P}, \mathbb{F}) -Brownian motion by Girsanov's theorem, and the dynamics of S is

$$dS_t = \text{diag}(S_t) \sigma dW_t^0.$$

Notice that $\mathbb{F}^S = \mathbb{F}^0$ the filtration generated by W^0 . We also denote by $L_t = 1/Z_t$ the (\mathbb{P}, \mathbb{F}) -martingale governed by

$$dL_t = L_t \sigma^{-1} \beta_t \cdot dW_t^0.$$

Next, we use Bayes formula and rewrite the gain (resp. cost) functionals of our two portfolio optimization problems as

$$\begin{aligned} J_1(\alpha) &= \mathbb{E}\left[\int_0^T L_t \alpha_t (S_t + f(\alpha_t)) dt + L_T \ell(Y_T)\right] \\ &= \mathbb{E}\left[\int_0^T \bar{L}_t^0 \alpha_t (S_t + f(\alpha_t)) dt + \bar{L}_T^0 \ell(Y_T)\right] \\ &= \mathbb{E}\left[\int_0^T \bar{L}_t^0 \alpha_t (\bar{S}_t^0 + f(\alpha_t)) dt + \bar{L}_T^0 \ell(\bar{Y}_T^0)\right] \\ J_2(\alpha) &= \mathbb{E}[L_T U(X_T)] = \mathbb{E}[\bar{L}_T^0 U(X_T)] = \mathbb{E}[\bar{L}_T^0 U(\bar{X}_T^0)] \end{aligned}$$

where $\bar{L}_t^0 = \mathbb{E}[L_t | W^0] = \int \ell \mathbb{P}_{L_t}^{W^0}(d\ell)$, $\bar{X}_t^0 = \mathbb{E}[X_t | W^0] = \int x \mathbb{P}_{X_t}^{W^0}(dx) = X_t$, $\bar{Y}_t^0 = \mathbb{E}[Y_t | W^0] = \int y \mathbb{P}_{Y_t}^{W^0}(dy) = Y_t$, $\bar{S}_t^0 = \mathbb{E}[S_t | W^0] = \int s \mathbb{P}_{S_t}^{W^0}(ds) = S_t$, and we used the law of conditional expectations and the fact that S , X and Y are \mathbb{F}^0 -adapted. This formulation of the functional J_1 (resp. J_2) fits into the MKV framework of Section 3.1 with state variables (X, L, β) (resp. (Y, S, L, β))

We now consider the particular case when β is an \mathcal{F}_0 -measurable random variable distributed according to some probability distribution $\nu(db)$: this corresponds to a Bayesian point of view when the agent's belief about the drift is modeled by a prior distribution. In this case, let us show how our partial observation problem can be embedded into a finite-dimensional full observation Markov control problem. Indeed, by noting that β is independent of the Brownian motion W^0 under \mathbb{P} , we have

$$\bar{L}_t^0 = \mathbb{E}\left[\exp\left(\sigma^{-1} \beta \cdot W_t^0 - \frac{1}{2} |\sigma^{-1} \beta|^2 t\right) | W^0\right] = F(t, W_t^0),$$

where

$$F(t, w) = \int \exp\left(\sigma^{-1} b \cdot w - \frac{1}{2} |\sigma^{-1} b|^2 t\right) \nu(db).$$

Hence, the functionals J_1 and J_2 can be written as

$$\begin{aligned} J_1(\alpha) &= \mathbb{E}\left[\int_0^T F(t, W_t^0) \alpha_t (S_t + f(\alpha_t)) dt + F(T, W_T^0) \ell(Y_T)\right] \\ J_2(\alpha) &= \mathbb{E}\left[F(T, W_T^0) U(X_T)\right]. \end{aligned}$$

We are then reduced to a $(\mathbb{P}, \mathbb{F}^0)$ -control problem with state variables (W^0, X) for problem (1)

and (W^0, S, Y) for problem (2) with dynamics under \mathbb{P} :

$$\begin{aligned} dS_t &= \text{diag}(S_t)\sigma dW_t^0, \quad S_0 = s_0 \in (\mathbb{R}_+)^d \\ dX_t &= \alpha_t^\top \sigma dW_t^0, \quad X_0 = 0 \\ dY_t &= \alpha_t dt, \quad Y_0 = y_0 \in \mathbb{R}_+. \end{aligned}$$

Remark 3.12. Another example of partial observation for the drift β is the case when it is modeled by a linear Gaussian process. This would lead to the well-known Kalman-Bucy filter, hence to a finite-dimensional control problem. However, for general unobserved drift process β , we fall into an infinite dimensional control problem involving the filter process.

3.4.2 Interbank systemic risk with partial observation

We consider the following model of systemic risk inspired by the model in [CFS15]. The log-monetary reserves of N banks lending to and borrowing from each other are governed by the system

$$dX_t^i = \frac{\kappa}{N} \sum_{j=1}^N (X_t^j - X_t^i) dt + \sigma X_t^i (\sqrt{1 - \rho^2} dW_t^i + \rho dW_t^0), \quad i = 1, \dots, N$$

where W^i , $i = 1, \dots, N$, are independent Brownian motions, representing the idiosyncratic risk of each bank, W^0 is a common noise independent of W^i , $\sigma > 0$ is given real parameter, $\rho \in [-1, 1]$, and where X_0^i , $i = 1, \dots, N$ are i.i.d.. The mean-reversion coefficient $\kappa > 0$ models the strength of interaction between the banks where bank i can lend to and borrow from bank j with an amount proportional to the difference between their reserves. In the asymptotic regime when $N \rightarrow \infty$, the theory of propagation of chaos implies that the reserve state X^i of individual banks become independent and identically distributed conditionally on the common noise W^0 , with a state governed by

$$dX_t = \kappa (\mathbb{E}[X_t|W^0] - X_t) dt + \sigma X_t (\sqrt{1 - \rho^2} dB_t + \rho dW_t^0)$$

for some Brownian motion B independent of W^0 .

Let us now consider a central bank, viewed as a social planner, who only observes the common noise and not the reserves of each bank, and can influence the strength of the interaction between the individual banks, through an \mathbb{F}^0 -adapted control process α_t . The reserve of the representative bank in the asymptotic regime is then driven by

$$dX_t = (\kappa + \alpha_t) (\mathbb{E}[X_t|W^0] - X_t) dt + \sigma X_t (\sqrt{1 - \rho^2} dB_t + \rho dW_t^0), \quad X_0 \sim X_0^1,$$

and we consider that the objective of the central bank is to minimize

$$J(\alpha) = \mathbb{E} \left[\int_0^T \left(\frac{1}{2} \alpha_t^2 + \frac{\eta}{2} (X_t - \mathbb{E}[X_t|W^0])^2 \right) dt + \frac{c}{2} (X_T - \mathbb{E}[X_T|W^0])^2 \right],$$

where $\eta > 0$ and $c > 0$ penalize the departure of the reserve from the average. This is a MKV control problem under partial observation, but notice that it does not belong to the class of linear quadratic (LQ) MKV problems due to the control α which appears in a multiplicative form with the state. However, it fits into our class of polynomial MKV problem, and can be embedded into standard control problem as follows: We set $\bar{X}_t = \mathbb{E}[X_t|W^0]$ and $Y_t = \mathbb{E}[(X_t - \bar{X}_t)^2|W^0]$. The

cost functional is then written as

$$J(\alpha) = \mathbb{E} \left[\int_0^T \left(\frac{1}{2} \alpha_t^2 + \frac{\eta}{2} Y_t \right) dt + \frac{c}{2} Y_T \right]$$

where the dynamics of \bar{X} and Y are governed by

$$\begin{aligned} d\bar{X}_t &= \sigma \rho \bar{X}_t dW_t^0, \quad \bar{X}_0 = x_0 = \mathbb{E}[X_0] \\ dY_t &= [(\sigma^2 - 2(\kappa + \alpha_t))Y_t + \sigma^2(1 - \rho^2)\bar{X}_t^2]dt + 2\rho\sigma Y_t dW_t^0, \quad Y_0 = \text{Var}(X_0). \end{aligned}$$

We have then reduced the problem to a $(\mathbb{P}, \mathbb{F}^0)$ -control problem in dimension two with state variables (\bar{X}, Y) , which is neither LQ, but can be solved numerically.

3.5 Numerical results

3.5.1 Portfolio Optimization

Let us now illustrate numerically the impact of uncertain Bayesian drift on the Portfolio Liquidation problem and the Portfolio Selection problem in dimension $d = 1$, by considering a Gaussian prior distribution $\beta \sim \nu = \mathcal{N}(b_0, \gamma_0^2)$, with $b_0 \in \mathbb{R}$ and $\gamma_0 > 0$. In this case, F is explicitly given by:

$$F(t, w) = \frac{\sigma}{\sqrt{\sigma^2 + \gamma_0^2 t}} \exp \left(\frac{1}{2(\sigma^2 + \gamma_0^2 t)} (-b_0^2 t + 2b_0 \sigma w + \gamma_0^2 w^2) \right).$$

1. Portfolio Liquidation Let us first consider the Portfolio Liquidation problem (1) with a linear price impact function $f(a) = \gamma a$, $\gamma > 0$, and a quadratic loss function $\ell(y) = \eta y^2$, $\eta > 0$. The optimal trading rate is given by (see [Pha16])

$$\alpha_t^* = -\frac{Y_t^*}{T-t+\gamma/\eta} + \frac{1}{2\gamma} \left(\frac{1}{T-t+\gamma/\eta} \int_t^T \mathbb{E}^0[S_u | \mathcal{F}_t^S] du - S_t \right)$$

where Y^* is the associated inventory with feedback control α^* : $dY_t^* = \alpha_t^* dt$, $Y_0^* = y_0$. Since we consider a Gaussian prior $\mathcal{N}(b_0, \gamma_0^2)$ for β , the optimal trading rate is explicitly given by

$$\alpha_t^* = -\frac{1}{T-t+\gamma/\eta} \left\{ Y_t^* + \frac{1}{2\gamma} \left[-\frac{1}{\gamma_0} \sqrt{\frac{\pi}{2}} e^{-\frac{b_0^2}{2\gamma_0^2}} \left(\text{erfi} \left(\frac{b_0 + \gamma_0^2(T-t)}{\sqrt{2}\gamma_0} \right) - \text{erfi} \left(\frac{b_0}{\sqrt{2}\gamma_0} \right) \right) + (T-t+\frac{\gamma}{\eta}) \right] S_t \right\},$$

where erfi is the imaginary error function, defined as:

$$\text{erfi}(x) = \frac{2}{\sqrt{\pi}} \int_0^x e^{t^2} dt.$$

Remark 3.13. In the particular case where the price process is a martingale, i.e., $b_0 = 0$, and in the limiting case when the penalty parameter η goes to infinity, corresponding to the final constraint $Y_T = 0$, we see that α_t^* converges to $-Y_t^*/(T-t)$, hence it becomes independent of the price process, and this leads to an explicit optimal inventory: $Y_t^* = y_0 \frac{T-t}{T}$ with constant trading rate $\alpha_t^* = -y_0/T$. We retrieve the well-known VWAP strategy obtained in [AC01].

We solve the problem numerically, taking $N = 100$ for the time discretization, and fixing the other parameters as follows: $\gamma=5$, $S_0=6$, $Y_0=1$, $\eta=100$ and $\sigma=0.4$. We run two sets of forward Monte Carlo simulations for $b_0 = 0.1$, $T = 1$ and $b_0 = -0.1$, $T = 0.5$ changing the value of γ_0 . We tested the Regress-Later Monte Carlo (RLMC), the Control Randomization (CR) and the

Quantization (Q) algorithms. In particular, we wanted to compare the performance of these algorithms with $(\alpha_{t_n}^*)_{n=0}^{N-1}$, where α^* , defined above, is the optimal strategy associated to the continuous-time Portfolio Liquidation problem. We refer to this discrete-time strategy as α^* (i.e., re-using the same notation), and we use Opt, or *continuous-time* optimal strategy when we want to stress the fact that this strategy is optimal for the *continuous-time* control problem, and not for the discrete time one. We also tested a benchmark strategy (Bench) which consists in liquidating the inventory at a constant rate $-y_0/T$. The test consisted in computing the estimates $\widehat{V}_{\Delta t}(t_0 = 0, S_0 = 6, Y_0 = 1)$ associated to the different algorithms.

We display the results obtained by the different algorithms in Table 3.1 and plot them in Figure 3.2. One can observe in Figure 3.2 that for $\Delta t = \frac{1}{100}$ the estimations $\widehat{V}_{\Delta t}(t_0 = 0, S_0 = 6, Y_0 = 1)$ of the value function $V_{\Delta t}(t_0 = 0, S_0 = 6, Y_0 = 1)$, provided by RLMC, CR or Q-based methods, are sometimes such that

$$\widehat{V}_{\Delta t}(t_0 = 0, S_0 = 6, Y_0 = 1) \leq \widehat{J}_{\Delta t}(t_0 = 0, S_0 = 6, Y_0 = 1, \alpha^*),$$

where $\widehat{J}_{\Delta t}(\cdot, \cdot, \cdot, \alpha^*)$ is a Monte Carlo estimate of $J_{\Delta t}(\cdot, \cdot, \cdot, \alpha^*)$ applying strategy $(\alpha_{t_n}^*)_{n=0}^{N-1}$ (see in Figure 3.2 the curve Opt). It means that RLMC, CR, or Q-based methods sometimes provide better estimations of the optimal strategy than α^* for the discrete time control problem. However, since under suitable conditions (see, e.g., [KLP15]), the optimal strategy for the discrete time control problem $\alpha_{\Delta t}^*$ converges toward α^* , i.e. we have $\alpha_{\Delta t}^* \xrightarrow{\Delta t \rightarrow 0} \alpha^*$, then it holds:

$$\widehat{J}_{\Delta t}(t_0 = 0, S_0 = 6, Y_0 = 1, \alpha^*) \xrightarrow{\Delta t \rightarrow 0} V(t_0 = 0, S_0 = 6, Y_0 = 1).$$

Figure 3.3 shows a sample of the inventory $(Y_t)_{t \in [0, T]}$ when the agent follows α^* and the Quantization algorithm. One can notice that given the chosen penalization parameters, it is optimal to short some stocks at terminal time. Finally, notice that the concavity of the curves comes from the fact that the running cost does not penalize the inventory. If so, we expect the curves of the inventory w.r.t. time to be convex, see, e.g., [GS13].

Discussion on the Regress Later and Control Randomization algorithms. The implementation of Regression Monte Carlo algorithms has required intense tuning and the use of the performance iteration technique introduced in Subsection 3.3.3 in order to obtain satisfactory results. Paramount is, in addition, the distribution chosen for the training points in Regress-Later and for the initial control in Control Randomization. The problem of finding the best set of data to provide to the backward procedure is similar in the two Regression Monte Carlo algorithms. However little study is available in the literature; for more details on this problem in the Regress-Later setting see [NMS17] and [BP17]. In the case of RL algorithm a training measure μ_n has been chosen in order to sufficiently explore the state space in the Y dimension, in particular we considered $\mu_n = \mathcal{U}[-0.5, 0.5 + \frac{T-t_n}{t_n}]$. Similarly for CR we seek a distribution of the random control such that the controlled process Y results in having a distribution similar to μ_n . In order to achieve such goal we follow the “explore first, exploit later” approach presented in Subsection 3.3.3 and use a perturbed version of the empirical distribution of the control (to avoid concentration of the training points) obtained at previous iteration of the method to determine the random control at next iteration of the method.

In order to choose the basis functions, we used the fact that we expect the value function to be convex in the Y dimension with minimum around the optimal inventory level and monotone in the S dimension. For RL algorithm we choose therefore the following set of basis functions: $\{s, y, y^2, sy, sy^2\}$, where we take the square function y^2 as a general approximator for convex functions around their minima (where we expect the measure μ_n to be concentrated). On the other hand, CR requires that we guess what the functional form of the conditional expectation of

the value function is with respect to the control process. Considering our argument on square function approximating general convex functions we choose to add the set $\{\alpha, \alpha^2, \alpha y, \alpha s\}$ to the set of basis functions used by RL.

Note that there is no need for time-consuming optimal control searching with such a choice of basis functions, as explained in Remark 3.11.

Finally note that we observed very high volatility in the quality of the policy estimated by control randomization. For this reason we estimated the policy 50 times, and report in Table 3.1 the results provided by the best performing one; increasing the number of training points further affects the variability only marginally.

Discussion on the Quantization algorithm implementation To numerically solve this example, we used the optimal grid of the Gaussian random variable with $L = 50$ points, denoted by Γ_L^ε , to define the grid⁴ $\Gamma_n^W = t_n \Gamma_L^\varepsilon$ that discretizes W_{t_n} , the Brownian motion at time t_n , and the grid $\Gamma_n^Y = Y_0 - \frac{t_n}{T} + t_n \Gamma_L^\varepsilon$ that discretizes Y_{t_n} , the inventory at time t_n , for $n = 0, \dots, N$. Note that Γ_n^Y , for $n = 0, \dots, N$, is centered at point $Y_0 - \frac{t_n}{T}$ because we guessed that the optimal liquidation rate was close to $\frac{Y_0}{T}$ (see Figure 3.3 to check that our guess is correct). We then considered the grid $\Gamma_n = \Gamma_n^W \times \Gamma_n^Y$ to discretize $Z_{t_n} = (W_{t_n}, Y_{t_n})$, $n = 0, \dots, N$.

We first tried to design a quantization algorithm using the following expression for the conditional expectation approximations:

$$\begin{aligned} \mathbb{E}_{n,(w,y)}^a \left[\widehat{V}_{\Delta t}^Q \left(t_{n+1}, \text{Proj}_{\Gamma_{n+1}^W} (W_{t_{n+1}}), \text{Proj}_{\Gamma_{n+1}^Y} (Y_{t_{n+1}}) \right) \right] \\ \approx \sum_{\ell=1}^L p_\ell \widehat{V}_{\Delta t}^Q \left(t_{n+1}, \text{Proj}_{\Gamma_{n+1}^W} (G_{\Delta t}((w, y), a, e_\ell)), \text{Proj}_{\Gamma_{n+1}^Y} (G_{\Delta t}((w, y), a, e_\ell)) \right), \quad (3.29) \\ \text{for } (w, y, a) \in \Gamma_n^W \times \Gamma_n^Y \times A, \end{aligned}$$

where the first and second components of the process $Z = (W, Y)$ are projected respectively on the grids Γ_n^W and Γ_n^Y ; and $\text{Proj}_{\Gamma_n^W}$ (resp. $\text{Proj}_{\Gamma_n^Y}$) stands for the Euclidean projection of the first (resp. second) component of $Z = (W, Y)$ on Γ_n^W (resp. Γ_n^Y).

This approximation belongs to the family of constant piecewise approximations, and is in the spirit of multidimensional component-wise-quantization methods already studied in the literature (see, e.g., [PS15]).

Unfortunately, as it can be seen in Figure 3.1, approximation (3.29) is discontinuous w.r.t. the control variable a in such a way that the optimal control searching task suffered from instability and inaccuracy, which implied bad value function estimations at time $n = 0, \dots, N - 1$. We thus had to use a better conditional expectation approximation.

We then decided to smooth the previous approximation of the conditional expectations w.r.t. the control variable by considering the following

$$\begin{aligned} \mathbb{E}_{n,(w,y)}^a \left[\widehat{V}_{\Delta t}^Q \left(t_{n+1}, \text{Proj}_{\Gamma_{n+1}^W} (W_{t_{n+1}}), \text{Proj}_{\Gamma_{n+1}^Y} (Y_{t_{n+1}}) \right) \right] \\ \approx \sum_{\ell=1}^L p_\ell \left[\lambda_a^{e_\ell, (w,y)} \widehat{V}_{\Delta t}^Q \left(t_{n+1}, \text{Proj}_{\Gamma_{n+1}^W} [G_{\Delta t}^w((w, y), a, e_\ell)], y_+ \right) \right. \\ \left. + (1 - \lambda_a^{e_\ell, (w,y)}) \widehat{V}_{\Delta t}^Q \left(t_{n+1}, \text{Proj}_{\Gamma_{n+1}^W} [G_{\Delta t}^w((w, y), a, e_\ell)], y_- \right) \right], \end{aligned}$$

⁴We use the notation $tB = \{tb, b \in B\}$, where $t \in \mathbb{R}$ and B is a set.

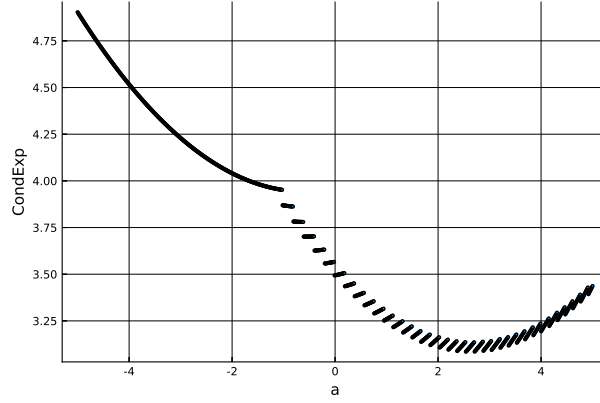


Figure 3.1 – Plot of the quantized-based piecewise-constant approximation of the conditional expectation CondExp:

$$a \mapsto \sum_{e \in \Gamma^\varepsilon} \mathbb{P}(\hat{\varepsilon} = e) \widehat{V}_{\Delta t}^Q \left(t_{n+1}, \text{Proj}_{\Gamma_{n+1}^W} \left(G_{\Delta t}((w, y), a, e) \right), \text{Proj}_{\Gamma_{n+1}^Y} \left(G_{\Delta t}((w, y), a, e) \right) \right).$$

We took $n = N - 1$, $w = 0$, and $y = -0.18$ to plot the curve. Observe that the approximation is discontinuous w.r.t. the control variable a in such a way that it makes the search of the minimizer of this function very difficult by usual (gradient descent-based) algorithms. Also, observe that the minimum of the function, which is actually equal to the estimation of the value function at time $N - 1$ at point $(w = 0, y = -0.18)$, suffers from inaccuracy.

where, in the spirit of the semi-linear approximation presented in Subsection 3.3.2, we have for all $\ell = 1, \dots, L$:

- $G_{\Delta t}^w((w, y), a, e_\ell)$ and $G_{\Delta t}^y((w, y), a, e_\ell)$ respectively stand for the first and the second component of $G_{\Delta t}((w, y), a, e_\ell)$, i.e., $G_{\Delta t}((w, y), a, e_\ell) = (G_{\Delta t}^w((w, y), a, e_\ell), G_{\Delta t}^y((w, y), a, e_\ell))$. See (3.21) for the definition of $G_{\Delta t}$.
- y_- and y_+ are the two closest states in Γ_{n+1}^Y from $G_{\Delta t}^y((w, y), a, e_\ell)$, such that $y_- < G_{\Delta t}^y((w, y), a, e_\ell) < y_+$ if such point exists; y_- and y_+ are equal to the closest state in Γ_{n+1}^Y from $G_{\Delta t}^y((w, y), a, e_\ell)$ otherwise.
- $\lambda_a^{e_\ell, (w, y)} = \frac{G_{\Delta t}^y((w, y), a, e_\ell) - y_-}{y_+ - y_-}$ in the first case of the definition of y_- and y_+ above; $\lambda_a^{e_\ell, (w, y)} = 1$ otherwise.

This approximation is a slight generalization (to dimension $d=2$) of the semi-linear approximation developed in (3.22). Its main interest lies in the continuity of the approximation w.r.t. the control variable a , which provides stability and accuracy to the usual (gradient descent-based) algorithms for the optimal controls searching, as can be seen on the numerical results (see, e.g., Table 3.1).

2. Portfolio Selection. Consider the Portfolio Selection problem with one risky asset. We choose a CARA utility function $U(x) = -\exp(-px)$, with $p > 0$. It has been shown in [GP16, Corollary 1] that the optimal portfolio strategy is explicitly given by

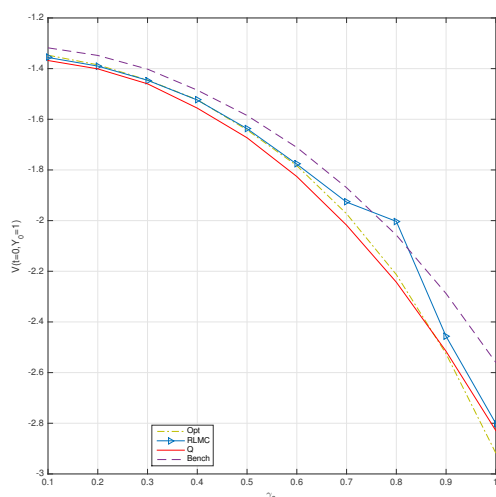
$$\alpha_t^* = \frac{\sigma^2 + \gamma_0^2 t}{\sigma^2 + \gamma_0^2 T} \frac{\hat{\beta}_t}{p\sigma^2}$$

where

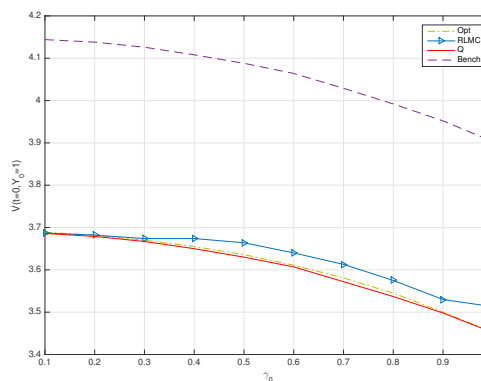
$$\hat{\beta}_t = \mathbb{E}^0[\beta | \mathcal{F}_t^S] = \frac{\sigma^2}{\sigma^2 + \gamma_0^2 t} b_0 + \frac{\gamma_0^2}{\sigma^2 + \gamma_0^2 t} \left(\ln \frac{S_t}{S_0} + \frac{1}{2} \sigma^2 t \right),$$

Table 3.1 – Portfolio Liquidation results. Estimations of the value functions at point $(s_0 = 6, y_0 = 1)$ and time 0 provided by different algorithms.

γ_0	$b_0 = 0.1, T = 1$					$b_0 = -0.1, T = 1/2$				
	Opt	RLMC	CR	Q	Bench	Opt	RLMC	CR	Q	Bench
0.1	-1.347	-1.356	-1.278	-1.368	-1.318	3.689	3.687	3.995	3.686	4.144
0.2	-1.385	-1.390	-1.283	-1.401	-1.348	3.682	3.682	3.847	3.679	4.138
0.3	-1.445	-1.446	-1.314	-1.460	-1.402	3.670	3.674	4.034	3.667	4.126
0.4	-1.523	-1.524	-1.323	-1.556	-1.485	3.655	3.674	4.128	3.650	4.108
0.5	-1.642	-1.637	-1.348	-1.673	-1.585	3.636	3.664	4.243	3.630	4.088
0.6	-1.783	-1.777	-1.425	-1.826	-1.711	3.611	3.640	4.386	3.607	4.064
0.7	-1.973	-1.927	-1.513	-2.018	-1.870	3.581	3.613	4.783	3.572	4.029
0.8	-2.213	-2.003	-1.637	-2.243	-2.057	3.545	3.575	5.142	3.537	3.992
0.9	-2.526	-2.457	-1.819	-2.516	-2.288	3.500	3.530	5.345	3.498	3.952
1	-2.918	-2.801	-1.806	-2.829	-2.560	3.453	3.513	6.765	3.452	3.903



$b_0 = 0.1$ and $T = 1$.



$b_0 = -0.1$ and $T = 0.5$.

Figure 3.2 – Results for the Portfolio Liquidation problem. Estimation of the value function at point $(s_0 = 6, y_0 = 1)$ at time 0 provided by different strategies w.r.t. γ_0 . We took $\gamma=5, S_0=6, Y_0=1, \eta=100$ and $\sigma=0.4$.

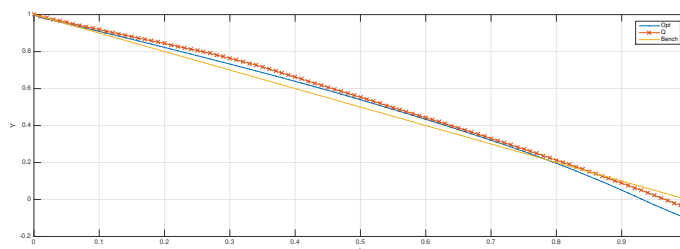


Figure 3.3 – Simulation of $(Y_t)_{t \in [0, T]}$ using the (continuous-time) optimal strategy (Opt), the (Q) estimated one, and the Benchmark strategy (Bench) to solve the Portfolio Liquidation problem. We took $T = 1, \sigma = 0.4, \gamma_0 = 1, b_0 = 0.1, S_0 = 6, Y_0 = 1, N = 100, \gamma = 5, \eta = 100$.

is the posterior mean of the drift (Bayesian learning on the drift), and the optimal performance by

$$J_2(\alpha^*) = -\exp \left[-p \left(x_0 + \frac{1}{2p} \left(\ln \left(\frac{\sigma^2 + \gamma_0^2 T}{\sigma^2} \right) - \frac{\gamma_0^2 T}{\sigma^2 + \gamma_0^2 T} \right) + \frac{b_0^2}{2p\sigma^2} \frac{\sigma^2 T}{\sigma^2 + \gamma_0^2 T} \right) \right].$$

The Portfolio Selection problem, even though in many aspects similar to the Portfolio Liquidation problem, is interesting in its own right because the control acts only on the variance of the controlled wealth process. We tested the Regress-Later Monte Carlo (RLMC), the Control Randomization (CR) and the Quantization (Q) algorithm on the Portfolio Selection problem. Similarly to what has been done for Portfolio Liquidation problem, we discretized time choosing $N = 100$ and solved the discrete time problem associated. We considered two set of experiments, $b_0 = 0.1, T = 1$ and $b_0 = -0.1, T = 0.5$, for different values of $\gamma_0 \in [0, 1], p = 1, \sigma = 0.4$. Given all these different parameters, we compared the performance of these algorithms with the one of the optimal strategy for the continuous-time problem α^* (Opt). The general test consists in computing a forward Monte Carlo with 500000 samples, following optimal strategy estimated using different strategies, to provide estimates of $V(t_0 = 0, X_0 = 0, W_0 = 0)$ the value function at time 0.

We present the results of our numerical experiments in Table 3.2. One can see that the Quantization algorithm performs similarly to the theoretical optimal strategy (Opt) for the continuous time problem, which can be interpreted as stability and accuracy of the Q algorithm, and also shows that the time discretization error is almost zero here.

We also present in Figure 3.4 a sample of the wealth of the agent following the optimal strategy and the (Q) estimated one. One can see that the strategies slightly differ when the drift is high, and remain the same when the drift is low. The small difference can be explained by the fact that the optimal strategy (Opt) is not optimal for the discrete time version of the problem.

Remarks on the Quantization algorithm. We designed the same Quantization algorithm as the one built to solve the Portfolio Liquidation problem. We nevertheless had to take a larger number of points in the grids to minimize the back-propagation of errors from the borders of the grids; and had to use the “explore first, exploit later” idea (see Subsection 3.3.3) to improve the results.

Remarks on the RL and CR algorithms implementation. When implementing Regression Monte Carlo algorithms, and choosing basis functions, the control on variance implies that low order polynomial can not be used alone, as they can easily cause the control to be bang-bang between the boundaries of its domain. Similarly, piecewise approximations are not very effective, as the dependence on the control is very weak, requiring a high number of local supports and making the computational complexity overwhelming. We tested both value and performance iteration and tried to employ different kinds of basis functions and training points. Unfortunately, both Regress-Later and Control Randomization do not cope well with controlling the dynamics of a process through the variance only. A tailor-made implementation of Regression Monte Carlo to deal with this kind of problems is outside the scope of this paper and further investigation will follow in future work. For now, we chose not to provide results based on RL and CR methods.

3.5.2 Interbank systemic risk

For this problem, in the absence of analytical solution, we decided to compare the estimations of the value function at time 0 provided by our algorithms with a numerical approximation based

Table 3.2 – Portfolio Selection results. Estimations of the value function at point $(x_0 = 0, S_0 = 6)$ time 0 using the *continuous-time* optimal strategy (Opt) and (Q) estimated optimal strategy.

γ_0	$b_0 = 0.1, T = 1$		$b_0 = -0.1, T = 0.5$	
	Opt	Q	Opt	Q
0.1	-0.985	-0.985	-0.992	-0.992
0.2	-0.982	-0.982	-0.991	-0.991
0.3	-0.973	-0.973	-0.988	-0.988
0.4	-0.954	-0.953	-0.981	-0.981
0.5	-0.927	-0.927	-0.969	-0.969
0.6	-0.896	-0.896	-0.952	-0.952
0.7	-0.863	-0.863	-0.932	-0.932
0.8	-0.830	-0.830	-0.910	-0.910
0.9	-0.797	-0.797	-0.886	-0.886
1	-0.767	-0.766	-0.863	-0.863

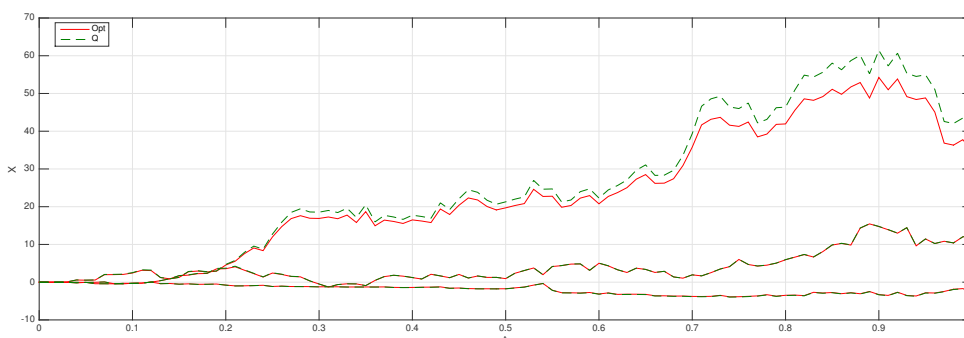


Figure 3.4 – 3 simulations of the agent's wealth $(X_t)_{t \in [0, T]}$ when the latter follows the continuous-time optimal strategy (Opt) and the (Q) estimated optimal strategy to solve the Portfolio Selection problem. We took $\sigma=0.4$, $T=1$, $P=0.1$, $\gamma_0=5$, $b_0=0,1$. One can see that the two strategies are the same when the drift is low; but Q performs slightly better than Opt when the drift is high, which is a time-discretization effect.

on finite difference scheme provided by Mathematica, of the solution to the 2-dimensional HJB equation associated to the systemic risk problem:

$$\begin{cases} \partial_t V + \frac{\eta}{2}y + \left((\sigma^2 - 2\kappa)y + \sigma^2(1 - \rho^2)x^2 \right) \partial_y V + \sup_{a \in A} \left[\frac{1}{2}a^2 - 2ay\partial_y V \right] \\ \quad + \frac{\sigma^2 \rho^2 x^2}{2} \partial_{xx}^2 V + 2\sigma^2 \rho^2 xy \partial_{xy} V + 2\sigma^2 \rho^2 y^2 \partial_{yy}^2 V = 0, & \text{for } (t, x, y) \in [0, T) \times \mathbb{R} \times \mathbb{R}_+, \\ V(T, x, y) = \frac{c}{2}y, & \forall (x, y) \in \mathbb{R} \times \mathbb{R}_+. \end{cases} \quad (3.30)$$

We refer to the solution of this partial differential equation (obtained using Mathematica using finite differences as explained below) as the Benchmark (or simply Bench) in the sequel.

We computed $\widehat{V}_{\Delta t}(t_0 = 0, x_0 = 10, y_0 = 0)$ using RL, CR and Q methods by considering a sample of size 500 000, and using the following parameters $T = 1$, $\sigma = 0.1$, $\kappa = 0.5$ and $X_0 = 10$. We recall that $\widehat{V}_{\Delta t}(t_0 = 0, x_0, y_0)$ is an estimation of $V(0, x_0 = 10, y_0 = 0)$, the value function at (x_0, y_0) and time 0 (see its definition on the last step of each pseudo-code presented in Subsection 3.3.6).

In Table 3.3 we display the numerical results of experiments run for two situations: we took $\eta = 10$, $c = 100$ and $\eta = 100$, $\rho = 0.5$ and vary the value of ρ in the first case, and vary the value of c in the second one. Plots of the two tables are also available in Figure 3.5. One can observe that the algorithms performs well. Mainly, Bench and Q provide slightly better results than the Regression Monte Carlo-based algorithms (the curves of Bench and Q are below those of the other two).

Figure 3.6 shows two examples of paths $(X_t)_{t \in [0, T]}$ controlled by RLMC (curve ‘‘RLMC’’), $(X_t)_{t \in [0, T]}$ naively controlled by $\alpha = 0$ (curve ‘‘uncontrolled’’), and the conditional expectation of X ($\bar{X}_t)_{t \in [0, T]}$ (curve ‘‘ $E(X|W)$ ’’). One can see in these two examples that the (RLMC estimated) optimal control is as follows:

- do nothing when the terminal time is far, i.e., take $\alpha = 0$, not to pay any running cost.
- catch \bar{X} when the terminal time is getting close, to minimize the terminal cost.

We finally present a sample of paths $(Y_t)_{t \in [0, T]}$ controlled by the decisions given by Q in Figure 3.7. One can see that the (Q estimated) optimal strategy minimizes the running cost first by letting Y grow; and deals with the terminal cost later by making Y small when the terminal time is approaching.

Remarks on the Regress Later and Control Randomization algorithms. For the implementation of the RL algorithm we decided to use polynomial basis functions up to degree 2. This choice allows us to compute the optimal control analytically as a function of the regression coefficients (see Remark 3.11). Compared to other optimization techniques, explicit expression allows for much faster and error-free computations (see Remark 3.11). For CR, we used basis functions up to degree 3 in all dimensions to obtain more stable results.

Regarding the choice of the training measure in RL, we employed marginal normal distributions on each dimension. As we know that the inventory dimension Y represents the conditional variance of the original process X , we centered the training distribution μ_n at zero but considered only training points $Y_n^m \geq 0$. In CR, on the other hand, we need to carefully choose the distribution of the random control so that the process Y does not become negative. Notice in fact that the Euler approximation, contrary to the original SDE describing Y , does not remain positive and we would therefore need to carefully choose a control to avoid driving Y negative. In order to achieve such goal, without having to worry too much about the control, we modified

the Euler approximation of (3.29) to feature a reflexive boundary at zero. Such features allow to train the estimated control policy to not overshoot when trying to drive the process Y to zero, without having Y to become negative.

Remarks on the Quantization algorithm. As stated above, it is straightforward that $Y > 0$ on $(0, T]$. However, the Euler scheme used to approximate the dynamics of Y does not prevent the associated process $(Y_{t_i})_{0 < i \leq N}$ to be non-positive. When implementing the Q algorithm for the systemic risk problem, we forced $(\text{Proj}_{\Gamma_i^Y}(Y_{t_i}))_{0 < i \leq N}$ to remain positive by simply choosing positive points for the grids Γ_i^Y that quantize the states of Y_{t_i} , at time t_i for $i = 0, \dots, N$.

Also, given the expression of the instantaneous and terminal reward, one can expect Y to stay close to 0, but we do not have any idea of how small Y should stay for the strategy to be optimal (cf. Figure 3.7 to see a posteriori where Y lies). To deal with this situation, we decided to adopt the “explore first, exploit later” procedure. First, we chose some random grids with a lot of points near 0 and computed the optimal strategy on these grids. Then, we ran forward Monte Carlo simulations and generated an empirical distribution of the quantized Y . Second, we build new grids of Quantization for Y by generating new points according to the empirical distribution that we got from in the previous step. Finally, we computed new (hopefully better) estimations of the optimal strategy by running the Q algorithm using the new grids. The Q strategy performed better after applying this step, but not significantly since our first naive guess for the grids (i.e., before bootstrapping) was already good enough.

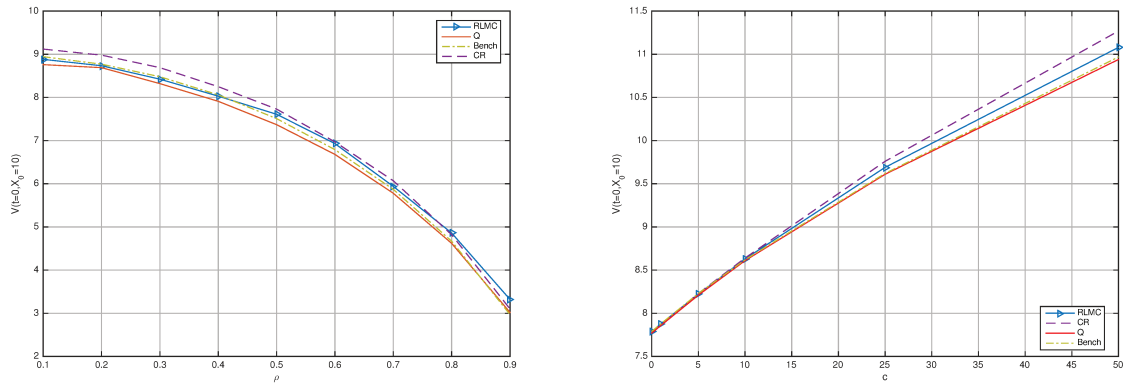
Details on the deterministic algorithm for the resolution of the HJB. We use the `NDSolve` function in Mathematica based on finite difference method to solve (3.30). Note that usually terminal and boundary conditions are required to get numerical results. The final condition: $V(T, x, y) = \frac{c}{2}y$ is already given by (3.30). However, the boundary conditions on $V(t, 0, y)$ and $V(t, x, 0)$ are missing, except the trivial condition consisting of $V(t, 0, 0) = 0$. We then provided the HJB without boundary conditions to the Mathematica function `NDSolve`, and let the latter add artificial boundary conditions by itself to output results.

ρ	RLMC	CR	Q	Bench	c	RLMC	CR	Q	Bench
0.1	8.88	9.12	8.76	8.94	0	7.79	7.78	7.77	7.79
0.2	8.73	8.98	8.69	8.77	1	7.88	7.87	7.86	7.88
0.3	8.42	8.69	8.32	8.48	5	8.22	8.23	8.21	8.23
0.4	8.02	8.25	7.91	8.06	10	8.63	8.64	8.61	8.62
0.5	7.61	7.73	7.37	7.51	25	9.69	9.76	9.61	9.62
0.6	6.93	6.97	6.68	6.79	50	11.08	11.27	10.94	10.97
0.7	5.94	6.07	5.78	5.87					
0.8	4.86	4.82	4.62	4.67					
0.9	3.32	3.10	3.02	2.97					

$\rho = 0.5$ and $\eta = 100$.

$c = 100$ and $\eta = 10$.

Table 3.3 – Results for the systemic risk problem. Estimations of the value function at point $(x_0 = 10)$ at time 0 provided by different strategies. We took $T = 1$, $N = 100$, $\sigma = 0.1$, $\kappa = 0.5$, $X_0 = 10$.



$c = 100$ and $\eta = 10$.

$\rho = 0.5$ and $\eta = 100$.

Figure 3.5 – Results for the systemic risk problem. Estimations of the value function at time 0 using different algorithms w.r.t. ρ and c . We took $T=1$, $N=100$, $\sigma=0.1$, $\kappa=0.5$, $x_0=10$.

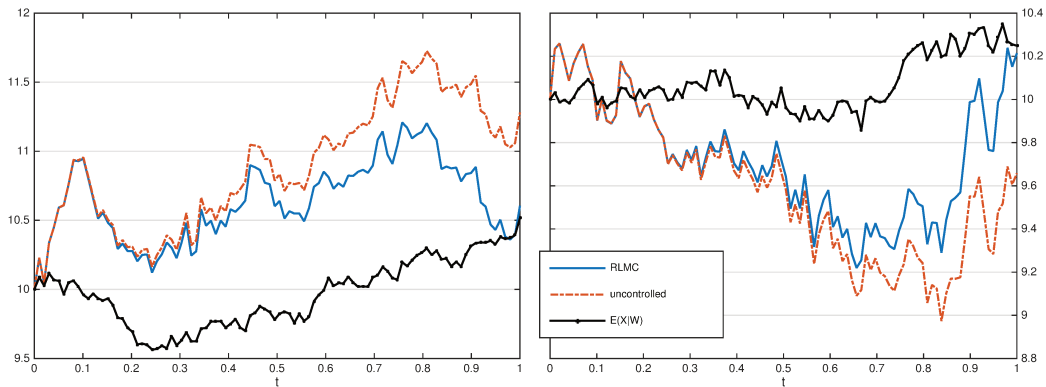


Figure 3.6 – Two realizations of $(X_t)_{t \in [0, T]}$ controlled by RLMC (curve “RLMC”), $(X_t)_{t \in [0, T]}$ naively controlled taken $\alpha = 0$ (curve “uncontrolled”), and \bar{X} (curve “ $E(X|W)$ ”). The optimal control for the systemic risk problem (computed by RLMC) is to do nothing at first, and catch \bar{X} when the terminal time is getting close.

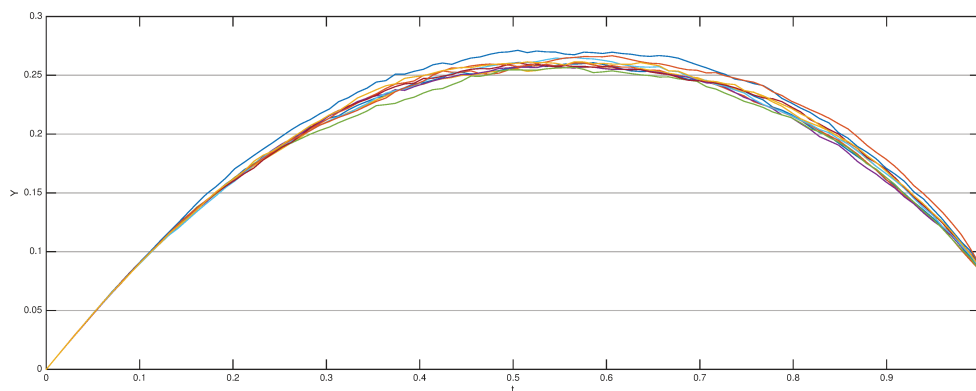


Figure 3.7 – Sample of $(Y_t)_{t \in [0, T]}$ controlled by Q. The (Q) estimated optimal control for the systemic risk problem is to initially let Y become large, and then reduce its value when the approaching the terminal time.

3.6 Discussion

In this work, we have investigated how to use probabilistic numerical methods for some classes of mean field control problem via Markovian embedding. We focused on two types of Regression Monte Carlo methods (namely, Regress-Later and Control Randomization) and Quantization. We have then presented three different examples of applications.

Regression Monte Carlo techniques. We found that the Regression Monte Carlo algorithms perform well in problems of control of the drift. In such problems, they are much faster than Quantization for similar precision. In particular, we noticed that Regress-Later is usually more reliable than Control Randomization; often the choice of a uniform distribution of the training points on an appropriate interval is sufficient to obtain high-quality estimations.

On the other hand Control Randomization is very sensitive to the choice of the distribution of the randomized control, and often a few iterations are necessary before finding a good control distribution. We have also tried to use the performance iteration or path recomputation method, but on the examples we considered, it was very time consuming and did not help much in terms of accuracy.

Despite the success of Regression Monte Carlo methods in problems with control on the drift, the example of Portfolio Selection highlighted a possible weakness of these algorithms. When the control acts on the variance only, we found difficult to make the numerical scheme converge to sensible results within the computational resources available. We realized that the study of these problems and the solution via Regression Monte Carlo methods is outside the scope of this paper. This is probably related to another limitation of this family of methods: the choice of the basis functions for the regression. Indeed, for some problems, a good basis might be very large or might require several steps of trials and errors.

Quantization techniques. Quantization-based method, on the other hand, provided very stable and accurate results. A first interesting and practical feature of the Q-algorithm is that regressing the value function using quantization-based methods is local. So, first, it can be easily parallelized to provide fast results, and, second, it is easy to check at which points of the grids the estimations suffer from instability and how to change the grid to fix the problem (basically, by adding more points where the estimations need to be improved).

Another interesting feature of the quantization methods is that, one can choose the grids on which to learn the value function. It is possible to exploit this feature in the case where one has, a priori, a rough idea of where the controlled process should be driven by the optimal strategy (see, e.g., the liquidation problem). In this case, one should build grids with many points located where the process is supposed to go.

When one has no guess of where the optimal process goes, it is always possible to use bootstrapping methods to build better grids iteratively, starting from a random guess for the grid (see, e.g., the systemic risk problem). In both cases, one has to be particularly careful with the borders of the grids that have been built. Indeed, the decisions computed by quantization-based methods at the borders might easily be wrong if the grids do not have a “good shape” at the borders. Unfortunately, the shape of the grid that should be used depends heavily on the problem under consideration.

Except in special cases, it is not possible to avoid deterministic algorithms (such as gradient descent methods or extensive search) to find the optimal action at each point of the grid. A smooth expression of the conditional expectations of the quantized processes is necessary for the deterministic algorithms find optimal strategy efficiently. Once again, the use of parallel

computing can alleviate the time-consuming task of searching for the optimal control at each point of the grids.

A non-intrusive stratified resampler for multi-factor models: Application in energy market

This chapter is based on the paper [GPZ18].

Abstract

We aim to solve dynamic programming equations (DPE) related to financial valuations in energy market. On the underlying asset, we consider that a calibrated model is not available and a limited sample from its historical data is accessible. We look for a non-intrusive method solving the DPE with empirical regression techniques, by suitable resampling of the historical data (and thus without calibration of the model). In power market, the forward contracts are driven by hidden factors modeled by Markov processes. Even though the DPE solution depends on these hidden factors, we come up with a resampling scheme using only the historical data of the observable log forward prices.

4.1 Introduction

Stochastic dynamic programming equations are related to the resolution of non-linear problems, e.g., the stochastic control (see [TV01, Pha05, BKS10]) or non-linear partial differential equations (see [BT04, Pha15, GT16]). These equations were studied in [MB46, Bel57, Put14] motivated by problem arising in almost all areas of science, from water reservoir management to finance.

We aim to solve dynamic programming equations (DPE) related to a financial application in energy market. For illustration, we can be concerned in the pricing of Bermudan or Swing options, where the underlying assets are power spot prices or forward contracts. We plan to build a non-intrusive method to solve those DPE using the observable data in the underlying market, without a full model calibration. We deal with a discrete-time framework, the DPE takes the following form:

$$Y_N = \tilde{g}_N(X_N),$$

$$Y_i = \mathbb{E}[\tilde{g}_i(Y_{i+1}, \dots, Y_N, X_i, \dots, X_N) \mid X_i], \quad i = N - 1, \dots, 0,$$

for some real functions \tilde{g}_N and \tilde{g}_i depending on the problem under consideration. Here, X is a Markov chain in \mathbb{R}^n , determined by the problem.

In the usual methodology, we observe a historical data O of size M , from which we extract the values of parameters of some model describing the stochastic evolution of X , and then we perform regression Monte Carlo (MC) algorithms (with many sampled points X , say N_{MC}) to obtain the regression functions \tilde{y}_i such that $Y_i = \tilde{y}_i(X_i)$ (see [LS01, Egl05]). In the current situation we are interested in, M is usually small compared to N_{MC} and thus, the statistical step gives a larger error than the empirical regression one. This is an incentive to perform a direct approach consisting in resampling (or bootstrapping) the observed data and obtain directly the regression functions, see Figure 4.1, this is inspired from [PBS01, GLZ18].

In energy market, the price at time t of a T -forward contract, denoted by $F(t, T)$, is usually modeled as driven in terms of hidden factors X , modeled by a Markov process (see [CMR09]). We suppose a dynamic for the hidden Markov processes X representing different characteristics of the forward price, for example, the short-term and equilibrium components (see [SS00]). We get from the market some observable data at different times t , which we denote by O ; typically O_t is the set of forward prices $F(t, T)$ or a transformation of them. Practical examples correspond to solving pricing problems when the time horizon T is one year. Regarding dates, we have daily historical data over last 20 years. However, energy data have a high seasonality component: for instance, the volatility in winter is quite different from volatility in summer (for instance as reference to seasonality effects in forward contracts, see [BG06]). Therefore, it is questionable to estimate winter volatilities using summer data. It is more reasonable to consider that we have $M = 20$ independent realizations of a one-year path.

When we rewrite the DPE in terms of the underlying $F(\cdot, T)$, which is also function of X , we get control or value functions depending on the hidden factors. Because these Markov process X are not observable, meaning we have no direct access to them, we can not apply a resampling method (as in [GLZ18]) that regenerates each one of those factor individually. To deal with this problem, we suggest a resampling method on the observable data O , by modifying the scheme of the aforementioned reference.

In view of applying a resampling method on observable data, the discrete dynamic programming equation (DDPE) will take another form:

$$Y_N = g_N(O_N),$$

$$Y_i = \mathbb{E}[g_i(Y_{i+1}, \dots, Y_N, O_i, \dots, O_N) \mid O_i], \quad i = N - 1, \dots, 0.$$

The aim is to compute the value function y_i such that $Y_i = y_i(O_i)$ (not as a function of X_i) using regression Monte Carlo methods. Here, X is a hidden process valued in \mathbb{R}^n and say that the

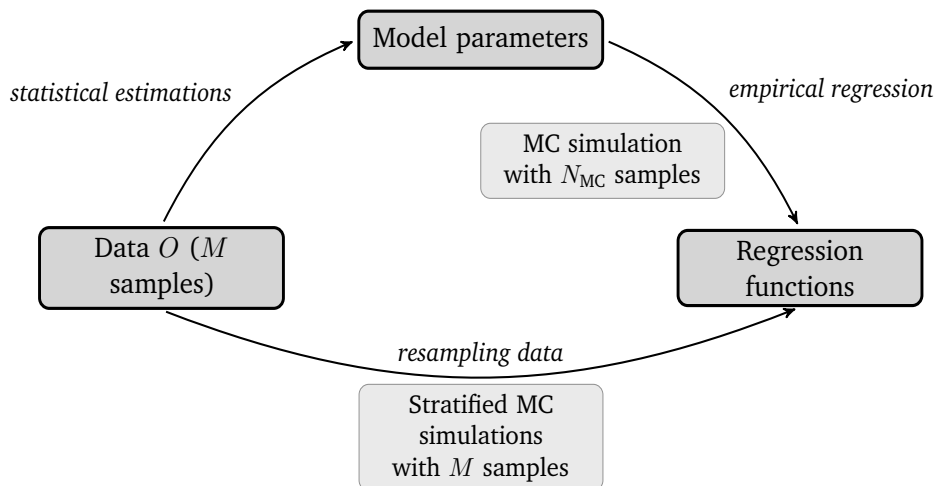


Figure 4.1 – Regression function reconstruction: Two different approaches

process O takes values in \mathbb{R}^d . For instance, O and X could represent the market log-prices and the factors on a multi-factor model. Usual modeling works give that, in the market, we take $d = 36$ different transformed prices $(O^{(1)}, \dots, O^{(36)})$, driven by $n = 3$ hidden factors $(X^{(1)}, X^{(2)}, X^{(3)})$ (see [KSB09], for a two factor model in electricity forward market).

Let us briefly recall the Regression Monte Carlo (RMC) methodology. Owing to the \mathbb{L}^2 -projection property of the conditional expectation, we have that y_i solves the minimization problem

$$\min_{\phi: \mathbb{R}^d \rightarrow \mathbb{R}} \mathbb{E} \left[\left| g_i(Y_{i+1}, \dots, Y_N, O_i, \dots, O_N) - \phi(O_i) \right|^2 \right], \quad i = N-1, \dots, 0,$$

where ϕ is a real measurable function on \mathbb{R}^d . The Regression Monte Carlo approach comprises simulating N_{MC} paths of observable variable O , say $O^{1:N_{\text{MC}}} := O^1, \dots, O^{N_{\text{MC}}}$ (input of the algorithm), and finding approximations $\hat{y}_i^{N_{\text{MC}}, \mathcal{L}}$ in a vector space of functions \mathcal{L} . The real function $\hat{y}_i^{N_{\text{MC}}, \mathcal{L}}$ (output of the algorithm) solves

$$\min_{\phi \in \mathcal{L}} \sum_{m=1}^{N_{\text{MC}}} \left| g_i(\hat{y}_{i+1}^{N_{\text{MC}}, \mathcal{L}}(O_{i+1}^m), \dots, \hat{y}_N^{N_{\text{MC}}, \mathcal{L}}(O_N^m), O_i^m, \dots, O_N^m) - \phi(O_i^m) \right|^2,$$

for every $i = N-1, \dots, 0$. But, to ensure accurate scheme, the RMC methods usually need that the number of simulations N_{MC} has to be much larger than the dimension of the vector space \mathcal{L} (number of coefficients). Note that the inputs are sampled from a model which is estimated from M data only: in our setting where M is small, model error may be a significant concern and we wish to better account for this.

In [GLZ18], the authors design a scheme where the N_{MC} simulations are replaced by the M observed data (and thus $M = N_{\text{MC}}$). To overcome the problem of numerous coefficients to compute (despite the small number of data), they combine this with a stratification strategy and local approximations.

This approach is called the Non-Intrusive Stratified Resampler (NISR) scheme. More precisely, this resampling method consists of subdividing \mathbb{R}^d in strata $(\mathcal{H}_k)_k$ and figuring out the approximations on each stratum \mathcal{H}_k using the root sample $O^{1:M}$. Due to the stratification part of the NISR scheme, we choose small stratum allowing to use a small dimensional approximation space \mathcal{L}_k , thus alleviating the large M constraint. The key point in the non-intrusive part of the scheme is to assume that only the structure of the model for O is known, but not the values of its parameters/coefficients.

In our setting, had we observed the factors X , we would have resampled it and used directly the NISR scheme. Unfortunately, we have only access to O which is a affine and partially unknown transformation of X : this is why we have to propose an extension of the native NISR scheme in a version which takes as inputs O . To make it possible, we consider that the factors model is described as multidimensional Ornstein-Uhlenbeck (OU) which we know the mean-reverting parameters, but not the volatilities/correlation.

We end this introduction with an outline of this chapter. The dynamics programming setting is discussed in Section 4.2 that leads to the NISR for regression Monte Carlo. Main notations, technical assumptions and model structures are also introduced. In Section 4.3, the algorithm to solve dynamic programming equation under standard conditions of accessibility to historical data is presented (see Algorithm 4). Section 4.4 is devoted to financial applications in energy market: the pricing of Bermudan and Swing options. Numerical results are in prospect.

4.2 Setting

4.2.1 Dynamic programming equation

In the following, we consider N fixed times t_i and denote by O_i the observable process O at time t_i . Our aim is to compute the following discrete dynamic programming equations (DPE) in general form:

$$\begin{aligned} Y_N &= g_N(O_N), \\ Y_i &= \mathbb{E}[g_i(Y_{i+1:N}, O_{i:N}) | O_i], \quad 0 \leq i \leq N-1, \end{aligned} \quad (4.1)$$

for some functions g_N and g_i . Here, $(O_i)_{0 \leq i \leq N}$ and $(Y_i)_{0 \leq i \leq N}$ are random processes taking values in \mathbb{R}^d and \mathbb{R} , respectively, and we use the short notation $O_{i:N} = (O_i, \dots, O_N)$.

We take a hidden Markov process X in \mathbb{R}^n such that the observable variable O in \mathbb{R}^d is described in terms of X as $O_i = \phi(X_i)$, thanks to some measurable function $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^d$. The modeling interpretation of X and O is that X drives the evolution of O ; only O can be observed in practice, see the subsequent Example 4.3. We assume that there are more observable variables than hidden variables, i.e., $d \geq n$. We also assume that O is a Markov chain.

Consider X_i denoting the process X at time t_i . Due the Markovian property of the data D , we write the previous system as follows

$$\begin{aligned} Y_N &= \tilde{g}_N(X_N), \\ Y_i &= \mathbb{E}[\tilde{g}_i(Y_{i+1:N}, X_{i:N}) | X_i], \end{aligned} \quad (4.2)$$

for some functions \tilde{g}_N and \tilde{g}_i depending on g_N , g_i and ϕ . Assuming that measurable function ϕ is not completely explicit, the composed functions \tilde{g}_N and \tilde{g}_i are neither.

Discussion. Up to this point, we have presented two different approaches depending on the underlying Markov process we choose O or X . If we apply the NISR scheme to

- Equation (4.1), the Markovian variable O is in dimension d , much larger than the effective Markovian dimension (i.e., n). We pay the price of a higher dimension. On the other hand, we may observe the data O solving a OU dynamics so that we can resample them easily (like in [GLZ18]).
- Equation (4.2), the Markovian variable X is in low dimension n (good) but the payoff depends on ϕ (that we do not know). Also we do not observe X directly, we need ϕ to get X from O . The problem is that ϕ depends on others unknown parameters such as correlation, volatilities (many parameters we do not know). This is why we follow the approach in (4.1).

From (4.1), we can deduce the existence of a measurable function y_i such that $Y_i = y_i(O_i)$. Therefore, our aim is to computation an approximation to the value functions y_i for all i using the NISR scheme. Hence, for technical reasons we prefer to assume bounded value functions as in [GLZ18, Assumption 2.1].

4.2.2 Modeling of the observable and hidden processes

We need to compute the coefficients of an approximation of y_i on a vector space using the samples of O (regression Monte Carlo). To invoke the NISR algorithm, we need to resample O and for this we specify the dynamics of O , X and the relation between O and X . First, essentially we will require $\phi : \mathbb{R}^n \rightarrow \mathbb{R}^d$ to be a linear transformation.

Assumption 4.1 (Model structure). *The observable data O_i is related to the hidden Markov state X_i as follows*

$$O_i = AX_i + b, \quad (4.3)$$

where A is a real matrix in $\mathbb{R}^{d \times n}$ and b is a vector in \mathbb{R}^d , with $d \geq n$. The matrix A is deterministic and **known**, while b is deterministic and unknown. The (column) rank of A is n .

Hidden model. We present the underlying structure of the hidden Markov chain X as discrete version of a Markov process with time independent parameters. Let $W = (W^1, \dots, W^n)$ be a n -dimensional Brownian motion on a probability space $(\Omega, \mathcal{F}, \mathbb{P})$ endowed with the natural (complete and right-continuous) filtration $\mathbb{F} = (\mathcal{F}_t)_{t \leq T}$ generated by W up to some fixed time horizon $T > 0$.

Assumption 4.2 (Hidden model). *We suppose that X is a n -dimensional OU process in \mathbb{R}^n given by*

$$dX_t = (\mu - \alpha X_t) dt + \sigma dW_t, \quad (4.4)$$

where μ is a drift vector in \mathbb{R}^n , α is a mean-reverting matrix in $\mathbb{R}^{n \times n}$, and σ is a diffusion matrix in $\mathbb{R}^{n \times n}$. We assume that α is **known**.

The value of n will depend on the model used to describe the commodity prices. Notice that X_i (random variable indexed by i) denotes the process X at time t_i , i.e., $X_i := X_{t_i}$. Below we show an example of a hidden multi-factor X in dimension $n = 2$.

Example 4.3 (Schwartz-Smith model, $n = 2$). Let (W_χ, W_ξ) be a couple of correlated Brownian motions given by

$$(W_\chi, W_\xi) = \left(W^1, \rho W^1 + \sqrt{1 - \rho^2} W^2 \right),$$

with correlation ρ . Let $X = (\chi, \xi)$ be the process satisfying (4.4) whose parameters μ, α, σ are given by

$$\mu = \begin{pmatrix} 0 \\ \mu_\xi \end{pmatrix}, \quad \alpha = \begin{pmatrix} \kappa & 0 \\ 0 & 0 \end{pmatrix}, \quad \sigma = \begin{pmatrix} \sigma_\chi & 0 \\ \rho\sigma_\xi & \sqrt{1 - \rho^2}\sigma_\xi \end{pmatrix}.$$

Those parameters are independent of time. Then the spot price S_t is modeled in terms of two correlated factors χ and ξ as follows

$$\log S_t = \chi_t + \xi_t.$$

Here, the processes χ and ξ represent short-term deviations and equilibrium levels respectively. After the change of probability to bring the process $X = (\chi, \xi)$ to a risk neutral measure, we express the T -forward price $F(t, T)$ as expectation of the future spot price S_T and obtain

$$\log F(t, T) = e^{-\kappa(T-t)} \chi_t + \xi_t + a(T - t),$$

where $a(\Delta\tau)$ is a function of the time to maturity $\Delta\tau = T - t$ given by

$$\begin{aligned} a(\Delta\tau) &= (\mu_\xi - \lambda_\xi) \Delta\tau - (1 - e^{-\kappa\Delta\tau}) \frac{\lambda_\chi}{\kappa} \\ &\quad + \frac{1}{2} \left((1 - e^{-2\kappa\Delta\tau}) \frac{\sigma_\kappa^2}{2\kappa} + \sigma_\xi^2 \Delta\tau + 2(1 - e^{-\kappa\Delta\tau}) \frac{\rho\sigma_\chi\sigma_\xi}{\kappa} \right). \end{aligned}$$

Here, λ_ξ and λ_χ are the risk premiums for the short-term deviations and equilibrium level, respectively. For more details, we leave the reader to see the so-called Schwartz-Smith model in [SS00].

Linear model structure. Regarding the Schwartz-Smith model in Example 4.3, we write model structure (see (4.3)) in continuous time

$$O_t = AX_t + b, \quad (4.5)$$

for a set of forward prices

$$F(t, t + \Delta\mathcal{T}) = \{F(t, t + \Delta\tau) : \Delta\tau \in \Delta\mathcal{T}\}$$

with time to maturities $\Delta\mathcal{T} = \{\Delta\tau_1, \dots, \Delta\tau_{36}\}$ as follows

$$O_t = \begin{pmatrix} \log F(t, t + \Delta\tau_1) \\ \log F(t, t + \Delta\tau_2) \\ \vdots \\ \log F(t, t + \Delta\tau_{36}) \end{pmatrix}, \quad A = \begin{pmatrix} e^{-\kappa\Delta\tau_1} & 1 \\ e^{-\kappa\Delta\tau_2} & 1 \\ \vdots & \vdots \\ e^{-\kappa\Delta\tau_{36}} & 1 \end{pmatrix}, \quad b = \begin{pmatrix} a(\Delta\tau_1) \\ a(\Delta\tau_2) \\ \vdots \\ a(\Delta\tau_{36}) \end{pmatrix} \quad (4.6)$$

and $X_t = (\xi_t, \chi_t)$ for some time t . In the previous example, we have assumed the existence of 36 forward contracts (with a sufficient liquidity).

According to Assumptions 4.1 and 4.2, we suppose the **knowledge of matrix A and α (i.e., of the parameter κ)**, but not of vector b (i.e., of the parameters $\mu_\xi, \lambda_\xi, \lambda_\chi, \sigma_\xi, \sigma_\chi$ and ρ). Otherwise, it would mean we estimate all parameters of the OU model. Because we are looking for a non-intrusive algorithm, the knowledge of all model parameters is not reasonable. Say differently, to solve the DPE (4.1) we will assume to know κ only, and we will use observations of O .

Observable model. From the relation between the observable data O and the hidden Markov state X , we now deduce the dynamics of O .

Proposition 4.4. *Under Assumption 4.1-4.2, we have*

$$O_t = e^{-\tilde{\alpha}(t-t_i)} O_i + \int_{t_i}^t e^{-\tilde{\alpha}(t-s)} (\tilde{\mu} ds + \tilde{\sigma} dW_s), \quad (4.7)$$

with

$$\tilde{\alpha} = A\alpha A^\dagger \in \mathbb{R}^{d \times d}, \quad \tilde{\mu} = A\mu + \tilde{\alpha}b \in \mathbb{R}^d, \quad \tilde{\sigma} = A\sigma \in \mathbb{R}^{d \times n},$$

and $A^\dagger \in \mathbb{R}^{n \times d}$ is the Moore-Penrose inverse of A (see the definition in [Pen55]).

Proof. From the assumption that the (column) rank of A is n , $A^\top A \in \mathbb{R}^{n \times n}$ is invertible. Therefore, we have the $n \times d$ matrix A^\dagger is the left inverse of A

$$A^\dagger = (A^\top A)^{-1} A^\top, \quad A^\dagger A = I, \quad (4.8)$$

where I is the $n \times n$ identity matrix.

From the relation in (4.5) and the left-inverse identity in (4.8), we obtain

$$X_t = A^\dagger(O_t - b). \quad (4.9)$$

By differentiating O in (4.5), by multiplying the SDE on X in (4.4) by A on the left-hand side and using the left-inverse identity in (4.8), we obtain that O solves also another OU process with

time independent parameters, after replacing the remaining X by the expression in (4.9)

$$\begin{aligned} dO_t &= (A\mu - A\alpha A^\dagger(O_t - b)) dt + A\sigma dW_t \\ &= (\tilde{\mu} - \tilde{\alpha}O_t) dt + \tilde{\sigma} dW_t. \end{aligned} \quad (4.10)$$

From the SDE in (4.10), we conclude the announced result. \square

The equation (4.7) shows how the dependence between O_t and O_i is parameterized by $\tilde{\alpha}$, and the random part independent of O_i . Notice that, in the case $d = n$ and $\text{rank}(A) = n$, we have the existence of A^{-1} and we simply obtain $\tilde{\alpha} = A\alpha A^{-1}$. In the Schwartz-Smith model, the knowledge of $\tilde{\alpha}$ is implied by the knowledge of the mean-reverting speed κ and the time to maturity $\Delta\tau$ of forward contracts $F(t, t + \Delta\tau)$.

4.2.3 Non-Intrusive resampling

For the regression Monte Carlo algorithm in the NISR framework, we use samples of O under the following assumption.

Assumption 4.5 (Observable data). *We have access to the observation of M paths of O , denoted by $O^{1:M} = \{O_i^m : 0 \leq i \leq N, 1 \leq m \leq M\}$. They are independent and called **root sample**.*

From the representation of the flow of the Markov chain O , we write

$$O_j^{i,z} = \Theta_{i,j}(z, U)$$

for different times $i \in \{1, \dots, N\}$ and points $z \in \mathbb{R}^d$, where U is a random vector and $\Theta_{i,j}$ is a deterministic function for pair (i, j) , called *flow functions* here.

We exhibit the flow function in the case of the Schwartz-Smith model.

Example 4.6 (Flow function in Schwartz-Smith model). Following the Assumptions 4.1-4.2 and the Example 4.3 (Schwartz-Smith model), we have obtained a model structured presented in (4.6). After computing the pseudo-inverse A^\dagger , we deduce that $\tilde{\alpha}$ is

$$\tilde{\alpha}_{i,j} = \kappa e^{-\kappa\Delta\tau_i} \frac{36e^{-\kappa\Delta\tau_j} - \sum_{j'=1}^{36} e^{-\kappa\Delta\tau_{j'}}}{36 \sum_{j'=1}^{36} e^{-2\kappa\Delta\tau_{j'}} - 2 \sum_{j'=1}^{36} e^{-\kappa\Delta\tau_{j'}}},$$

for all $1 \leq i, j \leq 36$. Therefore, knowing $\tilde{\alpha}$ is equivalent to known the mean-reverting speed κ and all time to maturities $(\Delta\tau_1, \dots, \Delta\tau_{36})$.

Given $N + 1$ times $\{t_i : 0 \leq i \leq N\}$, we observe a path $\{O_i : 0 \leq i \leq N\}$. Following the explicit solution of O in (4.7), we define the flow functions as

$$\Theta_{i,j}(z, U) = e^{-\tilde{\alpha}(t_j - t_i)z} + U_{i,j},$$

and we extract a random source U

$$U_{i,j} = O_j - e^{-\tilde{\alpha}(t_j - t_i)O_i},$$

for all $0 \leq i \leq j \leq N$.

Note that from Assumption 4.5, observing M independent (O^1, \dots, O^m) implies we can get M independent realizations (U^1, \dots, U^m) of U . In addition, owing to Assumptions 4.1 and 4.2, the flow function $\Theta_{i,j}$ is **known**.

Instead of identifying the distribution of observable data O , we simply assume to be able to retrieve the random sources that generated the historical data. In other words, we suppose the existence of a general structure of the Markov chain model, but not a complete knowledge of the model parameters. This is the non-intrusive characteristic of the NISR scheme.

4.3 The NISR scheme

In this section, we present the NISR-regression Monte Carlo algorithm assuming the complete access to the observable data. From the historical data, we construct on a root sample $O^{1:M}$ of size M (possibly small). Here we show how to recover the random sources generating the root sample, then be able to resampling M paths from any initial point z in \mathbb{R}^d . Assuming to know $\tilde{\alpha}$, we apply the Example 4.6 in general case.

4.3.1 Extraction

Pick one path $\{O_i^m : 0 \leq i \leq N\}$ of root sample $O^{1:M}$ along $N+1$ discrete times $\{t_i : 0 \leq i \leq N\}$. The explicit solution of O in (4.7) helps us to extract the random source U as follows

$$U_{i,j}^m := O_j^m - e^{-\tilde{\alpha}(t_j-t_i)} O_i^m, \quad \text{for all } 0 \leq i \leq j \leq N. \quad (4.11)$$

Notice that for each path $O_{1:N}^m$ we obtain a triangular matrix U^m . So the flow functions are

$$\Theta_{i,j}(z, U) := e^{-\tilde{\alpha}(t_j-t_i)} z + U_{i,j}, \quad \text{for all } 0 \leq i \leq j \leq N. \quad (4.12)$$

Here, with a random source U , we obtain another sample starting from a initial point $z \in \mathbb{R}^d$.

We call the mapping in (4.11) from the path space \mathbb{R}^{N+1} into random source space $\mathbb{R}^{\frac{(N+1)(N+2)}{2}}$ the *extractor function* Φ

$$U = \Phi(t_{0:N}, O_{0:N}) := \left(O_j - e^{-\tilde{\alpha}(t_j-t_i)} O_i \right)_{i,j:0 \leq i \leq j \leq N}. \quad (4.13)$$

In the following, $\Theta_{i,j}$ in (4.12) and Φ in (4.13) will be inputs of the NISR algorithm as well as the root sample $O^{1:M}$.

4.3.2 Stratification

Because the size M of the root sample is possibly small, we are not able to perform an accurate regression Monte Carlo method. Since the number of samples required for the empirical regression is much larger than sample extractable from the historical data, we will use the extracted random sources and resample them from different initial conditions (t_i, z) thanks to flow function in (4.12).

In fact, we need that M be greater the dimension of approximation spaces \mathcal{L} . For that, we resample the paths from distinct initial points in different regions of the space \mathbb{R}^d , then we perform a regression Monte Carlo on a low-dimensional approximation space. More precisely, we consider K strata $\{\mathcal{H}_k, 1 \leq k \leq K\}$ is a partition of the space \mathbb{R}^d , i.e.,

$$\mathcal{H}_k \cap \mathcal{H}_{k'} = \emptyset \text{ for } k \neq k', \quad \bigcup_{k=1}^K \mathcal{H}_k = \mathbb{R}^d.$$

and a probability measure ν such that its restriction ν_k on \mathcal{H}_k is denoted by

$$\nu_k(\mathrm{d}z) = \frac{\mathbf{1}_{\mathcal{H}_k}(z)}{\nu(\mathcal{H}_k)}\nu(\mathrm{d}z). \quad (4.14)$$

Here, we construct a resampling method using the probability measure ν and the random source U . To obtain a stability in the propagation of errors, we need a condition on the probability measure ν and the Markov chain O .

According to Proposition 3.2 in [GLZ18], we provide the probability measure ν that satisfying a stability condition for the OU processes, that we are resampling.

Proposition 4.7. *Let $p > 0$. Assume that $O_{i+1}^{i,z} = Dz + U_i$ for a diagonal invertible matrix $D = \text{diag}(D^1, \dots, D^d)$ with $\mathbb{E}[(1 + |U_i|)^{d(p+1)}] < +\infty$. Then the tensor-product Pareto-type distribution*

$$\nu(\mathrm{d}z) := \prod_{j=1}^d \frac{p}{2} (1 + |z_j|)^{-p-1} \mathrm{d}z \quad (4.15)$$

satisfies a **norm-stability condition**: there exists a constant $K \geq 1$ such that, for any $\varphi : \mathbb{R}^d \rightarrow \mathbb{R} \in \mathbb{L}^2(\nu)$, we have

$$\int_{\mathbb{R}^d} \mathbb{E}|\phi(O_{i+1}^{i,z})|^2 \nu(\mathrm{d}z) \leq K \int_{\mathbb{R}^d} |\phi(z)|^2 \nu(\mathrm{d}z),$$

for any $i \in \{0, \dots, N-1\}$.

The proof of Proposition 4.7 is given [GLZ18].

Now we present the stratified resampler using the distribution ν .

Definition 4.8 (Resampler). We define the M sample at time t_i until time t_N in the stratum \mathcal{H}_k as follows

- Let $V^{i,k,m}$ be a sample independent in i, k, m according to the distribution ν_k ; we set

$$O_i^{i,k,m} = V^{i,k,m}, \text{ for all } 1 \leq m \leq M.$$

- For every $j \in \{i+1, \dots, N\}$, we set

$$O_j^{i,k,m} = \Theta_{i,j}(O_i^{i,k,m}, U^m), \text{ for all } 1 \leq m \leq M,$$

where U^m is given by (4.11).

From Assumption 4.1, we have that random source $U^{1:M} = (U^1, \dots, U^M)$ are independent. We observe that $\{O_{i:N}^{i,k,m} : 1 \leq m \leq M\}$ are independent and identically distributed copies of $O_{i:N}$ starting at $O_i \sim \nu_k$, because

$$(O_j, i \leq j \leq N \mid O_i = z) \stackrel{d}{=} (\Theta_{i,j}(z, U), i \leq j \leq N).$$

4.3.3 NISR-regression Monte Carlo algorithm

By using M -sample given in Definition 4.8, we compute the approximation of the value function y_i according to Algorithm 4. Here, we consider a piecewise linear basis functions

$$\mathcal{L}_k = \text{span}\{\mathbf{1}_{\mathcal{H}_k}, z_1 \mathbf{1}_{\mathcal{H}_k}, \dots, z_d \mathbf{1}_{\mathcal{H}_k}\}.$$

In addition, we suppose the dimension of the space \mathcal{L}_k are independent from k , for the sake of simplicity. Indeed, we have $\dim(\mathcal{L}_k) = d + 1$.

In view of the following algorithm, we need to introduce the Ordinary Least Square (OLS) operator that approximates the function $C : \mathbb{R}^{d \times (N-i+1)} \rightarrow \mathbb{R}$ on the approximation space \mathcal{L}_k using the sample $O_{i:N}^{k,1:M}$. The operator **OLS** is defined as

$$\mathbf{OLS}(C, \mathcal{L}_k, O_{i:N}^{k,1:M}) = \arg \min_{\phi \in \mathcal{L}_k} \sum_{m=1}^M \left| C(O_{i:N}^{i,k,m}) - \phi(O_i^{i,k,m}) \right|^2.$$

The NISR-regression Monte Carlo algorithm is given in Algorithm 4. In the Algorithm 4, the root sample $O^{1:M}$ is an input to the Non-Intrusive Stratified Resampler given in Definition 4.8 through the random source $U^{1:M}$. Then, the existence of algorithm's approximation is conditioned to the fact that for any $i \in \{0, \dots, N\}$ the *observable* data O_i is indeed observed.

Algorithm 4 General NISR-regression Monte Carlo algorithm

Inputs:

- Extractor function Φ , Flow functions $\Theta_{i,j}$, resampling distribution ν ;
- Truncation operator $T_L(x) = -L \vee x \wedge L$ and regression function bound $|y_i|_\infty$.

- 1: Set $U^{1:M} = \Phi(t_{0:N}, O_{0:N}^{1:M})$.
- 2: Set $\hat{y}_N^M(\cdot) = g_N(\cdot)$.
- 3: **for** $i = N - 1$ to 0 **do**
- 4: **for** $k = 1$ to K **do**
- 5: Set $(V^{i,k,m})_{1 \leq m \leq M}$ according to ν_k (given in (4.14)).
- 6: Set $O_{i+1}^{i,k,m} = V^{i,k,m}$ and $O_{i+1}^{i,k,m} = \Theta_{i,j}(O_i^{i,k,m}, U^m)$ for $j \in \{i + 1, \dots, N\}$.
- 7: Set $C^M(z_{i:N}) = g_i(\hat{y}_{i+1}^M(z_{i+1}), \dots, \hat{y}_N^M(z_N), z_{i:N})$.
- 8: Compute $\phi^{M,k} = \mathbf{OLS}(C^M, \mathcal{L}_k, O_{i:N}^{i,k,1:M})$.
- 9: Set $\hat{y}_i^{M,k} = T_{|y_i|_\infty}(\phi^{M,k})$.
- 10: Set $\hat{y}_i^M(\cdot) = \sum_{k=1}^K \hat{y}_i^{M,k} \mathbf{1}_{\mathcal{H}_k}$.

Outputs: Approximation \hat{y}_i^M of regression function y_i .

Here, we state the convergence result for the Algorithm 4 (in the case of the one-step ahead DPE). It is expressed in term of the quadratic error of the best approximation of y_i on each stratum \mathcal{H}_k :

$$\mathcal{E}_{i,k} := \inf_{\phi \in \mathcal{L}_k} |\phi - y_i|_{\nu_k}^2,$$

where $|\phi|_{\nu_k}^2 := \int_{\mathbb{R}^d} |\phi(x)|^2 \nu_k(dx)$. In fact, the Theorem 4.9 gives us an upper bound for

$$\mathbb{E}[|\hat{y}_i^M - y_i|_\nu^2],$$

where $|\phi|_\nu^2 := \int_{\mathbb{R}^d} |\phi(x)|^2 \nu(dx)$. The previous expectation is also taken over the random sources $U^{1:M}$, making the quadratic error averaged on the root sample.

Theorem 4.9 (Converge result in the one-step ahead DPE case, see Corollary 3.1 in [GLZ18]). *Assume 4.1-4.2-4.5 and define \hat{y}_i^M as in Algorithm 4. Suppose each function g_i is Lipschitz w.r.t y_{i+1} (with Lipschitz constant L_{g_i}) and bounded at $y_{i+1} = 0$ w.r.t. $z_{i:N}$ by C_{g_i} . Let ν be the distribution given in Proposition 4.7. Then there are two constants $C_1(N)$ and $C_2(M)$ (with polynomial growth*

in N and log growth M , resp.) such that, for any $i \in \{0, \dots, N-1\}$,

$$\begin{aligned} \mathbb{E}[|\hat{y}_i^M - y_i|_\nu^2] &\leq C_1(N) \sum_{j=i}^{N-1} \left[\sum_{k=1}^K \nu(\mathcal{H}_k) \mathcal{E}_{j,k} \right. \\ &\quad \left. + \frac{1}{M} \left(C_2(M) (|y_i|_\infty^2 + L_{g_i}^2 |y_{i+1}|_\infty^2) + (d+1)(C_{g_i} + L_{g_i} |y_{i+1}|_\infty)^2 \right) \right]. \end{aligned}$$

The proof is given in [GLZ18, Section 3.3]. Usually, in Regression Monte Carlo methods, there is a competition between the approximation error and the statistical error. Indeed, while the dimension of the approximation space $K \dim(\mathcal{L}_k)$ goes to infinity, the statistical error (proportional to $K \dim(\mathcal{L}_k)/M$) explodes. In Theorem 4.9, the bias term (approximation error) is well controlled in K since the terms $\nu(\mathcal{H}_k)$ sum up to 1: we just obtain the average of the local quadratic errors $\mathcal{E}_{i,k}$. More striking is related to the variance term: it does not explode with the number K of strata. Therefore, taking a smaller stratum \mathcal{H}_k (which is interesting when we have a small root sample) does not make the statistical error increase. Briefly, the stratification improves the method by including several decoupled regression problems on a low-dimensional space.

4.4 Applications to Energy Market

In this section, we search for other applications to energy market. First, we look the valuation problem of Bermudan option whose the underlying processes are forward contracts in Subsection 4.4.1. Then, we explore the Gas storage valuation such as Swing option (multiple exercises rights) in Subsection 4.4.2. This study can be helpful for power producer willing to value its Swing contracts with its supplier.

4.4.1 Bermudan options

In this subsection, we consider the case where the observable data O_i throughout the dates $t_{0:N}$ is the log of forward prices $F(t_i, t_i + \Delta\tau)$ with time to maturity $\Delta\tau$.

Forward contracts. We suppose to have d different forward contracts, then O_i is given by

$$O_i = \begin{pmatrix} \log F(t_i, t_i + \Delta\tau_1) \\ \log F(t_i, t_i + \Delta\tau_2) \\ \vdots \\ \log F(t_i, t_i + \Delta\tau_d) \end{pmatrix}$$

for some collection of time to maturity $\Delta\mathcal{T} = \{\Delta\tau_j : 1 \leq j \leq d\}$.

In the actual market, the forward prices are characterized by their maturity dates T . By definition, whenever the current time t reaches the maturity T , the forward contract is equal to the spot S_T and after that it disappears. Notice then the actual forward data is defined for fixed maturity not for fixed time to maturity. We tackle this problem by performing a linear interpolation between the log forward prices $\log F(t_i, T_d(i, j))$ and $\log F(t_i, T_u(i, j))$, where $T_d(i, j)$ (resp. $T_u(i, j)$) is the largest (resp. the smallest) expiry date lesser (resp. greater) than the unobservable maturity $t_i + \Delta\tau_j$ for every $i \in \{0, \dots, N\}$ (see Definition 4.11).

Following the Assumption 4.1, we take the matrix A in (4.3) to be known and given by

$$A = \begin{pmatrix} e^{-\kappa_1 \Delta \tau_1} & \dots & e^{-\kappa_n \Delta \tau_1} \\ e^{-\kappa_1 \Delta \tau_2} & \dots & e^{-\kappa_n \Delta \tau_2} \\ \vdots & \vdots & \vdots \\ e^{-\kappa_1 \Delta \tau_d} & \dots & e^{-\kappa_n \Delta \tau_d} \end{pmatrix}.$$

For more details in energy market derivatives, see [Gem05] and [CS00] among others.

Option description. Let us describe the optimal stopping problem related the valuation of Bermudan options in Markovian framework. We take a set \mathcal{T}_N of $N + 1$ deterministic exercise times t_i and consider the optimal stopping rule associated: at final time T , the holder of the option exercises the option if it is in the money; at time t_i , the holder of the option compares the cash flow from immediate exercise with the expected cash flow from continuation. The value $V_i = V(t_i, O)$ of the Bermudan option at time t_i is given by the supremum over (integer-valued) stopping times τ in $\{i, \dots, N\}$ of the immediate reward g :

$$V(t_i, O) = \sup_{\tau \in \{i, \dots, N\}} \mathbb{E}[g(O_\tau) \mid O_i = O], \quad 0 \leq i \leq N - 1. \quad (4.16)$$

Following the dynamic programming principle (see [Kar88], for American options in continuous time and [Lam09] in discrete time), we have the existence of the optimal stopping time to (4.16):

$$\tau_i^* = \inf \left\{ j \geq i : V(t_j, O_j) = g(O_j) \right\},$$

corresponding to the optimal stopping strategy given the option have not been exercised until time t_i . Then, we rewrite the Bermudan option value $V_i = V(t_i, O)$ in terms of the continuation value $Y_i = Y(t_i, O)$ given by

$$V(t_i, O) = \max \{ g(O_i), Y(t_i, O) \}, \quad 0 \leq i \leq N - 1,$$

where

$$Y(t_i, O) = \mathbb{E}[V(t_{i+1}, O_{i+1}) \mid O_i = O], \quad 0 \leq i \leq N - 1$$

and $Y_N = 0$.

Regarding the discrete dynamic programming equation (4.1), the Bermudan option valuation corresponds to

$$g_N(O_N) = 0, \quad g_i(Y_{i+1}, O_{i+1}) = \max(g(O_{i+1}), Y(t_{i+1}, O_{i+1})), \quad (4.17)$$

similar to [TV01]. Here we set the continuation value as the expectation of the optimal value function at the next time, differently from the Longstaff-Schwartz approach proposed in [LS01].

Below we give some examples of reward (or payoff) functions g .

Example 4.10 (Options payoffs). Consider an option over $N = 12$ months such that exercise date t_i is the first day of the i -th month of the year.

1. The payoff function of a Bermudan Put on a 1 Month-Ahead (MAH) forward contract with strike price λ is given by

$$g(O_i) = (\lambda - \varphi(O_i))_+$$

with $\varphi(z) = e^z$ and $O_i = \log F(t_i, t_i + 1M)$.

2. The payoff of a Bermudan Exchange option on a 1 MAH forward contract against to 2 MAH one with strike price λ is

$$g(O_i) = (\varphi_1(O_i) - \lambda\varphi_2(O_i))_+$$

with $\varphi(z_1, z_2) = (e^{z_1}, e^{z_2})$ and $O_i = (\log F(t_i, t_i + 1M), \log F(t_i, t_i + 2M))^\top$.

3. The payoff of Bermudan Exchange Asian option on 1-6 MAH forward contracts versus 7-12 MAH ones with strike price λ is

$$g(O_i) = (\varphi_1(O_i) - \lambda\varphi_2(O_i))_+$$

with $\varphi(z_1, \dots, z_{12}) = (\frac{1}{6} \sum_{j=1}^6 e^{z_j}, \frac{1}{6} \sum_{j=7}^{12} e^{z_j})$ and $O_i = (\log F(t_i, t_i + 1M), \dots, \log F(t_i, t_i + 12M))^\top$.

For the previous payoff, we need to know at most 12 monthly forward contracts.

Complete and incomplete historical data. Regarding the Example 4.10, we require $(F(t_i, t_i + \Delta\tau_j))_{1 \leq j \leq n}$ to be observable in the market for every $i \in \{0, \dots, N\}$, which is summarized in the following case.

Case 1 (Complete data). For any $i \in \{0, \dots, N\}$ and $j \in \{1, \dots, n\}$, $t_i + \Delta\tau_j$ is the maturity time of a observable forward contract.

Therefore, the component $O_i^j = \log F(t_i, t_i + \Delta\tau_j)$ is directly extracted from the historical data without any interpolation. However, whenever the log forward price $\log F(t_i, t_i + \Delta\tau_j)$ is not available in the historical data, we perform an interpolation of log prices. This approach is summarized in the Case 2. In Figure 4.2, we see the French Power Base Futures Prices at July 2nd 2018 for different delivery periods. Those future contracts have fixed maturities. Moreover, we observe at this date

- The August 2018 futures is the most traded month contract;
- The Calendar 2022 and 1-Quarter 2020 futures are not traded yet;

implying that not every expiry dates are listed.

For that, we define the closest expiry times available in the market.

Definition 4.11. Let $\mathbf{T} = \{T_1, T_2, \dots\}$ be the collection of expiry time available in the market at every time t_i . For any time t , we consider

$$T_d(t) = \sup\{T \in \mathbf{T} : T \leq t\} \text{ and } T_u(t) = \inf\{T \in \mathbf{T} : T > t\}$$

as the greatest available expiry time less than time t and the least available expiry greater than time t , respectively.

Here, we interpolate those time to maturities which are not accessible by those which are.

Case 2 (Incomplete data). For any $i \in \{0, \dots, N\}$, the log forward price $\log F(t_i, t_i + \Delta\tau)$ is given by

$$\begin{aligned} \log F(t_i, t_i + \Delta\tau) &= \frac{T_u(t_i + \Delta\tau) - (t_i + \Delta\tau)}{T_u(t_i + \Delta\tau) - T_d(t_i + \Delta\tau)} \log F(t_i, T_d(t_i + \Delta\tau)) \\ &+ \frac{(t_i + \Delta\tau) - T_d(t_i + \Delta\tau)}{T_u(t_i + \Delta\tau) - T_d(t_i + \Delta\tau)} \log F(t_i, T_u(t_i + \Delta\tau)) \end{aligned} \quad (4.18)$$

for every time to maturity $\Delta\tau$ in $\Delta\mathcal{T} = \{\Delta\tau_1, \dots, \Delta\tau_d\}$.

Name	Best Bid	Best Ask	No. of Contracts	Last Price	Abs. Change	Last Time	Last Vol.	Settl. Price	Vol. Exchange	Vol. Trade Registration	Open Interest Prev. Day
Jul/18	-	-	-	-	-0.57	-	-	46.98	-	-	25,098
Aug/18	-	-	116	45.20	0.73	16:01	744	45.40	15,624	70,680	23,619
Sep/18	-	-	8	-	0.98	-	-	52.79	-	5,760	22,188
Oct/18	-	-	-	-	-0.10	-	-	58.80	-	-	5
Nov/18	-	-	-	-	1.03	-	-	64.44	-	-	-
Dec/18	-	-	-	-	-0.25	-	-	62.51	-	-	-

(a) Month futures with delivery period from July 2018 to December 2018.

Name	Best Bid	Best Ask	No. of Contracts	Last Price	Abs. Change	Last Time	Last Vol.	Settl. Price	Vol. Exchange	Vol. Trade Registration	Open Interest Prev. Day
4/18	-	-	52	61.95	0.28	15:52	4,418	61.89	26,508	88,360	19,950
1/19	-	-	20	61.05	0.45	16:45	10,795	61.14	43,180	-	814
2/19	-	-	10	39.00	0.50	10:55	10,920	39.16	10,920	10,920	1,260
3/19	-	-	-	-	0.41	-	-	40.10	-	-	1,240
4/19	-	-	8	54.75	0.37	15:40	11,045	54.69	17,672	-	221
1/20	-	-	-	-	-	-	-	-	-	-	-

(b) Quarter futures with delivery period from Quarter-4 2018 to Quarter-4 2019.

Name	Best Bid	Best Ask	No. of Contracts	Last Price	Abs. Change	Last Time	Last Vol.	Settl. Price	Vol. Exchange	Vol. Trade Registration	Open Interest Prev. Day
Cal-19	-	-	23	48.70	0.45	15:10	43,800	48.73	131,400	70,080	6,592
Cal-20	-	-	5	45.50	0.87	15:51	43,920	45.50	43,920	-	1,726
Cal-21	-	-	-	-	0.71	-	-	42.69	-	-	288
Cal-22	-	-	-	-	-	-	-	-	-	-	-
Cal-23	-	-	-	-	-	-	-	-	-	-	-

(c) Calendar futures with delivery period from 2019 to 2021.

Figure 4.2 – French Power Base Futures Contracts at July 2nd 2018 for different delivery periods (Month, Quarter, Year) and different expiry dates. Screen captured at July 10th 2018 (Source: <https://www.eex.com/en/market-data/power/futures/french-futures#!/2018/07/02>).

The forward prices $F(t_i, T_d(t_i + \Delta\tau))$ and $F(t_i, T_u(t_i + \Delta\tau))$ are those available in the market whose expiry times are the closest to $F(t_i, t_i + \Delta\tau)$. We illustrate below some examples of payoffs with incomplete data.

Example 4.12 (Incomplete data). Consider an option over $N = 24$ exercise dates t_{2i-1} with t_{2i} and is the 1st and the 15th day of the i -th month of the year, respectively. The payoff function of a Bermudan Exchange option on a 1 MAH forward contract with strike price λ is given by

$$g(O_i) = (\lambda - \varphi(O_i))_+$$

with $\varphi(z) = e^z$ and $O_i = \log F(t_i, t_i + 1M)$.

Notice the 1MAH forward contract is the contract with delivery date starting in the first next month. Therefore, this is a contract with fixed maturity and not fixed time to maturity.

NISR-regression Monte Carlo algorithm: Incomplete data In the case of incomplete data, we use the interpolation of log forward prices between different time to maturities given in (4.18).

Therefore, we modified the Algorithm 4 to solve the problem corresponding to the Bermudan option valuation in (4.17). The modified NISR-regression Monte Carlo algorithm uses the Ordinary Least Square operator that approximates the function $C : \mathbb{R}^{2d} \rightarrow \mathbb{R}$ on the approximation space \mathcal{L}_k using the sample $O_{i:i+1}^{k,1:M}$.

Up to now, we deal with Bermudan options which are in between American and European options in the sense of exercises dates. However, in this example, the contract buyer has only one exercise right. Now we are going extend to multiple exercise rights.

4.4.2 Swing options

In this subsection, we again consider that the observable data O is the log of forward prices $F(t_i, t_i + \Delta\tau)$. Then to recover the forward prices from the data O , we use the function $\varphi(z_1, \dots, z_d) = \sum_{i=1}^d \exp(z_i)/d$ (i.e., a basket of future contracts).

Option description. Now we describe the stochastic optimal control problem related to the valuation of Swing options. We consider a set $\mathcal{T}_N = \{t_0, \dots, t_i, \dots, t_N\}$ of $N + 1$ deterministic exercises dates. We give by q_i the local consumption at time t_i under (local) loading constraint:

$$q_i \in [q_{\min}, q_{\max}].$$

The cumulative volume purchased up time t_i is then denoted by $Q_i = \sum_{j=0}^{i-1} q_j$, which satisfies the (total) volume constraint:

$$Q_N \in [Q_{\min}, Q_{\max}].$$

At each time t_i , the buyer of the contract gains $q_i(\varphi(O_i) - \lambda)$ by purchasing forward contracts at price λ (we swing forward contracts against fixed prices). Moreover, whenever the total volume constraint is violated, the buyer must pay some penalties at maturity T . For example, these penalties can be linear in term of the forward contracts O_N at time t_N and in term of the under or over-consumption:

$$P_N(O_N, Q_N) = -\alpha_1 \varphi(O_N)(Q_N - Q_{\min})_- - \alpha_2 \varphi(O_N)(Q_N - Q_{\max})_+,$$

for some positives constants α_1 and α_2 .

Then, we consider V_i the market value of a variable volume Swing option at time t_i defined at the supremum over the control variable (consumption strategy) $q = (q_i)_{1 \leq i \leq N}$

$$V_i = \sup_{(q_j)_{i \leq j \leq N}} \mathbb{E}[CF_i(O_{i:N}, Q_{i:N}, q_{i:N}) \mid O_i, Q_i],$$

where CF_i is the future cash flow starting from date t_i are given by

$$CF_i(O_{i:N}, Q_{i:N}, q_{i:N}) = \sum_{j=i}^{N-1} q_j(\varphi(O_j) - \lambda) + P_N(O_N, Q_N).$$

Above, CF_i is the future returns starting from date t_i for a given consumption $(q_j)_{i \leq j \leq N}$ and underlying values $(O_j)_{i \leq j \leq N}$.

In the following, we present a version of the Swing option example given in [BBD⁺06, Case 1] with weekly exercise.

Example 4.13 (Weekly Swing options). Consider an option where the buyer has the right to purchase each first day of $N = 52$ weeks a gas quantity $q \in [q_{\min}, q_{\max}]$ at a fixed strike price λ . But he must purchase globally over the year a volume $Q \in [Q_{\min}, Q_{\max}]$. Here, we take the following parameters: $q_{\min} = 0$, $q_{\max} = 6$, $Q_{\min} = 1300$, $Q_{\max} = 1900$, $\lambda = 20$, $T = 1$ year with weekly exercise $N = 52$.

Writing the associated DDPE. The dynamic programming equation related to the Swing option valuation is described in terms of a Markov chain (O_t, Q_t) , which is controlled by the local consumption $q = (q_i)_{0 \leq i \leq N}$. Then the price $V_i = V(t_i, O, Q)$ is defined as the supremum of the expectation of CF_i over all admissible consumption q .

Following the dynamic programming principle (see [BS96], for more details in discrete-time stochastic control problem), we obtain the related DPE

$$V(t_N, O, Q) = P_N(O, Q),$$

$$V(t_i, O, Q) = \max_{q \in [q_{\min}, q_{\max}]} \left\{ q(\varphi(O) - \lambda) + \mathbb{E}[V(t_{i+1}, O_{i+1}, Q + q) \mid O_i = O] \right\}.$$

Here, we allow the local consumption to break the total volume constraint. Otherwise, we can simply take an infinity penalty (i.e., $\alpha_1 = \alpha_2 = +\infty$) which leads to the following constraint

$$q_{\min} \leq q \leq q_{\max}$$

$$Q_{\min} - (N - i)q_{\max} \leq Q + q \leq Q_{\max} - (N - i)q_{\min}.$$

This is a classical approach to solve swing option valuation problem (see [WW08], for finite element method).

Numerical results related to those applications are in prospectation. Please see the later version of [GPZ18] for further details.

Bibliography

- [AC01] Robert Almgren and Neil Chriss. Optimal execution of portfolio transactions. *Journal of Risk*, 3:5–40, 2001. (Cited on page 111)
- [AGPW18] Clémence Alasseur, Emmanuel Gobet, Isaque Pimentel, and Xavier Warin. Cost Management under asymmetric risk valuation. *Working paper*, 2018. (Cited on page 65)
- [AM11] Frédéric Abergel and Nicolas Millot. Nonquadratic local risk-minimization for hedging contingent claims in incomplete markets. *SIAM Journal on Financial Mathematics*, 2(1):342–356, 2011. (Cited on pages 3, 17, and 30)
- [BBD⁺06] Christophe Barrera-Esteve, Florent Bergeret, Charles Dossal, Emmanuel Gobet, Asma Meziou, Rémi Munos, and Damien Reboul-Salze. Numerical methods for the pricing of Swing options: a stochastic control approach. *Methodology and Computing in Applied Probability*, 8(4):517–540, 2006. (Cited on page 138)
- [Bel57] Richard Bellman. A Markovian decision process. *Journal of Mathematics and Mechanics*, 6(5):679–684, 1957. (Cited on page 123)
- [BG06] Svetlana Borovkova and Helyette Geman. Seasonal and stochastic effects in commodity forward curves. *Review of Derivatives Research*, 9(2):167–186, 2006. (Cited on page 124)
- [BHL⁺18] Alessandro Balata, Côme Huré, Mathieu Laurière, Huyên Pham, and Isaque Pimentel. A Class of Finite-Dimensional Numerically Solvable McKean-Vlasov Control Problems. *arXiv preprint arXiv:1803.00445*, 2018. (Cited on pages 10, 23, and 91)
- [BKS10] Denis Belomestny, Anastasia Kolodko, and John Schoenmakers. Regression methods for stochastic control problems and their convergence analysis. *SIAM Journal on Control and Optimization*, 48(5):3562–3588, 2010. (Cited on page 123)
- [BLPR17] Rainer Buckdahn, Juan Li, Shige Peng, and Catherine Rainer. Mean-field stochastic differential equations and associated pdes. *The Annals of Probability*, 45(2):824–878, 2017. (Cited on page 93)
- [BP17] Alessandro Balata and Jan Palczewski. Regress-Later Monte Carlo for optimal control of Markov processes. *arXiv preprint arXiv:1712.09705*, 2017. (Cited on pages 99 and 112)
- [BS96] Dimitir P Bertsekas and Steven Shreve. *Stochastic optimal control: the discrete-time case*. Athena Scientific, 1996. (Cited on page 138)

- [BT04] Bruno Bouchard and Nizar 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 page 123)
- [BW12] Bruno Bouchard and Xavier Warin. Monte-Carlo Valuation 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*, pages 215–255. Springer, Berlin, Heidelberg, 2012. (Cited on page 57)
- [Car10] Pierre Cardaliaguet. Notes on mean field games. Technical report, from P.-L. Lions’ lectures at Collège de France, 2010. (Cited on page 93)
- [CCD14] Jean-François Chassagneux, Dan Crisan, and François Delarue. A probabilistic approach to classical solutions of the master equation for large population equilibria. *arXiv preprint arXiv:1411.3009*, 2014. (Cited on page 93)
- [CDM17] Panagiotis Christodoulou, Nils Detering, and Thilo Meyer-Brandis. Quadratic hedging with multiple assets under illiquidity with applications in energy markets. *arXiv preprint arXiv:1705.06918*, 2017. (Cited on page 29)
- [CFS15] René Carmona, Jean-Pierre Fouque, and Li-Hsien Sun. Mean field games and systemic risk. *Communications in Mathematical Sciences*, 13(4):911–933, 2015. (Cited on page 110)
- [CMR09] Olivier Cappé, Eric Moulines, and Tobias Rydén. *Inference in Hidden Markov Models*. Springer-Verlag New York, 2009. (Cited on page 124)
- [Cre13] Stéphane Crepey. *Financial modeling: a Backward Stochastic Differential Equations perspective*. Springer. Springer-Verlag, 2013. (Cited on pages 3, 17, and 30)
- [CS00] Les Clewlow and Chris Strickland. *Energy Derivatives: Pricing and Risk Management*. Lacima Publications, 2000. (Cited on page 134)
- [Egl05] Daniel Egloff. Monte Carlo algorithms for optimal stopping and statistical learning. *The Annals of Applied Probability*, 15(2):1396–1432, 2005. (Cited on page 124)
- [EPQ97] Nicole El Karoui, Shige Peng, and Marie Claire Quenez. Backward Stochastic Differential Equations in finance. *Mathematical finance*, 7(1):1–71, 1997. (Cited on pages 30 and 31)
- [Fri08] Avner Friedman. *Partial Differential Equations of Parabolic type*. Courier Dover Publications, 2008. (Cited on page 63)
- [FS88] Hans Föllmer and Martin Schweizer. Hedging by sequential regression: An introduction to the mathematics of option trading. *Astin Bulletin*, 18(2):147–160, 1988. (Cited on page 30)
- [Gem05] Helyette Geman. *Commodities and commodity derivatives - Modeling and Pricing for Agriculturals, Metals and Energy*. Wiley Finance, 2005. (Cited on page 134)
- [GLW05] Emmanuel Gobet, Jean-Philippe Lemor, and Xavier 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 57)
- [GLW16] Hugo Gevret, Jerome Lelong, and Xavier Warin. STochastic OPTimization library in C++. Technical report, EDF Lab, 2016. (Cited on page 58)

- [GLZ18] Emmanuel Gobet, Gang Liu, and Jorge Zubelli. A Non-intrusive Stratified Resampler for regression Monte Carlo: application to solving nonlinear equations. *SIAM Journal on Numerical Analysis*, 56(1):50–77, 2018. (Cited on pages 10, 11, 24, 124, 125, 126, 131, 132, and 133)
- [GP16] Olivier Guéant and Jiang Pu. Portfolio choice, portfolio liquidation, and portfolio transition under drift uncertainty. *arXiv preprint arXiv:1611.07843*, 2016. (Cited on page 114)
- [GPW18] Emmanuel Gobet, Isaque Pimentel, and Xavier Warin. Option Valuation and Hedging using Asymmetric risk function: Asymptotic Optimality through Fully Nonlinear Partial Differential Equations. *HAL preprint hal-01761234*, 2018. (Cited on pages 4, 18, 29, 65, 67, 68, 69, and 72)
- [GPZ18] Emmanuel Gobet, Isaque Pimentel, and Jorge Zubelli. A Non-Intrusive Stratified Resampler for Multi-Factor Models: Application to the Pricing of Bermudan and Swing options. *Working paper*, 2018. (Cited on pages 13, 26, 123, and 138)
- [GS13] Jim Gatheral and Alexander Schied. Dynamical models of market impact and algorithms for order execution. In J-P. Fouque and J. Langsam, editors, *Handbook on Systemic Risk*, pages 579–599. Cambridge University Press, 2013. (Cited on page 112)
- [GT01] Emmanuel Gobet and Emmanuel Temam. Discrete time hedging errors for options with irregular payoffs. *Finance and Stochastics*, 5(3):357–367, 2001. (Cited on page 29)
- [GT16] Emmanuel Gobet and Plamen Turkedjiev. Linear regression MDP scheme for discrete backward stochastic differential equations under general conditions. *Mathematics of Computation*, 85(299):1359–1391, 2016. (Cited on pages 58 and 123)
- [GY04] Paul Glasserman and Bin Yu. Simulation for American options: Regression now or regression later? In H. Niederreiter, editor, *Monte Carlo and Quasi-Monte Carlo Methods 2002*, pages 213–226. Springer, Berlin, Heidelberg, 2004. (Cited on page 99)
- [Int16] International Energy Agency. *World Energy Outlook*. OECD IEA, Paris, 2016. ISBN: 978-92-64-26495-3. (Cited on page 66)
- [Kar88] Ioannis Karatzas. On the pricing of American options. *Applied mathematics and optimization*, 17(1):37–60, 1988. (Cited on page 134)
- [KLP14] Idris Kharroubi, Nicolas Langrené, and Huyên 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 page 98)
- [KLP15] Idris Kharroubi, Nicolas Langrené, and Huyên Pham. Discrete time approximation of fully nonlinear HJB equations via BSDEs with nonpositive jumps. *Annals of Applied Probability*, 25:2301–2338, 2015. (Cited on page 112)
- [KS98] Ioannis Karatzas and Steven Shreve. *Methods of mathematical finance*, volume 39. Springer Science & Business Media, 1998. (Cited on pages 2, 16, 29, and 31)
- [KSB09] Rüdiger Kiesel, Gero Schindlmayr, and Reik H Börger. A two-factor model for the electricity forward market. *Quantitative Finance*, 9(3):279–287, 2009. (Cited on page 125)

- [Lab17] Pierre-Henry Labordere. Deep Primal-Dual Algorithm for BSDEs: Applications of Machine Learning to CVA and IM. *SSRN*, 2017. (Cited on page 105)
- [Lam09] Damien Lamberton. Optimal stopping and american options. *Ljubljana Summer School on Financial Mathematics*, 2009. (Cited on page 134)
- [LGW06] Jean-Philippe Lemor, Emmanuel Gobet, and Xavier Warin. Rate of convergence of an empirical regression method for solving generalized backward stochastic differential equations. *Bernoulli*, 12(5):889–916, 2006. (Cited on page 58)
- [Lio12] Pierre-Louis Lions. Théorie des jeux de champ moyen et applications (mean field games). *Cours du College de France*, 2012. (Cited on page 93)
- [LM18] Michael Ludkovski and Aditya Maheshwari. Simulation methods for stochastic storage problems: A statistical learning perspective. *arXiv:1803.11309*, 2018. (Cited on page 102)
- [LS01] Francis Longstaff and Eduardo Schwartz. Valuing American options by simulation: a simple least-squares approach. *The review of financial studies*, 14(1):113–147, 2001. (Cited on pages 124 and 134)
- [MB46] Pierre Massé and Roger Boutteville. *Les réserves et la régulation de l'avenir dans la vie économique*. Hermann & Cie., 1946. (Cited on page 123)
- [NMS17] Selvaprabu Nadarajah, François Margot, and Nicola Secomandi. Comparison of least squares Monte Carlo methods with applications to energy real options. *European Journal of Operational Research*, 256(1):196–204, 2017. (Cited on pages 99 and 112)
- [PB04] Benoit Pochart and Jean-Philippe Bouchaud. Option pricing and hedging with minimum local expected shortfall. *Quantitative Finance*, 4(5):607–618, 2004. (Cited on page 30)
- [PBS01] Marc Potters, Jean-Philippe Bouchaud, and Dragan Sestovic. Hedged Monte-Carlo: low variance derivative pricing with objective probabilities. *Physica A: Statistical Mechanics and its Applications*, 289(3):517–525, 2001. (Cited on pages 57 and 124)
- [Pen55] Roger Penrose. A generalized inverse for matrices. In *Mathematical proceedings of the Cambridge philosophical society*, volume 51, pages 406–413. Cambridge University Press, 1955. (Cited on page 128)
- [Pen04] Shige Peng. Nonlinear Expectations, Nonlinear Evaluations and Risk Measures. In *Stochastic Methods in Finance*, pages 165–253. Springer, Berlin, Heidelberg, 2004. (Cited on page 30)
- [Pha00] Huyên Pham. Dynamic \mathbb{L}^p -hedging in discrete time under cone constraints. *SIAM Journal on Control and Optimization*, 38(3):665–682, 2000. (Cited on pages 3, 17, and 30)
- [Pha05] Huyên Pham. On some recent aspects of stochastic control and their applications. *Probability Surveys*, 2:506–549, 2005. (Cited on page 123)
- [Pha15] Huyên Pham. Feynman-Kac representation of fully nonlinear PDEs and applications. *Acta Mathematica Vietnamica*, 40(2):255–269, 2015. (Cited on page 123)
- [Pha16] Huyên Pham. Linear quadratic optimal control of conditional McKean-Vlasov equation with random coefficients and applications. *Probability, Uncertainty and Quantitative Risk*, 1(1):7, 2016. (Cited on page 111)

- [PPP04] Gilles Pagès, Huyên Pham, and Jacques Printems. Optimal quantization methods and applications to numerical problems in finance. *Handbook of computational and numerical methods in finance*, pages 253–297, 2004. (Cited on page 101)
- [PS15] Gilles Pagès and Abass Sagna. Markovian and product quantization of an \mathbb{R}^d -valued Euler scheme of a diffusion process with applications to finance. *pré-pub LPMA 1670*, 2015. (Cited on page 113)
- [Put14] Martin Puterman. *Markov decision processes: discrete stochastic dynamic programming*. John Wiley & Sons, 2014. (Cited on page 123)
- [PW17] Huyên Pham and Xiaoli Wei. Dynamic Programming for Optimal Control of Stochastic McKean–Vlasov Dynamics. *SIAM Journal on Control and Optimization*, 55(2):1069–1101, 2017. (Cited on pages 92 and 93)
- [Rog02] Leonard Cristopher Gordon Rogers. Monte Carlo Valuation of American Options. *Mathematical Finance*, 12:271–286, 2002. (Cited on page 105)
- [RY99] Daniel Revuz and Marc Yor. Continuous martingales and Brownian motion, vol. 293 of *Grundlehren der Mathematischen Wissenschaften*, 1999. (Cited on page 45)
- [Sch99] Martin 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. Cambridge University Press, 1999. (Cited on pages 2, 16, and 30)
- [SS00] Eduardo Schwartz and James Smith. Short-Term Variations Dynamics in and Long-Term Prices Commodity. *Management Science*, 46(7):893–911, 2000. (Cited on pages 124 and 127)
- [TV01] John Tsitsiklis and Benjamin Van Roy. Regression methods for pricing complex American-style options. *IEEE Transactions on Neural Networks*, 12(4):694–703, 2001. (Cited on pages 123 and 134)
- [WW08] Martina Wilhelm and Christoph Winter. Finite element valuation of Swing options. *Journal of Computational Finance*, 11(3):107, 2008. (Cited on page 138)

Titre : Valorisation optimale asymptotique avec risque asymétrique et applications en finance

Mots Clefs : Risque asymétrique, Optimalité asymptotique, Equations aux Dérivées Partielles Non-linéaires, Régression empirique, Marché d'électricité

Résumé : Dans cette thèse, nous étudions des problèmes de couverture et de valorisation d'options liés à une mesure de risque. Notre approche principale est l'utilisation d'une fonction de risque asymétrique et d'un cadre asymptotique dans lequel nous obtenons des solutions optimales à travers des équations aux dérivées partielles (EDP) non-linéaires.

En se concentrant sur la valorisation et la couverture des options européennes, nous considérons le problème de l'optimisation du risque résiduel généré par une couverture à temps discret en présence d'un critère asymétrique de risque. Au lieu d'analyser le comportement asymptotique de la solution du problème discret associé, nous étudions la mesure asymétrique du risque résiduel intégré dans un cadre Markovian. Dans ce contexte, nous montrons l'existence de cette mesure de risque asymptotique. Ainsi, nous décrivons une stratégie de couverture asymptotiquement optimale via la solution d'une EDP totalement nonlinéaire.

Ensuite, nous appliquons de cette méthode de couverture au problème de valorisation de la production d'une centrale. Puisque la centrale génère de coûts de maintenance qu'elle soit allumée ou non, nous nous sommes intéressés à la réduction du risque associé aux revenus incertains de cette centrale en se couvrant avec des contrats à terme. Nous étudions l'impact d'un coût de maintenance dépendant du prix d'électricité dans la stratégie couverture.

Title : Asymptotic optimal pricing with asymmetric risk and applications in finance

Keys words : Asymmetric Risk, Asymptotic optimality, Nonlinear Partial Differential Equations, Empirical regression, Electricity market

Abstract : In this thesis, we study several problems of hedging and pricing of options related to a risk measure. Our main approach is the use of an asymmetric risk function and an asymptotic framework in which we obtain optimal solutions through nonlinear partial differential equations (PDE).

Focusing on pricing and hedging European options, we consider the optimization problem of the residual risk generated by discrete-time hedging in the presence of an asymmetric risk criterion. Instead of analyzing the asymptotic behavior of the solution to the associated discrete problem, we study the integrated asymmetric measure of the residual risk in a Markovian framework. In this context, we show the existence of the asymptotic risk measure. Thus, we describe an asymptotically optimal hedging strategy via the solution to a fully nonlinear PDE.

Later, we apply this hedging method to the valuation problem of the power plant. Since the power plant generates maintenance costs whether it is on or off, we are interested in reducing the risk associated with its uncertain incomes by hedging with forwards contracts. We study the impact of a maintenance cost depending on the electricity price into the hedging strategy.

