HAL
open science

Modeling and optimal strategies in short-term energy markets

Laura Tinsi

To cite this version:

Laura Tinsi. Modeling and optimal strategies in short-term energy markets. Optimization and Control [math.OC]. Institut Polytechnique de Paris, 2021. English. NNT : 2021IPPAG005 . tel-03380450

HAL Id: tel-03380450

https://theses.hal.science/tel-03380450v1

Submitted on 15 Oct 2021

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

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

INSTITUT POLYTECHNIQUE DE PARIS

NNT : 2021IPPAG005

Thèse de doctorat

Modeling and optimal strategies in short-term energy markets

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à l'Ecole nationale de la statistique et de l'administration économique

École doctorale n°574 Ecole Doctorale de Mathématiques Hadamard (EDMH)
Spécialité de doctorat: Mathématiques appliquées

Thèse présentée et soutenue à Palaiseau, le 30 Septembre 2021, par

Laura Tinsi

Composition du Jury :

Gilles Pages Professeur, Sorbonne Université (LPSM)	Président
Almut Veraart Professeure, Imperial College London (Department of Mathematics)	Rapportrice
Huyên Pham Professeur, Université Paris 7 Diderot (LPMA)	Rapporteur
Marc Hoffmann Professeur, Université Paris-Dauphine PSL (CEREMADE)	Examinateur
Olivier Féron Ingénieur chercheur, EDFLab	Examinateur
Peter Tankov Professeur, ENSAE (CREST)	Directeur de thèse
Arnak Dalalyan Professeur, ENSAE (CREST)	Co-directeur de thèse

A ma mère, Corinne Mari.

Acknowledgements

J'aimerais tout d'abord remercier mes directeurs de thèse Peter Tankov et Arnak Dalalyan. Peter, par ta rigueur, ton écoute, ta disponibilité et ta constance, tu as su me tirer toujours plus vers le haut et m'as permis de m'accomplir dans mes recherches. Arnak, tu as su faire preuve d'une immense pédagogie pour m'initier à la recherche en statistique, et j'ai trouvé un grand intérêt à travailler avec toi sur ces sujets que tu maîtrises et connais si bien.

J'aimerais aussi remercier Olivier Féron, mon encadrant EDF. Olivier, ta connaissance des marchés de l'électricité m'a aidée à mieux comprendre la mécanique complexe qui permet à chacun de nous de profiter de cette richesse qu'est l'électricité, mais aussi à m'épanouir pleinement en donnant un sens réel à mes recherches. Ta bienveillance, ta psychologie, ton discernement et ton soutien ont été des éléments essentiels à l'aboutissement de cette thèse.

Un grand merci également à mes rapporteurs Almut Veraart et Huyên Pham, qui ont pris le temps de lire attentivement mon travail et d'en faire le rapport. Merci pour leurs commentaires et également pour leurs questions lors de la soutenance. Merci à mon examinateur Marc Hoffmann et au président du jury Gilles Pages pour leur implication, leurs questions, mais aussi leur sympathie durant la soutenance.

J'en profite également pour remercier le groupe R33 d'OSIRIS dans lequel j'ai eu la chance d'arriver en thèse, et que j'ai la chance de continuer à côtoyer. La bonne humeur qui y règne, tant à Saclay que lors d'afterworks parisiens, rend le travail agréable et épanouissant. Parmi les formidables personnes que j'ai pu rencontrer, j'aimerais remercier plus particulièrement Thomas, qui m'a donné l'opportunité de travailler sur des problèmes plus concrets de la R&D, tout en ponctuant nos échanges de blagues et jeux de mots toujours à propos. Pierre non plus n'est pas en reste tant pour les remerciements que je lui adresse que pour les blagues qu'il sait adapter en toute occasion! Merci mille fois Pierre pour ton coaching pointu lors de ma candidature à EDF. Une candidature durant laquelle je me suis sentie soutenue par Edouard et Clémence qui se sont montrés et se montrent toujours des managers bienveillants et à l'écoute. Enfin, j'aimerais dire un grand merci à Xavier, qui par son efficacité hors pair, même à distance pendant ces multiples confinements, a su trouver des solutions pour me permettre d'exploiter au mieux StOpt et sans lequel mon troisième papier aurait mis beaucoup plus de temps à voir le jour!

J'aimerais remercier les doctorants que j'ai pu côtoyer au cours de cette thèse. Ceux de l'ENSAE pour leur dynamisme et leur sympathie. Merci aussi aux doctorants d'EDF: à mon co-bureau Carl, à Margaux et à Maximilien avec lesquels les échanges sont toujours très intéressants. Merci aussi à Emma pour sa bonne humeur et ses conseils.

Présents bien avant ces trois belles années de thèse, j'aimerais maintenant remercier mes proches dont le soutien me porte depuis toujours.

Mes amis de l'ENSAE d'abord: Damien, Pierre, JC, Julien mais surtout Apolline et Marianne avec lesquelles j'ai découvert la vie à l'étranger et partagé d'inoubliables souvenirs à New York mais aussi à San Francisco et dans le fin fond des Etats-Unis.

Nicolas, un si long chemin parcouru depuis Dauphine. Un chemin géographique d'abord, puisque de place Dauphine nous avons progressivement migré vers le sud à Vanves puis à Saclay, allongeant les trajets mais en faisant d'autant plus d'occasions de discuter musique dans les embouteillages Porte d'Orléans. Une migration vers le sud qui ne s'arrête pas pour toi puisqu'elle te conduira bientôt vers l'Italie d'ailleurs. Un chemin intellectuel aussi qui nous a conduit aujourd'hui à soutenir nos thèses à quelques jours d'intervalle. Un chemin amical enfin, au travers des nombreux moments partagés à parler de la vie autour d'une bière. Alors merci pour tous ces chemins parcourus ensemble.

Merci à Marie et Marlène pour leur bonne humeur et leur joie de vivre, merci aux soirées et bons plans de Marie toujours à point nommé pour me changer les idées après une longue journée ponctuée d'équations.

Et j'aimerais dire, évidemment, l'incroyable chance que j'ai d'avoir des amis d'enfance aussi exceptionnels. Quentin, merci pour tes mots, tes réflexions qui m'ont depuis toujours permis de prendre de la hauteur. Lesia, Emmanuelle, Laure, Paola, Léa, Adriana, Cristiana et Andrea tous nos souvenirs, tous ces moments d'entraide, de joie et de partage sont un moteur pour moi. Je suis tellement heureuse de vous avoir dans ma vie et de tous les beaux moments à vivre ensemble qui nous attendent. Vous êtes des amies en or, je vous aime. Spéciale dédicace à Andrea, avec laquelle nous soutenons notre thèse à 2 semaines d'intervalle. Mac, je ne te remercierai jamais assez pour tous les fous rires, les debrief et les discussions aussi nombreuses que variées qui nous lient.

Armand, merci pour ton amour, ton incroyable gentillesse et pour ce quotidien que tu rends si beau. Je suis infiniment chanceuse de t'avoir à mes côtés. Merci aussi à ta famille, Domi, Vincent, tes frères, qui m'ont accueillie pendant ces longs mois de confinement en Mars 2020 et ont fait régner joie, bonne humeur (et barbecues!) durant cette période compliquée en tant que thésarde.

Merci à ma famille pour l'affection, le soutien et le réconfort qu'ils m'apportent. Boubou, merci pour tout l'amour dont tu m'as entourée depuis mon plus jeune âge. Merci pour ta gentillesse et ta bonne humeur constante qui contribuent pleinement à mon bonheur et à tous les beaux souvenirs que j'ai et que j'aurai à Cargèse. Merci pour la confiance que tu as en moi et en mes capacités, tu as toujours trouvé les mots pour me faire évoluer. Enfin, papa, cette thèse, bien que marquant l'aboutissement de mon parcours universitaire, n'est qu'une infime partie des remerciements que je devrais t'adresser. Merci pour ton indéfectible soutien, ta volonté de toujours me faire plaisir, ton éducation qui m'ont permis de devenir qui je suis aujourd'hui. Ton amour et ta présence ont été et seront toujours pour moi un pilier, un moteur et une joie. Tu

es un père extraordinaire.

Abstract

This thesis focuses on providing theoretical tools to help in the development and management of intermittent renewable energy in short term electricity markets.

In the first part, we develop a tractable equilibrium model for price formation in intraday electricity markets. For this, we propose a non cooperative game between several producers interacting in the market and facing an intermittent renewable production. Using stochastic control and game theory, we derive explicit optimal strategies for these producers as well as a closed form equilibrium price for different information structures and player characteristics. Our model allows to reproduce and explain the main stylized features of the intraday market such as the specific time dependence of volatility and the correlation between the price and the renewable production forecasts.

In the second part, we study dynamic probabilistic forecasts in the diffusion framework. We propose several stochastic differential equation models to capture the dynamic evolution of the uncertainty associated to a forecast, derive the associated predictive densities and calibrate the model on real meteorological data. We then apply it to the problem of a wind energy producer receiving sequential updates of the probabilistic forecasts of the wind speed used to predict her production and make trading decisions in the market. We show to what extent this method can outperform the use of point forecasts in decision-making processes.

Finally, in the last part, we propose to study the properties of aggregated shallow neural networks. We explore the PAC-Bayesian framework as an alternative to the classical empirical risk minimization approach. We focus on Gaussian priors and derive non-asymptotic risk bounds for the aggregated neural networks. These bounds yield minimax rates of estimation over Sobolev smoothness classes. This analysis also provides a theoretical basis for tuning the parameters and offers new perspectives for applications of aggregated neural networks to practical high dimensional problems increasingly present in energy decision problems involving renewables or storage.

Keywords: Intraday electricity market, market impact, renewable energy, mean field games, major player, probabilistic forecasting, ensemble forecasting, stochastic control, wind power trading, shallow neural networks, PAC Bayes, mirror averaging, exponentially weighted aggregate

Résumé

Cette thèse vise à fournir des outils théoriques pour soutenir le développement et la gestion des énergies renouvelables intermittentes sur les marchés court terme de l'électricité.

Dans la première partie, nous développons un modèle d'équilibre exploitable pour la formation des prix sur les marchés infrajournaliers de l'électricité. Pour cela, nous proposons un jeu non coopératif entre plusieurs producteurs interagissant sur le marché et faisant face à une production renouvelable intermittente. En utilisant la théorie des jeux et celle du contrôle stochastique, nous dérivons des stratégies optimales explicites pour ces producteurs ainsi qu'un prix d'équilibre en forme fermée pour différentes structures d'information et caractéristiques des joueurs. Notre modèle permet de reproduire et d'expliquer les principaux faits stylisés du marché intraday tels que la dépendance temporelle spécifique de la volatilité et la corrélation entre le prix et les prévisions de production renouvelable.

Dans la deuxième partie, nous étudions des prévisions probabilistes dynamiques sous la forme de processus de diffusion. Nous proposons plusieurs modèles d'équations différentielles stochastiques pour capturer l'évolution dynamique de l'incertitude associée à une prévision, nous dérivons les densités prédictives associées et nous calibrons le modèle sur des données météorologiques réelles. Nous l'appliquons ensuite au problème d'un producteur éolien recevant des mises à jour séquentielles des prévisions probabilistes de la vitesse du vent, utilisées pour prédire sa production, et prendre des décisions d'achat ou de vente sur le marché. Nous montrons dans quelle mesure cette méthode peut être avantageuse comparée à l'utilisation de prévisions ponctuelles dans les processus décisionnels.

Enfin, dans la dernière partie, nous proposons d'étudier les propriétés des réseaux de neurones peu profonds agrégés. Nous explorons le cadre PAC-Bayesien comme alternative à l'approche classique de minimisation du risque empirique. Nous nous concentrons sur les priors Gaussiens et dérivons des bornes de risque non asymptotiques pour les réseaux de neurones agrégés. Ces bornes atteignent des vitesses de convergence minimax pour l'estimation dans des espaces de Sobolev. Cette analyse fournit également une base théorique pour le réglage des paramètres et offre de nouvelles perspectives pour l'application des réseaux de neurones agrégés à des problèmes pratiques en grande dimension, de plus en plus présents dans les processus de décision liés à l'énergie et impliquant des moyens de production renouvelable ou du stockage.

Mots clés: Marché infrajournalier de l'électricité, impact de marché, énergie renouvelable, jeux à champ moyen, joueur majeur, prévisions probabilistes, prévisions d'ensemble, contrôle stochastique, trading de production éolienne, réseaux de neurones, PAC Bayes, moyenne miroir, agrégation à poids exponentiels

Contents

1 Introduction **1**

1.1 General context and motivation . . . 1

1.1.1 Energy and technology transition . . . 1

1.1.2 Successive short term electricity markets: a brief overview . . . 2

1.1.3 Main practical challenges of the thesis . . . 4

1.2 Theoretical background . . . 6

1.2.1 Stochastic games . . . 6

1.2.2 Probabilistic forecasts . . . 9

1.2.3 Neural networks and PAC-Bayesian inequalities . . . 12

1.3 Main contributions of the thesis . . . 15

1.3.1 Optimal strategies for renewable producers . . . 16

1.3.2 Price formation and main characteristics of the intraday market . . . 18

1.3.3 Probabilistic forecasts in sequential decision making problems . . . 20

1.3.4 Learning by aggregated shallow neural networks . . . 23

1.4 Outline of the thesis . . . 25

2 Introduction (Français) **27**

2.1 Context général et motivation . . . 27

2.1.1 Une transition énergétique et technologique . . . 27

2.1.2 Les marchés successifs court terme de l'électricité: un bref aperçu . . . 28

2.1.3 Les principaux enjeux pratiques de la thèse . . . 30

2.2 Cadre théorique . . . 32

2.2.1 Jeux stochastiques . . . 33

2.2.2 Prévisions probabilistes . . . 36

2.2.3 Réseaux de neurones et inégalités PAC-Bayesiennes . . . 39

2.3 Principales contributions de la thèse . . . 43

2.3.1 Stratégies optimales pour des producteurs renouvelables . . . 43

2.3.2 Formation des prix et principales caractéristiques du marché infrajournalier 46

2.3.3 Prévisions probabilistes dans des problèmes de décision séquentiels . . . 47

2.3.4 Apprentissage par réseaux de neurones agrégés à une couche cachée . . . 51

2.4 Structure de la thèse . . . 53

3 Price formation and optimal trading in intraday electricity markets **55**
3.1 Introduction . . . 56
3.2 Preliminaries on electricity markets . . . 59
3.3 Optimal trading strategies and equilibrium price . . . 60
3.3.1 N-player setting . . . 60
3.3.2 Mean-field game setting . . . 65
3.3.3 Relationship between N-player setting and MFG setting . . . 67
3.4 Intraday electricity prices: theoretical insights . . . 69
3.4.1 Price impact of forecast adjustments . . . 69
3.4.2 Volatility and Samuelson's effect . . . 70
3.4.3 Price-forecast covariance . . . 73
3.4.4 Trading costs . . . 73
3.5 Empirical results and numerical illustrations . . . 74
3.5.1 Stylized features of intraday electricity market prices . . . 74
3.5.2 Numerical illustration of our model . . . 77
3.6 Appendix . . . 83
3.6.1 Proof of Theorem 14 . . . 83
3.6.2 Proofs of Propositions 19 and 20 . . . 86

4 Price formation and optimal trading in intraday electricity markets with a major player **89**
4.1 Introduction . . . 90
4.2 Preliminaries . . . 92
4.3 A game of a major and minor agents . . . 95
4.4 Approximate Nash equilibrium in the N-player Stackelberg game . . . 100
4.5 Numerical illustration . . . 102
4.6 Conclusion . . . 106
4.7 Appendix . . . 107
4.7.1 Proofs of Lemma 29 and Proposition 30 . . . 107
4.7.2 Proof of Theorem 31 . . . 109
4.7.3 Proofs of Corollary 34 and Corollary 35 . . . 110
4.7.4 Proof of Proposition 38 . . . 110

5 Decision making with dynamic probabilistic forecasts **113**
5.1 Introduction . . . 114
5.2 Modeling probabilistic forecasts . . . 117
5.2.1 Forecast of a real-valued quantity . . . 118
5.2.2 Forecast of a positive quantity . . . 122
5.3 Fitting forecast models to data . . . 125
5.3.1 Presentation of the dataset . . . 125
5.3.2 Model calibration . . . 127

5.3.3 Numerical illustrations . . . 130
5.4 Application to wind power trading . . . 137
5.4.1 Description of the problem . . . 137
5.4.2 Numerical resolution . . . 139
5.5 Appendix . . . 143
5.5.1 Talagrand diagrams and PIT histograms for lead times 36h and 48h . . . 143

6 Risk bounds for aggregated shallow neural networks using Gaussian priors 145
6.1 Introduction . . . 146
6.2 Preliminaries and notations . . . 148
6.2.1 General framework . . . 149
6.2.2 PAC-Bayesian type upper bounds . . . 149
6.2.3 Shallow neural networks . . . 150
6.2.4 Spherical Gaussian prior distribution . . . 150
6.3 Examples of application . . . 151
6.3.1 Fixed design regression . . . 151
6.3.2 Random design regression . . . 151
6.3.3 Density estimation . . . 152
6.3.4 Classification for Φ-risk . . . 152
6.4 Risk bounds . . . 153
6.4.1 Risk bound for an arbitrary centered Gaussian prior . . . 154
6.4.2 Oracle inequality for optimized Gaussian prior . . . 155
6.5 Approximation bounds . . . 156
6.5.1 Bounds for sigmoidal activation functions . . . 157
6.5.2 Bounds for the ReLU activation function . . . 159
6.6 Worst-case risk bounds over smoothness classes . . . 161
6.6.1 Sigmoid activation functions . . . 161
6.6.2 ReLU activation function . . . 162
6.6.3 Related works for risk bound of neural networks . . . 164
6.7 Conclusion and outlook . . . 165
6.8 Appendix . . . 165
6.8.1 Some useful lemmas . . . 166
6.8.2 Proof of Proposition 47 . . . 169
6.8.3 Proof of Corollary 48 . . . 170
6.8.4 Proof of Lemma 51 . . . 172
6.8.5 Proof of Proposition 55 . . . 173

Chapter 1

Introduction

This thesis focuses on providing scientific tools to support the increasing penetration of intermittent renewable energy in the electric supply system, by trying to understand and manage its impact on short term electricity markets.

1.1 General context and motivation

1.1.1 Energy and technology transition

We are witnessing a surge of political actions on climate change these last years: the European Union climate and energy package (2008 and 2014), the Paris agreement (2015), the European Green deal (2020); and in France more particularly, the act on energy transition for green growth (2015), the climate plan (2017). The European countries are determined and ambitious: during the last decade, they all committed to national actions to increase the contribution of renewables in their energy mix, or limit carbon emission; some of them even would like to reach carbon neutrality by 2050. In France, to meet this goal, it is planned to take out of the market diesel and petrol powered cars by 2040, taking the side of a 100 % electric vehicle fleet. The electric car, spearhead of a carbon free system, also contributes to a higher and more uncertain electricity demand. The country will also support personal initiatives to consume self-produced energy, thanks to photo-voltaic systems for example. Concomitantly, a societal and technological transition emerges: we evolve towards a more energy and electricity dependent way of life. The rise of distance working these last two years is a very concrete example of this change: it made us more connected but also allowed for more flexibility giving birth to less predictable new consumption habits.

We witness the development of a fortunately greener, but also more uncertain energy mix, facing less predictable demand habits and a stronger dependence on the energy supply. A question arises then: how to make these two adverse trends sustainable ?

Smart grids for more system flexibility, progress in affordable means of storage, and short term electricity markets are key elements to tackle this challenge, but they require rigorous upstream practical and theoretical research. My thesis is a contribution to this research and mainly focuses on the management of intermittent renewable energy by and for renewable producers

using the short term electricity markets. This work addresses four main topics: modeling price formation in short term electricity markets, modeling renewable production forecasts, providing optimal strategies for renewable producers, and finally, coming up with quantitative and insightful guarantees for estimation by neural networks.

By way of a preamble, the next section is dedicated to a succinct description of short term electricity markets.

1.1.2 Successive short term electricity markets: a brief overview

Since the liberalization of the electricity markets in Europe and in France, initiated in the 90's, the non storable nature of electricity and the energy transition policy led to the creation of several new markets. Among them, we can distinguish the trading markets composed of the day ahead (also named "spot"), and the intraday market; as well as the ancillary service markets used to manage the reserves and support the balancing mechanism and its settlement. These short term markets constitute an important tool to manage the uncertainty and risks related to energy and technological transition. As long term energy markets such as the futures market, or over the counter forward contracts, appear to be essential to hedge long term contracts and scheduling, short term energy markets are crucial to ensure the constant balance between supply and demand of electricity.

Concerning the trading markets, several exchange platforms have emerged in Europe since 2001 with the creation of the day ahead market, followed by the intraday in 2010. These platforms can cover a country such as Powernext – initially created in France, or a specific area such as Northpool for Scandinavian countries or Apx-endex for England, Belgium and the Netherlands. These cross-border markets are especially important since they allow to smooth the available energy supply and demand between countries. As a consequence, a larger European market, EPEXSPOT, born from the merger of some of the above mentioned exchange platforms, allowing for national and cross-border trades, was created in 2008 and is today the leader of electricity markets in Europe [1]. Concerning the ancillary service markets and the adjustment system, the Electricity Balancing agreement (2017) of the EU regulation commission, established a common area and common rules on balancing settlement in Europe, where reserves are constituted by tender. These very recent evolutions demonstrate the still young age and ongoing expansion of these markets.

For each delivery period of the day, generally divided in hours, several actors are in charge of ensuring the energy equilibrium. The chronology of short term markets emphasizes their different roles: the producers and the electricity suppliers mainly intervene in the trading markets, before the delivery period; while the Transport System Operator (TSO) intervenes in real time, or a posteriori, in the management of reserves and balancing mechanisms. As they are of paramount interest for our work, in the sequel, we detail each step of these successive markets and illustrate some mechanisms with the example of a wind energy producer, *Producer A*, facing an uncertain

[1] All empirical analyses made in this thesis are based on EPEXSPOT intraday order books for France or Germany for winters of 2015 and 2017.

terminal production at a delivery time T, whose forecast is regularly updated. The following description is also summed up in a more synthetic way in Figure 1.1.

Day ahead/Spot market It is a blind auction market that opens one day before the delivery time and closes at 12 AM. Each delivery hour of the day after is a product of this market. For each of the 24 products, from the bid and ask volumes and prices, the supply and demand curves determine the market clearing price (MCP).
Producer A: she bids a certain volume of power X_0 that she commits to supply at the delivery time T corresponding to a certain hour of the day after. If the bid is accepted, for this hour and her volume, she gets paid at the MCP, noted S_0 in €/MWh and receives $X_0 S_0$.

Intraday electricity market It is a continuous trading market with an on going order book that matches the orders as they come (no clearing price). It opens at 3 PM the day before the delivery, closes 5 minutes or 15 minutes before the delivery time depending on the area, and, proposes hourly products – as in the day ahead market, half-hour products and even quarter-hour products.
Producer A: at some time $t \in [0, T]$, she receives an update of her production forecast denoted X_t, such that $\Delta X_t = X_0 - X_t > 0$, which means that she plans to produce less than what she committed to on the spot market. In reaction, she takes a position ϕ_t in the intraday market to adjust the forecast error ΔX_t. When the market closes, she has a position ϕ_T in the intraday market and a global engagement $X_0 - \phi_T$ of production, that must be compared to the real quantity X_T she will produce.

Ancillary services and balancing mechanism The balance responsible parties (BRPs) acting in electricity markets, are responsible for the balance between their supply and offer commitments. During the delivery hour, the TSO (or operators under contract with the TSO) is in charge of providing for excess/lack of supply. For this, primary and secondary reserves, called uniformly "Frequency Containment Reserve (FCR)" and "Frequency Restoration Reserve (FRR)", can be used automatically or semi-automatically to inject/remove power from the grid, at a quick time-response from real-time to less than 15 minutes. If the imbalance is too important or lasts to long, a slower and manually activated third reserve, named "Replacement Reserve (RR)", is used to restore or support the required level of FRR to be prepared for possible additional system imbalances. The global balancing process then requires reserves proposed by producers and/or consumers that must be paid for it. On the other hand, the BRPs are "penalized" with respect to their proper imbalance. The imbalance penalty is fixed by the TSO (generally a few hours after the delivery time) depending on the costs incurred by the TSO for managing the overal system imbalance.
Producer A: assume that her effective production is such that $X_0 - \phi_T > X_T$, she will then have to pay the penalty P computed by the TSO for the difference $X_0 - \phi_T - X_T$ at a global cost $(X_0 - \phi_T - X_T)P$.

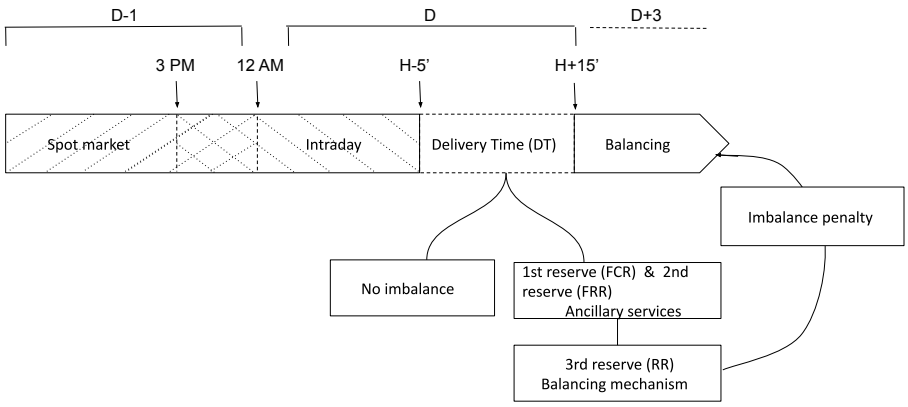

Figure 1.1: Schematic chronology of the short term electricity markets

In the sequel, we present the main challenges related to renewable energy management in short term electricity markets.

1.1.3 Main practical challenges of the thesis

From the previous section we understand why the rise of intermittent renewable energy and the emergence of new usages of electricity put a stronger emphasis on short term electricity markets. The increasing participation of renewable producers in the intraday market brings several challenges that need to be taken up in order to succeed a sustainable transition towards a carbon-free energy system. One consequence of the deeper penetration of renewables is the emergence of stylized features on intraday prices, such as the negative correlation between prices and renewable production forecasts (see Figure 1.2a). We refer the reader to Kiesel and Paraschiv (2017); Karanfil and Li (2017); Rowińska et al. (2018) for further empirical studies on short term electricity markets.

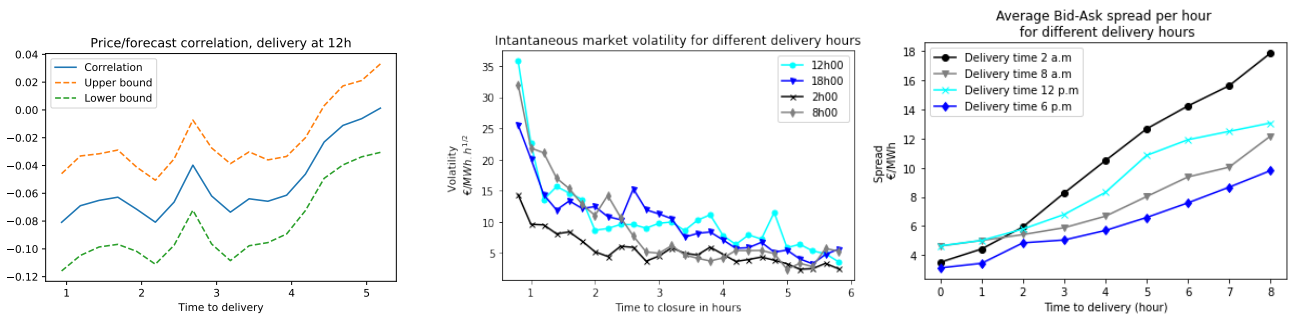

(a) Correlation between price and wind renewable production forecasts

(b) Empirical volatility

(c) Bid-ask spread

Figure 1.2: Empirical features for the German delivery zone in winters 2015 and 2017

Furthermore, to avoid costs caused by balancing penalties, most actors might tend to adjust their position until short time before the delivery. Apart from the correlation between price and production forecasts, this behavior may also explain other stylized features such as the increasing volatility Figure 1.2b, and the increasing liquidity Figure 1.2c, materialized by a decreasing bid-ask spread during the last trading hours. From the market characteristics, we figure out a negative feedback mechanism: the trading of numerous renewable producers impacts the price that in return impacts their revenues. These first observations trigger the two main aspirations of Chapter 3 and Chapter 4:

Problem 1. *Find strategies for producers interacting on the market that allow them to minimize their costs while taking into account their uncertain production and their market impact*

Problem 2. *Capture, using our model, how interactions of renewable producers may be a driving element in explaining the features observed in the intraday market*

Making this higher level of uncertainty economically sustainable in the long term is conditioned by three main issues. First, find convenient ways to model the market price and design optimal strategies for producers as we just saw. Second, in order to avoid extensive use of balancing mechanisms and make the market/energy mix more efficient, it is absolutely crucial to work on the ability of the production forecast to: (i) properly reflect the production to come, (ii) dynamically quantify the variations of the forecast uncertainty. Third, provide theoretical tools capable of ensuring good estimates in the increasingly complex problems that arise in the management of renewable generation.

For the second issue, the participation of renewable producers obviously depends on the vision they have of their production but is also highly driven by how trustworthy they consider it. In that respect, one needs to quantify the uncertainty related to their forecasts in a way that enables them to take well-informed decisions. This second observation drives us to problems that we will address in Chapter 5:

Problem 3. *Provide forecasts that allow to track dynamically the variations of the uncertainty based on real data*

Problem 4. *Be able to exploit the forecast as well as its uncertainty evolution in decision making processes*

Finally, the third line of research that may support the energy transition builds on the following observations. Prices and energy supply for close delivery hours are correlated and this effect might be strengthened by the future development of affordable means of storage. Therefore, modeling successive short term electricity markets, each involving at least 24 possibly correlated delivery products, and the same number of production forecasts, should be preferred as compared to the case of a specific delivery hour treated in Problems 1–4 of this thesis. However, this increases the complexity and the difficulty of estimating model parameters and making appropriate decisions. These last years, machine learning and, more specifically, neural networks gained popularity and are now often used for tackling these kind of high-dimensional

problems. However, theoretical guarantees justifying the successful empirical results of neural networks are still scarce. We believe that the evolution and perennial use of neural networks needs to go along with theoretical advances to legitimize the existing empirical evidence of their usefulness. This is why, in the last chapter, in contrast with the previous challenges stated, we will focus exclusively on the aforementioned theoretical analysis, leaving the application to energy markets to a future work:

Problem 5. *Exhibit risk bounds for learning by aggregated shallow neural networks*

Problem 6. *Provide practical indications on tuning parameters –e.g. weights initialization, size of the networks–, drawn from theoretical analysis*

To answer the questions raised in this section, we rely on mathematical tools, mainly coming from stochastic control theory, game theory and statistics. In the next section, we present them and provide some relevant literature references.

1.2 Theoretical background

In this section we present the main theoretical background of our developments. We do not intend to make an exhaustive review of the state of the art but simply mention the results that seem to be the most relevant to contextualize the present work. The scope of the thesis goes from deriving optimal trading strategies for renewable producers, to machine learning theory, through calibration and simulations. The two first chapters mainly focus on stochastic control problems, the third chapter is balanced between control and statistics and finally, the fourth chapter concentrates on statistics. We follow the organization of the chapters to present the theoretical tools.

1.2.1 Stochastic games

As mentioned in the previous section, we are interested in several renewable producers equipped with uncertain production forecasts that intend to minimize their cost by interacting on the market. In this setting, it is natural to consider them as the players of a non cooperative game. For this, the stochastic games offer a suitable framework: in the sequel we specify two main types of stochastic games that are of interest for the problems we will develop in the body of the thesis.

Optimal control Stochastic games are rooted in the stochastic control theory, whose value for modern financial mathematics is widely acknowledged. For a problem modeled by an objective function to minimize and a state space, including one or several controlled variables, one aims to exhibit, if it exists, a control that guarantees optimality. The notion of value function for a control problem stems from the dynamic programming principle. Heuristically, the Feynmac-Kac formula then bridges the gap between the dynamic programming principle and its infinitesimal counterpart, so that the solution usually takes the form of a partial differential equation (PDE):

the so-called Hamilton-Jacobi-Bellman (HJB) equation. The Pontryagin principle offers another approach in the form of a Backward Stochastic Differential Equation (BSDE). The latter might appear as more intuitive when thinking of the usual backward procedure employed to tackle dynamic programming problems such as the classical financial problem of replicating contingent claims. We refer the reader to Touzi (2012) for a comprehensive introduction to classical, or more sophisticated, stochastic control problems.

Optimal trading problems in financial markets for a single agent with price impact have been extensively studied (see *e.g.* Bank et al. (2017)); and in most cases, a model *à la* Almgren and Chriss (1999, 2001) is used, *i.e* with a linear price impact. This price maker setting has also been applied to energy markets: in the works of Aid et al. (2016); Tan and Tankov (2018), a producer facing an uncertain terminal production/demand has to handle her impact on the price. Our first area of research is very close to these works in spirit, but we consider this time a market with several agents, where the movement of the crowd can impact each of them.

Nash equilibrium This notion has been defined in the seminal paper of Nash (1950), and matches with the non cooperative game we aim to build, where the existence of an equilibrium in the market depends on the interactions of the participants. Assume that the game occurs from a date $t = 0$, to a date T just before the closure of the market. Consider N players, a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} := (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{P})$, and $\phi^i := (\phi^i_t)_{t\in[0,T]}$, $i = 1, \dots, N$, the $\mathbb{F}$-adapted strategic positions of the players in the market. Each player benefits from a reward given by the function $J^{N,i}$, depending on her strategy and the strategies of the other players. A Nash equilibrium is a state such that none of the players has an incentive to move from her position, that is, more formally:

Definition 7 (Nash Equilibrium)**.** *We say that $(\phi^{i*})_{i=1,\dots,N}$ is a Nash Equilibrium for the N-player game if it is a vector of admissible strategies – based on some admissibility criteria, and for each $i = 1, \dots, N$,*

$$J^{N,i}(\phi^i, \phi^{-i*}) \leq J^{N,i}(\phi^{i*}, \phi^{-i*}) \tag{1.1}$$

for any other admissible strategy ϕ^i, where $\phi^{-i} := (\phi^1, \dots, \phi^{i-1}, \phi^{i+1}, \dots, \phi^N)$.

Another interesting notion, and to which we shall come back later (see Section 1.3.1), is the one of ε-Nash equilibrium that consists in finding a set of admissible positions $\left(\phi^{i*}\right)_{i=1,\dots,N}$ that is ε-optimal in the sense of the following definition:

Definition 8 (ε-Nash equilibrium)**.** *We say that $(\phi^{i*})_{i=1,\dots,N}$ is an ε-Nash equilibrium for the N-player game if these strategies are admissible and for any other admissible strategy ϕ^i of the player i, $J^{N,i}(\phi^i, \phi^{-i*}) - \varepsilon \leq J^{N,i}(\phi^{i*}, \phi^{-i*})$, $i = 1, \dots, N$.*

Equilibria in stochastic games are not easy to tackle but the linear quadratic setting is a classical tool to obtain explicit solutions (Bouchard et al., 2018; Voß, 2019; Evangelista and Thamsten, 2020; Bank et al., 2021). However, in some cases, this simplified design is not

sufficient to exhibit an equilibrium. In game theory, one can distinguish the setting where all players have a complete information on the state and the strategy of the others, from the incomplete setting where a certain part of the strategic behavior of the other players is hidden. To relate this consideration with the initial problem of the renewable producers, note that their own production forecast may be either totally observed, partially observed, or not observed at all by the other actors; and the same remark could be made on the trade of the crowd that is, the aggregate position of all the players. The first case is usually easily tackled with Nash theory and one can exhibit the equilibrium in closed-form in the linear quadratic setting. On the contrary, the last cases generally cannot be explicitly solved using the present framework. To circumvent this shortcoming, we present tools coming from the mean field games theory in the next paragraph.

Mean field games They were introduced by Lasry and Lions (2007); Huang et al. (2006) and differ from the Nash theory in that they consider a game with infinitely many players, traditionally symmetric. For the Nash equilibrium case, the equilibrium results from the interactions of N players through their positions $(\phi^i)_{i=1,\dots,N}$. For example, in Chapter 3, agents will interact through the price impact expressed as a linear function of the aggregate position of the agents: $\frac{1}{N}\sum_{i=1}^N \phi_t^i$, $\forall t \in [0,T]$. It is convenient to understand the latter as an expectation with respect to the empirical distribution $\mu_t^N = \frac{1}{N}\sum_{i=1}^N \delta_{\phi_t^i}$, for which, heuristically, a natural extension in the mean field setting is the limiting flow of distribution $(\mu_t^\phi)_{t\in[0,T]}$ of the representative agent position $\phi := (\phi_t)_{t\in[0,T]}$. If we intend to develop a model where information is not totally shared by the agents, in addition to the filtration $\mathbb{F}$ to which all processes are adapted, we introduce a smaller filtration $\mathbb{F}^0$, representing the so-called common noise, and to which private information is not adapted. It follows that the mean field analogue of aggregate quantities in the N-agent problem are conditional expectations with respect to the common noise filtration $\mathbb{F}^0$. The game is now represented by the interaction of agents through the conditional distribution flow of the state process: $\mu_t^{\phi,0} := \mathcal{L}(\phi_t|\mathcal{F}_t^0)$. If we denote J^{MF} the objective function of the representative agent, we can define a mean field equilibrium similarly to the Nash one,

Definition 9 (mean field equilibrium)**.** *An admissible strategy ϕ^* is a mean field equilibrium if for any admissible strategy ϕ,*

$$J^{MF}(\phi, \bar{\phi}^*) \leq J^{MF}(\phi^*, \bar{\phi}^*),$$

with $\bar{\phi}_t = \mathbb{E}[\phi_t|\mathcal{F}_t^0] = \int_{\mathbb{R}} x\mu_t^{\phi,0}(dx)$.

To deepen the theoretical aspects, we refer the reader to Lasry and Lions (2007); Carmona and Delarue (2018); Cardaliaguet et al. (2019); Lacker (2020).

To tackle this kind of stochastic games, the methods are in the end, also very close to the ones evoked in the simpler framework of stochastic control. The seminal approach of Lasry and Lions (2007); Huang et al. (2006) relies on the HJB equation for the value function, coupled with a Fokker-Planck equation to formalize the mean field distribution through which players interact. Mirroring the stochastic maximum principle, the probabilistic approach of Carmona and

Delarue (2018) formalizes the problem through a McKean-Vlasov forward backward stochastic differential equation (FBSDE). The McKean-Vlasov extension of the classic FBSDE is introduced to manage the unknown mean field distribution flow $(\mu_t^{\phi})_{t\in[0,T]}$. The latter approach is closer in spirit to the method employed in the Chapters 3 and 4.

The way mean field games permit to simplify complex problems makes them very attractive in applications to finance (Casgrain and Jaimungal, 2018, 2020; Fu and Horst, 2020b; Fu et al., 2020; Fujii and Takahashi, 2020; Fu et al., 2021), to energy markets (Shrivats et al., 2020), or, to treat energy storage allocation problems (Alasseur et al., 2020). However, the classical setting of mean field games considers a class of symmetric agents. In Alasseur et al. (2020); Casgrain and Jaimungal (2018, 2020), they introduce heterogeneity through sub-populations with partial information and different beliefs. For the problem we address in the thesis, one can argue on which setting is the most relevant given real market conditions. A striking example of these conditions is illustrated by the French market where EDF has a dominant position for the production and supply of energy, which makes the company a major actor of the market compared to other producers. With the development of renewable energies – especially photovoltaic, an opposite phenomenon arises as the number of very small producers that are or may be one day connected to the market and the grid is increasing. In this case, heterogeneity comes from the size and thus the market power of the actors. Hopefully, this situation can be handled thanks to the special framework of mean field games with a major player also known as the Stackelberg mean field games (see *e.g.* Huang (2010); Nourian and Caines (2013); Bensoussan et al. (2016b); Carmona et al. (2016) for valuable insights on the topic, an application to financial markets can be found in Evangelista and Thamsten (2020)). Chapter 4 is dedicated to this heterogeneous setting, and its consequences on the strategic behavior of agents.

Stochastic games form a convenient tool to address the Problems 1–2. However, the gains of producers also crucially depend on the quality of the forecast they use to make their decisions on the market. This observation relates to Problems 3–4, for which we introduce in the next paragraph the notion of probabilistic forecasts, their most common approach and the calibration techniques they generally involve.

1.2.2 Probabilistic forecasts

Let assume one needs to forecast a random variable X at a given date t before its realization at a time $T > t$. The forecast can come in the form of a single 'point' information that is a view on the value of X at time t; or, it can also provide indications on how trustworthy is this information, which is the goal of probabilistic forecasts. As opposed to point forecasts, the latter enables one to take into account the uncertainty carried by a forecast, which permits to address Problem 3. Most of the time, probabilistic forecasts are presented in the form of confidence intervals or quantiles around the point forecast as it is the case in Figure 1.3. In the literature, a fair amount of the research has focused on how to model, calibrate and assess the quality of probabilistic forecasts.

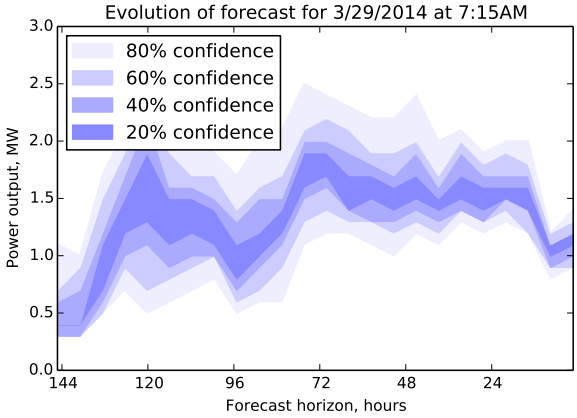

Figure 1.3: Evolution of the probabilistic forecast of power production of a wind plant as function of time, for a fixed production time

To estimate the entire predictive density is generally very costly since it is an infinite dimensional object. Yet, most of the time, the actual available information is low-dimensional. For this reason, a two-dimensional specification is sometimes chosen, based on two observable factors such as the mean and the variance of the random variable X given the information available at time t, that we can respectively denote m and V. The seminal papers of Gneiting et al. (2005, 2007) discuss the calibration and assessment of such models. For a given forecast horizon, they use a Gaussian distribution for the predictive density, $\mathcal{N}(m, V)$, where m, V have to be determined by calibration. Following this work, several authors developed more elaborate models using mixture of Gaussians (Wilks, 2002; Raftery et al., 2005), truncated distributions (Baran, 2014), log-normal ones (Baran and Lerch, 2015), mixtures of truncated and log-normal distributions (Baran and Lerch, 2016); for different purposes such as handling positive quantities or extreme value events. In these references, probabilistic forecasts are generally used for meteorological data. This makes it very appropriate to the context of renewable energy or electricity demand that largely depend on meteorological phenomena such as the wind or the temperature, on which we will focus in Chapter 5.

Calibration method Most of the time, meteorological data is represented in the form of ensemble forecasts. An example is displayed in Figure 1.4a for the wind speed: it shows an ensemble forecast for a forecast horizon of 48h, composed of $M = 50$ members that we will denote $X_1, \ldots, X_M$. The different scenarii evolve with the time horizon and the differences between them vary with time. For example, quite surprisingly, uncertainty 20 hours before the realization seems less prominent than uncertainty 8 hours before. Let us denote $\overline{X}_M = \frac{1}{M}\sum_{m=1}^{M} X_m$ and $S_M = \frac{1}{M-1}\sum_{m=1}^{M}(X_m - \overline{X}_M)^2$ the empirical mean and variance in the ensemble forecast.

A first naive attempt would be, in the case of the simple Gaussian model from Gneiting et al. (2005), to chose as a predictive density $\mathcal{N}(\overline{X}_M, S_M)$. However, members of the ensemble are obtained by running a deterministic forecasting model with perturbed initial conditions. Because of this, they do not represent the best approximation of the realization and the empirical

distribution of the ensemble members often suffers from a bias and underdispersion as the Figure 1.4b shows.

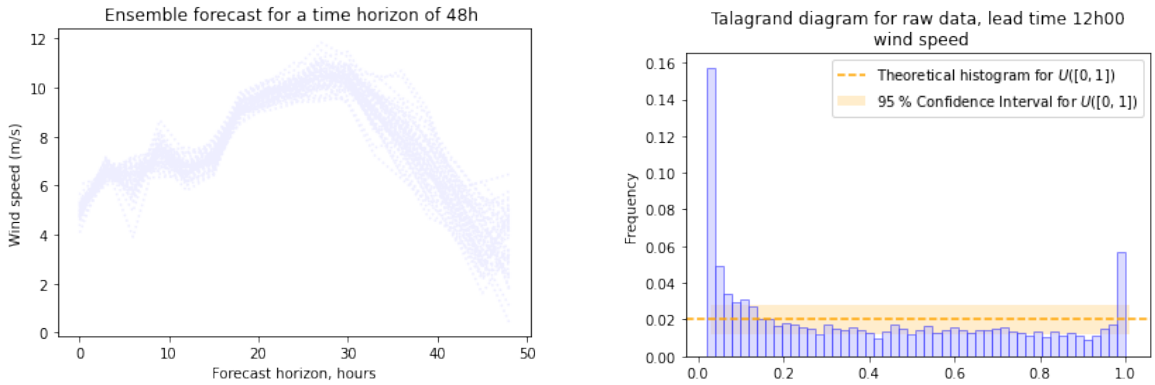

(a) Ensemble Forecasts

(b) Talagrand Diagram

Figure 1.4: An example of ensemble forecasts and the shortcomings associated to this method

The Talagrand diagram displayed in this figure is a tool to check the quality of a calibration and is analogous to a rank histogram $R(y, (X^m)_{m=1\ldots M})$ of the actual realization y within the ensemble forecast $(X^m)_{m=1\ldots M}$. In a perfect ensemble forecast, the histogram should be the one of a uniform distribution. In the present case, the histogram in Figure 1.4b, presents a U-shaped profile, which corresponds to under-dispersion of the ensemble and an asymmetric form that indicates a bias. To address this problem, depending on the type of predictive density used, several calibration methods exist. Among them, the most common ones are the Ensemble Model Output Statistic (EMOS) and the Bayesian Model Averaging (BMA). In this work, we will use a technique inspired by the EMOS post-processing method that simply amounts to estimate the distribution with a linear correction of the mean and variance – *e.g.* in Gneiting et al. (2005), they calibrate the normal distribution: $\mathcal{N}(a_0 + \sum_{m=1}^{M} a_m X_m, c + dS_M)$ where $a_0, \ldots, a_M, c, d$ are regression coefficients.

The diffusion setting In the works previously mentioned, there is no dynamical dimension in the models, which therefore do not account for temporal dependencies between the forecasts at different time horizons. This weakness can be addressed using the stochastic diffusion framework, which has been developed in the setting of wind speed forecasts (Iversen et al., 2016; Bensoussan and Brouste, 2016) and solar irradiance forecasts (Badosa et al., 2017). These works highlight how useful is this area of research for practical/industrial applications in the field of renewable energy. They mainly focus on modeling the evolution of point forecasts using stochastic differential equations (SDE), allowing to derive the probabilistic forecasts, often in the form of quantiles. Bensoussan and Brouste (2016) developed a Cox-Ingersoll-Ross (CIR) model for the wind speed which allows to derive a flow of predictive densities that can be properly calibrated. In Chapter 5, we propose several SDE models accounting for the dynamic evolution of the forecast and of its uncertainty. Our models also permit to derive predictive densities, for which we provide the following definition in the stochastic framework.

Definition 10. *Let $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ be a filtered probability space and fix a time horizon T. Then, a probabilistic forecast of an $\mathcal{F}_T$-measurable random variable $X \in \mathbb{R}^d$ is the conditional distribution p_t of X given $\mathcal{F}_t$. A dynamic model for a probabilistic forecast is then a flow of probability measures $(p_t)_{0 \leq t \leq T}$, which can be identified with the flow of conditional distributions of X.*

The diffusion setting also permits to integrate probabilistic forecasts in stochastic decision-making processes. We propose to address Problem 4 in Chapter 5, by the numerical resolution of a three dimensional sequential control problem. The Least Squares Monte Carlo regression we use only works well up to dimension 3 or 4, hence, the dimension becomes an obstacle for the development of more complex management problems. The need for new methods in learning problems has led to the development of machine learning and more specifically of neural networks, whose outstanding empirical results are not to be proven but often lack theoretical guarantees. In Chapter 6, we exhibit risk bounds for neural networks in the Probably Approximately Correct-Bayesian (PAC-Bayesian) setting that we overview in the next section.

1.2.3 Neural networks and PAC-Bayesian inequalities

The last part of this thesis is devoted to establishing performance guarantees for aggregated neural networks. As such, it does not directly relate to electricity markets. Nevertheless, advanced machine learning techniques have shown great empirical results for dealing with high dimensional problems and some recent work uses neural networks in stochastic control problems with highly appealing applications to energy storage (Bachouch et al., 2021) and risk management (Fécamp et al., 2019). Some of these works even provide convergence analysis for specific algorithms (see Huré et al. (2021) for dynamic programming, Becker et al. (2019, 2021) for optimal stopping problems and Huré et al. (2020); Germain et al. (2020) for BSDE approaches). Despite these remarkable advances, they often rely on the universal approximation theorem (Hornik et al., 1989) and discretization schemes (Bouchard and Touzi, 2004) to exhibit convergence guarantees, which do not highlight the finite-sample behavior of the neural networks and the impact of their structure on the test error. Establishing general theoretical guarantees coping with these shortcomings is an active area of research in machine learning, with the existing results mostly focusing on the properties of the neural network minimizing the training error.

Chapter 6 thus proposes to investigate risk bounds in the framework of aggregate predictors, whose application to neural networks is a very recent topic (Dziugaite and Roy, 2017; Neyshabur et al., 2017), and which have been far less studied than the standard empirical risk minimizer (ERM). During the last decades, the PAC-Bayesian theory has offered vast opportunities of progress on statistical risk guarantees for estimators defined as a "mixture" of a given family of weak estimators. In the rest of this subsection, we define more precisely the learning problem addressed in Chapter 6 and present some general elements on the PAC-Bayesian theory.

Some generalities on statistical learning and neural networks Let $(\mathcal{Z}, \mathscr{A})$ be a measurable space and $\mathbf{Z^n} = (Z_1, \ldots, Z_n) \in \mathcal{Z}^n$ be realizations drawn from an unknown distribution $\mathcal{P}$ on $(\mathcal{Z}^n, \mathscr{A}^{\otimes n})$. Define a set $\mathcal{X} \subset \mathbb{R}^{D_0}$, $D_0 \geq 1$, and let μ be a σ-finite measure on $\big(\mathcal{X}, \mathscr{B}(\mathcal{X})\big)$. We

consider the problem of estimating a measurable function $f_{\mathcal{P}} \in \mathcal{F} := \{f : \mathcal{X} \to \mathbb{R}^{D_L},\ D_L \in \mathbb{N}\}$, depending on the distribution $\mathcal{P}$. Generally, it is too costly to look for a function belonging to the entire set $\mathcal{F}$. Instead, it is appealing to look for an approximation chosen from a smaller parameterized subset of functions $\mathcal{F}_{\mathsf{W}} := \{f_{\boldsymbol{w}}, \boldsymbol{w} \in \mathsf{W}\} \subset \mathcal{F}$, indexed by elements of the measurable space $(\mathsf{W}, \mathscr{B}(\mathsf{W}))$, $\mathsf{W} \subset \mathbb{R}^d$. We also define a loss function $\ell : \mathcal{F} \times \mathcal{F} \mapsto \mathbb{R}_+$, such that $\ell(f_{\boldsymbol{w}}, f_{\mathcal{P}}) = \int_{\mathcal{X}} \mathrm{d}(f_{\boldsymbol{w}}(\boldsymbol{x}), f_{\mathcal{P}}(\boldsymbol{x}))\mu(d\boldsymbol{x})$ measures the performance of a function $f_{\boldsymbol{w}} \in \mathcal{F}_{\mathsf{W}}$, for a given distance d defined on $\mathbb{R}^{D_L}$.

Feedforward dense neural networks of depth $L \in \mathbb{N}$ are particular specifications of the subset $\mathcal{F}_{\mathsf{W}}$, where W defines the weights of the neural network. More precisely, for $\boldsymbol{w} = (\boldsymbol{w}_1, \boldsymbol{w}_2, \ldots, \boldsymbol{w}_L) \in \mathsf{W}$, the overall dimension of W is given by $d = D_0 D_1 + \cdots + D_{L-1} D_L$, where D_ℓ is the number of units in the layer ℓ; an element of $\mathcal{F}_{\mathsf{W}}$ is then defined by $f_{\boldsymbol{w}}(x) = \boldsymbol{w}_L^\top \sigma_{L-1}(\boldsymbol{w}_{L-1}^\top \sigma_{L-2}(\ldots \boldsymbol{w}_2^\top \sigma_1(\boldsymbol{w}_1^\top x)))$, where σ is an activation function and $\sigma_\ell : x \in \mathbb{R}^{D_\ell} \mapsto (\sigma(x_1), \ldots, \sigma(x_{D_\ell}))^\top$, $\ell = 1, \ldots, L-1$. Each w_ℓ is a $D_{\ell-1} \times D_\ell$ matrix. In this thesis, our goal is to derive theoretical guarantees in the simple tractable setting of shallow networks – with a single hidden layer; we refer the reader to Bishop (1994); Goodfellow et al. (2016) for an introduction and a comprehensive overview on more complex structures such as convolutional neural networks, LTST memory networks, GANs, VAEs, etc.

The first question that needs to be answered in such a setting is how well a shallow neural network approximates the function $f_{\mathcal{P}}$. The corresponding approximation error depends on the properties of the function $f_{\mathcal{P}}$ (*e.g.* Lipschitz or Hölder continuity, Sobolev smoothness), the size of the class $\mathcal{F}_{\mathsf{W}}$ and the type of activation function σ. The literature on this topic has developed around these main characteristics, such that for the approximation by shallow networks with sigmoid activation functions we can cite for example Mhaskar and Micchelli (1994); Petrushev (1998); Burger and Neubauer (2001); Cao et al. (2008); Costarelli and Spigler (2013a,b); Barron (1993); Delyon et al. (1995); Maiorov and Meir (2000); Maiorov (2006). In particular, Maiorov and Meir (2000); Maiorov (2006) consider the class of Sobolev functions $W_2^r([0,1]^{D_0})$ and obtain an approximation error of order $O(D_1^{-r/D_0})$. This result is particularly interesting in that it involves the number of units in the hidden layer D_1, the regularity r of the function $f_{\mathcal{P}}$ and the problem dimension D_0. In the case of ReLU activation functions, still for shallow networks, we may cite Petrushev (1998); Klusowski and Barron (2016b,a); Bach (2017); Xu (2020); Siegel and Xu (2020). Siegel and Xu (2020) prove that for the class of Barron spectral space $\mathscr{B}^s([0,1]^{D_0})$ – we refer the reader to Chapter 6 Section 6.5.2 expression (6.16) for a definition of this space – the error is of order $O(D_1^{-K})$ where K depends on s and the problem dimension D_0. These last years, the development of deep networks has triggered the research on approximation guarantees for these deeper networks with ReLU activation functions (Yarotsky, 2017, 2018; Yarotsky and Zhevnerchuk, 2019; Shen et al., 2019; Gühring et al., 2020; Lu et al., 2020).

It should be noted that in practice, the function $f_{\mathcal{P}}$ being unknown, one has no access to the best approximation. Instead, an estimator $\widehat{f}_n : \mathcal{Z}^n \mapsto \mathcal{F}_{\mathsf{W}}$ computed from the observed data $\mathbf{Z^n}$

should be used. The quality of such an estimator is evaluated by the risk defined by:

$$\mathbf{E}_{\mathcal{P}}[\ell(\widehat{f}_n(\mathbf{Z^n}), f_{\mathcal{P}})] = \int_{\mathcal{Z}^n} \ell(\widehat{f}_n(\mathbf{z}), f_{\mathcal{P}})\mathcal{P}(d\mathbf{z}). \tag{1.2}$$

To assess the risk of the neural networks, a literature on non asymptotic analysis of neural networks has emerged since the 90's, but until now it remains scarce. We can mention works for dense neural networks with sigmoid activation functions (Barron, 1994; McCaffrey and Gallant, 1994) and neural networks with ReLU activation functions (Klusowski and Barron, 2016b; Bach, 2017). Recent results on deep neural networks prove their ability to reach minimax optimal rates up to $\log(n)$ factors for functions with specific hierarchical architecture and sufficient regularity, for ReLU deep networks (Schmidt-Hieber, 2020) as well as sigmoid ones (Bauer and Kohler, 2019).

Most guarantees established in the literature for a neural network with a given architecture focus on the minimizer of the empirical risk (also known as training error):

$$\widehat{f}_n^{\mathsf{ERM}} \in \operatorname{argmin}_{\mathcal{F}_{\mathsf{W}}} \frac{1}{n}\sum_{i=1}^{n} L(f_{\boldsymbol{w}}, Z_i),$$

where $L(f_{\boldsymbol{w}}, Z_i)$ is a suitably chosen proxy for the loss $\ell(f_{\boldsymbol{w}}, f_{\mathcal{P}})$.

Nevertheless, one can tackle the estimation problem from another perspective than empirical risk minimization. Instead of trying to estimate the best possible parameter value $\boldsymbol{w}^*$ by a data-driven element $\widehat{\boldsymbol{w}}_n$ from W, one can try to combine several functions $f_{\boldsymbol{w}}$ to obtain a good approximation of $f_{\mathcal{P}}$. In a more general framework than the one of neural networks, Juditsky et al. (2008) show that for finite classes of predictors $\mathcal{F}_{\mathsf{W}}$, an empirical convex aggregate estimator over this class attains sharper risk bounds than the empirical risk minimizer could ever attain. This triggers interest on PAC-Bayesian theory and its application to neural networks.

The PAC-Bayesian theory Roughly speaking, this theory intends to derive the best of both worlds between the generalized Bayesian theory and Probably Approximately Correct bounds. Generalized Bayesian theory for learning relies on the classic Bayesian theory: assume that we have an a priori opinion on the predictor $f_{\boldsymbol{w}}$ encoded in a prior distribution π over W, then the posterior distribution is given by $\widehat{\pi}_n(\boldsymbol{w}|\mathbf{Z}^n) \propto \mathsf{Likelihood}(\mathbf{Z}^n|\boldsymbol{w})\pi(\boldsymbol{w})$. Generalized Bayesian learning amounts to replacing the likelihood by a loss functional $\mathscr{L}_{\boldsymbol{w},n}$, which measures the performance of a function $f_{\boldsymbol{w}}$ given $\mathbf{Z^n}$, such that $\widehat{\pi}_n(\boldsymbol{w}|\mathbf{Z^n}) \propto \mathscr{L}_{\boldsymbol{w},n}(\mathbf{Z^n})\pi(\boldsymbol{w})$. With the notion of posterior distribution on the space W, following the idea of an aggregate estimator introduced in the previous paragraph, we can then for example define the mean aggregate estimator by:

$$\widehat{f}_n = \int_{\mathsf{W}} f_{\boldsymbol{w}}\widehat{\pi}_n(d\boldsymbol{w}). \tag{1.3}$$

The second ingredient of PAC-Bayesian theory, the PAC bounds, were initiated by Valiant (1984) and generally take the form

$$\mathcal{P}\left(\mathbf{E}_{\mathcal{P}}[\ell(\widehat{f}_n, f_{\mathcal{P}})] - \mathbf{E}_{\mathcal{P}}[\ell(f_{\boldsymbol{w}^*}, f_{\mathcal{P}})] \leq \delta_n(\varepsilon)\right) \geq 1-\varepsilon,\ \varepsilon \in (0,1)$$

where δ_n is a well chosen strictly positive function.

These two concepts have been connected in the seminal papers of McAllester (1999a, 2003) in a data driven empirical bound, then extended by the work of Catoni (2007) that led to the first oracle inequality in this framework. Following these results, the works of Leung and Barron (2006) and Juditsky et al. (2008) initiated PAC-Bayesian bounds in expectation that have been further developed by Lecué (2007); Alquier and Lounici (2011); Dalalyan and Tsybakov (2012a); Rigollet and Tsybakov (2012). In this setting and under some assumptions, these works propose an aggregate estimator $\widehat{f}_n$ verifying the following risk bound:

$$\mathbf{E}_{\mathcal{P}}[\ell(\widehat{f}_n, f_{\mathcal{P}})] \leq C \inf_{p \in \mathcal{P}_W} \left\{ \int_{\mathsf{W}} \ell(f_{\boldsymbol{w}}, f_{\mathcal{P}})\, p(d\boldsymbol{w}) + \frac{\beta}{n} D_{\mathsf{KL}}(p||\pi) \right\}, \tag{1.4}$$

where $\mathcal{P}_{\mathsf{W}}$ is the space of all probability measures on W, C is a universal constant, π is a fixed prior distribution on W, $\beta > 0$ a temperature parameter and $\widehat{f}_n$ generally writes as (1.3).

In view of its very generic form, the inequality (1.4) applies to many frameworks. It allows for online learning, thanks to the Mirror Averaging (MA) procedure (see Yuditskii et al. (2005); Juditsky et al. (2008) for this topic and Dalalyan and Tsybakov (2012a) for the sparse prior case), or deterministic designs – useful in signal or image processing, with the use of Exponentially Weighted Aggregate estimator (EWA) (Dalalyan and Tsybakov, 2007, 2012b). We refer the reader to the comprehensive Primer on PAC-Bayesian learning of Guedj (2019) and the references therein for a more detailed introduction. The inequality in these settings, depending on some assumptions, applies to various kinds of statistical problems such as regression, classification, as well as density estimation. We refer the reader to Chapter 6 Section 6.3 for a more detailed description of the type of estimation problems encountered.

By adapting inequality (1.4) to the specific class $\mathcal{F}_{\mathsf{W}}$ of shallow neural networks, the work of Chapter 6 allows us to address Problems 5–6. As the last part of this introduction, the next section briefly presents the results of the following chapters, to give a general idea of our contribution.

1.3 Main contributions of the thesis

In this section we give a brief overview of the main developments of the thesis. The first two chapters focus on optimal trading strategies for renewable producers enhanced by a study on price formation and its features, Chapter 5 treats probabilistic forecasts and their application to wind energy management, and finally Chapter 6 provides insights on the parameterization and on the risk of aggregate shallow neural networks.

1.3.1 Optimal strategies for renewable producers

In the intraday electricity market, the price integrates the updates in demand/production forecasts faced by the agents. This specific characteristic induces market impact and possible strategic behaviors. In Chapter 3 and Chapter 4, we intend to explain this phenomenon and its consequences on prices through a simplified model of the successive short term markets presented in Section 1.1.2. We address Problem 1 by developing a tractable equilibrium model in a linear quadratic setting, that is able to provide an explicit price and explicit optimal trading strategies for producers facing an intermittent terminal production/demand. We study this game in different informational structures: in Chapter 3 we consider the complete and incomplete information setting, with identical agents; in Chapter 4, we introduce a major player impacting the strategical behavior of the crowd.

N-player setting We consider a fixed delivery period starting at date T. We assume there are N identical agents that share the same information on the market price and the production forecast, such that we can define a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} := (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{P})$ to which all the processes are adapted. The market clearing price of the day ahead market, S_0, occurs at date $t = 0 < T$ and we denote by ϕ_0^i the position of i-th agent, $i = 1, \dots, N$, in this market. Thereupon, the intraday market allows continuous trading up till the market closure just before T, and we denote ϕ_t^i the position of i-th player at time t, as well as $\bar{\phi}_t^N = \frac{1}{N}\sum_{i=1}^N \phi_t^i$ the aggregate position of all the agents. The intraday market price P_t^N is composed of a fundamental price S_t and a linear market impact (Almgren and Chriss, 1999, 2001) through which the players interact with one another:

$$P_t^N = S_t + a(\bar{\phi}_t^N - \bar{\phi}_0^N), \quad \forall t \in [0,T],\ a > 0. \tag{1.5}$$

When trading, i-th player incurs a cost given by

$$\dot{\phi}_t^i P_t^N + \frac{\alpha(t)}{2}\dot{\phi}_t^i\Big(\dot{\phi}_t^i + b\dot{\bar{\phi}}_t^{N,-i}\Big), \quad \forall t \in [0,T]$$

where $\dot{\bar{\phi}}_t^{N,-i} = \frac{1}{N-1}\sum_{j=1, j\neq i}^N \dot{\phi}_t^j$. Here the first term represents the actual cost of buying the electricity, and the second term represents the cost of trading, where $\alpha(.)$ is a continuous strictly positive function on $[0,T]$, reflecting the variation of market liquidity at the approach of the delivery date. The term $b\dot{\bar{\phi}}_t^{N,-i}$ with $b > 0$ represents the impact of the crowd trading direction on the cost of trading of i-th agent, which accounts for possible synchronization of the players. To define her optimal position on the market, each agent disposes of an uncertain production forecast of her terminal production X_T^i, given by the process $(X_t^i)_{t\in[0,T]}$. At the closure of the intraday market, if $\phi_T^i \neq X_T^i$, the agent must first purchase the missing amount or sell the extra amount of electricity at price S_T and in addition pay an imbalance penalty $\frac{\lambda}{2}(\phi_T^i - X_T^i)^2$. Finally, i-th player aims to maximize a linear quadratic objective function, which depends on

the strategy of the other players through the market price:

$$J^{N,i}(\phi^i,\phi^{-i}) := -\mathbb{E}\Big[\underbrace{\phi_0^i S_0}_{\text{Day ahead}} + \underbrace{\int_0^T \Big\{\dot{\phi}_t^i P_t^N + \frac{\alpha(t)}{2}\dot{\phi}_t^i\Big(\dot{\phi}_t^i + b\dot{\bar{\phi}}_t^{N,-i}\Big)\Big\}\,dt}_{\text{Intraday}}$$
$$\underbrace{-(\phi_T^i - X_T^i)S_T + \frac{\lambda}{2}(\phi_T^i - X_T^i)^2}_{\text{Balancing}}\Big], \tag{1.6}$$

where $\phi^{-i} := (\phi^1,\dots,\phi^{i-1},\phi^{i+1},\dots,\phi^N)$. In line with the Definition 7, we derive in Theorem 14 Chapter 3 a unique explicit Nash equilibrium without any specific assumptions on the dynamics of the price S.

From the expression of the strategies and the equilibrium price, we can infer some insights on how the fundamental price and the production forecasts impact the decisions of the renewable producers. To cite some examples, when S is assumed to be a martingale, the strategy does not depend on the price, the trades are only triggered by forecasts adjustments and the position on the day ahead market will be given by $\phi_0^i = X_0^i$. On the contrary, if we assume the price has a positive deterministic drift, players will choose $\phi_0^i < X_0^i$, but there will be a positive trend in the aggregate strategy in the intraday market such that the market impact will amplify the price trend. Besides, the additional cost related to crowd behavior, b, leads agents to trade less actively in response to common forecasts updates impacting all the X^i's, than in response to individual ones. Concerning the terminal penalty, it is modeled by the soft constraint λ because it is supposed to model the real imbalance penalty which is based on the difference between X_T^i and ϕ_T^i. Nevertheless, when $\lambda \to \infty$, the hard constraint condition is satisfied that is $X_T^i = \phi_T^i,\ i = 1,\dots,N$. Hence, depending on the imbalance settlement (simplified into a deterministic penalty), the model allows producers to decide to what extent they have to fulfil their day ahead / intraday commitment.

As already evoked, assuming all the market participants have a perfect knowledge of the aggregate forecasts might seem unrealistic. To address this, we provide a brief presentation of the incomplete information framework in the sequel and refer the reader to Chapter 3 Section 3.3.2 for further details.

Mean field extension In the mean field variant of the game, we consider the limiting case with infinitely many agents and allow for incomplete information, modeled by the filtration $\mathbb{F}$, to which all processes are adapted, and $\mathbb{F}^0$ that contains only the common noise – *i.e* restricted to the information present in the market, already introduced in Section 1.2.1. The position of the representative agent is given by ϕ_t and her production forecast is denoted X_t. The latter are only partially observed: the other market participants can only observe $\bar{\phi}_t = \mathbb{E}[\phi_t|\mathcal{F}_t^0]$ and $\overline{X}_t = \mathbb{E}[X_t|\mathcal{F}_t^0]$, and the market price P_t, given by $P_t = S_t + a(\bar{\phi}_t - \bar{\phi}_0)$. In this case, we also obtain an explicit mean field equilibrium in the sense of Definition 9, with an equilibrium price and strategies in closed form.

More importantly, we are able to bridge the gap between the two frameworks. In practice,

markets with an infinity of agents are not realistic and when facing a N-player market, it is useful to quantify how close the mean field equilibrium is to the 'true' N-player Nash equilibrium. With the Definition 8, we prove that the mean field equilibrium can provide strategies for the N-player game that form an ε-Nash equilibrium of the game with $\varepsilon = \frac{C}{N^{\frac{1}{2}}}$, where $C > 0$ is a constant. We also prove the reverse approximation that is, when the number of players increases, the Nash equilibrium converges to the limiting mean field one. Figure 1.5 illustrates this phenomenon when increasing the number of players from $N = 5$ to $N = 100$.

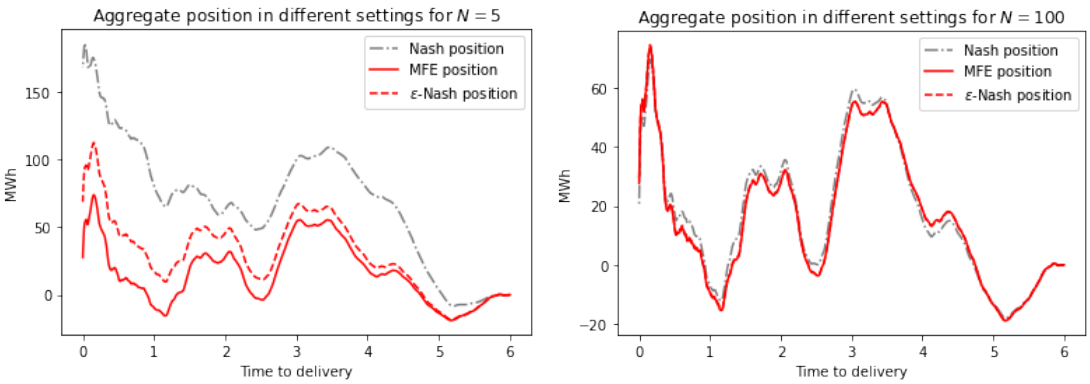

Figure 1.5: Aggregate position in different settings with $N = 5$ (left) and $N = 100$ (right) agents

Stackelberg game In Chapter 4, we develop a model in the mean field setting in the presence of a major agent with informational asymmetry. The infinitely many small and identical producers can observe the production forecasts of the major agent because of her size; but the reverse is not true: the major producer can only observe a common production forecast. While small producers have an informational advantage, the actions of the major producer can significantly impact the market. The small representative producer will have to strategically define her position with respect to the major player and vice versa. In this framework, we are still able to exhibit explicit optimal positions and a price equilibrium that allows to emphasize the impact of a major producer on the behavior of small ones. In addition, we can derive an ε-Nash equilibrium to adapt the strategies to a case where the number of players is finite.

Other insights on the strategic behavior of renewable producers are documented in Chapter 3 and Chapter 4. But, more importantly, these models lead to features of the market price that are in line with the empirical findings as explained in the next section.

1.3.2 Price formation and main characteristics of the intraday market

Volatility and Samuelson's effect In Figure 1.2b, Section 1.1.3, we displayed the empirical instantaneous volatility of the price estimated by kernel method (See Kristensen (2010) for the general methodology and Chapter 3 Section 3.5.1 for the setting of Figure 1.2b). We observe that the volatility increases as we approach the delivery time. This is a well documented feature of electricity futures and other futures markets (Jaeck and Lautier, 2016) known as the Samuelson effect. An important challenge of Chapters 3 and 4 is to reproduce and explain this stylized

feature. To investigate this phenomenon, we lead an analytic study on the closed-form expression of the equilibrium price and retrieve the additional volatility implied by the interaction between the agents thanks to the expression of the normalized price variance between small time steps h, such that, when $h \to 0$:

$$\frac{1}{h}\mathbb{E}[(P^N_{t+h} - P^N_t)^2] = \sigma_S^2 + h\mathcal{V}_t + O(h^2), \text{ with } \mathcal{V}_t = \frac{a^2\lambda^2\sigma_X^2}{\alpha\left(1+\frac{b}{2}\right)} f(t) \tag{1.7}$$

where σ_S^2, $\sigma_X^2 > 0$ are the constant variances of the fundamental price and the agents production forecasts, $\alpha > 0$ is a constant liquidity cost, and f is an increasing function of the time depending on the model parameters. Note that as $h \to 0$, the expression (1.7) converges to the variance of the fundamental price σ_S^2. However, an agent using volatility estimator with time step h on the price process, will find an extra variance of approximately $h\mathcal{V}_t$ (on average). Furthermore, we retrieve the Samuelson effect since f is increasing in t. This is confirmed by an empirical estimation of the instantaneous volatility, using exactly the same kernel method as in Figure 1.2b, on simulations of the model equilibrium price, as showed in Figure 1.6a. In addition, the explicit expression (1.7) brings information on the dependency of the volatility with respect to the various parameters of the model such as the imbalance penalty, the liquidity costs and the number of players. For more details on how volatility is impacted by these parameters, we refer the reader to Chapter 3 Section 3.4.2.

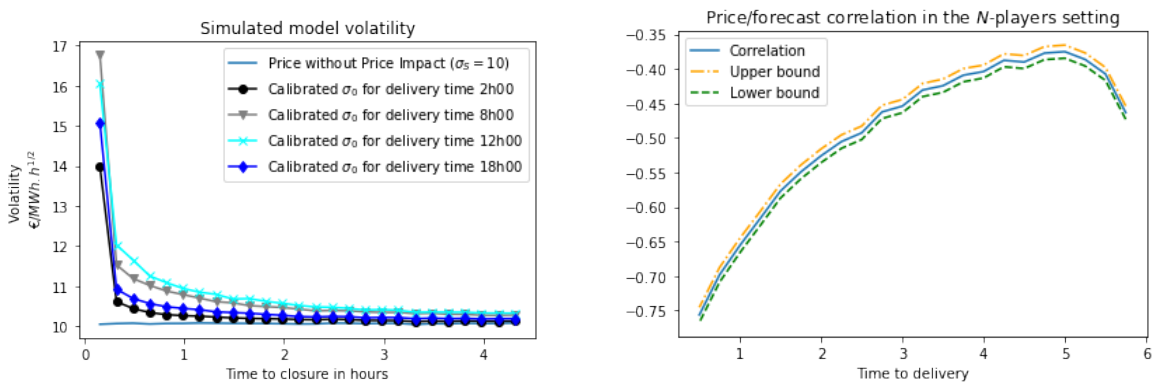

(a) Estimated volatility in the model

(b) Estimated correlation in the model

Figure 1.6: Estimated stylized features in the model

Correlation between renewable infeeds and market price Our model is also able to explain the empirically observed correlation between prices and production forecasts: the link between both quantities is also treated analytically in Section 3.4.3 Chapter 3 and we provide in Figure 1.6b an estimate of the correlation in the model from simulated equilibrium prices and renewable production forecasts. Correlation is negative and increases in absolute value as we approach the delivery time as it is the case in the empirical findings.

To conclude, Chapter 3 and 4 provide, to the best of our knowledge, the first equilibrium model capable of explaining the observed stylized features of the intraday electricity market. In

this way, it may offer a better understanding of the interactions in this market and help with the management of renewable energies. The recent paper of Aid et al. (2020), relates to our work. They developed a N-player heterogeneous complete information model where they exhibit an equilibrium in the intraday electricity market. Their model also reproduces the Samuelson's effect described by our model.

In the two first chapters of the thesis, the production forecasts are assumed to be point forecasts. However this notion is not sufficient to address Problem 3 and Problem 4. In the next section, we address the latter in the framework of probabilistic forecasts presented in Section 1.2.2.

1.3.3 Probabilistic forecasts in sequential decision making problems

In Chapter 5, we introduce a probabilistic forecast model in the diffusion setting to capture the time dependencies between different time horizons and provide an explicit flow of predictive densities that we can calibrate and exploit in stochastic decision making processes. To the best of our knowledge, the latter point has never been investigated yet. We develop four SDE models able to account for real valued quantities such as the temperature, or positive quantities such as the wind speed. For two of these models we describe a complete calibration procedure using ensemble forecasts from Paris area, of which an example is displayed in Figure 1.4a, and evaluate their accuracy. We then focus on a simple wind energy production management problem with a single agent where we use the calibrated probabilistic forecasts of the wind speed. In the sequel, we briefly present this example.

Density modeling We consider the filtered probability space $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$, a $\mathcal{F}_T$-measurable random variable m_T representing the wind speed, and refer to Definition 10 for the flow of predictive densities. For the predictive density at time $t \in [0, T]$, we choose a two-dimensional parametrization (m_t, V_t), where m_t is the conditional mean of m_T and V_t is a measure of the forecast uncertainty. More specifically, we assume that m_t and V_t are defined by the following Heston-like model:

$$\begin{aligned} \frac{dm_t}{m_t} &= \sqrt{V_t}\rho(T-t)dW_t \\ dV_t &= -V_t\rho^2(T-t)\left(1+\frac{b^2}{2}\right)dt + \sqrt{V_t}b\rho(T-t)dW_t'. \end{aligned} \tag{1.8}$$

where W, W' are independent Brownian motions, ρ is a positive deterministic function and $b > 0$ is a constant. From the results in Barndorff-Nielsen (1997), we derive that for m, V solution of (1.8), the conditional distribution of $\log m_T$ given $\mathcal{F}_t$ is the normal inverse Gaussian distribution on $\mathbb{R}$ with density

$$p_t(x) = \frac{\alpha\delta_t K_1\left(\alpha\sqrt{\delta_t^2 + (x-\mu_t)^2}\right)}{\pi\sqrt{\delta_t^2 + (x-\mu_t)^2}} e^{\delta_t\gamma + \beta(x-\mu_t)}, \forall t \in [0, T] \tag{1.9}$$

where the parameters are given by $\mu_t = \log m_t$, $\delta_t = \frac{V_t}{b}$, $\beta = -\frac{1}{2}$, $\alpha = \sqrt{\left(b^{-1} + \frac{b}{2}\right)^2 + \frac{1}{4}}$, $\gamma = b^{-1} + \frac{b}{2}$, and K is the modified Bessel function of the third kind. Following the ideas developed in Section 1.2.2, we then calibrate the density (1.9) at each lead time to use the diffusion model in a realistic wind energy management problem.

Calibration The data used to calibrate the model is composed of several ensemble forecasts that have a total time horizon of 48 hours and are updated every 12 hours. We thus calibrate the predictive density for time horizons 48h, 36h, 24h and 12h before the realization. As explained in Section 1.2.2, ensemble forecasts suffer from under-dispersion and bias. To correct these shortcomings, we employ a methodology inspired by EMOS for the different lead times. We refer the reader to Chapter 5 Section 5.3 for a detailed description of the method. Once this is done, we measure the quality of the calibration comparing with the raw ensemble data performances. We performed several investigations on the goodness of fit which demonstrate a clear improvement of the prediction ability compared to raw ensembles. Figure 1.7 shows the Probability Integral Transform (PIT) of the calibrated forecasts, which consists in the histogram of the calibrated predictive cumulative distribution function for the lead time 12h, evaluated at the point corresponding to the realized wind speed. If the predictive distribution is well calibrated, then the histogram should be close to the uniform distribution one. Comparing Figure 1.7 with the Talagrand diagram Figure 1.4b Section 1.2.2 we observe that the under-dispersion and bias have been corrected to a large extent.

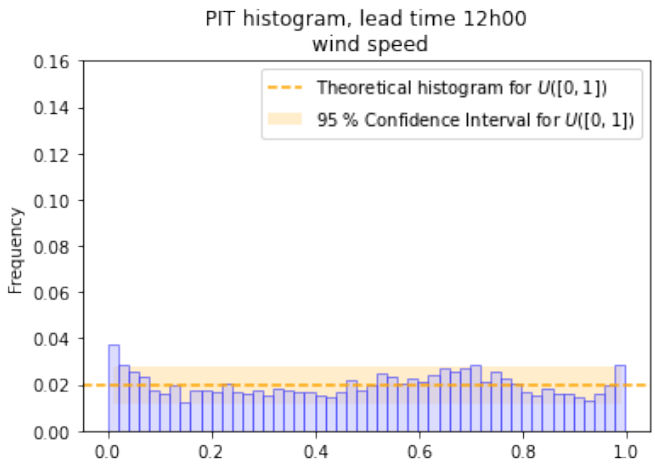

Figure 1.7: PIT histogram

Use of probabilistic forecasts for sequential decision making The core idea and advantage of the model we have developed is that it permits to exploit the information obtained from the evolution of the predictive density in a stochastic decision problem. To evaluate this advantage, we performed a study to determine the additional profit of a wind power producer using our model, compared to a standard model without an explicit dynamic description of the forecast uncertainty. To do so, we consider model A where the producer has access to probabilis-

tic forecasts and model B, which is based only on point forecasts of the wind speed. For a given delivery hour T, we denote the average wind speed during a small time interval around this date by m_T in both models. Their wind forecasts for this delivery time are regularly updated, denoted m_t at time $t < T$, and they deduce their production forecast updates from the wind forecast updates, thanks to a stylized production function f, such that the production forecast writes $f(m_t)$. They can adjust their position in the intraday market each time they receive new updates and must pay a terminal imbalance penalty. We denote by $t_0 < \cdots < t_{N-1}$ the discrete times at which the trades take place, $t_N = T$ being the time when the delivery starts. The fraction traded at the date t will have the price S_t, and we denote the total amount of power sold or bought up to date t by ϕ_t. Any power not sold in the intraday market prior to date T will be sold at date T at the balancing price denoted by S_T. In addition, balancing transactions are subject to imbalance penalty equal to a constant K times the volume of the transaction. The profits of the producers are thus given by

$$\begin{aligned}
&\underbrace{S_{t_0}\phi_{t_0} + \sum_{i=1}^{N-1} S_{t_i}(\phi_{t_i} - \phi_{t_{i-1}})}_{\text{Intraday market}} + \underbrace{S_T(f(m_T) - \phi_{t_{N-1}}) - K|f(m_T) - \phi_{t_{N-1}}|}_{\text{Imbalance payment}} \\
&= f(m_T)S_T - \sum_{i=0}^{N-1} \phi_{t_i}\Delta S_{t_i} - K|f(m_T) - \phi_{t_{N-1}}|,
\end{aligned}$$

where $\Delta S_{t_i} = S_{t_{i+1}} - S_{t_i}$. We assume that the producers aim to maximize the utility of profit at date T, that is, they solve the following control problem:

$$\max_{\phi := (\phi_{t_i})_{i=1}^{N-1}} \mathbb{E}\left[u\left(f(m_T)S_T - \sum_{i=0}^{N-1} \phi_{t_i}\Delta S_{t_i} - K|f(m_T) - \phi_{t_{N-1}}|\right)\right], \tag{1.10}$$

where u is a utility function. We then assess numerically their gains on simulated trajectories (of model A), for different parameterizations of the model, depending on the initial level of uncertainty at the opening of the market, and the drift μ_S of the price.

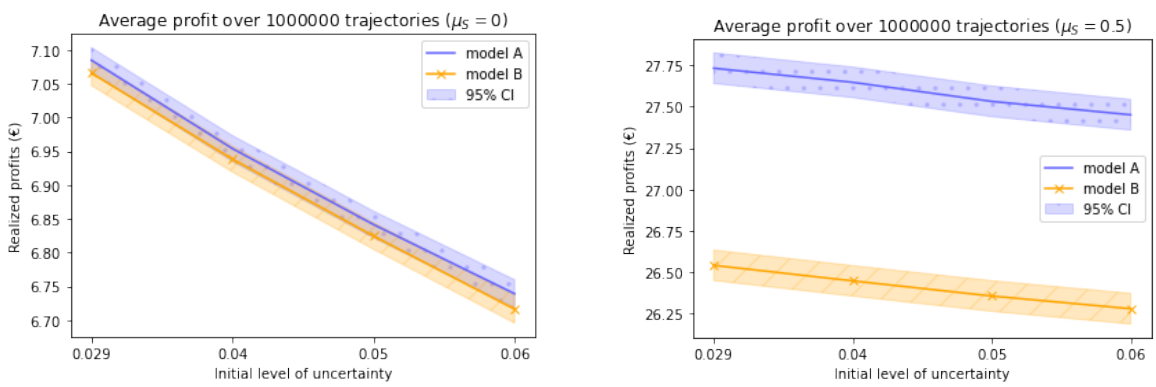

(a) $\mu_S = 0$ (b) $\mu_S = 0.5$

Figure 1.8: Profits in function of the initial wind speed forecasts uncertainty

The results in Figures 1.8a–1.8b show that the gains decrease when the initial uncertainty increases. In the case where the price is martingale the comparison of the results between the two models is not straightforward since the confidence intervals of the gains partially overlap; however, in the case where the price has a drift, gains are in average significantly higher for an agent using model A. From this study, we can conclude that this new method shows promising results and should be further investigated for practical uses.

In the next section, we present the main results we obtained while addressing Problems 5–6.

1.3.4 Learning by aggregated shallow neural networks

Chapter 6 presents the PAC-Bayesian bounds in the setting of shallow neural networks. Let $f_{\mathcal{P}}$ be the function we want to estimate and consider the framework introduced in Section 1.2.3 where $\mathcal{X} = [0,1]^{D_0}$, and the class $\mathcal{F}_{\mathsf{W}}$ is restricted to shallow neural networks, such that the set W of the weights of a neural network contains the weights $\boldsymbol{w}_1$ of the hidden layer, and the weights $\boldsymbol{w}_2$ of the output layer. Then $\boldsymbol{w} = (\boldsymbol{w}_1^\top, \boldsymbol{w}_2^\top)^\top$ belongs to $\mathbb{R}^d$ with $d = D_1(D_0 + D_2)$, $\boldsymbol{w}_1 \in \mathbb{R}^{D_0 D_1}$, $\boldsymbol{w}_2 \in \mathbb{R}^{D_1 D_2}$. The neural network parametrized by $\boldsymbol{w}$ thus writes:

$$f_{\boldsymbol{w}}(\boldsymbol{x}) = \boldsymbol{w}_2^\top \bar{\sigma}(\boldsymbol{w}_1^\top \boldsymbol{x}) \in \mathbb{R}^{D_2},\ \forall \boldsymbol{x} \in \mathcal{X} \quad \text{with} \quad \bar{\sigma} : \boldsymbol{x} \in \mathbb{R}^{D_1} \mapsto \begin{bmatrix} \sigma(x_1) \\ \dots \\ \sigma(x_{D_1}) \end{bmatrix} \in \mathbb{R}^{D_1}, \tag{1.11}$$

where $\sigma : \mathbb{R} \mapsto \mathbb{R}$ is an activation function.

We specify the loss function with the standard ℓ_2-norm for the PAC-Bayesian inequality (1.4), such that it rewrites:

$$\mathbf{E}_{\mathcal{P}}[\|\widehat{f}_n - f_{\mathcal{P}}\|^2_{\mathbb{L}_2(\mu)}] \leq C \inf_{p \in \mathcal{P}_{\mathsf{W}}} \left\{ \int_{\mathsf{W}} \|f_{\boldsymbol{w}} - f_{\mathcal{P}}\|^2_{\mathbb{L}_2(\mu)}\, p(d\boldsymbol{w}) + \frac{\beta}{n} D_{\mathsf{KL}}(p||\pi) \right\}. \tag{1.12}$$

Fixing the parameter $\bar{\boldsymbol{w}} \in \mathsf{W}$ and using triangle inequality, this yields:

$$\left(C^{-1}\, \mathbf{E}_{\mathcal{P}}\left[\|\widehat{f}_n - f_{\mathcal{P}}\|^2_{\mathbb{L}_2(\mu)}\right] \right)^{1/2} \leq \|f_{\bar{\boldsymbol{w}}} - f_{\mathcal{P}}\|_{\mathbb{L}_2(\mu)} + \mathrm{Rem}_n(\bar{\boldsymbol{w}})^{1/2}, \tag{1.13}$$

with the remainder term bounded by

$$\mathrm{Rem}_n(\bar{\boldsymbol{w}}) \leq \int_{\mathsf{W}} \|f_{\boldsymbol{w}} - f_{\bar{\boldsymbol{w}}}\|^2_{\mathbb{L}_2(\mu)}\, p(d\boldsymbol{w}) + \frac{\beta}{n} D_{\mathsf{KL}}(p||\pi), \tag{1.14}$$

where p is any distribution in $\mathcal{P}_{\mathsf{W}}$ and π is the prior distribution. Note that if we choose $\bar{\boldsymbol{w}} \in \mathrm{argmin}_{\boldsymbol{w} \in \mathsf{W}} \|f_{\boldsymbol{w}} - f_{\mathcal{P}}\|_{\mathbb{L}_2(\mu)}$, the right hand side of inequality (1.13) can be decomposed into the approximation error $\inf_{\boldsymbol{w} \in \mathsf{W}} \|f_{\boldsymbol{w}} - f_{\mathcal{P}}\|_{\mathbb{L}_2(\mu)}$ and the remainder term $\mathrm{Rem}_n(\bar{\boldsymbol{w}})$ assessing the estimation error. Another important remark, is that most of the time, in practice, neural networks are initialized with a Gaussian distribution. Following this idea, we set the prior distribution and the set $\mathcal{P}_{\mathsf{W}}$ as spherical Gaussian distributions.

Now that this general framework is set, our work in Chapter 6 consists in three main steps. First, find a 'good' measure p for the remainder term such that it shows good properties in

terms of tightness of the bound and allows to derive a tractable expression with respect to the model parameters. The last two steps consist in exploiting the tuning parameters of the neural network. The second step amounts to finding a suitable variance for the prior such that for Lipschitz activation functions– this includes sigmoid and ReLU, we derive an oracle bound from the remainder term (1.14), that highlights the roles of D_1 and n. The third step amounts to choosing an appropriate number D_1 of units in the hidden layer. From this stage, we restrict the class of functions $\mathcal{F}$ to real valued functions ($D_2 = 1$), use the results of Maiorov and Meir (2000) when $\mathcal{F}$ is the unit ball in the Sobolev space $W_2^r(\mathcal{X})$, and neural networks are equipped with sigmoid activation functions; and use the result of Siegel and Xu (2020), for Barron spectral space and sub-spaces of $W_2^r(\mathcal{X})$ –for a special choice of r, and ReLU neural networks. This permits to derive fully explicit oracle inequalities and, to cite an example, for the case of sigmoid activation functions, we obtain,

$$C^{-1}\mathbf{E}_{\mathcal{P}}\left[\|\widehat{f}_n - f_{\mathcal{P}}\|^2_{\mathbb{L}_2(\mu)}\right] \leq g(D_1)D_1^{-2r/D_0} + \bar{g}(n/d)\frac{D_1 D_0}{n}, \tag{1.15}$$

where g, $\bar{g}$ are at most logarithmic functions. The number of neurons D_1 has to be chosen carefully in order to control the balance between the approximation error and the estimation error, to guarantee the best bias-variance trade-off in terms of the sample size, as it is graphically illustrated for equation (1.15) in Figure 1.9.

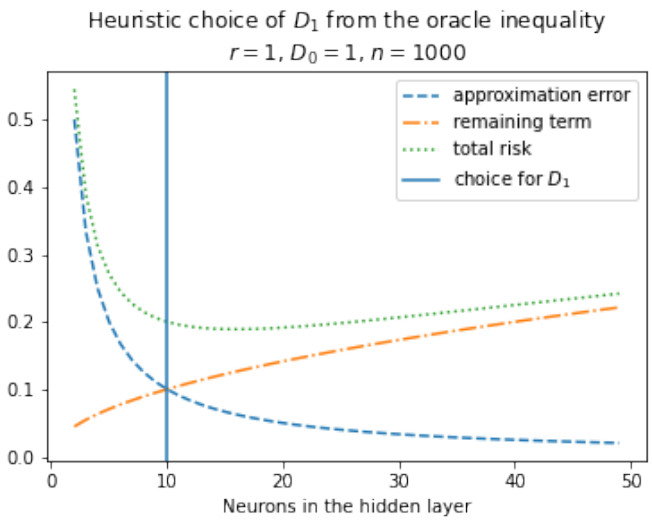

Figure 1.9: Balance between approximation error and estimation error

This has two main advantages: (i) obtain a precise non-asymptotic assessment of the risk (ii) provide indications, based on theoretical analysis on the weights initialization of the neural network and its architecture.

Our main contribution is to provide risk bounds for both sigmoid and ReLu activation functions, that, to the best of our knowledge, improve on the existing results for shallow neural networks minimizing the training error/empirical risk. In the case of the sigmoid activation functions it attains the minimax optimal rate $O(n^{-2r/(2r+D_0)})$ up to $\log(n)$ factors over the unit

ball in the Sobolev space $W_2^r(\mathcal{X})$. This result competes with deep neural networks, and shows that, roughly speaking, aggregation acts as an additional layer. For the case of ReLu, we obtain a first bound for functions of the Barron spectral space $\mathscr{B}^s(\mathcal{X})$ of order $O(n^{-2K/(2K+1)})$, where K depends on s and D_0. However it is not directly comparable with the sigmoid case since they do not deal with Sobolev spaces. Nevertheless, it is possible to bridge the gap between Barron space and Sobolev functions under some restrictions on r –we refer the reader to Section 6.6.2 of Chapter 6 for further details. In short, we obtain a risk bound of order $O(n^{-2\bar{r}/(2\bar{r}+D_0+1)})$, for any $\bar{r} < r$, over the unit ball in $W_2^r(\mathcal{X})$. The obtained bound is slightly worse than the minimax rate proved for deep networks: this is due to the absence, to the best of our knowledge, of sharp approximation properties for ReLU shallow neural networks in Sobolev spaces. Still, it improves existing bounds for ReLU shallow neural networks.

1.4 Outline of the thesis

The thesis is composed of 4 original chapters, based on the following works:

. Chapter 3: *Price formation and optimal trading in intraday electricity markets*, arXiv preprint arXiv:2009.04786. Joint work with Olivier Féron and Peter Tankov.

. Chapter 4: *Price formation and optimal trading in intraday electricity markets with a major player*, published in Risks 2020, 8(4), 133; Special Issue Stochastic Modeling and Pricing in Energy Markets. Joint work with Olivier Féron and Peter Tankov.

. Chapter 5: *Decision making with dynamic probabilistic forecasts*, arXiv preprint arXiv:2106.16047. Joint work with Peter Tankov.

. Chapter 6: *Risk bounds for aggregated shallow neural networks using Gaussian priors.* Joint work with Arnak Dalalyan.

Chapter 2

Introduction (Français)

Cette thèse vise à fournir des outils scientifiques pour soutenir la pénétration croissante des énergies renouvelables intermittentes dans le système de fourniture électrique, en essayant de comprendre et de gérer leur impact sur les marchés court terme de l'électricité.

2.1 Context général et motivation

2.1.1 Une transition énergétique et technologique

Nous assistons à une montée en puissance des actions politiques sur le thème du changement climatique ces dernières années: le paquet climat-énergie de l'Union européenne (2008 et 2014), l'accord de Paris (2015), le pacte vert pour l'Europe (2020); et en France plus particulièrement, la loi sur la transition énergétique pour la croissance verte (2015), le plan climat (2017). Les pays européens sont déterminés et ambitieux: au cours de la dernière décennie, ils se sont tous engagés dans des actions nationales visant à augmenter la contribution des énergies renouvelables dans leur mix énergétique, ou à limiter les émissions de carbone; certains d'entre eux souhaitent même atteindre la neutralité carbone d'ici 2050. En France, pour atteindre cet objectif, il est prévu de retirer du marché les voitures diesel et à essence d'ici 2040, au profit d'un parc automobile 100 % électrique. La voiture électrique, fer de lance d'un système décarboné, contribue également à une demande en électricité plus élevée et plus aléatoire. Le pays soutiendra également les initiatives personnelles de consommation d'énergie autoproduite, grâce à des systèmes photo-voltaïques par exemple. Parallèlement, une transition sociétale et technologique se dessine: nous évoluons vers un mode de vie plus dépendant en énergie et en électricité. L'essor du travail à distance ces deux dernières années est un exemple très concret de ce changement: il nous a rendu plus connectés mais a aussi permis plus de flexibilité donnant naissance à de nouvelles habitudes de consommation, moins prévisibles.

Nous assistons donc au développement d'un mix énergétique heureusement plus vert, mais aussi plus incertain, confronté à des habitudes de demande moins prévisibles et à une plus grande dépendance énergétique. Une question se pose alors: comment rendre ces deux tendances adverses durables ?

Les réseaux intelligents (plus communément appelés *smart grids*) visant à plus de flexibilité du système, le développement de moyens de stockage plus abordables, et les marchés court terme de l'électricité sont des éléments clés pour relever ce défi, mais ils nécessitent une recherche pratique et théorique rigoureuse en amont. Ma thèse est une contribution à ce travail de recherche et se concentre principalement sur la gestion des énergies renouvelables intermittentes par et pour les producteurs renouvelables utilisant les marchés court terme de l'électricité. Ce travail aborde quatre thèmes principaux: la modélisation de la formation des prix sur les marchés court terme de l'électricité, la modélisation des prévisions de production d'énergies renouvelables, l'obtention de stratégies optimales pour les producteurs d'énergies renouvelables, et enfin, l'apport de garanties quantitatives pertinentes pour l'estimation par réseaux de neurones.

En guise de préambule, la section suivante est consacrée à une description succincte des marchés court terme de l'électricité.

2.1.2 Les marchés successifs court terme de l'électricité: un bref aperçu

Depuis la libéralisation des marchés de l'électricité en Europe et en France, initiée dans les années 90, le caractère non stockable de l'électricité et la politique de transition énergétique ont conduit à la création de plusieurs nouveaux marchés. Parmi ceux-ci, on peut distinguer les marchés de *trading* composés du marché *day ahead* (également appelé *"spot"*) et du marché *intraday*; ainsi que les marchés de services auxiliaires utilisés pour gérer les réserves et assurer le mécanisme d'équilibrage et son règlement. Ces marchés à court terme constituent un outil important pour gérer l'incertitude et les risques liés à la transition énergétique et technologique. Alors que les marchés de l'énergie à long terme, tels que le marché des *futures* ou les contrats *forwards*, semblent essentiels pour couvrir les contrats à long terme et la planification de la production, les marchés court terme sont cruciaux pour assurer l'équilibre constant entre l'offre et la demande d'électricité.

Concernant les marchés de *trading*, plusieurs plateformes d'échange ont vu le jour en Europe depuis 2001 avec la création du marché *day ahead*, suivi de l'*intraday* en 2010. Ces plateformes peuvent concerner un pays comme Powernext -initialement pour la France, une zone géographique spécifique comme Northpool pour les pays scandinaves ou Apx-endex pour l'Angleterre, la Belgique et les Pays-Bas. Ces marchés transfrontaliers sont particulièrement importants car ils permettent de lisser l'énergie et la demande disponibles entre les pays. En conséquence, un marché européen plus large, EPEXSPOT, né de la fusion de certaines des plateformes d'échange mentionnées ci-dessus, permettant des échanges nationaux et transfrontaliers, a été créé en 2008 et est aujourd'hui le leader des marchés de l'électricité en Europe [1]. Concernant les marchés de services auxiliaires et le système d'ajustement, l'Electricity Balancing agreement (2017) de la commission de régulation de l'UE, a établi un espace commun et des règles communes sur le règlement des écarts en Europe, où les réserves sont constituées par appel d'offres. Ces évolutions très récentes démontrent du jeune âge et de l'expansion continue de ces marchés.

[1] Toutes les analyses empiriques faites dans cette thèse sont basées sur les carnets d'ordres intraday d'EPEXSPOT pour la France ou l'Allemagne pour les hivers 2015 et 2017.

Pour chaque période de livraison de la journée, généralement divisée en heures, plusieurs acteurs sont chargés d'assurer l'équilibre énergétique. La chronologie des marchés court terme met en évidence leurs différents rôles: les producteurs et les fournisseurs d'électricité interviennent principalement sur les marchés de *trading*, avant la période de livraison; tandis que le gestionnaire du réseau de transport (GRT) intervient en temps réel, ou a posteriori, sur la gestion des réserves et les mécanismes d'équilibrage. Comme ils sont d'un grand intérêt pour le présent travail de thèse, nous détaillons dans la suite chaque étape de ces marchés successifs et illustrons certains mécanismes avec l'exemple d'un producteur d'énergie éolienne, *Producteur A*, confronté à une production terminale incertaine pour une date de livraison T, dont la prévision est régulièrement mise à jour. La description suivante est également résumée de manière plus synthétique dans la Figure 2.1.

Marché day ahead/spot Régi par un mécanisme d'enchères à l'aveugle, il ouvre un jour avant l'heure de livraison et ferme à minuit. Chaque heure de livraison du jour suivant est un produit de ce marché. Pour chacun des 24 produits, à partir des volumes et des prix d'achats et de ventes proposés, les courbes d'offre et de demande déterminent le prix d'équilibre du marché.
Producteur A : elle offre un certain volume de puissance X_0 qu'elle s'engage à fournir à l'heure de livraison T correspondant à une certaine heure du jour suivant. Si l'offre est acceptée, pour cette heure et son volume, elle est payée au prix d'équilibre, noté S_0 en €/MWh et elle recevra alors $X_0 S_0$.

Marché intraday de l'électricité Il s'agit d'un marché de *trading* en continu doté d'un carnet d'ordres qui apparie les ordres au fur et à mesure qu'ils arrivent. Il ouvre à 15h la veille de la livraison, ferme 5 minutes ou 15 minutes avant l'heure de livraison selon les zones, et propose des produits horaires – comme le marché *day ahead*, des produits demi-horaires et même quart-horaires.
Producteur A : à une date $t \in [0, T]$, elle reçoit une mise à jour de sa prévision de production notée X_t, telle que $\Delta X_t = X_0 - X_t > 0$, signifiant qu'elle prévoit de produire moins que ce à quoi elle s'est engagée sur le marché spot. En réaction, elle prend une position ϕ_t sur le marché intraday pour ajuster l'erreur de prévision ΔX_t. A la clôture du marché, elle a une position ϕ_T sur le marché intraday et un engagement global $X_0 - \phi_T$ de production, qu'il faut comparer à la quantité réelle X_T qu'elle produira.

Services auxiliaires et mécanisme d'ajustement Les responsables d'équilibre agissant sur les marchés de l'électricité, sont responsables de l'équilibre entre leurs engagements de fourniture et d'offre. Pendant l'heure de livraison, le GRT (ou les opérateurs sous contrat avec le GRT) est chargé de pourvoir à l'excès/au manque d'offre. Pour cela, les réserves primaires et secondaires, désormais appelées uniformément "*Frequency Containment Reserve (FCR)*" et "*Frequency Restoration Reserve (FRR)*", peuvent être utilisées de manière automatique ou semi-automatique pour injecter/retirer de la puissance sur le réseau, dans un délai rapide - du temps réel à moins de 15 minutes. Si le déséquilibre est trop important ou dure trop longtemps, une

troisième réserve plus lente et activée manuellement, désormais appelée "*Replacement Reserve (RR)*", est utilisée pour rétablir ou soutenir le niveau requis de FRR afin de prévenir d'éventuels déséquilibres supplémentaires du système.

Le processus d'équilibrage global nécessite alors des réserves proposées par les producteurs et/ou les consommateurs qui doivent être payées pour cela. D'autre part, les responsables d'équilibre sont "pénalisés" par rapport à leur propre déséquilibre. La pénalité d'écart est fixée par le GRT (généralement quelques heures après l'heure de livraison) en fonction des coûts supportés par le GRT pour la gestion du déséquilibre global du système.
Producteur A : supposons que sa production effective soit telle que $X_0 - \phi_T > X_T$ *, elle devra alors payer la pénalité* P *calculée par le GRT pour la différence* $X_0 - \phi_T - X_T$ *à un coût global* $(X_0 - \phi_T - X_T)P$.

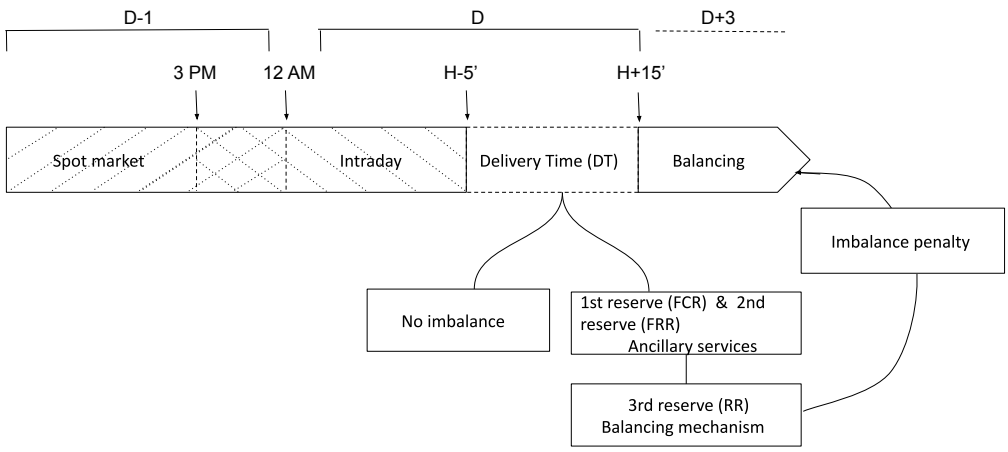

Figure 2.1: Chronologie schématique des marchés de l'électricité à court terme

Dans la suite, nous présentons les principaux défis liés à la gestion des énergies renouvelables sur les marchés court terme de l'électricité.

2.1.3 Les principaux enjeux pratiques de la thèse

Suite à la section précédente, nous comprenons pourquoi l'essor des énergies renouvelables intermittentes et l'émergence de nouveaux usages de l'électricité mettent davantage l'accent sur les marchés court terme de l'électricité. La participation croissante de producteurs d'énergies renouvelables sur le marché *intraday* entraîne plusieurs défis à relever pour réussir une transition durable vers un système énergétique décarboné. L'une des conséquences de la pénétration accrue des énergies renouvelables est l'émergence de faits stylisés sur les prix infrajournaliers, tels que la corrélation négative entre les prix et les prévisions de production renouvelable (voir la Figure 2.2a). Nous renvoyons le lecteur à Kiesel and Paraschiv (2017); Karanfil and Li (2017); Rowińska et al. (2018) pour d'autres études empiriques sur les marchés court terme de l'électricité.

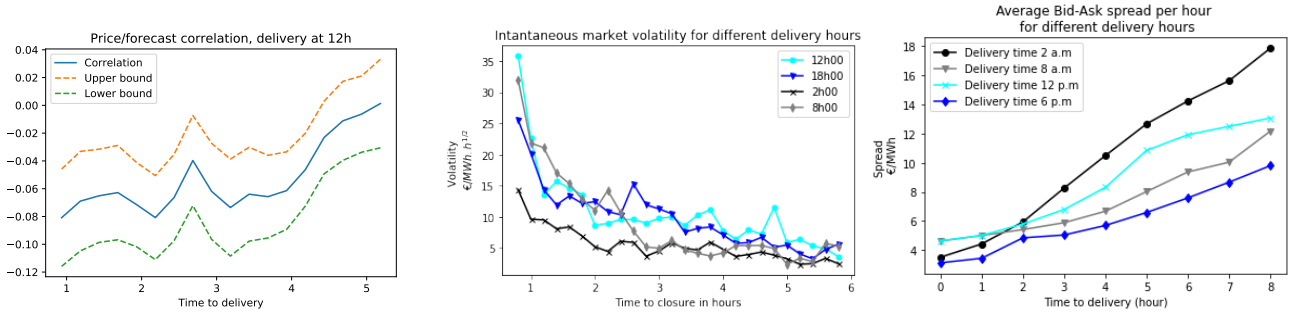

(a) Corrélation entre le prix et les prévisions de production d'énergie éolienne

(b) Volatilité empirique

(c) Ecarts de cotations

Figure 2.2: Caractéristiques empiriques de la zone de livraison allemande durant les hivers 2015 et 2017

En outre, pour éviter les coûts liés aux pénalités d'écart, la plupart des acteurs peuvent avoir tendance à ajuster leur position jusqu'à peu de temps avant la date de livraison. Outre la corrélation entre le prix et les prévisions de production, ce comportement peut également expliquer d'autres caractéristiques du marché telles que la volatilité croissante (Figure 2.2b), et la liquidité croissante (Figure 2.2c) matérialisée par des écarts de cotation décroissants au cours des dernières heures de *trading*. A partir des caractéristiques du marché, nous constatons un mécanisme de rétroaction négatif: les échanges de nombreux producteurs d'énergie renouvelable ont un impact sur le prix qui, en retour, a un impact sur leurs revenus. Ces premières observations suscitent les deux aspirations principales du Chapitre 3 et du Chapitre 4:

Problème 1. *Trouver des stratégies pour les producteurs interagissant sur le marché qui leur permettent de minimiser leurs coûts tout en prenant en compte la production incertaine à laquelle il font face et leur impact sur le marché.*

Problème 2. *Déterminer, à l'aide de notre modèle, comment les interactions des producteurs renouvelables peuvent être un élément clé pour expliquer les caractéristiques observées sur le marché intraday.*

Rendre ce niveau croissant d'incertitude économiquement viable à long terme est conditionné par trois enjeux. Premièrement, trouver des moyens concrets de modéliser le prix de marché et élaborer des stratégies optimales pour les producteurs, comme nous venons de le voir. Deuxièmement, afin d'éviter un recours massif aux mécanismes d'ajustement et de rendre le marché/mix énergétique plus efficace, il est absolument crucial de travailler sur la capacité des prévisions de production à: (*i*) refléter correctement la production à venir, (*ii*) quantifier de manière dynamique les variations de l'incertitude de la prévision. Troisièmement, fournir des outils théoriques capables de garantir de bonnes estimations dans les problèmes de plus en plus complexes qui se posent concernant la gestion de la production renouvelable.

En ce qui concerne le deuxième point, la participation des producteurs renouvelables dépend évidemment de la vision qu'ils ont de leur production, mais elle est aussi fortement influencée par la confiance qu'ils lui accordent. A cet égard, il est nécessaire de quantifier l'incertitude liée

à leurs prévisions de manière à leur permettre de prendre des décisions éclairées. Cette deuxième observation nous conduit aux problèmes que nous aborderons dans le Chapitre 5 :

Problème 3. *Fournir des prévisions qui permettent de suivre dynamiquement les variations de l'incertitude sur la base de données réelles.*

Problème 4. *Être capable d'exploiter la prévision ainsi que l'évolution de son incertitude dans des processus de prise de décision.*

Enfin, le troisième axe de recherche susceptible de soutenir la transition énergétique s'appuie sur les observations suivantes. Les prix et l'offre d'énergie pour des heures de livraison rapprochées sont corrélés et cet effet pourrait être renforcé par le développement futur de moyens de stockage plus abordables. Par conséquent, la modélisation des marchés court terme successifs, chacun impliquant au moins 24 produits de livraison éventuellement corrélés, et le même nombre de prévisions de production, devrait être privilégiée par rapport au cas d'une heure de livraison spécifique traité dans les Problèmes 1-4 de cette thèse. Cependant, cela augmente la complexité et la difficulté à estimer les paramètres du modèle ainsi qu'à prendre des décisions appropriées. Ces dernières années, le *machine learning* et, plus particulièrement, les réseaux de neurones ont gagné en popularité et sont maintenant fréquemment utilisés pour s'attaquer à ce type de problèmes en grande dimension. Cependant, les garanties théoriques justifiant le succès des résultats empiriques des réseaux de neurones sont encore rares. Nous pensons que l'évolution et l'utilisation pérenne des réseaux de neurones doit s'accompagner d'avancées théoriques pour légitimer leurs performances empiriques. C'est pourquoi, dans le dernier chapitre, contrairement aux précédents objectifs énoncés, nous nous concentrerons exclusivement sur l'analyse théorique susmentionnée, laissant l'application aux marchés de l'énergie à un travail ultérieur:

Problème 5. *Obtenir des bornes de risque pour l'apprentissage par réseaux de neurones agrégés peu profonds.*

Problème 6. *Fournir des indications pratiques sur les paramètres d'ajustement (initialisation des poids, taille des réseaux, etc.), tirées de l'analyse théorique.*

Pour répondre aux questions soulevées dans cette section, nous nous appuyons sur des outils mathématiques, provenant principalement de la théorie du contrôle stochastique, de la théorie des jeux et des statistiques. Dans la section suivante, nous les présentons et fournissons quelques références bibliographiques pertinentes.

2.2 Cadre théorique

Dans cette section, nous présentons les principaux fondements théoriques de nos développements. Ici, le but visé n'est pas de faire une revue exhaustive de l'état de l'art mais simplement de mentionner les résultats qui nous semblent les plus pertinents pour contextualiser le présent travail. Le champ d'application de la thèse s'étend de la dérivation de stratégies optimales pour les producteurs renouvelables, à la théorie de l'apprentissage automatique, en passant par la calibration

et la simulation. Les deux premiers chapitres se concentrent principalement sur les problèmes de contrôle stochastique, le troisième chapitre est partagé entre le contrôle et les statistiques et enfin, le quatrième chapitre se concentre sur les statistiques. Nous suivons l'organisation des chapitres pour présenter les outils théoriques.

2.2.1 Jeux stochastiques

Comme nous l'avons mentionné dans la section précédente, nous nous intéressons à plusieurs producteurs renouvelables dotés de prévisions de production incertaines et qui souhaitent minimiser leur coût en interagissant sur le marché. Dans ce contexte, il est naturel de les considérer comme les joueurs d'un jeu non coopératif. Pour cela, les jeux stochastiques offrent un cadre approprié: dans la suite, nous spécifions deux principaux types de jeux stochastiques d'intérêt pour les problèmes que nous développerons dans le corps de la thèse.

Contrôle optimal Les jeux stochastiques trouvent leurs racines dans la théorie du contrôle stochastique, dont la valeur pour les mathématiques financières modernes est largement reconnue. Pour un problème modélisé par une fonction objectif à minimiser et un espace d'état, comprenant une ou plusieurs variables contrôlées, on cherche à exhiber, s'il existe, un contrôle qui garantisse l'optimalité. La notion de fonction valeur pour un problème de contrôle découle du principe de la programmation dynamique. De manière heuristique, la formule de Feynmac-Kac comble alors le fossé entre le principe de programmation dynamique et sa contrepartie infinitésimale, de sorte que la solution prend généralement la forme d'une équation aux dérivées partielles (EDP): l'équation dite de Hamilton-Jacobi-Bellman (HJB). Le principe de Pontryagin offre une autre approche sous la forme d'une équation différentielle stochastique rétrograde (EDSR). Cette dernière peut sembler plus intuitive si l'on pense à la procédure rétrograde habituellement employée pour traiter les problèmes de programmation dynamique tels que les problèmes financiers classiques de réplication d'éventuelles créances futures. Nous référons le lecteur à Touzi (2012) pour une introduction complète aux problèmes classiques, ou plus sophistiqués, de contrôle stochastique.

Les problèmes de *trading* optimale sur les marchés financiers pour un agent unique avec impact sur les prix ont été largement étudiés (voir *e.g.* Bank et al. (2017)); et dans la plupart des cas, un modèle à la Almgren and Chriss (1999, 2001) est utilisé, *i.e* avec un impact de marché linéaire sur les prix. Ce cadre de *price maker* a également été appliqué aux marchés de l'énergie: dans les travaux de Aid et al. (2016); Tan and Tankov (2018), un producteur faisant face à une production/demande terminale incertaine doit gérer son impact sur le prix. Notre premier axe de recherche est très proche de ces travaux dans l'idée, mais nous considérons cette fois un marché avec plusieurs agents, où le mouvement de foule peut impacter chacun d'entre eux.

Équilibre de Nash Cette notion a été définie dans l'article fondateur de Nash (1950), et correspond au jeu non coopératif que nous voulons construire, où l'existence d'un équilibre sur le marché dépend des interactions des participants. Supposons que le jeu se déroule d'une date

$t = 0$, à une date T juste avant la fermeture du marché. Considérons N joueurs, un espace de probabilité filtré $(\Omega, \mathcal{F}, \mathbb{F} := (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{P})$, et $\phi^i := (\phi^i_t)_{t\in[0,T]}$, $i = 1, \dots, N$, les positions stratégiques $\mathbb{F}$-adaptées des joueurs sur le marché. Chaque joueur bénéficie d'une récompense donnée par la fonction $J^{N,i}$, dépendant de sa stratégie et des stratégies des autres joueurs. Un équilibre de Nash est un état tel qu'aucun des joueurs n'a d'incitation à changer de position, c'est-à-dire, plus formellement:

Définition 7 (Équilibre de Nash)**.** *$(\phi^{i*})_{i=1,\dots,N}$ est un équilibre de Nash pour le jeu à N joueurs si c'est un vecteur de stratégies admissibles – basé sur certains critères d'admissibilité, et pour chaque $i = 1, \dots, N$,*

$$J^{N,i}(\phi^i, \phi^{-i*}) \leq J^{N,i}(\phi^{i*}, \phi^{-i*}) \tag{2.1}$$

pour toute autre stratégie admissible ϕ^i, où $\phi^{-i} := (\phi^1, \dots, \phi^{i-1}, \phi^{i+1}, \dots, \phi^N)$.

Une autre notion intéressante, et sur laquelle nous reviendrons plus tard (voir la section 2.3.1), est celle d'un équilibre ε-Nash qui consiste à trouver un ensemble de positions admissibles $\left(\phi^{i*}\right)_{i=1,\dots,N}$ qui est ε-optimal au sens de la définition suivante :

Définition 8 (Equilibre ε-Nash)**.** *On dit que $(\phi^{i*})_{i=1,\dots,N}$ est un équilibre ε-Nash pour le jeu à N joueurs si ces stratégies sont admissibles et pour toute autre stratégie admissible ϕ^i du joueur i, $J^{N,i}(\phi^i, \phi^{-i*}) - \varepsilon \leq J^{N,i}(\phi^{i*}, \phi^{-i*})$, $i = 1, \dots, N$.*

Les équilibres dans les jeux stochastiques ne sont pas faciles à aborder mais le cadre linéaire quadratique est un outil classique pour obtenir des solutions explicites (Bouchard et al., 2018; Voß, 2019; Evangelista and Thamsten, 2020; Bank et al., 2021). Cependant, dans certains cas, ce schéma simplifié n'est toujours pas suffisant pour exhiber un équilibre. En théorie des jeux, on peut distinguer le cas où tous les joueurs ont une information complète sur l'état et la stratégie des autres joueurs, du cas incomplet où une certaine partie du comportement stratégique des autres joueurs est cachée. Cette considération fait écho au problème initial des producteurs renouvelables si l'on note que leurs propres prévisions de production peuvent être soit totalement observées, soit partiellement observées, soit pas du tout observées par les autres acteurs; et que la même remarque pourrait être faite sur le *trading* de la foule, c'est-à-dire la position agrégée de tous les acteurs. Le premier cas est généralement facilement abordé avec la théorie de Nash et l'on peut obtenir un équilibre en forme fermée dans le cadre linéaire quadratique. En revanche, les derniers cas ne peuvent généralement pas être résolus explicitement en utilisant le cadre actuel. Pour contourner ce problème, nous présentons dans le paragraphe suivant des outils issus de la théorie des jeux à champ moyen.

Jeux à champ moyen Ils ont été introduits par Lasry and Lions (2007); Huang et al. (2006) et diffèrent de la théorie de Nash en ce qu'ils considèrent un jeu avec une infinité de joueurs, traditionnellement symétriques. Dans le cas de l'équilibre de Nash, l'équilibre résulte des interactions de N joueurs à travers leurs positions $(\phi^i)_{i=1,\dots,N}$. Par exemple, dans le Chapitre 3, les

agents interagissent par le biais de l'impact des joueurs sur le prix exprimé comme une fonction linéaire de la position agrégée des agents : $\frac{1}{N}\sum_{i=1}^{N}\phi_t^i$, $\forall t \in [0,T]$. Il est intéressant de voir cette dernière comme une espérance par rapport à la distribution empirique $\mu_t^N = \frac{1}{N}\sum_{i=1}^{N}\delta_{\phi_t^i}$, pour laquelle, heuristiquement, une extension naturelle dans le cadre des jeux à champ moyen est le flux de la distribution limite $(\mu_t^\phi)_{t\in[0,T]}$ de la position de l'agent représentatif $\phi := (\phi_t)_{t\in[0,T]}$. Si nous voulons développer un modèle où l'information n'est pas totalement partagée par les agents, en plus de la filtration $\mathbb{F}$ à laquelle tous les processus sont adaptés, il faut introduire une plus petite filtration $\mathbb{F}^0$, représentant le bruit commun, à laquelle l'information privée n'est pas adaptée. Il s'ensuit que les analogues jeux à champ moyen des quantités agrégées dans le problème à N agents sont des espérances conditionnelles par rapport à la filtration du bruit commun $\mathbb{F}^0$. Le jeu est maintenant représenté par l'interaction des agents à travers le flux de la distribution conditionnelle du processus d'état : $\mu_t^{\phi,0} := \mathcal{L}(\phi_t|\mathcal{F}_t^0)$. Si nous désignons par J^{MF} la fonction objectif de l'agent représentatif, nous pouvons définir un équilibre de jeu à champ moyen de manière similaire à celui de Nash,

Définition 9 (Equilibre du jeu à champ moyen)**.** *Une stratégie admissible ϕ^* est un équilibre du jeu à champ moyen si pour toute stratégie admissible ϕ,*

$$J^{MF}(\phi, \bar{\phi}^*) \leq J^{MF}(\phi^*, \bar{\phi}^*),$$

avec $\bar{\phi}_t = \mathbb{E}[\phi_t|\mathcal{F}_t^0] = \int_{\mathbb{R}} x\mu_t^{\phi,0}(dx)$.

Pour approfondir les aspects théoriques, nous renvoyons le lecteur à Lasry and Lions (2007); Carmona and Delarue (2018); Cardaliaguet et al. (2019); Lacker (2020).

Pour aborder ce type de jeux stochastiques, les méthodes sont finalement aussi très proches de celles évoquées dans le cadre plus simple du contrôle stochastique. L'approche initiale de Lasry and Lions (2007); Huang et al. (2006) s'appuie sur l'équation de HJB de la fonction valeur, couplée à une équation de Fokker-Planck pour formaliser la distribution du champ moyen par laquelle les joueurs interagissent. Reflétant le principe du maximum stochastique, l'approche probabiliste de Carmona and Delarue (2018) formalise le problème par une équation différentielle stochastique progréssive rétrograde (EDSPR) de McKean-Vlasov. L'extension McKean-Vlasov de l'EDSPR classique est introduite pour gérer le flux de distribution inconnu du champ moyen $(\mu_t^\phi)_{t\in[0,T]}$. Cette dernière approche est plus proche, dans l'idée, de la méthode employée dans les Chapitres 3 et 4.

La façon dont les jeux à champ moyen permettent de simplifier des problèmes complexes les rend très attractifs dans des applications à la finance (Casgrain and Jaimungal, 2018, 2020; Fu et al., 2021; Fu and Horst, 2020b; Fu et al., 2020; Fujii and Takahashi, 2020), aux marchés de l'énergie (Shrivats et al., 2020), ou bien encore, pour traiter des problèmes d'allocation de stockage d'énergie (Alasseur et al., 2020). Cependant, le cadre classique des jeux à champ moyen considère une classe d'agents symétriques. Dans Alasseur et al. (2020); Casgrain and Jaimungal (2018, 2020), de l'hétérogénéité est introduite à travers des sous-populations disposant d'informations partielles et de croyances différentes. Pour le problème que nous abordons dans la thèse, le critère d'hétérogénéité sur lequel s'appuyer pour prendre en compte la réalité des

conditions du marché est à discuter. Un exemple frappant de ces conditions est illustré par le marché français où EDF a une position dominante pour la production et la fourniture d'énergie, ce qui fait de l'entreprise un acteur majeur du marché par rapport aux autres producteurs. Avec le développement des énergies renouvelables – notamment le photovoltaïque, un phénomène inverse se produit car le nombre de très petits producteurs qui sont ou pourraient être un jour connectés au marché et au réseau augmente. Dans ce cas, l'hétérogénéité provient de la taille et donc du pouvoir de marché des acteurs. Heureusement, cette situation peut être traitée grâce au cadre spécifique des jeux à champ moyen avec un acteur majeur, également connus sous le nom de jeux à champ moyen à la Stackelberg (voir *e.g.* Huang (2010); Nourian and Caines (2013); Bensoussan et al. (2016b); Carmona et al. (2016) pour plus de détails sur le sujet, une application aux marchés financiers peut être trouvée dans Evangelista and Thamsten (2020)). Le Chapitre 4 est consacré à ce cadre hétérogène, et à ses conséquences sur le comportement stratégique des agents.

Les jeux stochastiques constituent un outil pratique pour aborder les Problèmes 1–2. Cependant, les gains des producteurs dépendent aussi de manière importante de la qualité des prévisions qu'ils utilisent pour prendre leurs décisions sur le marché. Cette observation concerne les Problèmes 3–4, pour lesquels nous introduisons dans le paragraphe suivant la notion de prévisions probabilistes, leur approche la plus courante et les techniques de calibration qu'elles impliquent généralement.

2.2.2 Prévisions probabilistes

Supposons que l'on doive prévoir une variable aléatoire X à une date donnée t avant sa réalisation à un instant $T > t$. La prévision peut prendre la forme d'une information ponctuelle, c'est-à-dire d'une vision de la valeur de X à l'instant t; ou bien elle peut également fournir des indications sur la fiabilité de cette information, ce qui est l'essence des prévisions probabilistes. Par opposition aux prévisions ponctuelles, ces dernières permettent de prendre en compte l'incertitude portée par une prévision, ce qui constitue une réponse possible au Problème 3.

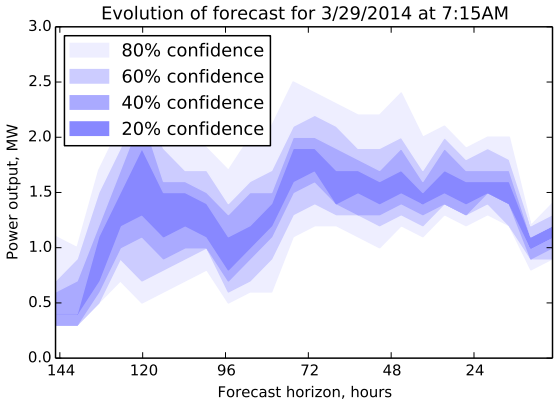

Figure 2.3: Évolution de la prévision probabiliste de la production électrique d'une centrale éolienne en fonction du temps, pour une date de production donnée

La plupart du temps, les prévisions probabilistes sont présentées sous la forme d'intervalles de confiance ou de quantiles autour de la prévision ponctuelle comme c'est le cas sur la Figure 2.3. Dans la littérature, une grande partie de la recherche s'est concentrée sur la façon de modéliser, de calibrer et d'évaluer la qualité des prévisions probabilistes.

L'estimation de la densité prédictive entière est généralement très coûteuse puisqu'il s'agit d'un objet de dimension infinie. Pourtant, la plupart du temps, l'information réellement disponible est de faible dimension. Pour cette raison, une spécification bidimensionnelle est parfois choisie, basée sur deux facteurs observables tels que la moyenne et la variance de la variable aléatoire X compte tenu de l'information disponible à l'instant t, que l'on peut désigner respectivement par m et V. Les articles novateurs de Gneiting et al. (2005, 2007) traitent de la calibration et de l'évaluation de tels modèles. Pour un horizon de prévision donné, ils utilisent une distribution Gaussienne pour la densité prédictive, $\mathcal{N}(m, V)$, où m, V doivent être déterminés par calibration. À la suite de ces travaux, plusieurs auteurs ont développé des modèles plus élaborés utilisant des mélanges Gaussiens (Wilks, 2002; Raftery et al., 2005), des distributions tronquées (Baran, 2014), des distributions log-normales (Baran and Lerch, 2015), des mélanges de distributions tronquées et log-normales (Baran and Lerch, 2016); pour différents objectifs tels que le traitement de quantités positives ou d'événements de valeurs extrêmes. Dans ces références, les prévisions probabilistes sont généralement utilisées pour des données météorologiques. Cela les rend spécialement appropriées au contexte des énergies renouvelables ou de la demande d'électricité qui dépendent largement de phénomènes météorologiques tels que le vent ou la température, sur lesquels nous nous concentrerons dans le Chapitre 5.

Méthode de calibration La plupart du temps, les données météorologiques sont représentées sous la forme de prévisions d'ensemble. La Figure 2.4a en donne un exemple pour la vitesse du vent: elle montre une prévision d'ensemble pour un horizon de prévision de 48h, composée de $M = 50$ membres que nous désignerons par $X_1, \ldots, X_M$. Les différents scenarii évoluent avec l'horizon de prévision et l'écart entre eux varient avec le temps. Par exemple, de manière assez surprenante, l'incertitude 20 heures avant la réalisation semble moins importante que l'incertitude 8 heures avant. Notons $X_M = \frac{1}{M}\sum_{m=1}^{M} X_m$ et $S_M = \frac{1}{M-1}\sum_{m=1}^{M}(X_m - \overline{X}_M)^2$ la moyenne et la variance empiriques de la prévision d'ensemble.

Une première tentative naïve serait, dans le cas du modèle Gaussien simple de Gneiting et al. (2005), de choisir comme densité prédictive $\mathcal{N}(\overline{X}_M, S_M)$. Cependant, les membres de l'ensemble sont obtenus en exécutant un modèle de prévision déterministe avec des conditions initiales perturbées. De ce fait, ils ne représentent pas la meilleure approximation de la réalisation et la distribution empirique des membres de l'ensemble souffre souvent d'un biais et d'une sous-dispersion comme le montre la Figure 2.4b.

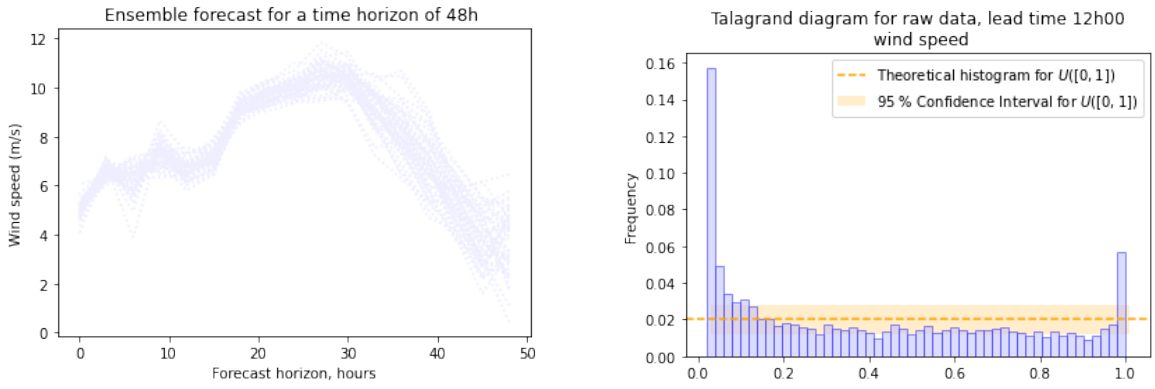

(a) Prévisions d'ensemble

(b) Diagramme de Talagrand

Figure 2.4: Un exemple de prévision d'ensemble et des possibles lacunes de ces prévisions

Le diagramme de Talagrand présenté dans cette figure est un outil permettant de vérifier la qualité d'une calibration et est analogue à un histogramme de rang $R(y, (X^m)_{m=1\ldots M})$ de la réalisation y dans les prévisions de l'ensemble $(X^m)_{m=1\ldots M}$. Dans une prévision d'ensemble parfaite, l'histogramme devrait être celui d'une distribution uniforme. Dans le cas présent, l'histogramme de la Figure 2.4b, présente un profil en U, qui correspond à une sous-dispersion de l'ensemble et une forme asymétrique qui indique un biais. Pour résoudre ce problème, en fonction du type de densité prédictive utilisé, plusieurs méthodes de calibration existent. Parmi elles, les plus courantes sont l'*Ensemble Model Output Statistic* (EMOS) et le *Bayesian Model Averaging* (BMA). Dans ce travail, nous utiliserons une technique inspirée de la méthode de traitement a posteriori EMOS qui revient simplement à estimer la distribution avec une correction linéaire de la moyenne et de la variance – *e. g.* dans Gneiting et al. (2005), la distribution normale est calibrée: $\mathcal{N}(a_0 + \sum_{m=1}^{M} a_m X_m, c + dS_M)$ avec $a_0, \ldots, a_M, c, d$ des coefficients de régression.

Le cadre de la diffusion Dans les travaux mentionnés précédemment, il n'y a pas de dimension dynamique dans les modèles, qui ne tiennent donc pas compte des dépendances temporelles entre les prévisions à différents horizons. Cette faiblesse peut être corrigée à l'aide d'une modélisation par diffusion stochastique, notamment développée dans le cadre de prévisions de vitesse du vent (Iversen et al., 2016; Bensoussan and Brouste, 2016) et de prévisions d'irradiance solaire (Badosa et al., 2017). Ces travaux soulignent l'utilité de ce champ de recherche pour les applications pratiques/industrielles dans le domaine des énergies renouvelables. Ils portent principalement sur la modélisation de l'évolution des prévisions ponctuelles à l'aide d'équations différentielles stochastiques (EDS), permettant de dériver des prévisions probabilistes, souvent sous la forme de quantiles. Bensoussan and Brouste (2016) ont développé un modèle Cox-Ingersoll-Ross (CIR) pour la vitesse du vent qui permet de dériver un flux de densités prédictives pouvant être correctement calibrées. Dans le Chapitre 5, nous proposons plusieurs modèles d'EDS rendant compte de l'évolution dynamique de la prévision et de son incertitude. Nos modèles permettent également de dériver des densités prédictives, dont nous donnons la définition suivante dans le cadre stochastique.

Définition 10. *Soit* $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ *un espace de probabilité filtré et* T *un horizon temporel fixé. Une prévision probabiliste d'une variable aléatoire mesurable* $X \in \mathbb{R}^d$ *est la distribution conditionnelle* p_t *de* X *étant donné* $\mathcal{F}_t$*. Un modèle dynamique pour une prévision probabiliste est alors un flux de mesures de probabilité* $(p_t)_{0 \leq t \leq T}$*, qui peut être identifié au flux des distributions conditionnelles de* X*.*

Le cadre diffusif permet également d'intégrer des prévisions probabilistes dans des problèmes de décision stochastiques. Dans le Chapitre 5, nous proposons de traiter le Problème 4 par la résolution numérique d'un problème de contrôle séquentiel tridimensionnel. La méthode de régression moindres carrés Monte Carlo que nous utilisons ne fonctionne bien que jusqu'en dimension 3 ou 4, la dimension devient donc un obstacle pour le développement de problèmes de gestion plus complexes. Le besoin de nouvelles méthodes pour les problèmes d'apprentissage et d'estimation a conduit au développement du *machine learning* et plus particulièrement des réseaux de neurones, dont les très bons résultats empiriques ne sont plus à prouver mais qui manquent cependant souvent de garanties théoriques. Dans le Chapitre 6, nous dérivons des bornes de risque pour les réseaux de neurones dans le cadre Probablement Approximativement Correct-Bayésien (PAC-Bayésien) que nous présentons dans la section suivante.

2.2.3 Réseaux de neurones et inégalités PAC-Bayesiennes

La dernière partie de cette thèse vise à obtenir des garanties sur la performance de réseaux de neurones agrégés. En tant que telle, elle n'est pas directement liée aux marchés de l'électricité. Néanmoins, les techniques avancées de *machine learning* ont montré d'excellents résultats empiriques pour traiter des problèmes en grande dimension, et de récents travaux utilisent les réseaux de neurones dans des applications à des problèmes de contrôle stochastique, très utiles pour le stockage d'énergie (Bachouch et al., 2021) et la gestion des risques (Fécamp et al., 2019). Certains de ces travaux fournissent même une analyse de la convergence pour des algorithmes spécifiques (voir Huré et al. (2021) pour la programmation dynamique, Becker et al. (2019, 2021) pour les problèmes d'arrêt optimal et Huré et al. (2020); Germain et al. (2020) pour les approches EDSR). Malgré ces avancées remarquables, les résultats s'appuient souvent sur le théorème d'approximation universelle (Hornik et al., 1989) et les schémas de discrétisation (Bouchard and Touzi, 2004) pour présenter des garanties de convergence, ce qui ne permet pas de mettre en évidence le comportement des réseaux de neurones avec un nombre fini d'observations, ou encore l'impact de leur structure sur l'erreur de test. L'établissement de garanties théoriques générales pour ce genre de problématique est un domaine de recherche actif en *machine learning*, les résultats existants se concentrant principalement sur les propriétés d'un réseau de neurones minimisant l'erreur d'apprentissage.

Le Chapitre 6 propose donc d'étudier des bornes de risque dans le cadre de prédicteurs agrégés, dont l'application aux réseaux de neurones est un sujet très récent (Dziugaite and Roy, 2017; Neyshabur et al., 2017), et bien moins étudié que le minimiseur du risque empirique (MRE) standard. Au cours des dernières décennies, la théorie PAC-Bayesienne a offert de nombreuses possibilités de progrès concernant les garanties du risque statistique pour les estimateurs définis

comme un "mélange" d'une famille donnée d'estimateurs faibles. Dans la suite de cette sous-section, nous définissons plus précisément le problème d'apprentissage abordé dans le Chapitre 6 et présentons quelques éléments généraux sur la théorie PAC-Bayesienne.

Quelques généralités sur l'apprentissage statistique et les réseaux de neurones Soit $(\mathcal{Z}, \mathscr{A})$ un espace mesurable et $\mathbf{Z^n} = (Z_1, \dots, Z_n) \in \mathcal{Z}^n$ des réalisations tirées d'une distribution inconnue $\mathcal{P}$ sur $(\mathcal{Z}^n, \mathscr{A}^{\otimes n})$. Considérons l'ensemble $\mathcal{X} \subset \mathbb{R}^{D_0}$, $D_0 \geq 1$, et soit μ une mesure σ-finie sur $\big(\mathcal{X}, \mathscr{B}(\mathcal{X})\big)$. Le problème considéré est l'estimation d'une fonction mesurable $f_{\mathcal{P}} \in \mathcal{F} := \{f : \mathcal{X} \to \mathbb{R}^{D_L},\ D_L \in \mathbb{N}\}$, dépendant de la distribution $\mathcal{P}$. Généralement, il est trop coûteux de chercher une fonction appartenant à l'ensemble entier $\mathcal{F}$. A la place, il est plus commode d'en chercher une approximation choisie à partir d'un sous-ensemble paramétré plus petit $\mathcal{F}_{\mathsf{W}} := \{f_{\boldsymbol{w}}, \boldsymbol{w} \in \mathsf{W}\} \subset \mathcal{F}$, indéxé par des éléments de l'espace mesurable $(\mathsf{W}, \mathscr{B}(\mathsf{W}))$, $\mathsf{W} \subset \mathbb{R}^d$. Nous définissons également une fonction de perte $\ell : \mathcal{F} \times \mathcal{F} \mapsto \mathbb{R}_+$, telle que $\ell(f_{\boldsymbol{w}}, f_{\mathcal{P}}) = \int_{\mathcal{X}} \mathrm{d}(f_{\boldsymbol{w}}(\boldsymbol{x}), f_{\mathcal{P}}(\boldsymbol{x}))\mu(d\boldsymbol{x})$ mesure la performance d'une fonction $f_{\boldsymbol{w}} \in \mathcal{F}_{\mathsf{W}}$, pour une distance donnée d définie sur $\mathbb{R}^{D_L}$.

Les réseaux de neurones denses de profondeur $L \in \mathbb{N}$ sont des spécifications particulières du sous-ensemble $\mathcal{F}_{\mathsf{W}}$, où W définit les poids du réseau. Plus précisément, pour les poids $\boldsymbol{w} = (\boldsymbol{w}_1, \boldsymbol{w}_2, \dots, \boldsymbol{w}_L) \in \mathsf{W}$, la dimension globale de W est donnée par $d = D_0 D_1 + \cdots + D_{L-1} D_L$, où D_ℓ est le nombre d'unités dans la couche ℓ; un élément de $\mathcal{F}_{\mathsf{W}}$ est alors défini par $f_{\boldsymbol{w}}(x) = \boldsymbol{w}_L^\top \sigma_{L-1}(\boldsymbol{w}_{L-1}^\top \sigma_{L-2}(\dots \boldsymbol{w}_2^\top \sigma_1(\boldsymbol{w}_1^\top x)))$, où σ est une fonction d'activation et $\sigma_\ell : x \in \mathbb{R}^{D_\ell} \mapsto (\sigma(x_1), \dots, \sigma(x_{D_\ell}))^\top$, ;$\ell = 1, \dots, L-1$. Chaque w_ℓ est une matrice $D_{\ell-1} \times D_\ell$. Dans cette thèse, notre objectif est d'obtenir des garanties théoriques dans le cadre simple et explicite des réseaux peu profonds – avec une seule couche cachée; nous renvoyons le lecteur à Bishop (1994); Goodfellow et al. (2016) pour une introduction et un aperçu complet sur des structures plus complexes telles que les réseaux de neurones convolutifs, les réseaux à mémoire *LTST*, les *GANs*, les *VAEs*, etc.

La première question à laquelle il faut répondre dans un tel contexte est de savoir dans quelle mesure un réseau de neurones peu profond approxime la fonction $f_{\mathcal{P}}$. L'erreur d'approximation correspondante dépend des propriétés de la fonction $f_{\mathcal{P}}$ (*e.g.* continuité de Lipschitz ou de Hölder, espace de Sobolev), de la taille de la classe $\mathcal{F}_{\mathsf{W}}$ et du type de fonction d'activation σ. La littérature sur ce sujet s'est développée autour de ces principales caractéristiques, de sorte que pour l'approximation par des réseaux peu profonds avec des fonctions d'activation sigmoïdes, on peut citer par exemple : Mhaskar and Micchelli (1994); Petrushev (1998); Burger and Neubauer (2001); Cao et al. (2008); Costarelli and Spigler (2013a,b); Barron (1993); Delyon et al. (1995); Maiorov and Meir (2000); Maiorov (2006). En particulier, Maiorov and Meir (2000); Maiorov (2006) considèrent la classe des fonctions de Sobolev $W_2^r([0,1]^{D_0})$ et obtiennent une erreur d'approximation d'ordre $O(D_1^{-r/D_0})$. Ce résultat est particulièrement intéressant car il fait intervenir le nombre d'unités dans la couche cachée D_1, la régularité r de la fonction $f_{\mathcal{P}}$ et la dimension du problème D_0. Dans le cas des fonctions d'activation ReLU, toujours pour les réseaux peu profonds, on peut citer Petrushev (1998); Bach (2017); Klusowski and Barron (2016b,a); Xu (2020); Siegel and Xu (2020). Siegel and Xu (2020) prouvent que pour la classe

des espaces spectraux de Barron $\mathscr{B}^s([0,1]^{D_0})$ – nous renvoyons le lecteur au Chapitre 6 Section 6.5.2 expression (6.16) pour une définition de cet espace – l'erreur est d'ordre $O(D_1^{-K})$ où K dépend de s et de la dimension du problème D_0. Ces dernières années, le développement du *deep learning* a encouragé la recherche de garanties d'approximation pour des réseaux plus profonds avec des fonctions d'activation ReLU: (Yarotsky, 2017, 2018; Yarotsky and Zhevnerchuk, 2019; Gühring et al., 2020; Lu et al., 2020; Shen et al., 2019).

Il faut noter qu'en pratique, la fonction $f_{\mathcal{P}}$ étant inconnue, la meilleure approximation n'est en conséquence pas accessible. A la place, un estimateur $\widehat{f}_n : \mathcal{Z}^n \mapsto \mathcal{F}_{\mathsf{W}}$ calculé à partir des données observées $\mathbf{Z^n}$ doit être utilisé. La qualité d'un tel estimateur est évaluée par le risque défini par

$$\mathbf{E}_{\mathcal{P}}[\ell(\widehat{f}_n(\mathbf{Z^n}), f_{\mathcal{P}})] = \int_{\mathcal{Z}^n} \ell(\widehat{f}_n(\mathbf{z}), f_{\mathcal{P}})\mathcal{P}(d\mathbf{z}). \tag{2.2}$$

Pour évaluer le risque des réseaux de neurones, une littérature sur l'analyse non asymptotique des réseaux de neurones a émergé depuis les années 90, mais a été peu enrichie jusqu'à présent. Nous pouvons citer des travaux pour les réseaux de neurones denses avec des fonctions d'activation sigmoïdes (Barron, 1994; McCaffrey and Gallant, 1994), et les réseaux de neurones avec des fonctions d'activation ReLU (Klusowski and Barron, 2016b; Bach, 2017). Des résultats récents sur les réseaux de neurones profonds démontrent leur capacité à atteindre des vitesses de convergence optimales minimax à des facteurs $\log(n)$ près, pour des fonctions avec une architecture hiérarchique spécifique et une régularité suffisante, dans le cadre des réseaux ReLU (Schmidt-Hieber, 2020) ainsi que sigmoïdes (Bauer and Kohler, 2019).

La plupart des garanties établies dans la littérature, pour un réseau de neurones avec une architecture donnée, se concentrent sur le minimiseur du risque empirique (également appelé erreur d'apprentissage):

$$\widehat{f}_n^{\mathsf{MRE}} \in \operatorname{argmin}_{\mathcal{F}_{\mathsf{W}}} \frac{1}{n}\sum_{i=1}^{n} L(f_{\boldsymbol{w}}, Z_i),$$

où $L(f_{\boldsymbol{w}}, Z_i)$ est un proxy convenable pour la perte $\ell(f_{\boldsymbol{w}}, f_{\mathcal{P}})$.

Néanmoins, on peut aborder le problème de l'estimation sous un autre angle que celui de la minimisation du risque empirique. Au lieu d'essayer d'estimer la meilleure valeur possible $\boldsymbol{w}^*$ du paramètre par un estimateur appartenant à W et se basant sur les données $\widehat{\boldsymbol{w}}_n$, on peut essayer de combiner plusieurs fonctions $f_{\boldsymbol{w}}$ pour obtenir une bonne approximation de $f_{\mathcal{P}}$. Dans un cadre plus général que celui des réseaux de neurones, Juditsky et al. (2008) montrent que pour des classes finies de prédicteurs $\mathcal{F}_{\mathsf{W}}$, un estimateur agrégé convexe empirique sur cette classe atteint des bornes de risque plus fines que celles que le minimiseur de risque empirique peut effectivement atteindre. Cette remarque suscite un intérêt pour la théorie PAC-Bayesienne et son application aux réseaux de neurones.

La théorie PAC-Bayesienne D'un point de vue informel, cette théorie vise à tirer le meilleur de la théorie Bayésienne généralisée et des bornes PAC. La théorie Bayésienne généralisée pour

l'apprentissage repose sur la théorie Bayésienne classique : supposons que l'on ait une opinion a priori sur le prédicteur $f_{\boldsymbol{w}}$ encodée au travers d'un prior π sur W, alors la distribution a posteriori est donnée par $\widehat{\pi}_n(\boldsymbol{w}|\mathbf{Z}^n) \propto \mathsf{Likelihood}(\mathbf{Z}^n|\boldsymbol{w})\pi(\boldsymbol{w})$. L'apprentissage Bayésien généralisé revient à remplacer la vraisemblance par une fonction de perte $\mathscr{L}_{\boldsymbol{w},n}$, qui mesure la performance d'une fonction $f_{\boldsymbol{w}}$ étant donné $\mathbf{Z^n}$, telle que $\widehat{\pi}_n(\boldsymbol{w}|\mathbf{Z^n}) \propto \mathscr{L}_{\boldsymbol{w},n}(\mathbf{Z^n})\pi(\boldsymbol{w})$. Avec la notion de posterior sur l'espace W, suivant l'idée d'estimateur agrégé introduite dans le paragraphe précédent, on peut alors par exemple définir l'estimateur moyenne agrégé par:

$$\widehat{f}_n = \int_{\mathsf{W}} f_{\boldsymbol{w}} \widehat{\pi}_n(d\boldsymbol{w}). \tag{2.3}$$

Le deuxième ingrédient de la théorie PAC-Bayésienne, les bornes PAC, ont été initiées par Valiant (1984) et prennent généralement la forme

$$\mathcal{P}\left(\mathbf{E}_{\mathcal{P}}[\ell(\widehat{f}_n, f_{\mathcal{P}})] - \mathbf{E}_{\mathcal{P}}[\ell(f_{\boldsymbol{w}^*}, f_{\mathcal{P}})] \leq \delta_n(\varepsilon)\right) \geq 1 - \varepsilon,\ \varepsilon \in (0,1)$$

où δ_n est une fonction strictement positive bien choisie.

Ces deux concepts ont été exploités conjointement dans les articles pionniers de McAllester (1999a, 2003) sous la forme d'une borne empirique basée sur les données, puis étendus par la suite au cadre des inégalités oracles grâce aux travaux de Catoni (2007). À la suite de ces résultats, les travaux de Leung and Barron (2006) et Juditsky et al. (2008) ont initié des bornes PAC-Bayesiennes en espérance, ensuite développées par Lecué (2007); Alquier and Lounici (2011); Dalalyan and Tsybakov (2012a); Rigollet and Tsybakov (2012). Dans ce cadre et sous certaines hypothèses, ces travaux proposent un estimateur agrégé $\widehat{f}_n$ vérifiant la borne de risque suivante :

$$\mathbf{E}_{\mathcal{P}}[\ell(\widehat{f}_n, f_{\mathcal{P}})] \leq C \inf_{p \in \mathcal{P}_W} \left\{ \int_{\mathsf{W}} \ell(f_{\boldsymbol{w}}, f_{\mathcal{P}})\, p(d\boldsymbol{w}) + \frac{\beta}{n} D_{\mathsf{KL}}(p||\pi) \right\}, \tag{2.4}$$

où $\mathcal{P}_{\mathsf{W}}$ est l'espace des mesures de probabilité sur W, C est une constante universelle, π est un prior donné dans W, $\beta > 0$ un paramètre de température et $\widehat{f}_n$ s'écrit généralement comme (2.3).

Compte tenu de sa forme très générique, l'inégalité (2.4) s'applique dans de nombreux cadres. Notamment pour l'apprentissage en ligne, grâce à la procédure *Mirror Averaging* (MA) (voir Yuditskii et al. (2005); Juditsky et al. (2008) pour ce sujet et Dalalyan and Tsybakov (2012a) pour le cas des priors *sparse*), ou dans des *designs* déterministes – utiles en traitement d'image ou de signal, avec l'utilisation d'agrégation à poids exponentiels (Dalalyan and Tsybakov, 2007, 2012b). Nous renvoyons le lecteur la très complète revue sur l'apprentissage PAC-Bayésien de Guedj (2019) et aux références qui y figurent pour une introduction plus détaillée. L'inégalité dans ces cadres, en fonction de certaines hypothèses, s'applique à divers types de problèmes statistiques tels que la régression, la classification, ainsi que l'estimation de densité. Nous renvoyons le lecteur au Chapitre 6 Section 6.3 pour une description plus détaillée du type de problèmes d'estimation rencontrés.

En adaptant l'inégalité (2.4) à la classe spécifique $\mathcal{F}_{\mathsf{W}}$ des réseaux de neurones peu profonds, les travaux du Chapitre 6 nous permettent de traiter les Problèmes 5–6. Comme dernière partie de cette introduction, la section suivante présente brièvement les résultats des chapitres suivants, pour donner une idée générale de notre contribution.

2.3 Principales contributions de la thèse

Dans cette section, nous donnons un bref aperçu des principaux développements de la thèse. Les deux premiers chapitres se concentrent sur les stratégies de *trading* optimales pour des producteurs renouvelables ainsi que sur une étude sur la formation du prix sur le marché *intraday* et ses caractéristiques, le Chapitre 5 traite des prévisions probabilistes et de leur application à la gestion de l'énergie éolienne, et enfin, le Chapitre 6 fournit des bornes de risque pour les réseaux de neurones agrégés peu profonds ainsi que des indications sur leur paramétrage.

2.3.1 Stratégies optimales pour des producteurs renouvelables

Sur le marché intraday de l'électricité, le prix intègre les mises à jour des prévisions de demande/production auxquelles sont confrontés les agents. Cette caractéristique spécifique induit un impact sur le marché et de possibles comportements stratégiques. Dans le Chapitre 3 et le Chapitre 4, nous nous proposons d'expliquer ce phénomène et ses conséquences sur les prix à travers un modèle simplifié des marchés court terme successifs présentés dans la section 2.1.2. Nous abordons le Problème 1 en développant un modèle d'équilibre traçable dans un cadre linéaire quadratique, capable de fournir un prix explicite et des stratégies de *trading* optimales explicites pour les producteurs confrontés à une production/demande terminale intermittente. Nous étudions ce jeu dans différentes structures d'information: dans le Chapitre 3, nous considérons les cadres de l'information complète et incomplète, avec des agents identiques; dans le Chapitre 4, nous introduisons un acteur majeur ayant un impact sur le comportement stratégique de la foule.

Cadre N-joueurs Nous considérons une période de livraison donnée commençant à la date T. Nous supposons qu'il existe N agents identiques qui partagent la même information sur le prix du marché et les prévisions de production, de sorte que nous pouvons définir un espace de probabilité filtré $(\Omega, \mathcal{F}, \mathbb{F} := (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{P})$ auquel tous les processus sont adaptés. Le prix du marché day ahead, S_0, est donné à la date $t = 0 < T$ et nous désignons par ϕ_0^i la position du i-ième agent, $i = 1, \dots, N$, sur ce marché. Suite à cela, le marché *intraday* permet de négocier en continu jusqu'à la fermeture du marché juste avant T, et nous désignons par ϕ_t^i la position du i-ième joueur au temps t, ainsi que $\overline{\phi}_t^N = \frac{1}{N}\sum_{i=1}^N \phi_t^i$ la position agrégée de tous les agents. Le prix du marché infrajournalier P_t^N est composé d'un prix fondamental S_t et d'un impact linéaire de marché (Almgren and Chriss, 1999, 2001) par lequel les acteurs interagissent entre eux:

$$P_t^N = S_t + a(\bar{\phi}_t^N - \bar{\phi}_0^N), \quad \forall t \in [0,T],\ a > 0. \tag{2.5}$$

Quand il trade, le joueur i subit un coût donné par

$$\dot{\phi}_t^i P_t^N + \frac{\alpha(t)}{2}\dot{\phi}_t^i\Big(\dot{\phi}_t^i + b\dot{\bar{\phi}}_t^{N,-i}\Big), \quad \forall t \in [0,T]$$

où $\dot{\bar{\phi}}_t^{N,-i} = \frac{1}{N-1}\sum_{j=1,j\neq i}^N \dot{\phi}_t^j$. Le premier terme représente le coût réel d'achat de l'électricité, et le second le coût de *trading*, où $\alpha(.)$ est une fonction continue strictement positive sur $[0,T]$, reflétant la variation de la liquidité du marché à l'approche de la date de livraison. Le terme $b\dot{\bar{\phi}}_t^{N,-i}$ avec $b > 0$ représente l'impact du sens de *trading* de la foule sur le coût de *trading* du i-ième agent, ce qui rend compte d'une éventuelle synchronisation des acteurs. Pour déterminer sa position optimale sur le marché, chaque agent dispose d'une prévision incertaine de sa production terminale X_T^i, donnée par le processus $(X_t^i)_{t\in[0,T]}$. A la clôture du marché intraday, si $\phi_T^i \neq X_T^i$, l'agent doit d'abord acheter la quantité manquante ou vendre la quantité supplémentaire d'électricité au prix S_T et, en plus, s'aquitter d'une pénalité de déséquilibre $\frac{\lambda}{2}(\phi_T^i - X_T^i)^2$. Enfin, le i-ème joueur vise à maximiser une fonction objectif linéaire quadratique, qui dépend de la stratégie des autres joueurs par le biais du prix du marché:

$$\begin{aligned} J^{N,i}(\phi^i,\phi^{-i}) := -\mathbb{E}\Big[& \underbrace{\phi_0^i S_0}_{\text{Day ahead}} + \underbrace{\int_0^T \Big\{\dot{\phi}_t^i P_t^N + \frac{\alpha(t)}{2}\dot{\phi}_t^i\Big(\dot{\phi}_t^i + b\dot{\bar{\phi}}_t^{N,-i}\Big)\Big\}\,dt}_{\text{Intraday}} \\ & \underbrace{-(\phi_T^i - X_T^i)S_T + \frac{\lambda}{2}(\phi_T^i - X_T^i)^2}_{\text{Règlement des écarts}}\Big], \end{aligned} \tag{2.6}$$

où $\phi^{-i} := (\phi^1,\dots,\phi^{i-1},\phi^{i+1},\dots,\phi^N)$. Dans le cadre de la Définition 7, nous dérivons dans le Théoreme 14 Chapitre 3 un équilibre de Nash unique et explicite sans hypothèses spécifiques sur la dynamique du prix S.

A partir de l'expression des stratégies et du prix d'équilibre, nous pouvons déduire de précieux éléments sur la façon dont le prix fondamental et les prévisions de production influencent les décisions des producteurs renouvelables. Pour citer quelques exemples, lorsque S est supposé être une martingale, la stratégie ne dépend pas du prix: les prises de positions sur le marché *intraday* ne sont déclenchées que par des ajustements dûs à des mises à jour de prévisions et la position sur le marché *day ahead* est donnée par $\phi_0^i = X_0^i$. Au contraire, si nous supposons que le prix a une dérive déterministe positive, les joueurs choisiront $\phi_0^i < X_0^i$, mais il y aura une tendance positive dans la stratégie agrégée sur le marché *intraday* de sorte que l'impact du marché amplifiera la tendance du prix. En outre, le coût supplémentaire lié au comportement de foule, b, conduit les agents à trader moins activement en réponse à des mises à jour de prévisions communes ayant un impact sur tous les X^i, qu'en réponse à des mises à jour individuelles. Concernant la pénalité terminale, elle est modélisée par la contrainte souple λ car elle est censée représenter la pénalité d'écart réelle qui est basée sur la différence entre X_T^i et ϕ_T^i. Néanmoins, lorsque $\lambda \to \infty$, la condition de contrainte dure est satisfaite, à savoir $X_T^i = \phi_T^i$, $i = 1,\dots,N$. Par conséquent, en fonction du règlement des écarts (simplifié en une pénalité déterministe), le modèle permet aux producteurs de décider dans quelle mesure ils doivent respecter leur engagement *day ahead*

/ *intraday*.

Comme nous l'avons déjà évoqué, supposer que tous les acteurs du marché ont une connaissance parfaite de la prévision aggrégée peut sembler irréaliste. Pour y remédier, nous présentons brièvement le cadre de l'information incomplète dans la suite et renvoyons le lecteur au Chapitre 3 Section 3.3.2 pour plus de détails.

Extension champ moyen Dans la variante à champ moyen du jeu, nous considérons le cas limite avec une infinité d'agents et envisageons un jeu où l'information est incomplète, modélisée par la filtration $\mathbb{F}$, à laquelle tous les processus sont adaptés, et $\mathbb{F}^0$ qui ne contient que le bruit commun – *i.e* restreint à l'information présente sur le marché, déjà introduite dans la section 2.2.1. La position de l'agent représentatif est donnée par ϕ_t et sa prévision de production est notée X_t. Ces dernières ne sont que partiellement observées: les autres participants du marché n'ont accès qu'à $\bar{\phi}_t = \mathbb{E}[\phi_t|\mathcal{F}_t^0]$ et $\overline{X}_t = \mathbb{E}[X_t|\mathcal{F}_t^0]$, et le prix du marché P_t, est donné par $P_t = S_t + a(\bar{\phi}_t - \bar{\phi}_0)$. Dans ce cas, nous obtenons également un équilibre explicite du jeu à champ moyen au sens de la Définition 9, avec un prix d'équilibre et des stratégies en forme fermée.

Autre fait remarquable, nous sommes en mesure de rapprocher les deux cadres (Nash et champ moyen). En pratique, les marchés avec une infinité d'agents ne sont pas vraiment réalistes et si nous sommes confrontés à un marché composé de N joueurs, il est utile de savoir quantifier dans quelle mesure l'équilibre du jeu à champ moyen est un proxy du "vrai" équilibre de Nash à N joueurs. Dans le cadre de la Définition 8, nous prouvons que l'équilibre du jeu à champ moyen peut fournir des stratégies pour le jeu à N joueurs qui forment un équilibre ε-Nash du jeu avec $\varepsilon = \frac{C}{N^{\frac{1}{2}}}$, où $C > 0$ est une constante. Nous prouvons également l'approximation inverse: lorsque le nombre de joueurs augmente, l'équilibre de Nash converge vers l'équilibre limite du jeu à champ moyen. La Figure 2.5 illustre ce phénomène lorsque le nombre de joueurs passe de $N = 5$ à $N = 100$.

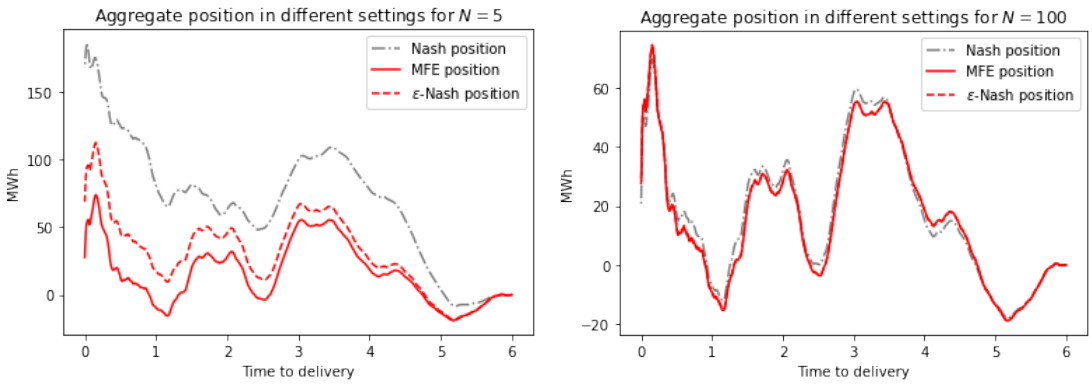

Figure 2.5: Positions agrégées dans différents cadres avec $N = 5$ (gauche) and $N = 100$ (droite) agents

Jeu à la Stackelberg Dans le Chapitre 4, nous développons un modèle dans le cadre du jeu à champ moyen en présence d'un agent majeur avec asymétrie d'information. Une infinité

de petits producteurs identiques peut observer les prévisions de production de l'agent majeur en raison de sa taille; mais l'inverse n'est pas vrai: le producteur majeur ne peut observer qu'une prévision de production commune tronquée. Si les petits producteurs ont un avantage informationnel, les actions du gros producteur peuvent avoir un impact significatif sur le marché. Le petit producteur représentatif devra alors définir stratégiquement sa position par rapport à l'acteur majeur et vice versa. Dans ce cadre, nous sommes également en mesure d'obtenir des positions optimales explicites et un équilibre de prix qui permet de mettre en avant l'impact d'un gros producteur sur le comportement de plus petits. En outre, nous pouvons dériver un équilibre ε-Nash pour adapter les stratégies à un cas où le nombre de joueurs serait fini.

D'autres informations sur le comportement stratégique des producteurs renouvelables sont présentées au Chapitre 3 et au Chapitre 4. Il est aussi important de mentionner que ces modèles conduisent à des caractéristiques du prix du marché qui sont en accord avec les résultats empiriques, comme expliqué dans la section suivante.

2.3.2 Formation des prix et principales caractéristiques du marché infrajournalier

Volatilité et effet Samuelson Dans la Figure 2.2b, Section 2.1.3, nous présentons la volatilité empirique instantanée du prix estimée par la méthode des noyaux (voir Kristensen (2010) pour la méthodologie générale et Chapitre 3 Section 3.5.1 pour le paramétrage de la Figure 2.2b). Nous observons que la volatilité augmente à mesure que l'on se rapproche de la date de livraison. Il s'agit d'une caractéristique bien documentée du marché des *futures* sur l'électricité et d'autres marchés à terme (Jaeck and Lautier, 2016) connue sous le nom d'effet Samuelson. Un défi important des Chapitres 3 et 4 est de reproduire et d'expliquer ce fait stylisé. Pour étudier ce phénomène, nous menons une étude analytique sur l'expression du prix d'équilibre afin d'expliciter la volatilité supplémentaire qu'implique l'interaction des agents grâce à l'expression de la variance normalisée du prix entre de petits pas de temps h, telle que, lorsque $h \to 0$:

$$\frac{1}{h}\mathbb{E}[(P^N_{t+h} - P^N_t)^2] = \sigma_S^2 + h\mathcal{V}_t + O(h^2), \ ; \text{with} \ ; \mathcal{V}_t = \frac{a^2\lambda^2\sigma_X^2}{\alpha\left(1+\frac{b}{2}\right)} f(t) \tag{2.7}$$

où σ_S^2, $\sigma_X^2 > 0$ sont les variances constantes du prix fondamental et des prévisions de production des agents, $\alpha > 0$ est un coût de liquidité constant, et f est une fonction croissante du temps dépendant des paramètres du modèle. Notons que lorsque $h \to 0$, l'expression (2.7) converge vers la variance du prix fondamental σ_S^2. Cependant, un agent utilisant un estimateur de volatilité du prix avec un pas de temps h, trouvera une variance supplémentaire d'environ $h\mathcal{V}_t$ (en moyenne). De plus, nous retrouvons l'effet Samuelson puisque f est croissante en t. Une estimation empirique de la volatilité instantanée sur des simulations du prix d'équilibre du modèle, en utilisant exactement la même méthode de noyaux que dans la Figure 2.2b, confirme ce phénomène comme indiqué dans la Figure 2.6a. De plus, l'expression explicite (2.7) apporte des informations sur la dépendance de la volatilité aux différents paramètres du modèle tels que la pénalité d'écart, les coûts de liquidité et le nombre de joueurs. Pour plus de détails sur la

façon dont la volatilité est influencée par ces paramètres, nous renvoyons le lecteur au Chapitre 3 Section 3.4.2.

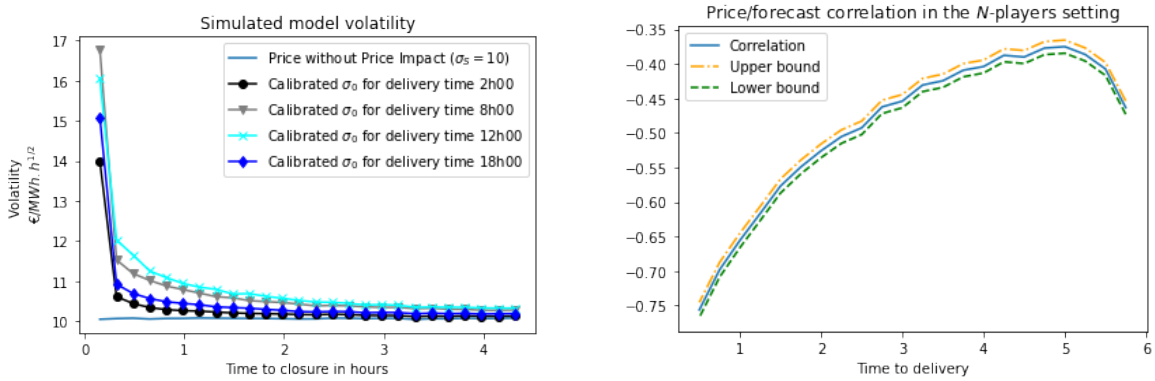

(a) Volatilité estimée dans le modèle

(b) Correlation estimée dans le modèle

Figure 2.6: Estimation des caractéristiques du prix dans le modèle

Corrélation entre les prévisions de production renouvelable et le prix du marché Notre modèle est également capable d'expliquer la corrélation observée empiriquement entre les prix et les prévisions de production: le lien entre ces deux quantités est également traité analytiquement dans la Section 3.4.3 Chapitre 3 et nous fournissons, Figure 2.6b, une estimation de la corrélation dans le modèle à partir du prix d'équilibre simulé et des prévisions de production renouvelable. La corrélation est négative et augmente en valeur absolue à mesure que l'on se rapproche de la date de livraison, à l'instar des résultats empiriques.

Pour conclure, les Chapitres 3 et 4 fournissent, à notre connaissance, le premier modèle d'équilibre capable d'expliquer les caractéristiques stylisées du marché *intraday* de l'électricité observées empiriquement. Ainsi, il peut offrir une meilleure compréhension des interactions sur ce marché et aider à la gestion des énergies renouvelables. L'article récent de Aid et al. (2020), est proche, dans l'idée, de notre travail. Ils ont développé un modèle à information complète où les N joueurs ont des caractéristiques hétérogènes et ils obtiennent un équilibre sur le marché *intraday* de l'électricité. Leur modèle reproduit également l'effet de Samuelson décrit par notre modèle.

Dans les deux premiers chapitres de la thèse, les prévisions de production sont supposées être des prévisions ponctuelles. Cependant, cette notion n'est pas suffisante pour répondre au Problème 3 et au Problème 4. Dans la section suivante, nous abordons ces derniers dans le cadre des prévisions probabilistes présentées dans la Section 2.2.2.

2.3.3 Prévisions probabilistes dans des problèmes de décision séquentiels

Dans le Chapitre 5, nous introduisons un modèle de prévisions probabilistes dans le cadre diffusif pour capturer les dépendances temporelles entre différents horizons de prévision et fournir un flux explicite de densités prédictives qu'il est possible de calibrer et d'exploiter ensuite dans

des processus de prise de décision stochastiques. À notre connaissance, ce dernier point n'a encore jamais été étudié. Nous développons quatre modèles d'EDS capables de prendre en compte des quantités à valeurs réelles comme la température, ou des quantités positives comme la vitesse du vent. Pour deux de ces modèles, nous décrivons une procédure de calibration complète en utilisant des prévisions d'ensemble de la région parisienne, dont un exemple est présenté dans la Figure 2.4a, et nous évaluons leur précision. Nous nous concentrons ensuite sur un problème simple de gestion de la production d'énergie éolienne avec un seul agent où les prévisions probabilistes calibrées de la vitesse du vent sont utilisées. Dans la suite, nous présentons brièvement cet exemple.

Modélisation de la densité Nous considérons un espace de probabilité filtré $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$, une variable aléatoire m_T $\mathcal{F}_T$-mesurable représentant la vitesse du vent, et nous nous référons à la Définition 10 pour le flux de densités prédictives. Pour la densité prédictive au temps $t \in [0, T]$, nous choisissons une paramétrisation bidimensionnelle (m_t, V_t), où m_t est l'espérance conditionnelle de m_T et V_t est une mesure de l'incertitude de la prévision. Plus précisément, nous supposons que m_t et V_t sont définis par le modèle de type Heston suivant:

$$\begin{aligned} \frac{dm_t}{m_t} &= \sqrt{V_t}\rho(T-t)dW_t \\ dV_t &= -V_t\rho^2(T-t)\left(1+\frac{b^2}{2}\right)dt + \sqrt{V_t}b\rho(T-t)dW_t'. \end{aligned} \tag{2.8}$$

où W, W' sont des mouvements browniens indépendants, ρ est une fonction déterministe positive et $b > 0$ est une constante. À partir des résultats obtenus dans Barndorff-Nielsen (1997), nous déduisons que pour m, V solution de (2.8), la distribution conditionnelle de $\log m_T$ étant donné $\mathcal{F}_t$ est la distribution normale inverse Gaussienne sur $\mathbb{R}$ qui a pour densité

$$p_t(x) = \frac{\alpha\delta_t K_1\left(\alpha\sqrt{\delta_t^2 + (x-\mu_t)^2}\right)}{\pi\sqrt{\delta_t^2 + (x-\mu_t)^2}} e^{\delta_t\gamma + \beta(x-\mu_t)}, \ \forall t \in [0, T] \tag{2.9}$$

où les paramètres sont donnés par $\mu_t = \log m_t$, $\delta_t = \frac{V_t}{b}$, $\beta = -\frac{1}{2}$, $\alpha = \sqrt{\left(b^{-1} + \frac{b}{2}\right)^2 + \frac{1}{4}}$, $\gamma = b^{-1} + \frac{b}{2}$, et K est la fonction de Bessel modifiée de la troisième espèce. En suivant les idées développées dans la Section 2.2.2, nous calibrons la densité (2.9) à chaque horizon de prévision pour ensuite utiliser le modèle de diffusion dans un problème réaliste de gestion d'énergie éolienne.

Calibration Les données utilisées pour calibrer le modèle sont composées de plusieurs prévisions d'ensemble qui sont mises à jour toutes les 12 heures et dont l'horizon temporel total est de 48 heures. Nous calibrons donc la densité prédictive pour les horizons 48h, 36h, 24h et 12h avant la réalisation. Comme expliqué dans la section 2.2.2, les prévisions d'ensemble souffrent de sous-dispersion et de biais. Pour corriger ces défauts, nous utilisons une méthodologie inspirée d'*EMOS* pour les différents horizons. Nous renvoyons le lecteur au Chapitre 5 Section 5.3 pour

une description détaillée de la méthode. Une fois cette étape réalisée, la qualité de la calibration est mesurée en la comparant aux performances des données d'ensemble brutes. Nous avons réalisé plusieurs études sur la qualité de la calibration qui démontrent une nette amélioration de la capacité de prédiction par rapport aux ensembles bruts. La Figure 2.7 montre la transformation intégrale de la probabilité (TIP) des prévisions calibrées, qui consiste en un histogramme de la fonction de répartition prédictive calibrée pour l'horizon de prévision 12h, évaluée au point correspondant à la vitesse du vent réalisée. Si la distribution prédictive est bien calibrée, alors l'histogramme devrait être proche de celui de la distribution uniforme. En comparant la Figure 2.7 avec le diagramme de Talagrand Figure 2.4b Section 2.2.2 nous observons que la sous-dispersion et le biais ont été corrigés dans une large mesure.

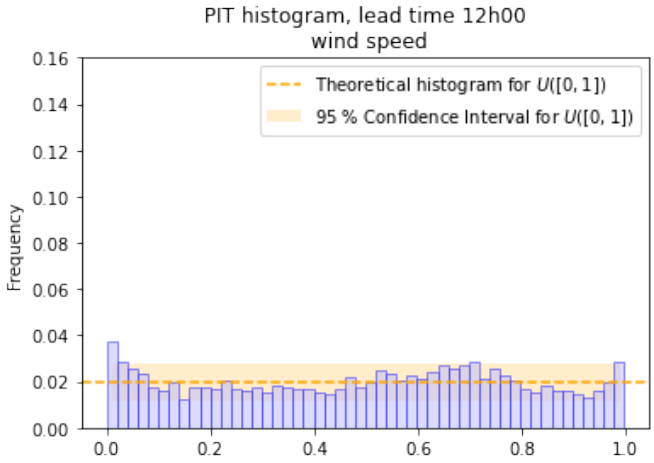

Figure 2.7: Histogramme TIP

Utilisation des prévisions probabilistes pour la prise de décision séquentielle L'idée centrale et l'avantage du modèle que nous avons développé est qu'il permet d'exploiter l'information sur l'évolution de la densité prédictive dans un problème de décision stochastique. Pour évaluer cet avantage, nous avons réalisé une étude visant à déterminer le profit supplémentaire d'un producteur éolien utilisant notre modèle, par rapport à un modèle standard sans description dynamique explicite de l'incertitude de la prévision. Pour ce faire, nous considérons le modèle A où le producteur a accès à des prévisions probabilistes et le modèle B, qui est basé uniquement sur des prévisions ponctuelles de la vitesse du vent. Pour une heure de livraison donnée T, nous désignons la vitesse moyenne du vent pendant un petit intervalle de temps autour de cette date par m_T. Leurs prévisions de vent pour cette heure de livraison sont régulièrement mises à jour, notées m_t à l'instant $t < T$, et ils déduisent les mises à jour de leurs prévisions de production des mises à jour des prévisions de la vitesse du vent, grâce à une fonction de production stylisée f, telle que la prévision de production est donnée par $f(m_t)$. Ils peuvent ajuster leur position sur le marché *intraday* chaque fois qu'ils reçoivent de nouvelles mises à jour et doivent payer une pénalité d'écart terminale. Nous désignons par $t_0 < \cdots < t_{N-1}$ les moments discrets auxquels de nouvelles prises de position ont lieu, $t_N = T$ étant le moment où la livraison commence. La

fraction échangée à la date t aura le prix S_t, et nous désignons la quantité totale de puissance vendue ou achetée jusqu'à la date t par ϕ_t. Toute puissance non vendue sur le marché infra-journalier avant la date T sera vendue à la même date au prix des écarts désigné par S_T. De plus, les transactions d'équilibrage sont soumises à une pénalité d'écart égale à une constante K multipliée par le volume de la transaction. Les profits des producteurs sont donc donnés par

$$\begin{aligned}
&\underbrace{S_{t_0}\phi_{t_0} + \sum_{i=1}^{N-1} S_{t_i}(\phi_{t_i} - \phi_{t_{i-1}})}_{\text{marché infrajournalier}} + \underbrace{S_T(f(m_T) - \phi_{t_{N-1}}) - K|f(m_T) - \phi_{t_{N-1}}|}_{\text{règlement des écarts}} \\
&= f(m_T)S_T - \sum_{i=0}^{N-1} \phi_{t_i}\Delta S_{t_i} - K|f(m_T) - \phi_{t_{N-1}}|,
\end{aligned}$$

où $\Delta S_{t_i} = S_{t_{i+1}} - S_{t_i}$. Nous supposons que les producteurs souhaitent maximiser l'utilité de leur profit à la date T, c'est-à-dire qu'ils résolvent le problème de contrôle suivant :

$$\max_{\phi:=(\phi_{t_i})_{i=1}^{N-1}} \mathbb{E}\left[u\left(f(m_T)S_T - \sum_{i=0}^{N-1} \phi_{t_i}\Delta S_{t_i} - K|f(m_T) - \phi_{t_{N-1}}|\right)\right], \tag{2.10}$$

où u est une fonction d'utilité. Nous évaluons ensuite numériquement leur gain sur des trajectoires simulées (du modèle A), pour différentes paramétrisations du modèle, en fonction du niveau d'incertitude initial à l'ouverture du marché, et de la dérive μ_S du prix.

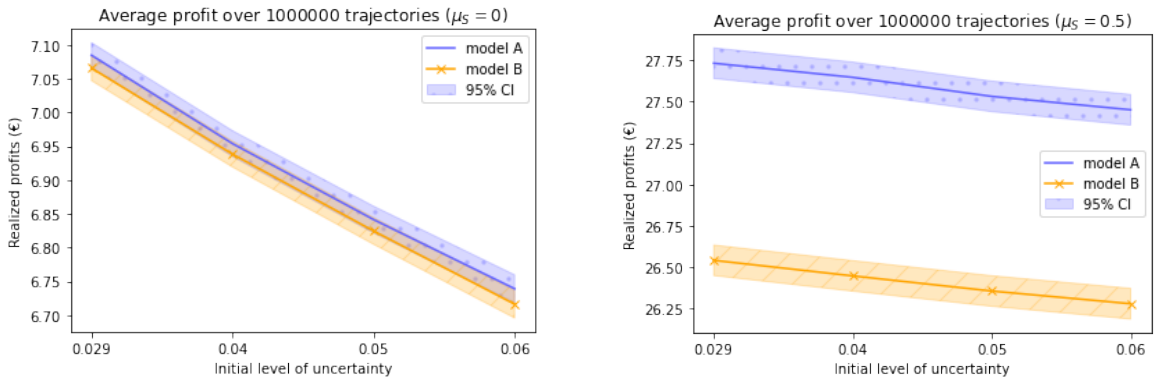

(a) $\mu_S = 0$ (b) $\mu_S = 0.5$

Figure 2.8: Profits en fonction du niveau initial d'incertitude de la prévision de la vitesse du vent

Les résultats des Figures 2.8a–2.8b montrent que les gains diminuent lorsque l'incertitude initiale augmente. Dans le cas où le prix est une martingale, une différence de gain entre les deux modèles ne peut être clairement établie puisque les intervalles de confiance des gains se chevauchent partiellement; cependant, dans le cas où le prix a une dérive, les gains sont en moyenne significativement plus élevés pour un agent utilisant le modèle A. De cette étude, nous pouvons conclure que cette nouvelle méthode montre des résultats prometteurs et devrait être étudiée plus en profondeur pour des utilisations pratiques.

Dans la section suivante, nous présentons les principaux résultats obtenus lors du traitement des Problèmes 5–6.

2.3.4 Apprentissage par réseaux de neurones agrégés à une couche cachée

Le Chapitre 6 fournit des bornes PAC-Bayésiennes dans le cadre des réseaux de neurones peu profonds. Soit $f_{\mathcal{P}}$ la fonction que nous voulons estimer et considérons le cadre introduit dans la Section 2.2.3 où $\mathcal{X} = [0,1]^{D_0}$, et la classe $\mathcal{F}_{\mathsf{W}}$ est limitée aux réseaux de neurones à une couche cachée, tel que l'ensemble W des poids d'un réseau de neurones contient les poids $\boldsymbol{w}_1$ de la couche cachée, et les poids $\boldsymbol{w}_2$ de la couche de sortie. Alors $\boldsymbol{w} = (\boldsymbol{w}_1^\top, \boldsymbol{w}_2^\top)^\top$ appartient à $\mathbb{R}^d$ avec $d = D_1(D_0 + D_2)$, $\boldsymbol{w}_1 \in \mathbb{R}^{D_0 D_1}$, $\boldsymbol{w}_2 \in \mathbb{R}^{D_1 D_2}$. Le réseau de neurones paramétré par $\boldsymbol{w}$ s'écrit donc :

$$f_{\boldsymbol{w}}(\boldsymbol{x}) = \boldsymbol{w}_2^\top \bar{\sigma}(\boldsymbol{w}_1^\top \boldsymbol{x}) \in \mathbb{R}^{D_2},\ \forall \boldsymbol{x} \in \mathcal{X} \quad \text{with} \quad \bar{\sigma} : \boldsymbol{x} \in \mathbb{R}^{D_1} \mapsto \begin{bmatrix} \sigma(x_1) \\ \dots \\ \sigma(x_{D_1}) \end{bmatrix} \in \mathbb{R}^{D_1}, \tag{2.11}$$

où $\sigma : \mathbb{R} \mapsto \mathbb{R}$ est une fonction d'activation.

Nous spécifions la fonction de perte avec la norme standard ℓ_2 pour l'inégalité PAC-Bayésienne (2.4), telle qu'elle se réécrit:

$$\mathbf{E}_{\mathcal{P}}[\|\widehat{f}_n - f_{\mathcal{P}}\|^2_{\mathbb{L}_2(\mu)}] \leq C \inf_{p \in \mathcal{P}_{\mathsf{W}}} \left\{ \int_{\mathsf{W}} \|f_{\boldsymbol{w}} - f_{\mathcal{P}}\|^2_{\mathbb{L}_2(\mu)}\, p(d\boldsymbol{w}) + \frac{\beta}{n} D_{\mathsf{KL}}(p||\pi) \right\}. \tag{2.12}$$

En fixant le paramètre $\bar{\boldsymbol{w}} \in \mathsf{W}$ et en utilisant l'inégalité triangulaire, cela donne:

$$\left(C^{-1}\, \mathbf{E}_{\mathcal{P}}\left[\|\widehat{f}_n - f_{\mathcal{P}}\|^2_{\mathbb{L}_2(\mu)}\right] \right)^{1/2} \leq \|f_{\bar{\boldsymbol{w}}} - f_{\mathcal{P}}\|_{\mathbb{L}_2(\mu)} + \mathrm{Rem}_n(\bar{\boldsymbol{w}})^{1/2}, \tag{2.13}$$

avec le terme résiduel borné par

$$\mathrm{Rem}_n(\bar{\boldsymbol{w}}) \leq \int_{\mathsf{W}} \|f_{\boldsymbol{w}} - f_{\bar{\boldsymbol{w}}}\|^2_{\mathbb{L}_2(\mu)}\, p(d\boldsymbol{w}) + \frac{\beta}{n} D_{\mathsf{KL}}(p||\pi), \tag{2.14}$$

où p est une distribution quelconque dans $\mathcal{P}_{\mathsf{W}}$ et π est le prior. Notons que si nous choisissons $\bar{\boldsymbol{w}} \in \mathrm{argmin}_{\boldsymbol{w} \in \mathsf{W}} \|f_{\boldsymbol{w}} - f_{\mathcal{P}}\|_{\mathbb{L}_2(\mu)}$, le côté droit de l'inégalité (2.13) peut être décomposé en l'erreur d'approximation $\inf_{\boldsymbol{w} \in \mathsf{W}} \|f_{\boldsymbol{w}} - f_{\mathcal{P}}\|_{\mathbb{L}_2(\mu)}$ et le terme résiduel $\mathrm{Rem}_n(\bar{\boldsymbol{w}})$ évaluant l'erreur d'estimation. Une autre remarque importante est que la plupart du temps, en pratique, les réseaux de neurones sont initialisés avec une distribution Gaussienne. Suivant cette idée, nous définissons le prior et l'ensemble $\mathcal{P}_{\mathsf{W}}$ comme des distributions Gaussiennes sphériques.

Une fois ce cadre général établi, notre travail dans le Chapitre 6 consiste en trois étapes principales. Tout d'abord, il faut trouver une "bonne" mesure p pour le terme résiduel, de sorte qu'elle présente de bonnes propriétés en termes de finesse de la borne et permette de dériver une expression qui dépende des paramètres du modèle. Les deux dernières étapes consistent à exploiter les réglages possibles pour les paramètres du réseau de neurones. La deuxième étape consiste à trouver une variance appropriée pour le prior de sorte que, pour les fonctions

d'activation Lipschitz (y compris sigmoïde et ReLU), nous dérivons une borne oracle à partir du terme résiduel (2.14), qui met en évidence les rôles de D_1 et n. La troisième étape consiste à choisir un nombre approprié D_1 d'unités dans la couche cachée. A partir de cette étape, nous restreignons la classe des fonctions $\mathcal{F}$ aux fonctions à valeur réelle ($D_2 = 1$), utilisons les résultats de Maiorov and Meir (2000) lorsque $\mathcal{F}$ est la boule unité dans l'espace de Sobolev $W_2^r(\mathcal{X})$, et que les réseaux de neurones sont équipés de fonctions d'activation sigmoïdes; et utilisons le résultat de Siegel and Xu (2020), pour l'espace spectral de Barron et les sous-espaces de $W_2^r(\mathcal{X})$ –pour un choix spécial de r, avec des réseaux de neurones ReLU. Ceci permet de dériver des inégalités oracles explicites et, pour citer un exemple, dans le cas de fonctions d'activation sigmoïdes, nous obtenons,

$$C^{-1}\mathbf{E}_{\mathcal{P}}\left[\|\widehat{f}_n - f_{\mathcal{P}}\|^2_{\mathbb{L}_2(\mu)}\right] \leq g(D_1)D_1^{-2r/D_0} + \bar{g}(n/d)\frac{D_1 D_0}{n}, \tag{2.15}$$

où g, $\bar{g}$ sont des fonctions à croissance au plus logarithmique. Le nombre de neurones D_1 doit être choisi avec soin afin de contrôler l'équilibre entre l'erreur d'approximation et l'erreur d'estimation, pour garantir le meilleur compromis biais-variance en fonction de la taille de l'échantillon, comme cela est illustré graphiquement pour l'équation (2.15) dans la Figure 2.9.

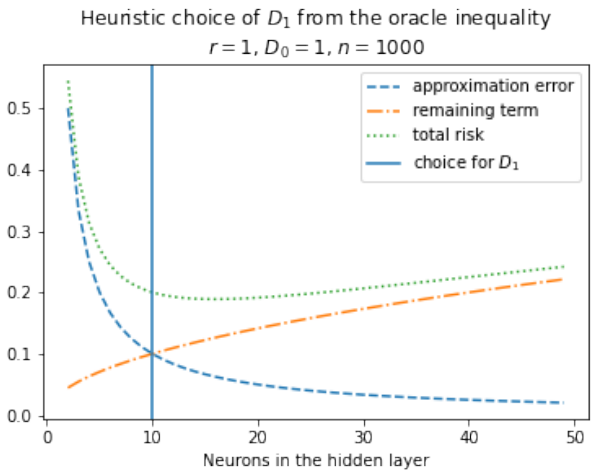

Figure 2.9: Equilibre entre l'erreur d'approximation et l'erreur d'estimation

Ceci présente deux avantages principaux : (i) obtenir une évaluation non-asymptotique précise du risque ; (ii) fournir des indications, basées sur une analyse théorique, sur l'initialisation des poids du réseau de neurones et son architecture.

Notre principale contribution consiste à fournir des bornes de risque pour les fonctions d'activation sigmoïde et ReLu, qui, à notre connaissance, améliorent les résultats existants pour les réseaux de neurones à une couche cachée minimisant l'erreur d'apprentissage/le risque empirique. Dans le cas des fonctions d'activation sigmoïde, nous atteignons la vitesse optimale minimax $O(n^{-2r/(2r+D_0)})$ à des facteurs $\log(n)$ près sur la boule unité dans l'espace de Sobolev $W_2^r(\mathcal{X})$. Ce résultat fait concurrence aux les réseaux de neurones profonds, et montre, de manière informelle, que l'agrégation agit comme une couche supplémentaire. Pour le cas de ReLu, nous

obtenons une première borne pour les fonctions de l'espace spectral de Barron $\mathscr{B}^s(\mathcal{X})$ d'ordre $O(n^{-2K/(2K+1)})$, où K dépend de s et D_0. Cependant, ce résultat n'est pas directement comparable au cas sigmoïde puisqu'il ne s'agit pas d'espaces de Sobolev. Néanmoins, il est possible de combler le fossé entre l'espace de Barron et les fonctions de Sobolev sous certaines restrictions sur r – nous renvoyons le lecteur à la Section 6.6.2 du Chapitre 6 pour plus de détails. En résumé, nous obtenons une borne de risque d'ordre $O(n^{-2\bar{r}/(2\bar{r}+D_0+1)})$, pour tout $\bar{r} < r$, sur la boule unitaire dans $W_2^r(\mathcal{X})$. La borne obtenue est légèrement moins bien que la vitesse minimax prouvée pour les réseaux profonds: ceci est dû à l'absence, à notre connaissance, de propriétés d'approximation plus fines pour les réseaux de neurones ReLU à une couche cachée dans les espaces de Sobolev. Néanmoins, elle améliore les résultats existants pour les réseaux de neurones ReLU à une couche cachée.

2.4 Structure de la thèse

Cette thèse est composée de 4 chapitres originaux basés sur les travaux suivants

. Chapitre 3: *Price formation and optimal trading in intraday electricity markets*, arXiv preprint arXiv:2009.04786. Travail en collaboration Olivier Féron et Peter Tankov.

. Chapitre 4: *Price formation and optimal trading in intraday electricity markets with a major player*, publié dans Risks 2020, 8(4), 133; Special Issue Stochastic Modeling and Pricing in Energy Markets. Travail en collaboration avec Olivier Féron et Peter Tankov.

. Chapitre 5: *Decision making with dynamic probabilistic forecasts*, arXiv preprint arXiv:2106.16047. Travail en collaboration avec Peter Tankov.

. Chapitre 6: *Risk bounds for aggregated shallow neural networks using Gaussian priors.* Travail en collaboration avec Arnak Dalalyan.

Chapter 3

Price formation and optimal trading in intraday electricity markets

Abstract

We develop a tractable equilibrium model for price formation in intraday electricity markets in the presence of intermittent renewable generation. Using stochastic control theory, we identify the optimal strategies of agents with market impact and exhibit the Nash equilibrium in closed form for a finite number of agents as well as in the asymptotic framework of mean field games. Our model reproduces the empirical features of intraday market prices, such as increasing price volatility at the approach of the delivery date and the correlation between price and renewable infeed forecasts, and relates these features with market characteristics like liquidity, number of agents, and imbalance penalty.

3.1 Introduction

The electricity markets around the world are undergoing a major transformation driven by the transition towards a carbon-free energy system. The increasing penetration of intermittent renewables puts a stronger emphasis on short-term electricity trading and balancing. The intraday electricity markets are increasingly used by the renewable producers to compensate forecast errors. This improves market liquidity and at the same time creates feedback effects of the renewable generation on the market price, leading to increased price volatility and negative correlations between renewable infeed and prices. These effects have an adverse impact on the revenues of renewable producers. They are already significant in countries with high renewable penetration and will become even more important as new renewable capacity comes online. A better understanding of the impact of intermittent renewable generation on intraday electricity market prices and trading volumes is therefore needed to ensure the long-term economic sustainability of the renewable energy production.

In this paper, we build an equilibrium model for the intraday electricity market, aiming to understand the price formation and identify the optimal strategies for market participants in the setting where both the strategies of the agents and the demand or generation forecasts may affect market prices. We consider an intraday electricity market, where the participants optimize their revenues based on imperfect forecasts of terminal demand or production. We place ourselves in the standard linear-quadratic setting with quadratic trading costs and linear market impact. The actions of each agent therefore impact market prices, leading to a stochastic game where players interact through the market price. We exhibit a closed-form Nash equilibrium for this game, and provide explicit formulas for the market price and the strategies of the agents under two different settings:

- the setting of N identical agents, having complete information about the forecasts of the other agents,
- the setting of an infinite number of identical small agents (the mean field), where each agent only observes the aggregate forecast as well as its own forecast.

We then show by theoretical analysis and through numerical simulations that our model reproduces the stylized features of the market price, which we document empirically. In particular,

- the market price becomes more volatile at the approach of the delivery time, a phenomenon known as Samuelson's effect in the empirical literature on futures markets;
- the market price exhibits negative correlation with the total renewable infeed forecast, which grows in absolute value at the approach of the delivery time.

Furthermore, our model provides direct quantitative links between market characteristics and market price features, as well as the gain of individual agents. For instance,

- observed price volatility increases for higher imbalance penalties which force the agents to follow the forecasts more closely;

- observed price volatility increases for lower instantaneous trading costs, which allow agents to trade more actively;
- increased competition (greater number of agents in the market) limits profit opportunities for individual agents and leads to lower price volatility.

Correlations between renewable infeed and intraday market prices have been studied empirically by a number of authors. Kiesel and Paraschiv (2017) perform an econometric analysis of the German intraday market and show that a deeper penetration of renewable energies increases market liquidity and price-infeed correlations. The wind power output forecast errors thus turn out to be of paramount importance in explaining the price differences between the day ahead and intraday prices. Karanfil and Li (2017) draw similar conclusions from an empirical study of the Danish market, and exhibit the impact of renewable energies on prices, bid-ask spread and volatility. Rowińska et al. (2018) establish a negative correlation between the wind energy penetration and the day ahead market prices. Jónsson et al. (2010) show that in addition to creating a negative correlation between the renewable infeed and spot prices, a deeper penetration of the intermittent energies significantly modifies the distribution of spot prices.

Optimal strategies in the intraday market for a single wind energy producer have also been the object of studies both in the price-taker and price-maker context. In the price-taker setting, Garnier and Madlener (2015) solve a discrete-time optimal trading problem to arbitrate between immediate and delayed trading when price and production forecasts are uncertain. In Morales et al. (2010), they consider a multimarket setting to derive an optimal bidding strategy for a wind energy producer in the day ahead and adjustment markets, while minimizing the cost incurred in the balancing market. Discrete decisions are taken for each delivery period, considering a finite number of probable scenarii. This approach has been enhanced by Zugno et al. (2013b), where the wind energy producer is now price maker in the balancing market. Following the same framework, Delikaraoglou et al. (2015) formulate a problem where the renewable producer is price maker in both the day ahead and balancing markets and assess the relevance of strategic behavior in the context of high renewable penetration and varying flexible capacities. Still in the price-maker setting, continuous-time approaches have also been developed. Aid et al. (2016), consider the optimal trading rate and power generation of a thermal producer when the residual demand at the terminal date is random. In the same trend, Tan and Tankov (2018) develop an optimal trading model for a wind energy producer. They quantify the evolution of forecast uncertainty at the approach of the delivery time, and exhibit optimal strategies depending on forecast updates.

In our study, the uncertain renewable production is also a source of randomness, and the producers' trading decisions impact the market. Unlike the previous papers on electricity markets, we consider the equilibrium setting with many agents and determine the market price as the result of their interaction. Explicit results for dynamic equilibria are often difficult to obtain. In particular, Nash equilibria often lead to systems of coupled partial differential equations. However, the linear-quadratic setting and in particular the Almgren-Chriss framework of linear market impact and quadratic trading costs has become a standard toolbox allowing many

authors to obtain the explicit form of equilibrium price under different market designs.

In the N-agent setting, Bouchard et al. (2018), study the equilibrium returns in a market with mean-variance optimizing investors under quadratic transactions costs. Closer to our model, Voß (2019) considers a game of two agents in the Almgren-Chriss framework, interacting through the market impact function, where each agent aims to follow a target as in the single-agent model of Bank et al. (2017). In Bank et al. (2021), they apply the Almgren-Chriss framework to the study of liquidity dynamics in OTC dealer markets. Evangelista and Thamsten (2020) consider liquidation games in a finite population of agents with information asymmetry.

The problem of finding the equilibrium may be simplified further by assuming a continuum of agents and using the mean field game approach. Fu et al. (2021) consider the optimal liquidation mean-field game in the generalized Almgren-Chriss framework and obtain the optimal strategies and the equilibrium price as a solution to a linear forward-backward stochastic differential equation (FBSDE) with a singular terminal condition. Fu and Horst (2020a) extend these results to a leader-follower setting using the theory of mean-field games with a major player and Fu et al. (2020) extend the framework with a self-exciting order flow. Fujii and Takahashi (2020) find an equilibrium price under market clearing conditions under quadratic trading costs. Casgrain and Jaimungal (2020, 2018) used the Almgren and Chriss framework in the mean field setting to deal with heterogeneous sub-populations of agents with distinct filtration and/or different beliefs for each sub-population. Shrivats et al. (2020) recently applied the theoretical setting developed in Casgrain and Jaimungal (2020) to the case of trading in solar renewable energy certificate markets. Finally, while this paper was under review, a different equilibrium price model for electricity market was proposed in Aid et al. (2020), where in particular the Samuelson's effect is explained through heterogeneity among agents.

Following these authors, our paper is based on the Almgren-Chriss toolbox, which we use to study market impact, trading strategies and equilibrium prices in intraday electricity markets. Electricity markets are very different from traditional stock/futures markets. The most fundamental difference is the predictability of prices: since electricity is non storable, shifts in demand and supply forecasts are reflected in the price. Other features, such as the Samuelson effect (growth of volatility at the approach of the delivery date) are specific to futures markets in general. Our paper is an attempt to model these features of electricity markets in an endogenous way, by relating price formation to production and demand forecasts. The classical Almgren-Chriss setting, extended with production and demand forecasts provides a simple and tractable framework to take into account the different features of electricity markets: market impact; low liquidity, which improves at the approach of the delivery date etc. We justify the use of this framework a posteriori by showing that the price obtained with our model reproduces the main observed empirical features of intraday electricity prices.

The main difference of our framework with the existing research, motivated by the predictability of electricity prices, is the presence of forecast processes, which determine the terminal constraint on the strategy of each agent. We consider deterministic market impact and trading cost parameters, which enables us to determine the trading strategies and the equilibrium price in explicit form, under very general assumptions on the price process and the forecast processes.

In particular, unlike the above quoted papers, the fundamental price process is not assumed to be a martingale. The explicit form of the equilibrium price enables us to carry out a theoretical study of various characteristics of electricity markets, such as observed volatility, price-forecast correlation, market impact of forecast adjustments and trading costs. All these quantities are also determined in explicit form. As a final contribution, in the last section of the paper we perform an empirical analysis of intraday electricity markets using order book data, and show empirically that the qualitative features of electricity markets are reproduced by our model.

The paper is structured as follows. Section 3.2 describes the market and introduces our modeling framework. In Section 3.3.1 we place ourselves in a setting with a finite number of agents, where all agents observe the forecasts of the other agents. In Section 3.3.2 we consider the mean field game, where agents only observe their individual forecasts and the common information. To make a connection between the N-agent setting and the mean-field game setting, we show in Section 3.3.3 that (i) the N-player equilibrium converges to the mean-field equilibrium as $N \to \infty$, and (ii) an ε-Nash equilibrium for the N-player problem may be constructed from the mean-field equilibrium. In Section 3.4 we use the results of Section 3.3.1 to analyze theoretically the properties of equilibrium price in electricity markets. Finally, in Section 3.5 we perform an empirical analysis of intraday electricity prices and confront it with the theoretical results obtained in the preceding sections.

3.2 Preliminaries on electricity markets

In this paper we consider a short-term electricity market, populated with small agents with identical characteristics. These agents face uncertain demand or supply at some future time, and use the electricity market to manage the associated risk. While our primary interest is to study the impact of increasing renewable penetration on intraday market prices, the market participants may in principle represent both renewable producers with uncertain generation forecasts and industrial consumers / utilities with uncertain demand. To simplify the language and notation, in the sequel, unless specified otherwise, we will refer to forecasts of all agents as demand forecasts (if the agent is a producer, its demand forecast will therefore be negative).

In most countries, the short-term electricity markets have the following structure (the specific times correspond to the EPEX Spot/Intraday market):

- The day-ahead market is a one-off trading venue, where the agents may make bids until 12PM (noon) on the day preceding the delivery day. At 12:55 the price is fixed using the merit order mechanism and the market clears. A major part of the electricity production is sold on the day-ahead market.

- At 3PM on the day preceding the delivery day, the intraday market opens, allowing continuous trading for each quarter-hour of the delivery day. The intraday market has higher trading cost than the day-ahead market, and is mainly used by market participants to adjust their day-ahead positions following forecast updates.

- 15 minutes before delivery the intraday trading for the given delivery period closes. At this point, negative production imbalances must be compensated to the market operator at the 'high imbalance settlement price', which is higher than the last intraday price, and for positive production imbalances, the producer is compensated at the 'low imbalance settlement price', which is lower than the intraday market price[1]. In addition, high imbalances carry a reputational cost for the market participant. To avoid paying the imbalance penalties, the aggregate position (day-ahead plus intraday) held by the agent at the delivery time must therefore be equal to the realized demand.

To represent this market structure in a simplified way, we consider a fixed delivery period, starting at time T, and assume that the day-ahead market allows agents to trade instantaneously, without transaction costs, at time $t = 0$, at price denoted by S_0. Then, between $t = 0$ and $t = T$, the agents may trade in the intraday market, at price $(P_t)_{0\leq t\leq T}$, which contains a market impact component, and subject to transaction costs. Finally, at time T, if there is an imbalance, the agents must purchase the missing amount /sell the extra amount of electricity at price S_T without transaction cost, and in addition, pay a penalty depending on the absolute value of the imbalance. In the following section we provide details of the model and compute explicitly the optimal strategies of the agents and the equilibrium intraday market price.

3.3 Optimal trading strategies and equilibrium price

In this section we introduce our model of electricity market and derive explicit expressions for the equilibrium price and optimal equilibrium trading strategies of the agents. We consider both the N-agent setting (Section 3.3.1) and the mean-field game setting (Section 3.3.2). Section 3.3.3 clarifies the relationship between the equilibrium strategies and prices in the N-agent market and those of the mean field game limit.

3.3.1 N-player setting

In this section we assume that in the market there are N identical agents, and we denote by ϕ_t^i the position of i-th agent at time t. As is common in optimal execution literature, we assume that the position of i-th agent is an absolutely continuous process, and we define the *rate of trading* $\dot{\phi}_t^i$. We introduce a filtered probability space $(\Omega, \mathcal{F}, \mathbb{F} := (\mathcal{F}_t)_{t\in[0,T]}, \mathbb{P})$ to which all processes are adapted, and which models the information available in the market to all the agents. The position of the i-th agent at time t is given by $\phi_t^i = \phi_0^i + \int_0^t \dot{\phi}_s^i ds$ with $\phi_0^i \in \mathcal{F}_0$ denoting the position of the agent in the day-ahead market. The fundamental electricity price process is denoted by $(S_t)_{t\in[0,T]}$, where S_0 corresponds to the day-ahead market price and S_t for $0 < t < T$ denotes the intraday market price net of the price impact component. The intraday market price with the price impact component is denoted by $(P_t^N)_{t\in[0,T]}$. The strategies of the

[1]See `www.services-rte.com/en/learn-more-about-our-services/becoming-a-balance-responsible-party/Imbalance-settlement-price.html` for details

agents impact the market price P_t^N as follows:

$$P_t^N = S_t + a(\bar{\phi}_t^N - \bar{\phi}_0^N), \quad \forall t \in [0, T], \tag{3.1}$$

where $\bar{\phi}_t^N = \frac{1}{N}\sum_{i=1}^N \phi_t^i$ is the average position of the agents and a is a constant. The parameter N describes the size of the market (number of agents), it is therefore natural that the trading strategy of each agent has an effect of order of $1/N$ on the market price. The permanent component of the price impact of trades in our model is thus linear, which is the only shape compatible with the absence of arbitrage, see Huberman and Stanzl (2004); Gatheral (2010). On the other hand, the transient component of market impact is not modelled directly. Literature on market microstructure mostly shows that metaorders have a concave transient impact on prices (see Bershova and Rakhlin (2013), Bacry et al. (2015), Bucci et al. (2020) and Bouchaud (2010)). However, for the sake of simplicity and in order to obtain an analytical solution for our model, we choose a linear impact function as in the seminal papers by Almgren and Chriss (1999, 2001) and many other more recent papers, including Aid et al. (2016) in the context of electricity markets. The transient component of the market impact is taken into account indirectly, via a trading cost penalty.

The agents trading in the market at time t incur an instantaneous cost,

$$\dot{\phi}_t^i P_t^N + \frac{\alpha(t)}{2}\dot{\phi}_t^i\Big(\dot{\phi}_t^i + b\dot{\bar{\phi}}_t^{N,-i}\Big), \quad \forall t \in [0, T]$$

for the i-th agent where $\dot{\bar{\phi}}_t^{N,-i} = \frac{1}{N-1}\sum_{j=1, j\neq i}^N \dot{\phi}_t^j$. Here the first term represents the actual cost of buying the electricity, and the second term represents the cost of trading, where $\alpha(.)$ is a continuous strictly positive function on $[0, T]$ reflecting the variation of market liquidity at the approach of the delivery date. The term $b\dot{\bar{\phi}}_t^{N,-i}$ with $b > 0$ represents the impact of the crowd trading direction on the cost of trading of a single agent, which accounts for possible synchronization of the agents. The instantaneous cost paid by each agent is thus independent of the size of the market. This corresponds to a market where immediately available liquidity (market depth) is low (thus even a minor agent has to pay order book costs) but the order book is resilient (thus the trade of a minor agent only has a lasting impact of order of $1/N$ on the price). This is consistent with recent empirical and theoretical studies of order book dynamics, for example, according to Donier et al. (2015), while the total daily volume exchange on a typical stock is around 1/200th of its market capitalization, the volume present in the order book at any given time is 1000 times smaller than this.

Each agent i has a demand forecast X_t^i and aims to maximize her gain from trading in the market under the volume constraint $\phi_T^i = X_T^i$. More precisely, whenever $\phi_T^i \neq X_T^i$, the agent must first purchase the missing amount or sell the extra amount of electricity at price S_T and in addition pay an imbalance penalty $\frac{\lambda}{2}(\phi_T^i - X_T^i)^2$. The actual imbalance mechanism of electricity markets boils down to applying a L^1 penalty function to the terminal imbalance; however, large imbalances may also create a reputational damage to the producer, thus a quadratic penalty, which penalizes large imbalances more strongly, appears appropriate. On the other hand, the

'hard constraint' may be recovered from our results by making the penalty parameter λ tend to infinity.

Our main results hold true under the following assumption.

Assumption 11. *The process S is $\mathbb{F}$-adapted and satisfies*

$$\mathbb{E}[\sup_{0\leq t\leq T} S_t^2] < \infty. \tag{3.2}$$

and the processes $(X^i)_{i=1}^N$ are square integrable $\mathbb{F}$-martingales.

Considering the demand forecast as a martingale is natural since it is the best estimate at time t of what the demand will be at the delivery time T given our current knowledge $\mathcal{F}_t$.

Definition 12 (Admissible strategy)**.** *We say that the strategy $(\phi_t^i)_{t\in[0,T]}$ of the i-th agent is admissible if $\phi_0^i \in \mathcal{F}_0$, the process $(\dot{\phi}_t)_{t\in[0,T]}$ is $\mathbb{F}$-adapted and*

$$\mathbb{E}\Big[(\phi_0^i)^2 + \int_0^T (\dot{\phi}_t^i)^2 dt\Big] < \infty.$$

Following the discussion above, the objective function maximized by agent i is written as follows:

$$\begin{aligned} J^{N,i}(\phi^i, \phi^{-i}) := -\mathbb{E}\Big[& \underbrace{\phi_0^i S_0}_{\text{Day ahead}} + \underbrace{\int_0^T \left\{ \frac{\alpha(t)}{2} \dot{\phi}_t^i \Big(\dot{\phi}_t^i + b\dot{\bar{\phi}}_t^{N,-i} \Big) + \dot{\phi}_t^i P_t^N \right\} dt}_{\text{Intraday}} \\ & \underbrace{-(\phi_T^i - X_T^i)S_T + \frac{\lambda}{2}(\phi_T^i - X_T^i)^2}_{\text{Balancing}} \Big], \end{aligned} \tag{3.3}$$

where $\phi^{-i} := (\phi^1, \ldots, \phi^{i-1}, \phi^{i+1}, \ldots, \phi^N)$ is the vector of positions of all agents except the i-th one. Here, the first term corresponds to the day-ahead market transaction, the integral term corresponds to the cost of purchasing electricity in the intraday market, and the term in the second line corresponds to the imbalance payment.

Because of the price impact, each agent's gain is affected by the decisions of others and we thus face a non-cooperative game. The optimal strategy of each player depends on the other players' actions and we want to describe the resulting dynamical equilibrium, which we define formally below.

Definition 13 (Nash Equilibrium)**.** *We say that $(\phi_t^{i*})_{t\in[0,T]}^{i=1\ldots N}$ is a Nash Equilibrium for the N-player game if it is a vector of admissible strategies, and for each $i = 1, \ldots, N$,*

$$J^{N,i}(\phi^i, \phi^{-i*}) \leq J^{N,i}(\phi^{i*}, \phi^{-i*}) \tag{3.4}$$

for any other admissible strategy ϕ^i.

The following theorem characterizes explicitly the Nash equilibrium of the N-player game. In the theorem and its proof, we denote the average forecast process by $\overline{X}_t^N := \frac{1}{N}\sum_{i=1}^N X_t^i$ and

use the following shorthand notation.

$$
\begin{aligned}
\Delta^N_{s,t} &:= \int_s^t \frac{\eta^N_{u,t}}{\alpha(u)\big(1+\frac{b}{2}\big)}du && \text{with} \quad \eta^N_{s,t} = e^{-\int_s^t \frac{(N-1)a}{N\alpha(u)\left(1+\frac{b}{2}\right)}du} \\
\widetilde{\Delta}^N_{s,t} &:= \int_s^t \frac{\tilde{\eta}^N_{u,t}}{\alpha(u)}du, && \text{with} \quad \tilde{\eta}^N_{s,t} = e^{\int_s^t \frac{a}{N\alpha(u)}du} \\
I^N_t &:= \int_0^t \frac{\eta^N_{s,t}}{\alpha(s)\big(1+\frac{b}{2}\big)}S_s ds, && \widetilde{I}^N_t := \mathbb{E}\left[\int_0^T \frac{\eta^N_{s,T}}{\alpha(s)\big(1+\frac{b}{2}\big)}S_s ds \Big| \mathcal{F}_t\right], \\
\widetilde{S}_t &:= \mathbb{E}[S_T|\mathcal{F}_t], && \check{X}^i_t = X^i_t - \overline{X}^N_t. \qquad (3.5)
\end{aligned}
$$

The proof of this theorem can be found in Section 3.6.1.

Theorem 14. *Under Assumption 11, the unique Nash equilibrium in the N-player game is given by*

$$
\begin{aligned}
\phi^{i*}_t = X^i_0 &+ \frac{1+\frac{a}{N}\Delta^N_{0,t}}{1+\frac{a}{N}\Delta^N_{0,T}}(\widetilde{I}^N_0 - \widetilde{S}_0\Delta^N_{0,T}) - (I^N_t - \Delta^N_{0,t}\widetilde{S}_0) \\
&+ \int_0^t \Delta^N_{s,t}\frac{\left(\frac{a}{N}+\lambda\right)d\widetilde{I}^N_s + \lambda d\overline{X}^N_s + d\widetilde{S}_s}{1+\left(\frac{a}{N}+\lambda\right)\Delta^N_{s,T}} + \int_0^t \widetilde{\Delta}^N_{s,t}\frac{\lambda d\check{X}^i_s}{1+\left(\frac{a}{N}+\lambda\right)\widetilde{\Delta}^N_{s,T}}.
\end{aligned} \tag{3.6}
$$

The equilibrium price has the following form:

$$
\begin{aligned}
P^N_t = S_t &+ a\frac{\frac{a}{N}\Delta^N_{0,t}}{1+\frac{a}{N}\Delta^N_{0,T}}(\widetilde{I}^N_0 - \widetilde{S}_0\Delta^N_{0,T}) - a(I^N_t - \Delta^N_{0,t}\widetilde{S}_0) \\
&+ a\int_0^t \Delta^N_{s,t}\frac{\left(\frac{a}{N}+\lambda\right)d\widetilde{I}^N_s + \lambda d\overline{X}^N_s + d\widetilde{S}_s}{1+\left(\frac{a}{N}+\lambda\right)\Delta^N_{s,T}}.
\end{aligned} \tag{3.7}
$$

Discussion The day-ahead market position of i-th agent is given by

$$
\phi^{i*}_0 = X^i_0 + \frac{\widetilde{I}^N_0 - \widetilde{S}_0\Delta^N_{0,T}}{1+\frac{a}{N}\Delta^N_{0,T}}.
$$

The agents, therefore, trade in the day-ahead market based on their forecasts at time 0 and apply a correction for the potential fundamental price trend, which disappears if the fundamental price is a martingale.

For nonzero trading costs, the strategies of the agents and thus the price impact have a finite variation. Hence, the price impact component does not directly induce additional volatility which may be a weakness of the model. However, the drift $\dot{\bar{\phi}}^N$ is stochastic and thus creates additional price variations, making the effective observed volatility larger. We will investigate this phenomenon in more details in Section 3.4.2.

The aggregate intraday market strategy $\bar{\phi}^{N*}$ (given by equation (3.27)) and, consequently, the equilibrium price have a complex structure because of the generality of our setting; in particular the fundamental price process $(S_t)_{0\le t\le T}$ is only assumed to be square integrable.

Under more stringent assumptions, important simplifications can be obtained, as the following examples illustrate.

- Assume that the fundamental price process S is a martingale. Then, $\widetilde{I}_t^N = -\int_0^t S_s d\Delta_{s,T}^N + S_t\Delta_{t,T}^N$ and $d\tilde{I}_t^N = \Delta_{t,T}^N dS_t$. Substituting this into (3.7), after cancellations, we find that aggregate strategy does not depend on the fundamental price:

$$\bar{\phi}_t^{N*} = \overline{X}_0^N + \int_0^t \frac{\Delta_{s,t}^N \lambda d\overline{X}_s^N}{1 + \left(\frac{a}{N} + \lambda\right)\Delta_{s,T}^N}.$$

  In the absence of price trend, the trades are therefore only provoked by forecast adjustments.

- Assume now that the fundamental price contains a martingale component M and a deterministic component A: $S_t = A_t + M_t$. The above argument shows that the aggregate strategy does not depend on the martingale part M. Thus, we can assume that S_t is deterministic, which means that $\widetilde{I}_t^N$ and $\widetilde{S}_t$ are constant, and the aggregate strategy becomes

$$\begin{aligned}\bar{\phi}_t^{N*} = \overline{X}_0^N &+ \frac{1 + \frac{a}{N}\Delta_{0,t}^N}{1 + \frac{a}{N}\Delta_{0,T}^N} \int_0^T \frac{\eta_{s,T}^N}{\alpha(s)\left(1 + \frac{b}{2}\right)}(A_s - A_T)ds \\ &- \int_0^t \frac{\eta_{s,t}^N}{\alpha(s)\left(1 + \frac{b}{2}\right)}(A_s - A_T)ds + \int_0^t \frac{\Delta_{s,t}^N \lambda d\overline{X}_s^N}{1 + \left(\frac{a}{N} + \lambda\right)\Delta_{s,T}^N}\end{aligned}$$

  If there is a positive trend in the fundamental price (that is, A is increasing), then the day-ahead position will be below the demand forecast, but there will be a positive trend in the aggregate strategy: the overall price trend will be amplified by the market impact component.

- Consider now the limiting case of infinite penalty: $\lambda \to \infty$. Then, using the dominated convergence as needed, we see that the aggregate strategy satisfies:

$$\begin{aligned}\lim_{\lambda\to\infty} \bar{\phi}_t^{N*} = \overline{X}_0^N &+ \frac{1 + \frac{a}{N}\Delta_{0,t}^N}{1 + \frac{a}{N}\Delta_{0,T}^N}(\widetilde{I}_0^N - \widetilde{S}_0\Delta_{0,T}^N) \\ &- (I_t^N - \Delta_{0,t}^N\widetilde{S}_0) + \int_0^t \Delta_{s,t}^N \frac{d\widetilde{I}_s^N + d\overline{X}_s^N}{\Delta_{s,T}^N}\end{aligned}$$

- Finally, let us compute the form of trading strategy in the limit of zero trading costs. To this end, we assume in addition that the fundamental price process S has a left limit at every point. Fixing $s < t \in [0, T]$, we have:

$$\Delta_{s,t}^N = \frac{N}{a(N-1)}\left(1 - e^{-\int_s^t \frac{a(N-1)}{\alpha(l)\left(1+\frac{b}{2}\right)N}dl}\right) \longrightarrow \frac{N}{a(N-1)} := \Delta^*$$

  as $\alpha(t) \to 0$ uniformly in t. From the left limit property of S, it is easy to see that $I_t^N \to$

$\Delta^* S_t$ almost surely, for every t. For similar reasons, using the dominated convergence theorem, $\widetilde{I}_t^N \to \Delta^* \widetilde{S}_t$. Finally,

$$\lim_{\|\alpha\|\to 0} \bar{\phi}_t^{N*} = \overline{X}_0^N + \frac{N}{a(N-1)}(\widetilde{S}_t - S_t) + \frac{\lambda}{a+\lambda}(\overline{X}_t^N - \overline{X}_0^N)$$

Thus, in the absence of trading costs, for $N \geq 2$, the aggregate equilibrium strategy is well defined, and the gain of each agent remains bounded in expectation. This is in contrast with the single-agent case, where the gain may be arbitrarily large, unless the fundamental price process is a martingale. Indeed, in the single-agent case, without transaction costs the objective function writes:

$$J^{1,i}(\phi) = \mathbb{E}\left[\int_0^T \phi_t dS_t - \frac{a}{2}(\phi_T - \phi_0)^2 - \frac{\lambda}{2}(X_T - \phi_T)^2 - X_T S_T\right],$$

and it is clear that unless the fundamental price process is a martingale, this expression can be made arbitrarily large. This means that the "price of anarchy" in this model is infinite: if the agents chose the same strategy, they could have all obtained an infinite gain, but competition between agents limits everybody's gain to a finite value.

- Finally, coming back to the form of the individual agent's strategy ϕ_t^{i*}, we see that the dependence on the trading cost α is different for the common part of the strategy and the individual part of the strategy (the last term of the formula). While the common part of the strategy depends on the "effective trading cost" $\alpha(1 + b/2)$, taking into account the crowd behavior, the individual part of the strategy depends only on α. We conclude that due to additional costs related to crowd behavior of agents, the agents trade less actively in response to common forecast updates than in response to individual forecast updates.

3.3.2 Mean-field game setting

In this section, we place ourselves in the mean field game limit, that is, we assume the number of agents in the market, N tends to infinity, while the strategy of each agent remains finite. We then consider a generic agent and denote by $X := (X_t)_{t\in[0,T]}$ the demand forecast of this agent, by ϕ the agent's position and by $\mathbb{F}$ the filtration which contains the information available to this agent. In addition we introduce a smaller filtration, containing the common noise and denoted by $\mathbb{F}^0$. This filtration contains the information about the fundamental price and potentially some information about the demand forecast but, in general, not the full individual demand forecast of the generic agent. We decompose the individual demand forecast as follows: $X_t = \overline{X}_t + \check{X}_t$, where $\overline{X}_t = \mathbb{E}\left[X_t|\mathcal{F}_t^0\right]$ is common for all agents (it can be seen as a national demand forecast). In this mean field game setting, the average quantities of the N-agent problem are replaced with conditional expectations with respect to the common noise filtration $\mathbb{F}^0$.

For any $\mathbb{F}$-adapted process $(\zeta_t)_{t\in[0,T]}$, we will denote $\bar{\zeta}_t = \mathbb{E}[\zeta_t|\mathcal{F}_t^0] = \int_{\mathbb{R}} x\mu_t^\zeta(dx)$ where: $\mu_t^\zeta := \mathcal{L}(\zeta_t|\mathcal{F}_t^0)$. The game is now represented by the interaction of agents through the conditional

distribution flow $\mu_t^\phi := \mathcal{L}(\phi_t|\mathcal{F}_t^0)$ of the state process. The price impact function, defined in the previous section as an expectation with respect to the empirical measure, is now an integral with respect to the measure flow:

$$P_t = S_t + a(\bar{\phi}_t - \bar{\phi}_0). \tag{3.8}$$

Each individual agent now has a negligible impact on the price, but the aggregate position of all agents has a nonzero impact. Thus, in the mean-field game setting, we consider that the market is very large compared to the size of the individual agent, but the immediately available liquidity in the order book is small, so that even a minor agent pays a non-zero trading cost.

The objective function for the generic agent is

$$\begin{aligned} J^{MF}(\phi, \bar{\phi}) := -\mathbb{E}\Big[\phi_0 S_0 + \int_0^T \frac{\alpha(t)}{2}\dot{\phi}_t\Big(\dot{\phi}_t + b\dot{\bar{\phi}}_t\Big) + \dot{\phi}_t(S_t + a(\bar{\phi}_t - \bar{\phi}_0))dt \\ -(\phi_T - X_T)S_T + \frac{\lambda}{2}(\phi_T - X_T)^2\Big]. \end{aligned} \tag{3.9}$$

As in the previous section, each agent maximizes this functional over the set of strategies satisfying Definition 12.

We now define the mean field equilibrium.

Definition 15 (mean field equilibrium)**.** *An admissible strategy $\phi^* := (\phi_t^*)_{t\in[0,T]}$ is a mean field equilibrium if for any admissible strategy ϕ,*

$$J^{MF}(\phi, \bar{\phi}^*) \leq J^{MF}(\phi^*, \bar{\phi}^*).$$

In this section, we make the following assumption.

Assumption 16.

- *The process S is adapted to the filtration $\mathbb{F}^0$ and satisfies (3.2).*
- *The process X is a square integrable martingale with respect to the filtration $\mathbb{F}$.*
- *The process $\overline{X}$ defined by $\overline{X}_t := \mathbb{E}[X_t|\mathcal{F}_t^0]$ for $0 \leq t \leq T$ is a square integrable martingale with respect to the filtration $\mathbb{F}$.*

Note that if X is an $\mathbb{F}$-martingale, then $\overline{X}$ is by construction an $\mathbb{F}^0$-martingale, but it may not necessarily be a martingale in the larger filtration $\mathbb{F}$.

The following theorem characterizes the mean field equilibrium in our setting. The statement of the theorem appears similar to that of Theorem 14, modulo replacing $\overline{X}^N$ with $\overline{X}$ and making N tend to infinity. However, the computation of the strategy and the market price in the N-player setting requires the knowledge of the sum of forecasts of all agents whereas in the mean-field setting one needs to know the conditional expectation of the agent's forecast with respect to the 'common knowledge' filtration. Thus, the theoretical price given by this theorem can be computed by the regulator, and the strategy of this theorem can be computed by an individual player, both of which do not have the complete information about the forecasts of other players.

In the theorem and its proof, we use the following shorthand notation.

$$\begin{aligned}
\Delta_{s,t} &:= \int_s^t \frac{\eta_{u,t}}{\alpha(u)\big(1+\frac{b}{2}\big)}du && \text{with} \quad \eta_{s,t} = e^{-\int_s^t \frac{a}{\alpha(u)\left(1+\frac{b}{2}\right)}du}, \\
I_t &:= \int_0^t \frac{\eta_{s,t}}{\alpha(s)\big(1+\frac{b}{2}\big)} S_s ds, && \widetilde{I}_t := \mathbb{E}\left[\int_0^T \frac{\eta_{s,T}}{\alpha(s)\big(1+\frac{b}{2}\big)} S_s ds \middle| \mathcal{F}_t\right], \\
\text{and} \quad \widetilde{\Delta}_{s,t} &:= \int_s^t \alpha^{-1}(u)du. && \qquad (3.10)
\end{aligned}$$

Theorem 17. *Under Assumption 16, the unique mean field equilibrium strategy is given by*

$$\begin{aligned}
\phi_t^* &= X_0 + (\widetilde{I}_0 - \widetilde{S}_0\Delta_{0,T}) - (I_t - \Delta_{0,t}\widetilde{S}_0) \\
&\quad + \int_0^t \Delta_{s,t}\frac{\lambda d\widetilde{I}_s + \lambda d\overline{X}_s + d\widetilde{S}_t}{1+\lambda\Delta_{s,T}} + \int_0^t \widetilde{\Delta}_{s,t}\frac{\lambda d\check{X}_s}{1+\lambda\widetilde{\Delta}_{s,T}}. \qquad (3.11)
\end{aligned}$$

The equilibrium price has the following form:

$$P_t = S_t - a(I_t - \Delta_{0,t}\widetilde{S}_0) + a\int_0^t \Delta_{s,t}\frac{\lambda d\widetilde{I}_s + \lambda d\overline{X}_s + d\widetilde{S}_t}{1+\lambda\Delta_{s,T}}. \qquad (3.12)$$

The proof of Theorem 17 follows the lines of that of Theorem 14 with some adjustments, and is thus omitted to save space.

3.3.3 Relationship between N-player setting and MFG setting

In this section, we study the relationship between the equilibrium strategies and prices in the N-agent market and those of the mean field game limit, and prove the following results.

- The market price and the agent's strategy in the N-agent model converge to their respective mean field values as $N \to \infty$. This shows that to understand the behavior of agents and prices in the realistic N-agent market, one can use the mean-field game model, which does not require the knowledge of individual forecasts, but only that of the common information filtration.

- An approximate equilibrium (ε-Nash equilibrium) in the N-player setting may be constructed from the MFG solution. In other words, an agent trading in the N-agent market may construct a strategy whose gain is sufficiently close to the optimal equilibrium gain using the mean-field game solution, which does not require the knowledge of the private forecasts of the other agents.

To address these questions, we need to make more precise assumptions on the probabilistic setup of the problem. In particular, since we would like to study the convergence of the N-agent problem as $N \to \infty$, we consider an infinity of agents. In addition, all N-agent problems and the mean field problem must be defined on the same probability space.

Assumption 18.

- *The process S adapted to the filtration $\mathbb{F}^0$ and satisfies (3.2).*
- *The processes $(X^i)_{i=1}^{\infty}$ are square integrable $\mathbb{F}$-martingales.*
- *There exists a square intergrable $\mathbb{F}$-martingale $\overline{X}$, such that for all $i \geq 1$, and all $t \in [0,T]$, almost surely, $\mathbb{E}[X_t^i|\mathcal{F}_t^0] = \overline{X}_t$.*
- *The processes $(\check{X}^i)_{i=1}^{\infty}$ defined by $\check{X}_t^i = X_t^i - \overline{X}_t$ for $t \in [0,T]$, are orthogonal square integrable $\mathbb{F}$-martingales, such that the expectation $\mathbb{E}[(\check{X}_T^i)^2]$ does not depend on i.*

Let us fix $N < \infty$, and consider a market with N agents. For a given $i \leq N$, we may define the "mean-field" strategy for the i-th agent as follows.

$$\begin{aligned}\phi_t^{MF,i*} &= X_0^i + (\widetilde{I}_0 - \widetilde{S}_0\Delta_{0,T}) - (I_t - \Delta_{0,t}\widetilde{S}_0) \\ &\quad + \int_0^t \Delta_{s,t}\frac{\lambda d\widetilde{I}_s + \lambda d\overline{X}_s + d\widetilde{S}_t}{1+\lambda\Delta_{s,T}} + \int_0^t \widetilde{\Delta}_{s,t}\frac{\lambda d\check{X}_s^i}{1+\lambda\widetilde{\Delta}_{s,T}}\end{aligned} \tag{3.13}$$

Unlike the true optimal strategy of the i-th agent, this strategy is computed using only the common information and the individual information of the i-th agent, it does not require the knowledge of the private forecasts of the other agents. Moreover, this strategy does not depend on N. The following two results show that, on the one hand, the true optimal strategy of the i-th agent in the N-player game converges to this mean-field strategy as $N \to \infty$, and on the other hand, that this mean-field strategy, if used by all agents in the N-player game, constitutes an ε-Nash equilibrium. Proofs of these results can be found in Appendix 3.6.2.

Proposition 19. *Let Assumption 18 holds true, and let ϕ^{i*} denote the optimal position of the i-th agent in the N-player setting, given by (3.6), and by $\phi^{MF,i*}$ the optimal position in the mean field setting, given by (3.13). Then, for all $N \geq 1$, the differences between the strategy of a single agent, the aggregate strategy and the equilibrium price in the N-agent model and the corresponding quantities in the mean-field model can be bounded as follows.*

$$\begin{aligned}&\sup_{0\leq t\leq T}\mathbb{E}[(\phi_t^{i*} - \phi_t^{MF,i*})^2] + \sup_{0\leq t\leq T}\mathbb{E}[(\overline{\phi}_t^{N*} - \overline{\phi}_t^{*})^2] + \sup_{0\leq t\leq T}\mathbb{E}[(P_t^N - P_t)^2] \\ &\qquad + \sup_{0\leq t\leq T}\mathbb{E}[(\dot{\overline{\phi}}_t^{N*} - \dot{\overline{\phi}}_t^{*})^2] \leq \frac{C}{N^2}\mathbb{E}[\sup_{0\leq t\leq T} S_t^2] + \frac{C}{N^2}\mathbb{E}[(\overline{X}_T)^2] + \frac{C}{N}\mathbb{E}[(\check{X}_T^i)^2],\end{aligned}$$

where the constant C depends only on the coefficients α, b, λ and a.

Proposition 20. *Under Assumption 18, consider the vector of admissible strategies for the N-player game defined by equation (3.13) for $i = 1, \ldots, N$. Then, there is a constant $C < \infty$ which does not depend on N, such that for any other vector of admissible strategies $(\phi_t^i)_{t\in[0,T]}^{i=1\ldots,N}$ for the N-player game,*

$$J^{N,i}(\phi^i, \phi^{MF,-i*}) - \frac{C}{N^{\frac{1}{2}}} \leq J^{N,i}(\phi^{MF,i*}, \phi^{MF,-i*}), \quad \forall i \in \{1, ..., N\},\ \forall t \in [0,T].$$

In other words, the vector of strategies $(\phi_t^{MF,i*})_{t\in[0,T]}^{i=1\ldots,N}$ *is an* ε*-Nash equilibrium for the* N*-player game with* $\varepsilon = \frac{C}{N^{\frac{1}{2}}}$.

3.4 Intraday electricity prices: theoretical insights

In this section we show through theoretical analysis how the main empirically documented features of electricity markets appear naturally as the result of our model. Empirical illustrations of these features are provided in the following section. We analyze the effect of market structure (number of participants, terminal penalty, trading costs and market impact parameters) on the overall costs/gains of participants as well as on the aggregate market parameters such as price volatility, the correlation between forecast and price and the impact of forecast adjustments on market prices. We consider the N-agent framework and make the following additional assumptions to simplify computations.

- The fundamental price process S is a martingale orthogonal to the forecast processes of the agents, with $\langle S\rangle_t = \sigma_S^2 t$.
- The trading cost parameter α is constant.
- The forecast processes of agents satisfy

$$\langle \overline{X}^N\rangle_t = \sigma_X^2 t, \quad \langle \check{X}^i\rangle_t = \check{\sigma}_X^2 t, \quad \text{and} \quad \langle \overline{X}^N, \check{X}^i\rangle_t = 0 \quad \forall\, i,$$

  for some constants σ_X and $\check{\sigma}_X$, where $\check{X}_t^i = X_t^i - \overline{X}_t^N$.

In addition, to make the notation more compact throughout this section we write $\tilde{\alpha} := \alpha(1+b/2)$.

Under these assumptions, the coefficients $\eta_{s,t}^N$ and $\Delta_{s,t}^N$ depend only on $t-s$ and not on s and t separately. We shall therefore write them as η_{t-s}^N and Δ_{t-s}^N, and similarly for the other coefficients, from now and until the end of this section. The aggregate position of N agents in equilibrium therefore writes:

$$\bar{\phi}_t^N = \overline{X}_0^N + \int_0^t \Delta_{t-s}^N \frac{\lambda d\overline{X}_s^N}{1 + \left(\frac{a}{N} + \lambda\right)\Delta_{T-s}^N}.$$

3.4.1 Price impact of forecast adjustments

Due to the non-storability of electricity, the prices of this commodity are strongly affected by demand and supply shocks. While for regular commodities these shocks may be compensated by changes in reserves, for electricity this is not possible. As a result, supply shocks caused, for instance, by the power plant breakdowns, and demand shocks often caused by weather forecast changes have a lasting impact on the price. In our model, the market impact of demand/supply shocks can be represented through a jump in the forecast process. An idiosyncratic supply shock may correspond to a jump in the individual forecast process X^i, while a generalized demand shock caused by weather forecast update may correspond to a jump in the aggregate forecast

process $\overline{X}^N$. Consider for instance a jump $\Delta\overline{X}^N$ in the aggregate forecast occurring at time t^*. The impact of this jump on the aggregate strategy $\bar{\phi}_t^N$, defined as the difference of the strategies with and without the forecast adjustment, is given by

$$\delta\bar{\phi}_t^N = \mathbf{1}_{t\geq t^*}\frac{\lambda\Delta_{t-t^*}^N\Delta\overline{X}^N}{1+\left(\frac{a}{N}+\lambda\right)\Delta_{T-t^*}^N}.$$

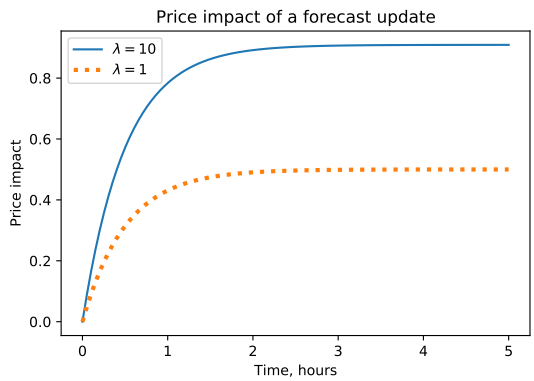

Figure 3.1: Price impact of an aggregate demand forecast adjustment at $t = 0$ (delivery time is $t = 5$). Parameter values: $a = 1$, $N = 100$, $\alpha = 0.5$, $b = 0$, forecast adjustment size: 1 MW.

For example, after a positive demand shock, the agents will need to purchase the missing electricity in the market, creating a permanent price impact given by a positive increasing continuous function of time (see Figure 3.1).

3.4.2 Volatility and Samuelson's effect

We have seen that since the strategy $\bar{\phi}^N$ is differentiable, the quadratic variation of the equilibrium price P_t^N coincides with the quadratic variation of the fundamental price. However, the actual observed volatility, which is estimated from discretely observed prices, may be different. The standard estimator of integrated variance by discrete quadratic variation over the interval $[t, t+h]$ with M steps is given by

$$Q^M(t, t+h) = \sum_{i=0}^{M-1}(P_{t+\frac{i+1}{M}h} - P_{t+\frac{i}{M}h})^2$$

To focus on the average behavior of volatility rather than on individual random trajectories, we consider the expectation of this estimator. Finally, to estimate the expected instantaneous variance, it seems natural to consider this estimator over one time step and normalize it by the step size. Thus, the expression

$$\frac{1}{h}\mathbb{E}[(P_{t+h}^N - P_t^N)^2]$$

represents the average instantaneous price variance, estimated over time step h.

The following lemma quantifies the behavior of this expression for small values of h.

Lemma 21. *As $h \to 0$, the equilibrium price satisfies*

$$\frac{1}{h}\mathbb{E}[(P^N_{t+h} - P^N_t)^2] = \sigma_S^2 + h\mathcal{V}_t + O(h^2), \tag{3.14}$$

where

$$\mathcal{V}_t = \frac{a^2\lambda^2\sigma_X^2}{\tilde{\alpha}^2}\int_0^t \frac{(\eta^N_{t-s})^2}{(1+(\frac{a}{N}+\lambda)\Delta^N_{T-s})^2}ds. \tag{3.15}$$

Proof. The expected squared change in the price process satisfies:

$$\mathbb{E}[(P^N_{t+h} - P^N_t)^2] = \sigma_S^2 h + 2a\mathbb{E}[(S_{t+h} - S_t)(\bar{\phi}^N_{t+h} - \bar{\phi}^N_t)] + a^2\mathbb{E}[(\bar{\phi}^N_{t+h} - \bar{\phi}^N_t)^2]$$

Since the fundamental process S is orthogonal to the forecast process, the second term in the right-hand side above is zero. The third term satisfies:

$$\begin{aligned}\mathbb{E}[(\bar{\phi}^N_{t+h} - \bar{\phi}^N_t)^2] = \sigma_X^2\lambda^2 &\int_0^t \frac{(\Delta^N_{t-s} - \Delta^N_{t+h-s})^2}{(1+(\frac{a}{N}+\lambda)\Delta^N_{T-s})^2}ds \\ &+ \sigma_X^2\lambda^2\int_t^{t+h} \frac{(\Delta^N_{t+h-s})^2}{(1+(\frac{a}{N}+\lambda)\Delta^N_{T-s})^2}ds.\end{aligned}$$

From the explicit form of Δ^N_t, it is clear that the second term above is of order of $O(h^3)$, and the first term equals

$$h^2\sigma_X^2\lambda^2\int_0^t \frac{((\Delta^N_{t-s})')^2}{(1+(\frac{a}{N}+\lambda)\Delta^N_{T-s})^2}ds,$$

up to terms of order of h^3. □

As expected, as $h \to 0$, the expression (3.14) converges to the variance of the fundamental price σ_S^2. However, an agent using volatility estimator with time step h on the fundamental price process, will find an extra variance of approximately $h\mathcal{V}_t$ (on average). For a given fixed time step, the function $\mathcal{V}$ can thus be used as a proxy of the additional volatility of the equilibrium prices.

In this section we draw conclusions about the behavior of price volatility by analyzing this proxy, and in Section 3.5 we will show in numerical examples that the actual volatility, estimated from discrete observations of simulated market price exhibits similar behavior.

- First of all, since Δ^N_t is increasing in t, the observed price volatility increases at the approach of the delivery date in our model (see Figure 3.2). This phenomenon, well documented in electricity futures and other futures markets Jaeck and Lautier (2016) is known as the Samuelson effect and we also illustrate it empirically in Section 3.5.

- The observed price volatility is increasing in λ: stronger imbalance penalties lead to higher

volatility in the intraday market (see Figure 3.2). Moreover,

$$\lim_{\lambda\to\infty} \mathcal{V}_t = \frac{a^2\sigma_X^2}{\tilde{\alpha}^2}\int_0^t \frac{(\eta_{t-s}^N)^2}{(\Delta_{T-s}^N)^2}ds,$$

and the latter expression explodes for $t \to T$. Thus, we conclude that the Samuelson effect is also stronger for higher imbalance penalties.

- In the small liquidity cost regime ($\tilde{\alpha} \to 0$), for $0 < t < T$, $\Delta_{t-s}^N \to \frac{N}{(N-1)a}$ uniformly on $s \in [0, t]$. Therefore, for $N > 1$,

$$\mathcal{V}_t \sim \frac{a^2\lambda^2\sigma_X^2}{\tilde{\alpha}^2}\frac{(a+\lambda)^2N^2}{a^2(N-1)^2}\int_0^t (\eta_{t-s}^N)^2 ds \sim \frac{\lambda^2\sigma_X^2}{\tilde{\alpha}}\frac{(a+\lambda)^2}{2a}\frac{N}{N-1}$$

  This shows that with decreasing trading costs extra variance of the equilibrium price grows like $\frac{1}{\tilde{\alpha}}$. Lower transaction costs allow the agents to follow the forecasts more closely, leading to a higher volatility of the aggregate position and of the market price. On the other hand, since the function $N \mapsto \frac{N}{N-1}$ is decreasing in N, we conclude that price volatility in the small liquidity cost regime is decreasing with the number of agents: in our model, competition between agents increases market frictions and leads to reduced volatility.

- In the large liquidity cost regime ($\tilde{\alpha} \to \infty$), $\eta_t^N \to 1$ and $\Delta_t^N \sim \frac{t}{\tilde{\alpha}}$, so that

$$\mathcal{V}_t \sim \frac{a^2\lambda^2\sigma_X^2 t}{\tilde{\alpha}^2}.$$

  Higher liquidity costs decrease the trading rate of agents and lead to a lower overall market volatility.

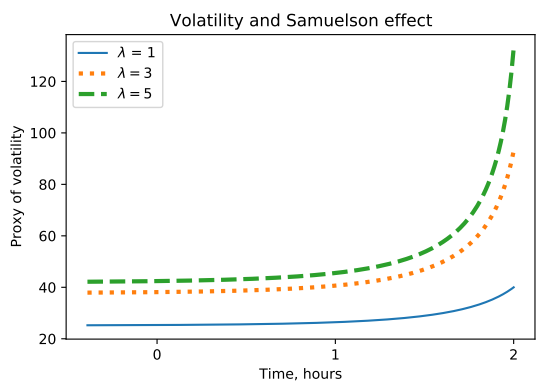

Figure 3.2: Proxy of expected observed volatility as function of time (delivery time is $t = 2$). Parameter values: $a = 1$, $N = 100$, $\alpha = 0.5$, $b = 0$, $\sigma_X = 50$.

3.4.3 Price-forecast covariance

To understand how the forecast updates influence prices, we compute the covariance of the increment of the aggregate strategy over an interval of length h with the increment of the aggregate forecast over the same interval. Using the explicit form of the strategy, we easily obtain,

$$\mathrm{Cov}[\bar{\phi}^N_{t+h} - \bar{\phi}^N_t, \overline{X}^N_{t+h} - \overline{X}^N_t] = \lambda(\sigma_X)^2 \int_t^{t+h} \frac{\Delta^N_{t+h-s}}{1 + (\frac{a}{N} + \lambda)\Delta^N_{T-s}} ds.$$

From this expression, we conclude that the covariance of equilibrium price with forecast updates increases when the terminal penalty λ increases, and when the time t approaches the delivery date.

3.4.4 Trading costs

In our model, the agents face three types of costs: the trading costs, the market impact costs, and the balancing costs. Using the martingale property of S and other assumptions of this section, these costs are evaluated as follows:

$$\begin{aligned}
C^{N,i}_{tra} &= \mathbb{E}\left[\int_0^T \frac{\alpha}{2}\dot{\phi}^i_t\Big(\dot{\phi}^i_t + b\dot{\bar{\phi}}^{N,-i}_t\Big)dt\right] \\
&= \frac{\alpha(1+b))}{2}\mathbb{E}\left[\int_0^T (\dot{\bar{\phi}}^N_t)^2 dt\right] + \frac{\alpha}{2}\frac{N-1-b}{N-1}\mathbb{E}\left[\int_0^T (\dot{\check{\phi}}^N_t)^2 dt\right] \\
C^{N,i}_{imp} &= \mathbb{E}\left[\int_0^T a\dot{\phi}^i_t(\bar{\phi}^N_t - \bar{\phi}^N_0)dt\right] = \frac{a}{2}\mathbb{E}\left[(\bar{\phi}^N_T - \bar{\phi}^N_0)^2\right] \\
C^{N,i}_{bal} &= \frac{\lambda}{2}\mathbb{E}\left[(\phi^i_T - X^i_T)^2\right] = \frac{\lambda}{2}\mathbb{E}\left[(\bar{\phi}^N_T - \overline{X}^N_T)^2\right] + \frac{\lambda}{2}\mathbb{E}\left[(\check{\phi}^i_T - \check{X}^i_T)^2\right]
\end{aligned}$$

After some tedious but straightforward computations, these costs are found to have the following integral form:

$$\begin{aligned}
C^{N,i}_{tra} &= \frac{1+b}{1+\frac{b}{2}}\frac{\lambda^2\sigma_X^2}{4}\int_0^T dt \frac{(1+\eta^N_{T-t})\Delta^N_{T-t}}{(1+(\frac{a}{N}+\lambda)\Delta^N_{T-t})^2} \\
&\quad + \frac{N-1-b}{N-1}\frac{\lambda^2\check{\sigma}_X^2}{4}\int_0^T dt \frac{(\tilde{\eta}^N_{T-t}+1)\widetilde{\Delta}^N_{T-t}}{(1+(\frac{a}{N}+\lambda)\widetilde{\Delta}^N_{T-t})^2} \\
C^{N,i}_{imp} &= \frac{a\sigma_X^2\lambda^2}{2}\int_0^T \frac{(\Delta^N_{T-t})^2}{(1+(\frac{a}{N}+\lambda)\Delta^N_{T-t})^2}dt \\
C^{N,i}_{bal} &= \frac{\lambda\sigma_X^2}{2}\int_0^T \frac{(1+\frac{a}{N}\Delta^N_{T-t})^2}{(1+(\frac{a}{N}+\lambda)\Delta^N_{T-t})^2}dt + \frac{\lambda\check{\sigma}_X^2}{2}\int_0^T \frac{(1+\frac{a}{N}\widetilde{\Delta}^N_{T-t})^2}{(1+(\frac{a}{N}+\lambda)\widetilde{\Delta}^N_{T-t})^2}dt,
\end{aligned}$$

which leads to the following conclusions:

- Trading costs are proportional to forecast variances: more precise forecasts lead to lower trading costs. However, while the trading costs and the balancing costs depend both on

the volatility of aggregate forecast and that of the individual forecast, the market impact costs only depend on the volatility of the aggregate forecast. Thus, an agent who has a better individual forecast will pay lower trading and balancing costs but the same market impact costs.

- Since both $C_{tra}^{N,i}$ and $C_{imp}^{N,i}$ are increasing in λ, stronger imbalance penalties lead to higher trading and market impact costs. The balancing cost $C_{bal}^{N,i}$ is increasing in λ for small values of λ, but may become decreasing for large λ. When $\lambda \to \infty$, the market impact costs and the balancing costs remain bounded, however it can be shown that the single agent trading cost tends to $+\infty$ at the rate of $\log \lambda$, thus very high imbalance penalties lead to prohibitive trading costs and are therefore detrimental for market liquidity.

- In the case of small liquidity costs ($\alpha \to 0$), each component of the cost converges to a nonzero limit:

$$
\begin{aligned}
C_{tra}^{N,i} &\to \frac{1+b}{1+\frac{b}{2}} \frac{\lambda^2 \sigma_X^2 T}{4} \frac{(N-1)a}{N(a+\lambda)^2} + \frac{N-1-b}{N-1} \frac{\lambda^2 \check{\sigma}_X^2 T}{4} \frac{aNT}{(a+\lambda N)^2} \\
C_{imp}^{N,i} &\to \frac{a\sigma_X^2 \lambda^2 T}{2(a+\lambda^2)} \\
C_{bal}^{N,i} &\to \frac{\lambda \sigma_X^2 T}{2} \frac{a^2}{(a+\lambda)^2} + \frac{\lambda \check{\sigma}_X^2 T}{2} \frac{a^2}{(a+\lambda N)^2}.
\end{aligned}
$$

As the cost per trade decreases, the agents trade more actively so that the overall trading cost does not tend to zero.

3.5 Empirical results and numerical illustrations

In this section our objective is to analyze the empirically observed features of intraday market prices, demonstrate that these features are reproduced by our model, and illustrate other properties of our model, such as the convergence of the N-agent model to the mean-field limit, with numerical examples.

3.5.1 Stylized features of intraday electricity market prices

A brief description of our dataset To compute the empirical price analyzed in the following sections, we used the limit order book data provided by EPEX electricity market for the Germany delivery zone for the 1st quarter of 2015 and January 2017. Although in the market it is possible to trade in quarter-hours, in this study we focus on the full hours only. The dataset contains full information about sell and buy orders recorded on any given day, whether they result in a transaction or not. From this data we reconstruct the state of the order book, which allows us in turn to derive the mid-quote price and the bid-ask spread.

Market liquidity In Figure 3.3, we plot the distribution of the times of orders and transactions as function of time to delivery computed over all orders and transactions in February

2015. We observe that the liquidity starts to appear only 5-6 hours before delivery, and grows very quickly at the approach of the delivery date.

Figure 3.3: Distribution of times orders and transactions in February 2015

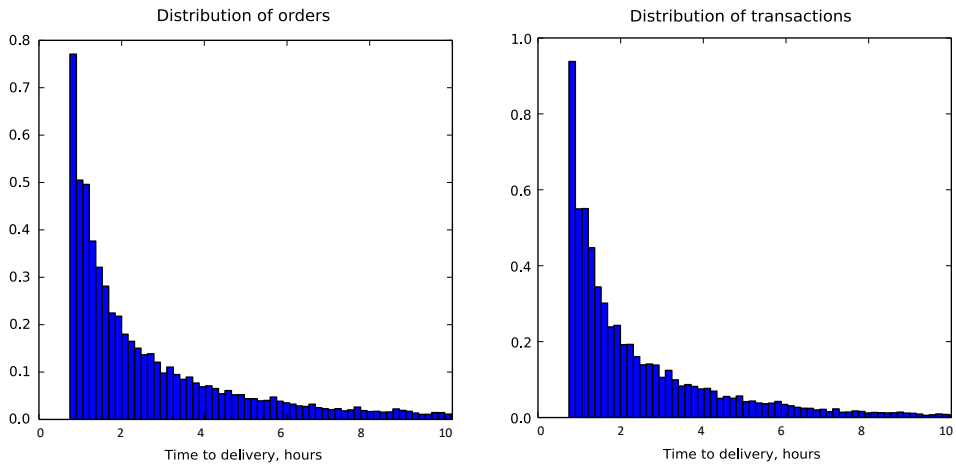

We also performed an estimation of the bid-ask spread in the German intraday electricity market for January 2017 for different delivery times. For each delivery time, Figure 3.4, shows that the spread averaged over each hour and over all days in January 2017 decreases as we approach the end of the trading period. This is consistent with the assumption that the market is used by the renewable energy producers to adjust their positions when precise forecasts become available.

Figure 3.4: Average spread per hour as a function of time to delivery over January 2017

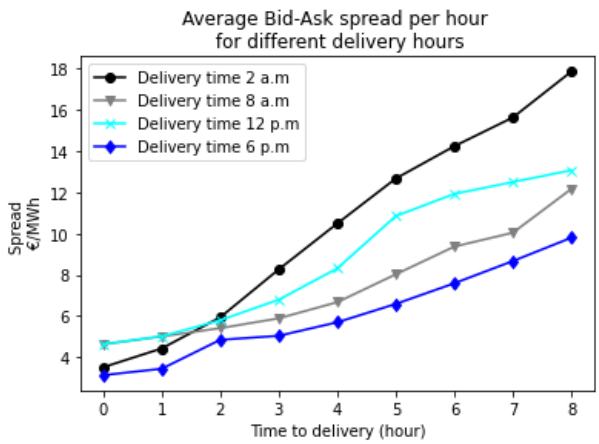

Price volatility To estimate the empirical volatility, we consider mid-quote prices reconstructed from the limit order book data of the Germany delivery zone for January 2017, as explained above. The mid-quote price was computed on a uniform grid with a time step of 1 minute. In January 2017 the market was already relatively liquid: the average number of daily price changes for a given delivery hour varied between approximately 3400 for the least liquid

delivery hour (2 AM to approximately 5800 for the most liquid delivery hour (6 PM). Given that, as we observed above, liquidity is concentrated in the last 5-6 hours, a one-minute interval during this time contains many price changes and the market microstructure effects are limited.

The observed midquote price is denoted $(\tilde{P}_t)_{t\in[0,T]}$. We denote by n the number of observations in the data of January 2017 and by $\{t_0, \ldots, t_i, \ldots, t_n\}$ the (uniform) time grid over which the observations are available. In contrast with the integrated volatility whose estimator is generally given by $\widehat{\int_0^T \sigma_s^2 ds} = \sum_{i=1}^n \Delta\tilde{P}_{t_{i-1}}^2$, estimating the instantaneous volatility is less straightforward. Following Kristensen (2010), we use a kernel-based non parametric estimator of the instantaneous volatility:

$$\widehat{\sigma}_t^2 = \frac{\sum_{i=1}^n K_h(t_{i-1} - t)\Delta\tilde{P}_{t_{i-1}}^2}{\sum_{i=1}^n K_h(t_{i-1} - t)(t_i - t_{i-1})}, \tag{3.16}$$

where $K(.)$ is the Epanechnikov kernel: $K(x) = \frac{3}{4}(1 - x^2)\mathbb{1}_{[-1,1]}(x)$ and $K_h(x) = \frac{1}{h}K(\frac{x}{h})$. The parameter h was taken equal to 0.08 hour ($\approx$ 5 minutes) after performing some cross-validation analyses and sensitivity tests. The paths of the estimated volatility as function of time to delivery for different delivery hours are given in Figure 3.5. We observe that the volatility increases as delivery time draws near and market participants trade more actively, giving an empirical evidence of the presence of the Samuelson effect in electricity market.

Figure 3.5: Instantaneous market volatility for different delivery hours

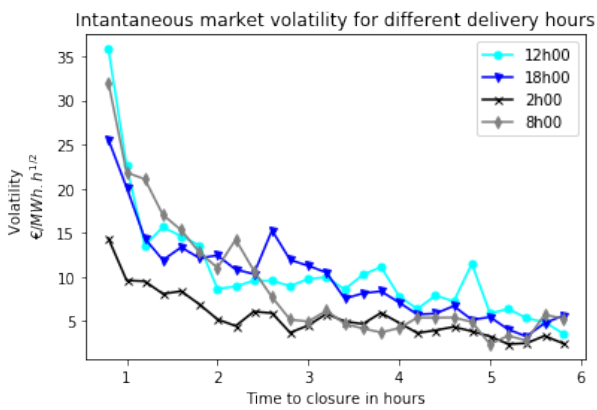

Correlation between price and renewable indeed forecasts We finally study the empirical correlation between the intraday market prices and the renewable wind production forecasts. Unlike the rest of the paper, here we use actual wind infeed forecasts, not the demand forecasts. To compute empirical correlation estimates, we use the limit order book data from the intraday EPEX market of the first three months of 2015 for the Germany delivery zone, from which, as before, we compute the mid-quote prices. The production forecasts correspond to the same period and are updated every 15 minutes for each delivery hour. In Figure 3.6, we plot the correlation between the increments of the market price and the increments of the production forecasts for the delivery time 12h (averaged over 90 days in the dataset), together with the

2-standard deviation bounds. To match the forecast update frequency, the mid-quote price is also sampled at 15-minute intervals here.

Figure 3.6: Correlation between the market price increments and the renewable production forecast increments for the German delivery zone in winter 2015

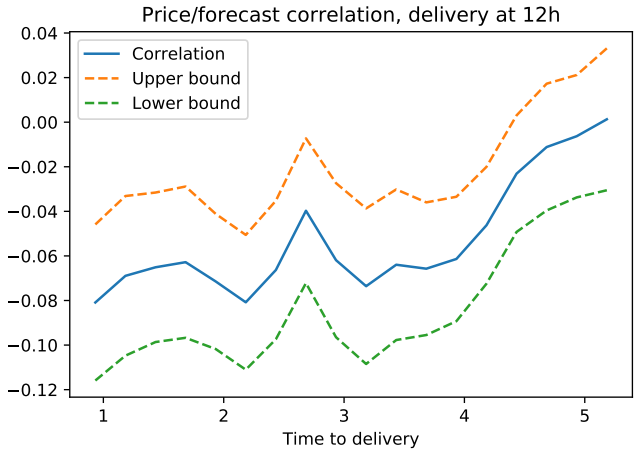

We find that the correlation between the price increments and those of the production forecast is negative and increases in absolute value as we approach the delivery time.

3.5.2 Numerical illustration of our model

Model specification We now define the dynamics for the fundamental price and for the demand forecasts used in the simulations. We also give the chosen values of the different parameters. Our objective here is to illustrate the features of the model and show that it reproduces the stylized facts of the market prices. Therefore, the majority of the parameters are not precisely estimated, but are given plausible values.

The evolution of the fundamental price is described as follows:

$$dS_t = \sigma_S dW_t \tag{3.17}$$

where σ_S is a constant and $(W_t)_{t\in[0,T]}$ is Brownian motion. We also assume that the liquidity function $\alpha(.)$ is given by

$$\alpha(t) \quad = \quad \alpha \times (T-t) + \beta, \quad \forall t \in [0,T] \tag{3.18}$$

where α and β are strictly positive constants. The liquidity function is decreasing with time. This assumption relies on the fact that, as we observed in Section 3.5.1, the market becomes more liquid as we get closer to the delivery time and it is less costly to trade when the market is liquid.

To simulate demand forecasts, we assume the following dynamics:

$$d\bar{X}_t^N = \sigma_X d\bar{B}_t \tag{3.19}$$

$$d\check{X}_t^i = \check{\sigma}_X dB_t^i, \qquad i \in \{1, \dots, N\} \tag{3.20}$$

where σ_X and $\check{\sigma}_X$ are constants and $(\bar{B}_t)_{t\in[0,T]}$, $(B_t^i)_{t\in[0,T]}$ are independent Brownian motions, also independent from $(W_t)_{t\in[0,T]}$.

In this illustration, we choose the same parameters for the dynamics of the common and the individual demand forecasts (that is, $\sigma_X = \check{\sigma}_X$). The common volatility is calibrated to wind energy forecasts in Germany over January 2015 during the last quotation hour, by using the classical volatility estimator

$$\sigma_X = \check{\sigma}_X = \frac{\sqrt{\Delta t}}{n'-1} \sum_{i=1}^{n'} Y_i^2 \tag{3.21}$$

with Δt the time step between two observations, $Y_i = X_{t_i} - X_{t_{i-1}}$ the increment between two successive observations and n' the total number of observed increments. As the forecasts are updated every 15 minutes, there are three increments during the last trading hour, available on each day from the 3rd of January to the 31th of January. Thus, for each delivery hour we dispose of $n' = 87$ increments points to estimate the volatility. The volatility, as well as the other model parameters are specified in Table 3.1.

Parameter	Value	Parameter	Value
S_0	40 €/MWh	a	1 €/MWh^2
σ_S	10 €/MWh·h$^{1/2}$	λ	100 €/MWh^2
$\overline{X}_0$, $\check{X}_0^i$	0 MWh	N	100
σ_X,$\check{\sigma}_X$	73 MWh/h$^{1/2}$	α	0.14 €/MW^2·h
b	0	β	0.06 €/MW^2

Table 3.1: Parameters of the model

Price trajectories. In Figure 3.7, we plot a simulated trajectory of the fundamental price S starting six hours before the delivery time (corresponding to $t = 0$), up to the time T of delivery, together with the market price P associated with the different settings studied in this paper: the N-player Nash equilibrium with $N = 100$ players, the mean field and the ϵ-Nash equilibrium. Graphs were all simulated with the same demand forecasts, initial values, volatilities and parameters as specified in Table 3.1.

In all settings, the model reflects the price impact of the positions taken by the agents. This price impact is influenced by the market price and the demand forecasts. If agents anticipate to have overestimated the demand (negative values of the demand process), there is an excess of supply in the market, thus the price impact is negative and the market price decreases. On the contrary, if they anticipate to have underestimated the demand (positive values of the demand forecast process), there is a lack of supply and the market price increases.

In Figure 3.7, the parameters are as specified in Table 3.1, with $b = 0$. To emphasize the

Figure 3.7: Model price trajectories (left) associated to a given common demand forecast trajectory (right) in different settings

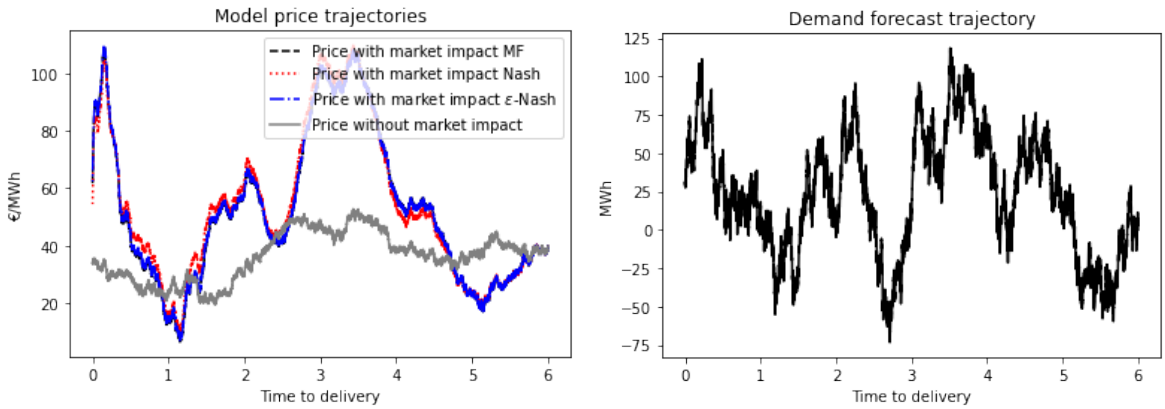

impact of crowd behavior on the cost of trading and the strategy, we display in Figure 3.8 the trajectory of the mean field at the equilibrium for different values of the parameter b. This is sufficient to capture the effect of a possible synchronization between agents since, from the discussion in Section 3.3.1, only the common part of the strategy is impacted by the effective cost $\alpha(1 + b/2)$.

Figure 3.8: Mean-field trajectories and demand forecasts as function of the impact of the crowd on the cost of trading

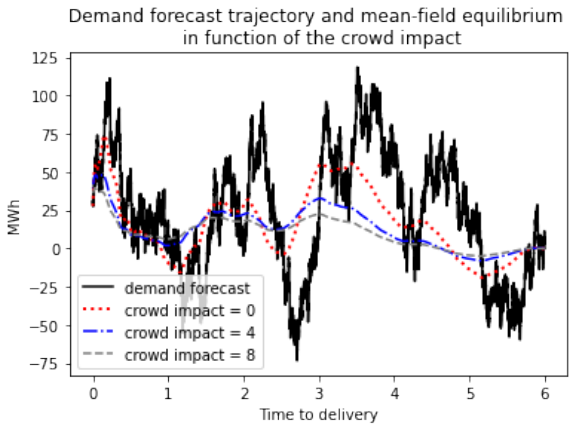

As we saw in Section 3.3.1, higher crowd trading parameter b leads to increased trading costs and therefore reduced volatility of the aggregate strategy, which therefore follows the forecast updates less closely.

Volatility and correlation In this paragraph we compare the price volatility and the correlation between price and renewable infeed forecasts in our model with the empirical ones. We have already seen through theoretical analysis in Section 3.4.2 that our model reproduces the observed features of the volatility; the goal of this paragraph is to confirm this using simulated prices. We once again highlight the fact that the market impact can induce an increase in the price variations but no changes in the quadratic variation since the price impact, though it is

stochastic, has a finite quadratic variation. However, the volatility estimated from discrete price observations, which is the only quantity relevant in practice, does increase in our model, as we shall see below.

We focus on hourly products and on several different delivery hours: 2 AM, 8 AM, 12 PM and 6 PM to include both peak (high electricity demand) and off-peak (low electricity demand) times. The volatility of the fundamental price S is assumed to be constant, ($\sigma_S = 10$ €/MWh·h$^{1/2}$) to ensure that the observed volatility changes are only due to the stochastic drift of the market price, i.e., the aggregate trading rates of the agents. The volatility of the production forecasts for the different delivery hours has been calibrated using the estimator defined in (3.21) and is shown in Table 3.2.

Hour	Volatility (MWh/h$^{1/2}$)
2h00	67
8h00	81
12h00	73
18h00	73

Table 3.2: Calibrated volatility of the production forecast for different delivery hours

During peak hours, both market activity and liquidity are higher. To account for this phenomenon in our model, we chose different levels of the liquidity coefficients α and β defined in (3.18) and presented in Table 3.3.

Hours \ Coefficients	α (€/h.MW^2)	β (€/MW^2)
2h00	0.24	0.10
8h00	0.10	0.04
12h00	0.06	0.02
18h00	0.14	0.06

Table 3.3: Liquidity coefficients used for different delivery hours

Since calibrating the model to market data is not the purpose of this study, we chose plausible values for these coefficients in an ad hoc manner with lower trading costs corresponding to delivery hours for which the market is more liquid. All other model parameters are specified in Table 3.1.

Figure 3.9 shows the estimated volatility of the simulated model price P in the Nash N-player game setting with $N = 100$, averaged over 1000 simulations. The volatility was computed using the estimator (3.16), with the same window width and time step as in the empirical analysis. From this graph we can see that the model is able to reproduce the increasing shape of the empirical market price volatility at the approach of the delivery time, and that it captures the different levels of volatility corresponding to the different delivery hours.

Figure 3.9: Simulated model volatility for different delivery hours

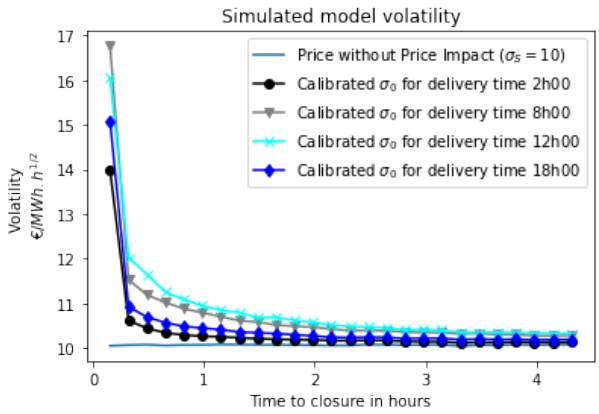

Correlation between price and renewable infeed An important stylized feature of intraday market prices, observed empirically in Kiesel and Paraschiv (2017) is the correlation between the price and the renewable production forecasts. Figure 3.10 plots the correlation between 15-minute increments of the simulated market price and the 15-minute increments of the simulated renewable production forecasts as function of time. For each time step, the correlation $\rho_t = \text{corr}(\Delta Y_t, \Delta P_t)$ is computed by Monte Carlo using the following estimator:

$$\widehat{\rho}_t = \frac{\sum_{k=1}^{N_{sim}} (\Delta Y_t^k - \overline{\Delta Y}_t)(\Delta P_t^k - \overline{\Delta P}_t)}{\sqrt{\sum_{k=1}^{N_{sim}} (\Delta Y_t^k - \overline{\Delta Y}_t)^2 \sum_{k=1}^{N_{sim}} (\Delta P_t^k - \overline{\Delta P}_t)^2}},$$

with N_{sim} stands for number of simulations (we considered $N_{sim} = 50000$), $\Delta Y_t^k = -(\overline{X}_{t+dt}^{N,k} - \overline{X}_t^{N,k})$, $\Delta P_t^k = P_{t+dt}^{N,k} - P_t^{N,k}$ and $N = 100$. Notice that we use the minus sign in front of the forecast increment to plot the correlation of production forecasts, whereas $\overline{X}$ stands for the demand forecast.

Figure 3.10: Correlation between the simulated market price increments and the renewable production forecast increments in the model during the last six hours of trading

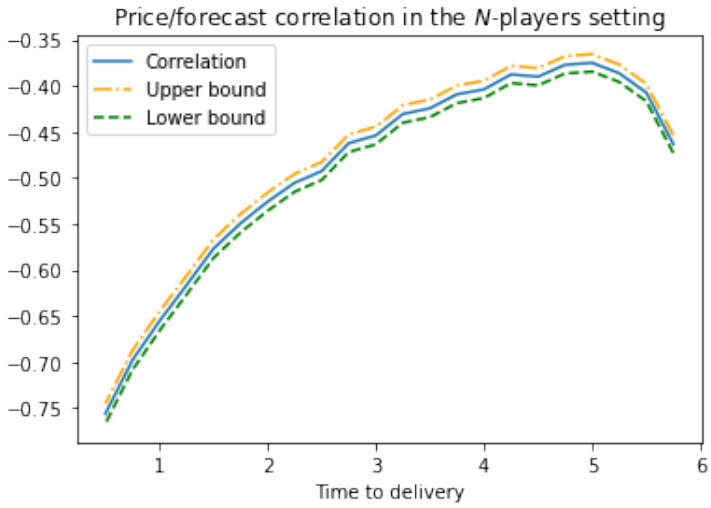

We first note that the correlation is negative: an expected increase of the renewable production is correlated to a decrease in the market price and an expected lack of renewable production is correlated to an increase in the price. As we get closer to the delivery date, the agents trade more actively as new forecast information becomes available, and the market price becomes more strongly dependent on the forecast updates. The model outputs qualitatively match the results observed empirically. However, the strength of the correlation seems to be greater in the model than in reality. This can be explained by the fact that the model does not take into account other renewable means of production such as the solar energy. The slight increase of the correlation for longer times to delivery (the right-hand side of the graph) may be explained by the fact that the correlation is computed as the ratio of the covariance to the square root of the product of variances. While both quantities decrease for longer times to delivery, the denominator may decrease faster, explaining the slight increase in the correlation values.

Convergence and approximations In Figure 3.11 we plot the mean field position, the aggregate N-player Nash equilibrium position and the aggregate position for the ϵ-Nash equilibrium (respectively given by Theorem 14, Theorem 17 and Proposition 20) for a model with $N = 5$ players and $N = 100$ players. The trajectories were computed with the same simulated fundamental price, common production forecast and parameters as the Figure 3.7 above, over the 6 hours preceding the delivery time. The left graph ($N = 5$) shows a big difference between the Nash equilibrium and ϵ-Nash approximation on one hand, and the mean field on the other hand. This is explained by the individual production forecast taken into account in the Nash and ϵ-Nash equilibria. When we consider a larger number of players, $N = 100$, the three position trajectories are much closer to each other. This confirms the asymptotic convergence to the mean field discussed in Section 3.3.3 for the N-player Nash equilibrium and ϵ-Nash equilibrium.

Figure 3.11: Aggregate position in different settings with $N = 5$ (left) and $N = 100$ (right) agents

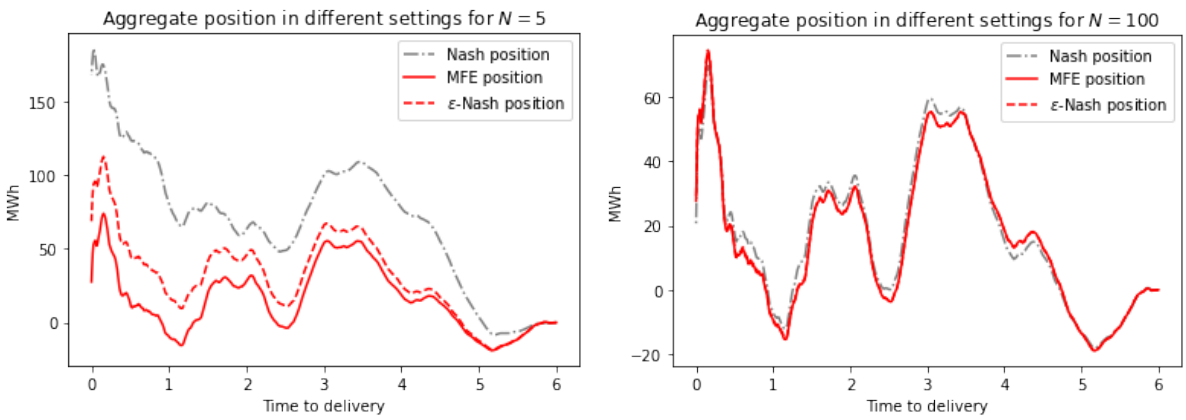

3.6 Appendix

3.6.1 Proof of Theorem 14

Proof. Step 1. First order condition of optimality for a single agent. In this step, we are going to show that for fixed ϕ^{-i*}, the strategy ϕ^{i*} satisfies (3.4) if and only if $\mathbb{E}[\phi_T^{i*} - X_T^i|\mathcal{F}_0] = 0$ and there exists a square integrable $\mathbb{F}$-martingale Y^i such that, almost surely,

$$\begin{aligned} &Y_t^i + \alpha(t)\Big\{\dot{\phi}_t^{i*}\Big(1 - \frac{b_N}{N}\Big) + b_N\dot{\bar{\phi}}_t^{N*}\Big\} \\ &\qquad + S_t + a(\bar{\phi}_t^{N*} - \bar{\phi}_0^{N*}) - \frac{a}{N}(\phi_t^{i*} - \phi_0^{i*}) = 0,\ 0 \le t \le T, \\ &Y_T^i = \frac{a}{N}(\phi_T^{i*} - \phi_0^{i*}) + \lambda(\phi_T^{i*} - X_T^i) - S_T, \end{aligned} \tag{3.22}$$

with the shorthand notation $b_N = \frac{N}{N-1}\frac{b}{2}$. Assume that ϕ^{i*} satisfies (3.4). Then, for any adapted square integrable process $(\nu_t)_{0\le t\le T}$, and for any $\delta \in \mathcal{F}_0$,

$$J^{N,i}(\phi^{i*} + \delta + \int_0^{\cdot} \nu_s ds, \phi^{-i*}) \le J^{N,i}(\phi^{i*}, \phi^{-i*}).$$

Developing the expressions, this is equivalent to

$$\begin{aligned} &\mathbb{E}\Big[\int_0^T \nu_t\Big\{\alpha(t)\Big[\dot{\phi}_t^{i*}\Big(1 - \frac{b_N}{N}\Big) + b_N\dot{\bar{\phi}}_t^{N*}\Big] + S_t + a(\bar{\phi}_t^{N*} - \bar{\phi}_0^{N*}) - \frac{a}{N}(\phi_t^{i*} - \phi_0^{i*})\Big\}dt \\ &+ \Big(\frac{a}{N}(\phi_T^{i*} - \phi_0^{i*}) + \lambda(\phi_T^{i*} - X_T^i) - S_T\Big)\int_0^T \nu_t dt + \lambda(\phi_T^{i*} - X_T^i)\delta\Big] \\ &+ \mathbb{E}\Big[\frac{1}{2}\int_0^T \alpha(t)\nu_t^2 dt + \Big(\frac{a}{N} + \frac{\lambda}{2}\Big)\Big(\int_0^T \nu_t dt\Big)^2 + \frac{\lambda}{2}\delta^2\Big] \ge 0, \end{aligned}$$

and since ν and δ are arbitrary, we see that optimality is equivalent to

$$\begin{aligned} &\mathbb{E}\Big[\int_0^T \nu_t\Big\{\alpha(t)\Big[\dot{\phi}_t^{i*}\Big(1 - \frac{b_N}{N}\Big) + b_N\dot{\bar{\phi}}_t^{N*}\Big] + S_t + a(\bar{\phi}_t^{N*} - \bar{\phi}_0^{N*}) - \frac{a}{N}(\phi_t^{i*} - \phi_0^{i*})\Big\}dt \\ &+ \Big(\frac{a}{N}(\phi_T^{i*} - \phi_0^{i*}) + \lambda(\phi_T^{i*} - X_T^i) - S_T\Big)\int_0^T \nu_t dt\Big] = 0, \end{aligned} \tag{3.23}$$

for any adapted square integrable ν, together with the condition that $\mathbb{E}[\phi_T^{i*} - X_T^i|\mathcal{F}_0] = 0$. Now, assume that Y^i is a square integrable martingale satisfying (3.22). Then, by integration by parts, the expression in the previous line equals

$$\mathbb{E}\Big[-\int_0^T \nu_t Y_t dt + Y_T\int_0^T \nu_t dt\Big] = \mathbb{E}\Big[\int_0^T \Big(\int_0^t \nu_s ds\Big)dY_t\Big] = 0,$$

and we see that the optimality condition is satisfied. Conversely, assume that (3.23) is satisfied for any adapted square integrable process ν, and let Y^i be a martingale such that

$$Y_T^i = \frac{a}{N}(\phi_T^{i*} - \phi_0^{i*}) + \lambda(\phi_T^{i*} - X_T^i) - S_T.$$

Then, by integration by parts, (3.23) is equivalent to

$$\mathbb{E}\Big[\int_0^T \nu_t\Big\{\alpha(t)\Big[\dot{\phi}_t^{i*}\Big(1-\frac{b_N}{N}\Big)+b_N\dot{\bar{\phi}}_t^{N*}\Big] + S_t + a(\bar{\phi}_t^{N*}-\bar{\phi}_0^{N*}) - \frac{a}{N}(\phi_t^{i*}-\phi_0^{i*}) + Y_t^i\Big\}dt\Big] = 0,$$

and since ν is arbitrary, we see that (3.22) is satisfied.

Step 2. Computing the average position. Let $(\phi^{i*})_{i=1,\dots,N}$ be a Nash equilibrium. We have seen that this is equivalent to (3.22) together with the condition that $\mathbb{E}[\phi_T^{i*}-X_T^i|\mathcal{F}_0]=0$ for $i=1,\dots,N$. Summing up these expressions for $i=1,\dots,N$ and denoting $\overline{Y}_t^N=\frac{1}{N}\sum_{i=1}^N Y_t^i$, we get

$$\begin{aligned}
&\overline{Y}_t^N + \alpha(t)\Big(1+\frac{b}{2}\Big)\dot{\bar{\phi}}_t^{N*} + S_t + a\frac{N-1}{N}(\bar{\phi}_t^{N*}-\bar{\phi}_0^{N*}) = 0, \quad 0\le t\le T,\\
&\overline{Y}_T^N = \frac{a}{N}(\bar{\phi}_T^{N*}-\bar{\phi}_0^{N*}) + \lambda(\bar{\phi}_T^{N*}-\overline{X}_T^N) - S_T,\\
&\mathbb{E}[\bar{\phi}_T^{N*}-\overline{X}_T^N|\mathcal{F}_0] = 0.
\end{aligned}$$

The first equation can be solved explicitly for $\bar{\phi}^{N*}$:

$$\bar{\phi}_t^{N*} = \bar{\phi}_0^{N*} - \int_0^t \eta_{s,t}^N \frac{\overline{Y}_s^N + S_s}{\alpha(s)\big(1+\frac{b}{2}\big)}ds. \tag{3.24}$$

Denoting $\widehat{\phi}_t := \bar{\phi}_t^{N*} + I_t^N$, we obtain simplified equations:

$$\begin{aligned}
\widehat{\phi}_t &= \bar{\phi}_0^{N*} - \int_0^t \eta_{s,t}^N \frac{\overline{Y}_s^N}{\alpha(s)\big(1+\frac{b}{2}\big)}ds,\\
\overline{Y}_T^N &= (\frac{a}{N}+\lambda)(\widehat{\phi}_T - I_T^N) - \frac{a}{N}\bar{\phi}_0^{N*} - \lambda\overline{X}_T^N - S_T.
\end{aligned}$$

Substituting $\widehat{\phi}_T$ into the second equation and taking the expectation, we obtain another linear equation, this time for $\overline{Y}_t^N$:

$$\begin{aligned}
\overline{Y}_T^N &= -(\frac{a}{N}+\lambda)\int_0^T \eta_{s,T}^N \frac{\overline{Y}_s^N}{\alpha(s)\big(1+\frac{b}{2}\big)}ds - (\frac{a}{N}+\lambda)I_T - \lambda\overline{X}_T^N + \lambda\bar{\phi}_0^{N*} - S_T,\\
\overline{Y}_t^N &= -\Big(\frac{a}{N}+\lambda\Big)\int_0^t \overline{Y}_s^N \frac{\eta_{s,T}^N}{\alpha(s)\big(1+\frac{b}{2}\big)}ds - \Big(\frac{a}{N}+\lambda\Big)\Delta_{t,T}^N\overline{Y}_t^N\\
&\quad - \Big(\frac{a}{N}+\lambda\Big)\widetilde{I}_t^N - \lambda\overline{X}_t^N + \lambda\bar{\phi}_0^{N*} - \widetilde{S}_t.
\end{aligned}$$

By integration by parts, this is equivalent to

$$\begin{aligned}\overline{Y}_t^N =& -\left(\frac{a}{N}+\lambda\right)\int_0^t \Delta_{s,T}^N d\overline{Y}_s^N - \left(\frac{a}{N}+\lambda\right)\Delta_{0,T}^N \overline{Y}_0^N \\ & -\left(\frac{a}{N}+\lambda\right)\widetilde{I}_t - \lambda \overline{X}_t^N + \lambda \bar{\phi}_0^{N*} - \widetilde{S}_t.\end{aligned}$$

Taking $t=0$, we get:

$$\overline{Y}_0^N = \frac{-\left(\frac{a}{N}+\lambda\right)\widetilde{I}_0^N - \lambda(\overline{X}_0^N - \bar{\phi}_0^{N*}) - \widetilde{S}_0}{1+\left(\frac{a}{N}+\lambda\right)\Delta_{0,T}^N}$$

On the other hand, in differential form,

$$\left\{1+\left(\frac{a}{N}+\lambda\right)\Delta_{t,T}^N\right\} d\overline{Y}_t^N = -\left(\frac{a}{N}+\lambda\right) d\widetilde{I}_t^N - \lambda d\overline{X}_t^N - d\widetilde{S}_t,$$

which is solved therefore explicitly by

$$\overline{Y}_t^N = \frac{-\left(\frac{a}{N}+\lambda\right)\widetilde{I}_0^N - \lambda(\overline{X}_0^N - \bar{\phi}_0^{N*}) - \widetilde{S}_0}{1+\left(\frac{a}{N}+\lambda\right)\Delta_{0,T}^N} - \int_0^t \frac{\left(\frac{a}{N}+\lambda\right) d\widetilde{I}_s^N + \lambda d\overline{X}_s^N + d\widetilde{S}_t}{1+\left(\frac{a}{N}+\lambda\right)\Delta_{s,T}^N} \tag{3.25}$$

Finally

$$\begin{aligned}\bar{\phi}_t^{N*} &= \bar{\phi}_0^{N*} - I_t^N + \int_0^t \overline{Y}_s^N d\Delta_{s,t}^N = \bar{\phi}_0^{N*} - I_t^N - \overline{Y}_0^N \Delta_{0,t}^N - \int_0^t \Delta_{s,t}^N d\overline{Y}_s^N \\ &= \bar{\phi}_0^{N*} - I_t^N + \Delta_{0,t}^N \frac{\left(\frac{a}{N}+\lambda\right)\widetilde{I}_0^N + \lambda(\overline{X}_0^N - \bar{\phi}_0^{N*}) + \widetilde{S}_0}{1+\left(\frac{a}{N}+\lambda\right)\Delta_{0,T}^N} \\ &+ \int_0^t \Delta_{s,t}^N \frac{\left(\frac{a}{N}+\lambda\right) d\widetilde{I}_s^N + \lambda d\overline{X}_s^N + d\widetilde{S}_t}{1+\left(\frac{a}{N}+\lambda\right)\Delta_{s,T}^N}.\end{aligned} \tag{3.26}$$

It remains to compute $\bar{\phi}_0^{N*}$ from the condition $\mathbb{E}[\bar{\phi}_T^{N*} - \overline{X}_T^N | \mathcal{F}_0] = 0$. Substituting this into the above expression, we find

$$\bar{\phi}_0^{N*} = \overline{X}_0^N + \frac{\widetilde{I}_0^N - \Delta_{0,T}^N \widetilde{S}_0}{1+\frac{a}{N}\Delta_{0,T}^N}$$

so that

$$\begin{aligned}\bar{\phi}_t^{N*} &= \overline{X}_0^N + \frac{1+\frac{a}{N}\Delta_{0,t}^N}{1+\frac{a}{N}\Delta_{0,T}^N}(\widetilde{I}_0^N - \Delta_{0,T}^N \widetilde{S}_0) \\ &\quad - (I_t^N - \Delta_{0,t}^N \widetilde{S}_0) + \int_0^t \Delta_{s,t}^N \frac{\left(\frac{a}{N}+\lambda\right) d\widetilde{I}_s^N + \lambda d\overline{X}_s^N + d\widetilde{S}_t}{1+\left(\frac{a}{N}+\lambda\right)\Delta_{s,T}^N}.\end{aligned} \tag{3.27}$$

<u>Step 3: computing the position of the agent.</u> Let $\check{\phi}_t^{i*} := \phi_t^{i*} - \bar{\phi}_t^{N*}$, $\check{X}_t^i = X_t^i - \overline{X}_t^N$ and $\check{Y}_t^i := Y_t^i - \overline{Y}_t^N$. Then, $\check{Y}^i$ is an $\mathbb{F}$-martingale and satisfies

$$\check{Y}_T^i = \frac{a}{N}\check{\phi}_T^{i*} + \lambda(\check{\phi}_T^{i*} - \check{X}_T^i), \qquad \check{Y}_t^i = -\alpha(t)\dot{\check{\phi}}_t^{i*} + \frac{a}{N}\check{\phi}_t^{i*}.$$

together with the additional condition $\mathbb{E}[\check{\phi}_T^{i*} - \check{X}_T^i|\mathcal{F}_t] = 0$. Similarly to the second part, this system admits an explicit solution:

$$\check{Y}_t^i = -\frac{\lambda(\check{X}_0^i - \check{\phi}_0^{i*})}{1 + \left(\frac{a}{N} + \lambda\right)\widetilde{\Delta}_{0,T}^N} - \int_0^t \frac{\lambda d\check{X}_s^i}{1 + \left(\frac{a}{N} + \lambda\right)\widetilde{\Delta}_{s,T}^N}.$$

and

$$\check{\phi}_t^{i*} = \check{X}_0^i + \int_0^t \widetilde{\Delta}_{s,t}^N \frac{\lambda d\check{X}_s^i}{1 + \left(\frac{a}{N} + \lambda\right)\widetilde{\Delta}_{s,T}^N}.$$

□

Remark 22. *The optimal strategy is obtained by solving a series of linear equations from an equivalent characterization of optimality. Since all equations admit unique solutions, the equilibrium strategy is unique. This is a consequence of the strict concavity of the objective function in our linear quadratic setting.*

3.6.2 Proofs of Propositions 19 and 20

Proof of Proposition 19. For all $t \in [0, T]$, we define:

$$g_{s,t}^N = \frac{\Delta_{s,t}^N}{(1 + (\frac{a}{N} + \lambda)\Delta_{s,T}^N)}, \quad g_{s,t} = \frac{\Delta_{s,t}}{(1 + \lambda\Delta_{s,T})},$$
$$\tilde{g}_{s,t}^N = \frac{\widetilde{\Delta}_{s,t}^N}{(1 + (\frac{a}{N} + \lambda)\widetilde{\Delta}_{s,T}^N)}, \quad \tilde{g}_{s,t} = \frac{\widetilde{\Delta}_{s,t}}{(1 + \lambda\widetilde{\Delta}_{s,T})},$$

so that, for some constant C depending only on the parameters a, b, α and λ, but not on other ingredients of the model,

$$|\Delta_{s,t}^N - \Delta_{s,t}| + |g_{s,t}^N - g_{s,t}| + |\tilde{g}_{s,t}^N - \tilde{g}_{s,t}| \leq \frac{C}{N}.$$

for all $s, t \in [0, T]$. Now, let us consider the optimal strategies $(\phi_t^{i*})_{t\in[0,T]}$ and $(\phi_t^{MF,i*})_{t\in[0,T]}$ of the generic agent i respectively in the N-player setting and the mean field setting. Fix $t \in [0, T]$. Then,

$$\begin{aligned}
\phi_t^{i*} - \phi_t^{MF,i*} &= \frac{a}{N}\frac{\Delta_{0,t}^N - \Delta_{0,T}^N}{1 + \frac{a}{N}\Delta_{0,T}^N}(\widetilde{I}_0^N - \widetilde{S}_0\Delta_{0,T}^N) + \widetilde{I}_0^N - \widetilde{I}_0 + \widetilde{S}_0(\Delta_{0,T} - \Delta_{0,T}^N) \\
&- I_t^N + I_t + \widetilde{S}_0(\Delta_{0,t}^N - \Delta_{0,t}) + \int_0^t (g_{s,t}^N - g_{s,t})\left\{\left(\frac{a}{N} + \lambda\right) d\widetilde{I}_s^N + \lambda d\overline{X}_s^N + d\widetilde{S}_s\right\} \\
&+ \lambda\int_0^t g_t(s)d(\widetilde{I}_s^N - \widetilde{I}_s + \overline{X}_s^N - \overline{X}_s) \\
&+ \int_0^t \lambda(\tilde{g}_{s,t}^N - \tilde{g}_{s,t})d(X_s^i - \overline{X}_s^N) + \lambda\int_0^t \tilde{g}_{s,t}d(\overline{X}_s - \overline{X}_s^N)
\end{aligned}$$

Therefore, for some constant C depending only on the parameters a, b, α and λ, but not on other ingredients of the model,

$$\begin{aligned}
\mathbb{E}[(\phi_t^{i*} - \phi_t^{MF,i*})^2] &\leq \frac{C}{N^2}\mathbb{E}[(\widetilde{I}_0^N)^2] + \frac{C}{N^2}\mathbb{E}[(\widetilde{S}_0)^2] + \mathbb{E}[(\widetilde{I}_0^N - \widetilde{I}_0)^2] + \mathbb{E}[(I_t^N - I_t)^2] \\
&+ \frac{C}{N^2}\mathbb{E}[(\widetilde{I}_t^N)^2] + C\mathbb{E}[(\widetilde{I}_t^N - \widetilde{I}_t)^2] + \frac{C}{N^2}\mathbb{E}[\widetilde{S}_t^2] \\
&+ \frac{C}{N^2}\mathbb{E}[(\overline{X}_t^N)^2] + C\mathbb{E}[(\overline{X}_t^N - \overline{X}_t)^2] + \frac{C}{N^2}\mathbb{E}[(\check{X}_t^i)^2] \\
&\leq \mathbb{E}[(I_t^N - I_t)^2] + \frac{C}{N^2}\mathbb{E}[(I_T^N)^2] + C\mathbb{E}[(I_T^N - I_T)^2] \\
&+ \frac{C}{N^2}\mathbb{E}[S_T^2] + \frac{C}{N^2}\mathbb{E}[(\overline{X}_t)^2] + \frac{C}{N^2}\sum_{i=1}^N \mathbb{E}[(\check{X}_t^i)^2] \\
&\leq \frac{C}{N^2}\mathbb{E}[\sup_{0\leq s\leq T} S_s^2] + \frac{C}{N^2}\mathbb{E}[(\overline{X}_t)^2] + \frac{C}{N}\mathbb{E}[(\check{X}_t^i)^2].
\end{aligned}$$

where the estimate for the first line above is obtained through Jensen's inequality. The other estimates of the proposition are obtained in a similar way.

Proof of Proposition 20. To lighten notation, and since we now have only one strategy, we omit in this proof the superscript MF in the candidate strategy $\phi^{MF,i*}$. The "distance to optimality" for this strategy is estimated as follows.

$$\begin{aligned}
&J^{N,i}(\phi^i, \phi^{-i*}) - J^{N,i}(\phi^{i*}, \phi^{-i*}) \\
&= J^{N,i}(\phi^i, \phi^{-i*}) - J^{MF}(\phi^i, \bar{\phi}^*) + J^{MF}(\phi^i, \bar{\phi}^*) - J^{MF}(\phi^{i*}, \bar{\phi}^*) \\
&\quad + J^{MF}(\phi^{i^*}, \bar{\phi}^*) - J^{N,i}(\phi^{i*}, \phi^{-i*}) \\
&\leq J^{N,i}(\phi^i, \phi^{-i*}) - J^{MF}(\phi^i, \bar{\phi}^*) + J^{MF}(\phi^{i^*}, \bar{\phi}^*) - J^{N,i}(\phi^{i*}, \phi^{-i*}).
\end{aligned}$$

The second difference is estimated as follows:

$$\begin{aligned}
&J^{MF}(\phi^{i^*}, \bar{\phi}^*) - J^{N,i}(\phi^{i*}, \phi^{-i*}) \\
&= -\mathbb{E}\Big[\int_0^T \dot{\phi}_t^{i*}\Big\{a\big(\bar{\phi}_t^* - \bar{\phi}_t^{N*} - \bar{\phi}_0^* + \bar{\phi}_0^{N*}\big) + \frac{\alpha(t)b}{2}\big(\dot{\bar{\phi}}_t^* - \dot{\bar{\phi}}_t^{N,-i*}\big)\Big\}dt\Big]
\end{aligned} \tag{3.28}$$

Since

$$\dot{\bar{\phi}}_t^* - \dot{\bar{\phi}}_t^{N,-i*} = \dot{\bar{\phi}}_t^* - \dot{\bar{\phi}}_t^{N*} + \frac{\dot{\phi}^{i*} - \dot{\bar{\phi}}^{N*}}{N-1}, \tag{3.29}$$

it follows from the Cauchy-Schwarz inequality and Proposition 19 that

$$|J^{MF}(\phi^{i^*}, \bar{\phi}^*) - J^{N,i}(\phi^{i*}, \phi^{-i*})| \leq \frac{C}{\sqrt{N}}. \tag{3.30}$$

The first difference admits the following estimate.

$$
\begin{aligned}
&J^{N,i}(\phi^i,\phi^{-i*}) - J^{MF}(\phi^i,\bar{\phi}^*) = a\mathbb{E}\Big[\int_0^T \dot{\phi}_t^i(\bar{\phi}_t^* - \bar{\phi}_t^{N*} - \bar{\phi}_0^* + \bar{\phi}_0^{N*})dt\Big] \\
&- \frac{a}{N}\mathbb{E}\Big[\int_0^T \dot{\phi}_t^i(\phi_t^i - \phi_t^{i*} - \phi_0^i + \phi_0^{i*})dt\Big] + b\mathbb{E}\Big[\int_0^T \alpha(t)\dot{\phi}_t^i(\dot{\bar{\phi}}_t^* - \dot{\bar{\phi}}_t^{N,-i*})dt\Big] \\
&\le \mathbb{E}\Big[\int_0^T (\dot{\phi}_t^i)^2 dt\Big]^{\frac{1}{2}}\Big\{a\mathbb{E}\Big[\int_0^T (\bar{\phi}_t^* - \bar{\phi}_t^{N*} - \bar{\phi}_0^* + \bar{\phi}_0^{N*})^2 dt\Big]^{\frac{1}{2}} \\
&+ \frac{a}{N}\mathbb{E}\Big[\int_0^T (\phi_t^{i*} - \phi_0^{i*})^2 dt\Big]^{\frac{1}{2}} + \frac{b}{N-1}\mathbb{E}\Big[\int_0^T \alpha^2(t)(\dot{\phi}_t^{i*} - \dot{\bar{\phi}}^{N*})^2 dt\Big]^{\frac{1}{2}} \\
&+ b\mathbb{E}\Big[\int_0^T \alpha^2(t)(\dot{\bar{\phi}}_t^* - \dot{\bar{\phi}}_t^{N,*})^2\Big]^{\frac{1}{2}}\Big\} \le \frac{C_0}{\sqrt{N}}\mathbb{E}\Big[\int_0^T (\dot{\phi}_t^i)^2 dt\Big]^{\frac{1}{2}}
\end{aligned}
\tag{3.31}
$$

for some constant $C_0 < \infty$, which does not depend on ϕ^i, in view of Proposition 19 and (3.29). On the other hand, the following estimate also holds true.

$$
\begin{aligned}
&J^{N,i}(\phi^i,\phi^{-i*}) = -\mathbb{E}\Big[\int_0^T \Big\{\frac{\alpha(t)}{2}\dot{\phi}_t^i\Big(\dot{\phi}_t^i + b\dot{\bar{\phi}}_t^{N,-i*}\Big) + \dot{\phi}_t^i(S_t + a(\bar{\phi}_t^{N*} - \bar{\phi}_0^{N*})\Big\}dt \\
&+ \frac{a}{N}\int_0^T \dot{\phi}_t^i(\phi_t^i - \phi_0^i - \phi_t^{i*} + \phi_0^{i*})dt + \phi_0^i S_0 - (\phi_T^i - X_T^i)S_T + \frac{\lambda}{2}(\phi_T^i - X_T^i)^2\Big], \\
&\le -\frac{\bar{\alpha}}{2}\mathbb{E}\Big[\int_0^T (\dot{\phi}_t^i)^2 dt\Big] - \frac{\lambda}{2}\mathbb{E}[(\phi_T^i)^2] + C_1\mathbb{E}\Big[\int_0^T (\dot{\phi}_t^i)^2 dt\Big]^{\frac{1}{2}} + C_2\mathbb{E}[(\phi_T^i)^2]^{\frac{1}{2}} + C_3,
\end{aligned}
\tag{3.32}
$$

where $\bar{\bar{\alpha}} = \max_{0\le t\le T}\alpha(t)$ and

$$
\begin{aligned}
C_1 &= \frac{b}{2}\mathbb{E}\Big[\int_0^T \alpha^2(t)(\dot{\bar{\phi}}_t^{N,-i*})^2 dt\Big]^{\frac{1}{2}} + a\mathbb{E}\Big[\int_0^T (\bar{\phi}_t^{N*} - \bar{\phi}_0^{N*})^2 dt\Big]^{\frac{1}{2}} \\
&+ \frac{a}{N}\mathbb{E}\Big[\int_0^T (\bar{\phi}_t^{i*} - \bar{\phi}_0^{i*})^2 dt\Big]^{\frac{1}{2}} + \mathbb{E}\Big[\int_0^T (S_t - S_0)^2\Big]^{\frac{1}{2}}, \\
C_2 &= \mathbb{E}[(S_T - S_0)^2]^{\frac{1}{2}} + 2\mathbb{E}[(X_T^i)^2]^{\frac{1}{2}}, \qquad C_3 = |\mathbb{E}[X_T^i S_T]|.
\end{aligned}
$$

Thus, there exists a constant $C^* < \infty$, which does not depend on ϕ^i, such that if

$$
\mathbb{E}\Big[\int_0^T (\dot{\phi}_t^i)^2 dt\Big] + \mathbb{E}[(\phi_T^i)^2] > C^*,
$$

then

$$
J^{N,i}(\phi^i,\phi^{-i*}) - J^{N,i}(\phi^{i*},\phi^{-i*}) < 0.
$$

Therefore, from (3.30) and (3.31) it follows that for any admissible strategy ϕ^i,

$$
\begin{aligned}
&J^{N,i}(\phi^i,\phi^{-i*}) - J^{N,i}(\phi^{i*},\phi^{-i*}) \\
&\le \mathbf{1}_{\mathbb{E}[\int_0^T (\dot{\phi}_t^i)^2 dt] + \mathbb{E}[(\phi_T^i)^2] \le C^*}\Big\{\frac{C}{\sqrt{N}} + \frac{C_0}{\sqrt{N}}\mathbb{E}[\int_0^T (\dot{\phi}_t^i)^2 dt]^{\frac{1}{2}}\Big\} \\
&\le \frac{C}{\sqrt{N}} + \frac{C_0\sqrt{C^*}}{\sqrt{N}}.
\end{aligned}
$$

Chapter 4

Price formation and optimal trading in intraday electricity markets with a major player

Abstract

We study price formation in intraday electricity markets in the presence of intermittent renewable generation. We consider the setting where a major producer may interact strategically with a large number of small producers. Using stochastic control theory we identify the optimal strategies of agents with market impact and exhibit the Nash equilibrium in closed form in the asymptotic framework of mean field games with a major player. This is a companion paper to Féron et al. (2020), where a similar model is developed in the setting of identical agents.

4.1 Introduction

The structure of electricity markets around the world has been profoundly transformed by the push towards liberalization in the late 90s and more recently by the massive arrival of renewable energy production. Distribution has been separated from production, and whereas in the past, a single producer could own the entire generation capacity of a given country or region, now a patchwork of small, often renewable, generators competes with a big historical producer.

The aim of this paper is to develop an equilibrium model for intraday electricity markets where a big producer with a significant market share competes with a large number of small renewable producers. Both the large producer and the small producers use the intraday markets to compensate their production and demand forecast errors, creating feedback effects on the market price. The large producer can act strategically, anticipating the impact of its decisions on the market prices and thus on the behavior of the small agents. The small agents are not strategic, and each one has a negligible effect on the market, however the behavior of all small agents taken together has a significant market impact. The large player has the first-mover advantage but does not observe the forecast of the minor players. These in turn have the information advantage since they observe the forecast of the major player as well as their own forecast. This leads to a stochastic leader-follower game where players interact through the market price. We place ourselves in the linear-quadratic setting, exhibit the unique Nash equilibrium for this game in closed form in the framework of mean field games with a major player, and provide explicit formulas for the market price and the strategies of the agents. For a game with a finite number of players, we show how an ε-Nash equilibrium can be constructed from the mean field game solution.

This paper is a companion paper to Féron et al. (2020), where a similar model is developed for the case of identical agents with symmetric interactions, and we refer the readers to that paper for a detailed review of literature on the stochastic and econometric modeling of intraday electricity markets. Here we simply mention that a similar linear-quadratic setting with linear market impact, has been used to determine optimal strategies for a single energy producer by Aid et al. (2016) and Tan and Tankov (2018), while Bouchard et al. (2018) found an analytic expression for the equilibrium price in a linear-quadratic model of the stock market with symmetric interactions and perfect information. We also mention the recent paper of Aid et al. (2020) where an equilibrium in complete information setting for a finite number of agents is derived in the intraday electricity market. This paper is close in spirit to the complete information framework of Féron et al. (2020), but allows to treat the case of heterogeneous agents in conditions of uncertain production with possible outages and uncertain demand. However, the complete information setting, where each agent observes all other agents' forecasts does not seem realistic in electricity markets. The incomplete information setting, where each agent only observes its own forecast and the aggregate forecast, may not be tractable for a finite number of agents. Nevertheless, in Féron et al. (2020), it has been shown that explicit solutions may be found in the mean field limit, where the number of agents is sent to infinity, and the influence of every single agent on the entire market becomes negligible. The use of mean field theory for stochastic control with

partial information has also recently been proposed by Bensoussan and Yam (2019), in a formal fashion, to solve the associated Zakai equation.

The mean field games (MFG) are stochastic differential games with infinitely many players and symmetric interactions. The seminal papers of Lasry and Lions (2007) and Huang et al. (2006) characterized the Nash equilibrium in this framework through a coupled system of a Hamilton-Jacobi-Bellman (HJB) and a Fokker-Planck (FP) equation. Carmona and Delarue (2018) developed an alternative probabilistic approach inspired by the Pontryagin principle and related the mean field game solution to a McKean-Vlasov Forward Backward Stochastic Differential Equation (FBSDE). The asymptotic results obtained in mean field games can be used to construct approximate equilibria (ε-Nash equilibria) for games with a finite number of players. Alternatively, equilibria of N-player games can be shown to converge to the corresponding weak mean field equilibria Lacker (2020).

While the original MFG setting involves symmetric agents, Huang (2010) introduced linear-quadratic mean field games with a major player. Nourian and Caines (2013) developed this approach in a general framework. In both papers, the mean field is exogenous to the actions of the major player. In contrast to these two papers, Bensoussan et al. (2016a), and Carmona et al. (2016), considered the endogenous case where the major player can influence the mean field. In Bensoussan et al. (2016a), this leads to a leader-follower setting, also known as Stackelberg game. The authors derived a HJB equation and a FP equation to characterize the solution in the general case, while the linear quadratic setting was tackled with a stochastic maximum principle approach. More recently, Lasry and Lions (2018), introduced a master equation accounting for this kind of major player model. Cardaliaguet et al. (2020) showed that the two previous approaches (Lasry and Lions (2018) and Carmona et al. (2016)) lead to the same Nash equilibria.

Financial markets and energy systems with many small interacting agents are a natural domain of applications of MFG. Casgrain and Jaimungal (2020) applied the MFG theory to optimal trade execution with price impact and terminal inventory liquidation condition, Fujii and Takahashi (2020), used this theory to find an equilibrium price under market clearing conditions. In Casgrain and Jaimungal (2020); Fujii and Takahashi (2020), the authors used the extended mean field setting to deal with heterogeneous sub-populations of agents and incomplete information for Casgrain and Jaimungal (2020). Alasseur et al. (2020) developed a model for the optimal management of energy storage and distribution in a smart grid system through an extended MFG. Shrivats, Firoozi and Jaimungal Shrivats et al. (2020) recently applied the theoretical setting developed in Casgrain and Jaimungal (2020) to the case of trading in solar renewable energy certificate markets. Financial markets with a major player, leader-follower interactions and terminal inventory constraint were recently analyzed in Fu and Horst (2020a); Evangelista and Thamsten (2020). In Fu and Horst (2020a), the authors consider a Brownian filtration, impose a zero terminal inventory constraint and characterize the equilibrium in terms of a McKean-Vlasov FBSDE. In Evangelista and Thamsten (2020), the authors study a market with a finite number of small players and a major player with first-mover advantage and information asymmetry, and characterize the solution in terms of a McKean-Vlasov FBSDE in a more general setting than that of Fu and Horst (2020a).

Among the cited papers, our methods and findings are closest in spirit to Bensoussan et al. (2016a); Evangelista and Thamsten (2020); Fu and Horst (2020a). The main novelty of our paper is the application of the linear quadratic MFG with a major player to the analysis of electricity markets in the presence of renewable production, however we make a number of contributions to the mathematical theory as well. Compared to the article Bensoussan et al. (2016a), which of course solves a more general problem, without focusing on a specific application, our paper allows a much more general dynamics for the driving processes (general semimartingales) and does not require an a priori bound on the strategies to prove the existence of the Nash equilibrium in the presence of a major player. Unlike the articles Evangelista and Thamsten (2020); Fu and Horst (2020a), which also study leader-follower games in financial markets, we consider a stochastic terminal constraint, characterize the equilibrium in explicit form, and show how an ε-Nash equilibrium for the finite-player game may be constructed from a mean field game solution.

The paper is organized as follows. In section 2 we introduce the model and briefly recall the mean field game solution obtained in Féron et al. (2020) in the case of identical agents. In section 3, we present the main results of this paper in the setting allowing for the presence of a major player, whose influence on the market is not negligible in the MFG limit. In section 4, we show how the limiting MFG solution may be used to construct an approximate Nash equilibrium in a Stackelberg game with one major player and N minor players. Finally, in section 5, equilibrium price trajectories, and the effect of market parameters on the price characteristics are illustrated with simulated data.

4.2 Preliminaries

In this paper, we place ourselves in the intraday market for a given delivery hour starting at time T, where time 0 corresponds to the opening time of the market (in EPEX Intraday this happens at 3PM on the previous day). In reality, trading stops a few minutes before delivery time (e.g. 5 minutes for Germany). However, for the sake of simplicity we assume that market participants can trade during the entire period $[0, T]$. In the market, there are agents (producers or consumers) who are assumed to have taken a position in the day-ahead market and use the intraday market to manage the volume risk associated to the imperfect demand/production forecast. These forecasts represent the best estimate of the additional demand compared to the position taken by the agent in the spot market: to avoid imbalance penalties, the intraday position of the agent at the delivery date must therefore be equal to the realized demand, or, in other words, the last observed value of the demand forecast.

We consider the case of a Stackelberg game where an agent called "major agent" faces a large number of smaller agents called "minor agents". We directly place ourselves in the setting of mean field games with a major player, that is, we assume that the number of small agents in the market is infinite, and the influence of each small agent on the market is negligible. The aggregate impact of the minor agents on the market is therefore modelled through a mean field.

Each agent observes the common national demand forecast, and the demand forecast of the major player. In addition, the small agents also observe their individual demand forecasts, which

are not observed by the other agents. The common filtration of the market thus contains the information about the forecast of the major player and the common part of the forecasts of the minor players, but the small agents benefit from a private information advantage compared to the major player.

The demand forecast process and the position of the generic minor agent are given, respectively, by $X := (X_t)_{0\leq t\leq T}$ and $\phi := (\phi_t)_{0\leq t\leq T}$, while the forecast process and the position of the major agent are given, respectively, by $(X_t^0)_{t\in[0,T]}$ and $(\phi_t^0)_{t\in[0,T]}$. Note that the position and forecast of the minor and major agents are not expressed in the same units. Indeed, in the mean field game limit considered in this paper, we assume that the market is very large, so that the position of every minor agent compared to the market size is negligible, but the major agent takes up a nonzero share of the market, so that ϕ^0 and X^0 denote the position and forecast of the major agent normalized by the market size.

We denote by $\mathbb{F}$ the filtration which contains all information available to the generic minor agent and by $\mathbb{F}^0$ the filtration which contains all information available to the major agent. This filtration contains the information about the fundamental price, the information about the demand forecast of the major agent, and potentially some information about the demand forecast of the generic minor agent (the common noise) but, in general, not the full individual demand forecast of the generic agent.

Throughout the paper and for any $\mathbb{F}$-adapted process $(\zeta_t)_{t\in[0,T]}$, we will denote $\bar{\zeta}_t = \mathbb{E}[\zeta_t|\mathcal{F}_t^0] = \int_{\mathbb{R}} x\mu_t^\zeta(dx)$ where: $\mu_t^\zeta := \mathcal{L}(\zeta_t|\mathcal{F}_t^0)$. In view of the convergence results of Proposition 9 and Proposition 10 in Féron et al. (2020), the (normalized) aggregate position of all minor agents is given by the expectation of ϕ with respect to the common noise: $\bar{\phi}_t = \mathbb{E}[\phi_t|\mathcal{F}_t^0]$.

We assume that the market price $(P_t)_{t\in[0,T]}$ is given by the fundamental price $(S_t)_{t\in[0,T]}$ plus a weighted combination of the aggregate position of the minor agents and the position of the major agent:

$$P_t = S_t + a\bar{\phi}_t + a^0\phi_t^0, \quad \forall t \in [0,T]$$

where a, a^0 are positive weights, which reflect the size of the major agent relative to the combined size of all minor agents and the overall strength of the market impact. Thus, the impact of each minor agent on the entire market is negligible, but the aggregate position of all minor agents, and the position of the major agent both have a nonzero impact.

We say that the strategy of the generic minor agent $(\dot{\phi}_t)_{t\in[0,T]}$ is admissible if it is $\mathbb{F}$-adapted and square integrable. Similarly, the strategy of the major agent $(\dot{\phi}_t^0)_{t\in[0,T]}$ is admissible if it is $\mathbb{F}^0$-adapted and square integrable. The instantaneous cost of trading for the major agent and for the generic minor agent are defined, respectively, by:

$$\dot{\phi}_t^0 P_t + \frac{\alpha_0(t)}{2}(\dot{\phi}_t^0)^2, \quad \text{and} \quad \dot{\phi}_t P_t + \frac{\alpha(t)}{2}(\dot{\phi}_t)^2, \quad \forall t \in [0,T] \tag{4.1}$$

In both instantaneous costs, the first term represents the actual cost of buying the electricity, and the second term represents the cost of trading, where $\alpha(.)$ and $\alpha_0(.)$ are continuous strictly positive functions on [0, T] reflecting the variation of market liquidity at the approach of the

delivery date.

The objective function of the minor agent has the following form:

$$J^{MF}(\phi, \bar{\phi}, \phi^0) := -\mathbb{E}\left[\int_0^T \frac{\alpha(t)}{2}\dot{\phi}_t^2 + (S_t + a\bar{\phi}_t + a^0\phi_t^0)\dot{\phi}_t dt + \frac{\lambda}{2}(\phi_T - X_T)^2\right], \tag{4.2}$$

while the objective function of the major agent writes,

$$J^{MF,0}(\phi^0, \bar{\phi}) := -\mathbb{E}\left[\int_0^T \frac{\alpha_0(t)}{2}\dot{\phi}_t^{0\,2} + (S_t + a\bar{\phi}_t + a^0\phi_t^0)\dot{\phi}_t^0 dt + \frac{\lambda_0}{2}(\phi_T^0 - X_T^0)^2\right]. \tag{4.3}$$

Note that this formulation implies (as it is the case in real markets) that the major agent pays a much lower trading cost per unit traded and a much lower imbalance penalty than the minor agents. Indeed, if the major agent paid the same quadratic cost/penalty as the minor agents, since the position of the major agent is very large, the quadratic trading cost/penalty would grow much faster than the linear part (the middle term in the formula), and the limiting formula would be degenerate, in the sense that the trading strategy would be independent from the price. To obtain a nondegenerate expression in terms of the normalized trading strategy of the major agent, we must therefore assume that the actual trading cost and penalty are also renormalized. The quantities $\alpha_0(t)$ and λ_0 are thus different from $\alpha(t)$ and λ since they are of different nature: $\alpha(t)$ and λ apply to the actual strategy of the generic agent, while $\alpha_0(t)$ and λ_0 apply to the normalized strategy of the major agent. The different nature of trading costs for minor and major agents is confirmed by other authors Donier et al. (2015): while the minor agents post their orders immediately in the order book, the major agent splits its orders into many small chunks to minimize trading costs.

To close this introductory section, we briefly recall one of the main results (Theorem 7) from Féron et al. (2020), which characterizes the mean field equilibrium in the setting of identical agents, in other words, we assume that $a_0 = 0$ until the end of this section.

Definition 23 (mean field equilibrium)**.** *An admissible strategy $\dot{\phi}^* := (\dot{\phi}_t^*)_{t\in[0,T]}$ is a mean field equilibrium in the setting of identical agents if it maximizes the functional (4.2) with $a_0 = 0$ and satisfies $\bar{\phi} = \bar{\phi}^*$.*

We make the following assumption.

Assumption 24.

- *The process S is square integrable and adapted to the filtration $\mathbb{F}^0$.*
- *The process X is a square integrable martingale with respect to the filtration $\mathbb{F}$.*
- *The process $\overline{X}$ defined by $\overline{X}_t := \mathbb{E}[X_t|\mathcal{F}_t^0]$ for $0 \le t \le T$ is a square integrable martingale with respect to the filtration $\mathbb{F}$.*

Note that if X is an $\mathbb{F}$-martingale, then $\overline{X}$ is by construction an $\mathbb{F}^0$-martingale, but it may not necessarily be a martingale in the larger filtration $\mathbb{F}$.

The following theorem characterizes the mean field equilibrium in the identical agent setting. In the theorem, we decompose the individual demand forecast as follows: $X_t = \overline{X}_t + \check{X}_t$, where $\overline{X}_t = \mathbb{E}\left[X_t|\mathcal{F}_t^0\right]$, and we use the following shorthand notation:

$$\Delta_{s,t} := \int_s^t \frac{\eta(u,t)}{\alpha(u)} du \quad \text{with} \quad \eta(s,t) = e^{-\int_s^t \frac{a}{\alpha(u)} du} \quad \text{and} \quad \widetilde{\Delta}_{s,t} := \int_s^t \alpha^{-1}(u) du$$
$$I_t := \int_0^t \frac{\eta(s,t)}{\alpha(s)} S_s ds, \quad \widetilde{I}_t := \mathbb{E}\left[\int_0^T \frac{\eta(s,T)}{\alpha(s)} S_s ds \Big| \mathcal{F}_t\right]. \tag{4.4}$$

Theorem 25. *Under Assumption 24, the unique mean field equilibrium in the setting of identical agents is given by*

$$\phi_t^* = -I_t + \lambda \left[\Delta_{0,t} \frac{\widetilde{I}_0 + \overline{X}_0}{1+\lambda\Delta_{0,T}} + \int_0^t \Delta_{s,t} \frac{d\widetilde{I}_s + d\overline{X}_s}{1+\lambda\Delta_{s,T}}\right. \tag{4.5}$$
$$\left. + \widetilde{\Delta}_{0,t} \frac{\check{X}_0}{1+\lambda\widetilde{\Delta}_{0,T}} + \int_0^t \widetilde{\Delta}_{s,t} \frac{d\check{X}_s}{1+\lambda\widetilde{\Delta}_{s,T}}\right].$$

The equilibrium price has the following form:

$$P_t = S_t - aI_t + a\lambda \left[\Delta_{0,t} \frac{\widetilde{I}_0 + \overline{X}_0}{1+\lambda\Delta_{0,T}} + \int_0^t \Delta_{s,t} \frac{d\widetilde{I}_s + d\overline{X}_s}{1+\lambda\Delta_{s,T}}\right]. \tag{4.6}$$

4.3 A game of a major and minor agents

In this section we proceed to characterize the Nash equilibrium in the Stackelberg mean field game with a major player. Since a single minor agent has an infinitesimal impact on the market and cannot influence the mean field or the strategy of the major agent, the problem of the generic minor agent is to maximize $J^{MF}(\phi, \bar{\phi}, \phi^0)$ for fixed $\bar{\phi}$ and ϕ^0. On the other hand, by modifying her strategy ϕ^0, the major agent may influence the strategies of the minor agents, and thus also the mean field $\bar{\phi}$. This leads to the following definition of mean field equilibrium. As in the preceding section, the "consistency condition" in this definition simply translates the fact that the aggregate position $\bar{\phi}$ of all minor agents is given by the expectation of the representative agent strategy ϕ with respect to the common noise.

Definition 26 (Stackelberg mean field equilibrium)**.** *We call the triple ϕ^*, $\bar{\phi}^*$, ϕ^{0*} Stackelberg mean field equilibrium for the game with a major and minor players if the following holds:*

i. ϕ^* and ϕ^{0*} are admissible strategies for, respectively, the representative minor and the major players, the consistency condition $\bar{\phi}_t^* = \mathbb{E}[\phi_t^*|\mathcal{F}_t^0]$ is satisfied for all $t \in [0,T]$ and for any other admissible strategy for the representative minor player ϕ,

$$J^{MF}(\phi, \bar{\phi}^*, \phi^{0*}) \leq J^{MF}(\phi^*, \bar{\phi}^*, \phi^{0*})$$

ii. For any other triple $(\phi, \bar{\phi}, \phi^0)$ satisfying condition i.,

$$J^{MF,0}(\phi^0, \bar{\phi}) \leq J^{MF,0}(\phi^{0*}, \bar{\phi}^*). \tag{4.7}$$

Assumption 27. *In addition to Assumption 24, we also assume that the process X^0 is a square integrable martingale with respect to the filtration $\mathbb{F}^0$.*

We start with the characterization of the optimal strategy for the minor agent.

Proposition 28 (Minor representative agent)**.** *Let $\bar{\phi}$ and ϕ^0 be fixed. The minor agent strategy ϕ maximizes (4.2) over the set of admissible strategies if and only if:*

$$\dot{\phi}_t = -\frac{Y_t + S_t + a\overline{\phi}_t + a^0\phi^0_t}{\alpha(t)}, \quad \forall t \in [0, T], \tag{4.8}$$

where Y is a $\mathbb{F}$-martingale satisfying $Y_T = \lambda(\phi_T - X_T)$.

Proof. The proof follows from the first step of the proof of Theorem 25 (see Theorem 7 in Féron et al. (2020)) taking $\widetilde{S} = S + a_0\phi^0$ as fundamental price instead of S. □

The problem of the major agent is more complex since the minor agents observe her actions and modify their strategies accordingly, which means that the mean field $\bar{\phi}$ depends on the major agent's strategy ϕ^0, and the problem of the major agent effectively becomes a stochastic control problem. We start with a reformulation of the definition of Stackelberg equilibrium in terms of $\bar{\phi}$ and ϕ^0 only.

Lemma 29. *Let $(\bar{\phi}^*, \phi^{0*})$ be $\mathbb{F}^0$-adapted square integrable processes. There exists ϕ^* such that $(\phi^*, \bar{\phi}^*, \phi^{0*})$ is a Stackelberg mean field equilibrium if and only if the couple $(\bar{\phi}^*, \phi^{0*})$ satisfies the following conditions:*

i. For every $\mathbb{F}^0$-adapted square integrable process ν,

$$\mathbb{E}\left[\int_0^T \nu_t\{\alpha(t)\dot{\bar{\phi}}^*_t + S_t + a_0\phi^{0*}_t + a\bar{\phi}^*_t\}dt + \lambda(\bar{\phi}^*_T - \overline{X}_T)\int_0^T \nu_t dt\right] = 0.$$

ii. For every other couple $(\phi^0, \bar{\phi})$ satisfying the condition i, the inequality (4.7) holds true.

The proof of this lemma is postponed to Appendix 4.7.1.

The following proposition provides a martingale characterization of the Stackelberg mean field equilibrium.

Proposition 30. *Let $(\bar{\phi}^*, \phi^{0*})$ be $\mathbb{F}^0$-adapted square integrable processes. There exists ϕ^* such that $(\phi^*, \bar{\phi}^*, \phi^{0*})$ is a Stackelberg mean field equilibrium if and only if*

$$\dot{\phi}^{0*}_t = -\frac{M^0_t + S_t + a\overline{\phi}_t - a^0 N_t}{\alpha_0(t)}, \quad \forall t \in [0, T], \tag{4.9}$$

where M^0 is an $\mathbb{F}^0$-martingale and N is an absolutely continuous $\mathbb{F}^0$-adapted process, and there exists an $\mathbb{F}^0$-martingale M, and an $\mathbb{F}^0$-martingale $\overline{Y}$ such that the following system of equations is satisfied:

$$\begin{cases} M_T^0 = a^0 N_T + a^0 \phi_T^0 + \lambda_0(\phi_T^0 - X_T^0) \\ M_t - a\phi_t^0 + \alpha(t)\dot{N}_t - aN_t = 0, \qquad M_T = a\phi_T^0 + (a+\lambda)N_T \\ \overline{Y}_t + \alpha(t)\dot{\bar{\phi}}_t + S_t + a\bar{\phi}_t + a^0\phi_t^0 = 0, \qquad \overline{Y}_T = \lambda(\bar{\phi}_T - \overline{X}_T) \end{cases} \tag{4.10}$$

The proof of this proposition is postponed to Appendix 4.7.1.

From Proposition 30, it follows that the existence and uniqueness of the Stackelberg equilibrium reduces to the existence and uniqueness of the solution, if any, of the linear system of coupled BSDEs equation (4.10).

The following theorem provides an explicit characterization of the equilibrium in the Stackelberg setting.

Theorem 31 (Explicit solution)**.** *Let $\Xi_t = (\phi_t^0, N_t, \bar{\phi}_t)'$. The unique equilibrium of the mean field game with a major agent is characterized by the following differential equation:*

$$B(t)^{-1}A\,\Xi_t + \dot{\Xi}_t = -\begin{pmatrix} \alpha_0(t)^{-1}(M_t^0 + S_t) \\ \alpha(t)^{-1}M_t \\ \alpha(t)^{-1}(\overline{Y}_t + S_t) \end{pmatrix}, \tag{4.11}$$

where N is a $\mathbb{F}^0$-adapted process with $N_0 = 0$ and M^0, M, $\overline{Y}$ are $\mathbb{F}^0$-martingales that satisfy:

$$\begin{cases} M_T^0 = a^0 N_T + a^0 \phi_T^0 + \lambda_0(\phi_T^0 - X_T^0) \\ M_T = a\phi_T^0 + (a+\lambda)N_T \\ \overline{Y}_T = \lambda(\bar{\phi}_T - \overline{X}_T) \end{cases} \tag{4.12}$$

and

$$A = \begin{pmatrix} 0 & -a^0 & a \\ -a & -a & 0 \\ a^0 & 0 & a \end{pmatrix}, \qquad B(t) = \begin{pmatrix} \alpha_0(t) & 0 & 0 \\ 0 & \alpha(t) & 0 \\ 0 & 0 & \alpha(t) \end{pmatrix}.$$

Denoting by $\Phi(t)$ the fundamental matrix solution of the equation $B(t)^{-1}A\,\Xi_t + \dot{\Xi}_t = 0$, the solution is given in integral form by the following expression:

$$\Xi_t = \Upsilon_t - \Pi_{0,t}(I + D\Pi_{0,T})^{-1}(D\widetilde{\Upsilon}_0 - \Lambda\mathcal{X}_0) - \int_0^t \Pi_{s,t}(I + D\Pi_{s,T})^{-1}(Dd\widetilde{\Upsilon}_s - \Lambda d\mathcal{X}_s). \tag{4.13}$$

where,

$$\Upsilon_t := -\Phi(t)\int_0^t \Phi(s)^{-1}\begin{pmatrix} \alpha_0(s)^{-1}S_s \\ 0 \\ \alpha(s)^{-1}S_s \end{pmatrix} ds, \quad \mathcal{X}_s := \begin{pmatrix} X_s^0 \\ 0 \\ \overline{X}_s \end{pmatrix}, \quad D = \begin{pmatrix} a^0 + \lambda_0 & a^0 & 0 \\ a & a+\lambda & 0 \\ 0 & 0 & \lambda \end{pmatrix},$$

$$\Lambda = \begin{pmatrix} \lambda_0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & \lambda \end{pmatrix}, \quad \Pi_{s,t} := \Phi(t) \int_s^t \Phi(u)^{-1} B(u)^{-1} du \quad \textit{and} \quad \widetilde{\Upsilon}_t = \mathbb{E}[\Upsilon_T | \mathcal{F}_t].$$

The proof of this theorem can be found in Appendix 4.7.2.

Remark 32. *These results can be generalized to the setting of several major agents, interacting with the mean-field of the minor agents, provided that each major agent observes the individual forecasts of the other agents, but not those of the minor agents. In this case, the constrained optimization problem of one major agent becomes a constrained game between several major agents. Since the setting remains linear-quadratic, one will still be able to obtain an explicit solution, at the price of more tedious computations.*

Remark 33. *If $\alpha_0(t) = c\alpha(t)$ for some constant c, the fundamental matrix solution is given explicitly by*

$$\Phi(t) = \exp\left(-\int_0^t B(s)^{-1} A ds\right)$$

Let us make some comments on how minor and major player strategies change when the parameters of the model vary. First, when the major player has no price impact, $a_0 = 0$, we recover the homogeneous mean field setting optimal strategy for the minor player from (4.10) and (4.12):

$$\bar{\phi}_t = \int_0^t \frac{\eta(s,t)}{\alpha(s)} (\overline{Y}_s + S_s) ds,$$

where $\overline{Y}$ satisfies the equation:

$$\begin{cases} \overline{Y}_t + \alpha(t)\dot{\bar{\phi}}_t^* + S_t + a\bar{\phi}_t^* = 0 \\ \overline{Y}_T = -\lambda(\bar{\phi}_T^* - \overline{X}_T). \end{cases}$$

Second, we explore the limiting behavior of the optimal strategies for the major agent and for the mean field in various limiting cases. In this corollary, we use the notation of Theorem 31. The proof of this corollary can be found in Appendix 4.7.3.

Corollary 34. *i. Assume that the fundamental price process S is a martingale. Then the equilibrium mean field position of minor agents, and the position of the major agent satisfy*

$$\Xi_t = -\Pi_{0,t}(I + D\Pi_{0,T})^{-1} \begin{pmatrix} S_0 - \lambda_0 X_0^0 \\ 0 \\ S_0 - \lambda \overline{X}_0 \end{pmatrix} - \int_0^t \Pi_{s,t}(I + D\Pi_{s,T})^{-1} d \begin{pmatrix} S_s - \lambda_0 X_s^0 \\ 0 \\ S_s - \lambda \overline{X}_s \end{pmatrix}.$$

ii. In the limit of infinite terminal penalty (when $\lambda, \lambda_0 \to \infty$) the equilibrium mean field position of minor agents, and the position of the major agent satisfy,

$$\Xi_t \to \Upsilon_t - \Pi_{0,t}\Pi_{0,T}^{-1}(\widetilde{\Upsilon}_0 - D_\infty \mathcal{X}_0) - \int_0^t \Pi_{s,t}\Pi_{s,T}^{-1}(d\widetilde{\Upsilon}_s - D_\infty d\mathcal{X}_s),$$

almost surely for all $t \in [0,T]$, where

$$D_\infty = \begin{pmatrix} 1 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & 1 \end{pmatrix}.$$

When the fundamental price process S is a martingale, in the limit of infinite terminal penalty, the strategies do not depend on the fundamental price and we have,

$$\Xi_t \to \Pi_{0,t}\,\Pi_{0,T}^{-1} \begin{pmatrix} X_0^0 \\ 0 \\ \overline{X}_0 \end{pmatrix} + \int_0^t \Pi_{s,t}\,\Pi_{s,T}^{-1}\,d \begin{pmatrix} X_s^0 \\ 0 \\ \overline{X}_s \end{pmatrix}.$$

iii. In the absence of terminal penalties (when $\lambda = \lambda_0 = 0$), the equilibrium mean field position of minor agents, and the position of the major agent satisfy,

$$\Xi_t = \Upsilon_t - \Pi_{0,t}(I + D_0\Pi_{0,T})^{-1}D_0\widetilde{\Upsilon}_0 - \int_0^t \Pi_{s,t}(I + D_0\Pi_{s,T})^{-1}D_0 d\widetilde{\Upsilon}_s. \tag{4.14}$$

Interestingly, when the players don't have a terminal penalty ($\lambda = \lambda_0 = 0$), the equilibrium positions of the agents in equation (4.14) still contain forward looking terms, which were absent in the case of the mean field game with identical players (see Equation (4.5) with $\lambda = 0$). The presence of these terms is due to the strategic interaction of the major player with the mean field of small agents.

In the limit of zero trading costs, the gain of the major player remains bounded in expectation, however, contrary to the case of identical players, the optimal strategy of the major agent cannot be determined uniquely from the optimization problem. Indeed, assuming that the trading cost for minor agents is zero, the equilibrium price (computed from Equation (12) in Féron et al. (2020) with $N \to \infty$) is given by

$$P_t = S_t + a_0\phi_t^0 + a\bar{\phi}_t = \frac{\lambda}{a+\lambda}(a\overline{X}_t + \mathbb{E}[S_T|\mathcal{F}_t] + a_0\mathbb{E}[\phi_T^0|\mathcal{F}_t]).$$

Substituting this expression into the optimization problem for the major player, we need to minimize the following functional:

$$\begin{aligned} &\mathbb{E}\left[\frac{\lambda}{a+\lambda}\int_0^T \dot{\phi}_t^0(a\overline{X}_t + \mathbb{E}[S_T|\mathcal{F}_t] + a_0\mathbb{E}[\phi_T^0|\mathcal{F}_t])dt + \frac{\lambda_0}{2}(\phi_T^0 - X_T^0)^2\right] \\ &= \mathbb{E}\left[\frac{\lambda}{a+\lambda}(a\phi_T^0\overline{X}_T + \phi_T^0 S_T + a_0(\phi_T^0)^2) + \frac{\lambda_0}{2}(\phi_T^0 - X_T^0)^2\right], \end{aligned}$$

where the equality follows, in particular, from the martingale property of $\overline{X}_T$. Since the expression to be minimized only depends on the terminal value ϕ_T^0 of the major agent's position, any strategy with the optimal terminal value will satisfy the condition of optimality: the Stackelberg

equilibrium will not be unique in this case.

To finish this section, we provide the explicit form of the strategy of the minor agents. The proof of this result can be found in Appendix 4.7.3.

Corollary 35 (Minor agent strategy)**.** *Under Assumption 27, the optimal generic minor agent position ϕ^* is given by:*

$$\phi_t^* = \int_0^t \widetilde{\Delta}_{s,t} \frac{\lambda d\check{X}_s}{1+\lambda\widetilde{\Delta}_{s,T}} + \widetilde{\Delta}_{0,t} \frac{\lambda \check{X}_0}{1+\lambda\widetilde{\Delta}_{0,T}} + \bar{\phi}_t^*$$

where $\bar{\phi}^*$ is the optimal aggregate position of the minor agents, as given by Theorem 31.

4.4 Approximate Nash equilibrium in the N-player Stackelberg game

In this section, we derive the ϵ-Nash approximation for the Stackelberg game. In the present leader-follower setting, we allow the minor agents to change their strategies when the major agent deviates from her optimal one.

Since we would like to study the rate of convergence as $N \to \infty$, we assume that there is a major player and an infinity of minor players replacing the generic agent. Their demand forecasts are respectively given by X_t^i, $i = 0, \dots, \infty$, $t \in [0, T]$. The private demand forecasts of all agents are defined on the same probability space. We therefore impose the following assumption.

Assumption 36.

- *The process S is square integrable and adapted to the filtration $\mathbb{F}^0$.*
- *The demand forecast X^0 of the major agent is a square integrable $\mathbb{F}^0$-martingale.*
- *The processes $(X^i)_{i=1}^\infty$ are square integrable $\mathbb{F}$-martingales.*
- *There exists a square intergrable $\mathbb{F}$-martingale $\overline{X}$, such that for all $i \geq 1$, and all $t \in [0, T]$, almost surely, $\mathbb{E}[X_t^i | \mathcal{F}_t^0] = \overline{X}_t$.*
- *The processes $(\check{X}^i)_{i=1}^\infty$ defined by $\check{X}_t^i = X_t^i - \overline{X}_t$ for $t \in [0, T]$, are orthogonal square integrable $\mathbb{F}$-martingales, such that the expectation $\mathbb{E}[(\check{X}_T^i)^2]$ does not depend on i.*

The strategy $(\dot{\phi}^i)$ of agent $i = 1, \dots, \infty$ is said to be admissible if it is $\mathbb{F}$-adapted and square integrable; the strategy $(\dot{\phi}^0)$ of the major agent is admissible if it is $\mathbb{F}^0$-adapted and square integrable. For a fixed $N \geq 1$, we denote:

$$P^N(\phi_t^0, ..., \phi_t^N) = S_t + a\overline{\phi}_t^N + a^0\phi_t^0$$
$$P^{MF}(\phi_t^0, \bar{\phi}_t) = S_t + a\bar{\phi}_t + a^0\phi_t^0,$$

where $\bar{\phi}_t^N = \frac{1}{N}\sum_{i=1}^N \phi_t^i$ is the average position of the minor agents. And we define in the N-player game, the objective functions for the major agent:

$$J^{N,0}(\phi^0, \phi^{-0}) := -\mathbb{E}\left[\int_0^T \left\{\frac{\alpha_0(t)}{2}(\dot{\phi}_t^0)^2 + \dot{\phi}_t^0 P^N(\phi_t^0, \ldots, \phi_t^N)\right\} dt + \frac{\lambda_0}{2}(\phi_T^0 - X_T^0)^2\right], \quad (4.15)$$

and for the minor agents $i = 1, \ldots, N$:

$$J^{N,i}(\phi^i, \phi^{-i}) := -\mathbb{E}\left[\int_0^T \left\{\frac{\alpha(t)}{2}(\dot{\phi}_t^i)^2 + \dot{\phi}_t^i P^N(\phi_t^0, \ldots, \phi_t^N)\right\} dt + \frac{\lambda}{2}(\phi_T^i - X_T^i)^2\right], \quad (4.16)$$

as well as the objective function for the minor agents $i = 1, \ldots, N$, in the mean field setting:

$$J^{MF}(\phi^i, \bar{\phi}, \phi^0) := -\mathbb{E}\left[\int_0^T \left\{\frac{\alpha(t)}{2}(\dot{\phi}_t^i)^2 + \dot{\phi}_t^i P^{MF}(\phi_t^0, \bar{\phi}_t)\right\} dt + \frac{\lambda}{2}(\phi_T^i - X_T^i)^2\right]. \quad (4.17)$$

We next provide a definition of the ϵ-Nash equilibrium in the present Stackelberg setting. As mentioned above, the deviations of the major and minor agents must be treated differently: when the major agent deviates, we allow the minor agents to adjust their strategies to respond optimally to the new strategy of the major agent. We say that the minor agent strategies $\phi^1, \ldots, \phi^N$ are an optimal response to the major agent strategy ϕ^0 if for every $i = 1, \ldots, N$ and for every admissible minor agent strategy $\tilde{\phi}^i$,

$$J^{N,i}(\tilde{\phi}^i, \phi^{-i}) \leq J^{N,i}(\phi^i, \phi^{-i}).$$

Definition 37 (Stackelberg ε-Nash equilibrium)**.** *We say that $(\phi_t^{i*})_{t\in[0,T], 0\leq i\leq N}$ is an ϵ-Nash equilibrium for the N-player game if these strategies are admissible and the following holds.*

i. ***Deviation of a minor player:*** *For any other admissible strategy ϕ^i for the minor player i, $i = 1, \ldots, N$,*

$$J^{N,i}(\phi^i, \phi^{-i*}) - \varepsilon \leq J^{N,i}(\phi^{i*}, \phi^{-i*}).$$

ii. ***Deviation of the major player:*** *For any other set of admissible strategies $(\phi^i), i = 0, \ldots, N$, such that $\phi^1, \ldots, \phi^N$ are optimal responses of minor players to the major player strategy ϕ^0, we have,*

$$J^{N,0}(\phi^0, \phi^{-0}) - \varepsilon \leq J^{N,0}(\phi^{0*}, \phi^{-0*}).$$

Our definition of ε-Nash equilibrium is different from the one in Carmona et al. (2016): while the latter paper assumes that the major player deviates from her strategy unilaterally (see Definition 4.2), we allow the minor players to respond to the deviation of the major player, in agreement with the leader-follower nature of the game. In addition, in Carmona et al. (2016), an a priori bound on the L^p-norm of the new strategy of the major agent is required to establish Theorem 4.1, whereas no such bound is needed in our setting.

Proposition 38. *Assume that the strategies of the N minor agents are given by*

$$\phi_t^{i*} = \int_0^t \widetilde{\Delta}_{s,t} \frac{\lambda d\check{X}_s^i}{1+\lambda\widetilde{\Delta}_{s,T}} + \widetilde{\Delta}_{0,t} \frac{\lambda \check{X}_0^i}{1+\lambda\widetilde{\Delta}_{0,T}} + \overline{\phi}_t^*,$$

where $\overline{\phi}^*$ is the third component of the mean field equilibrium defined in Theorem 31. Assume that the strategy of the major agent is given by Theorem 31 as well. Let assumption 36 hold true. Then there exists a constant $C < \infty$, which does not depend on N, such that these strategies form an ε-Nash equilibrium of the N-player game with $\varepsilon = \frac{C}{N^{1/2}}$.

The proof of this proposition is postponed to Appendix 4.7.4.

Remark 39. *The ε-Nash equilibrium described in Proposition 38 approximates the N-player equilibrium in the complete information setting (where every player observes the others' actions), but its implementation for each agent only requires the knowledge of the common information $\mathbb{F}^0$ as well as the agent's individual forecast.*

4.5 Numerical illustration

In this section, our objective is to illustrate the theoretical results presented in sections 4.3 and 4.4 with numerical simulations. We analyze the role of the major producer in the market and its impact on price characteristics such as volatility and price-forecast correlation, and compare this situation to the homogeneous agent setting studied in Féron et al. (2020). Some comparisons with the empirical market characteristics are also performed, but we refer the reader to Féron et al. (2020) and other papers cited therein for a more detailed description of intraday electricity markets and their empirical features. As empirical analysis in Féron et al. (2020) are led on actual wind infeed forecasts, we will consider production forecasts here instead of demand forecasts as it is the case in the rest of the paper. Throughout, we consider that the production forecasts are the differences between actual production forecasts and the agents' positions in the market at time 0. Therefore, the initial values X_0^i, $i = 0, \dots, N$ will be set to 0.

Model specification We now define the dynamics for the fundamental price and for the production forecasts used in the simulations and specify the parameter values. The objective being to illustrate the model, the majority of the parameters are not precisely estimated, but are given ad hoc plausible values.

The evolution of the fundamental price is described as follows:

$$dS_t = \sigma^S dW_t \tag{4.18}$$

where σ^s is a constant and $(W_t)_{t\in[0,T]}$ is Brownian motion. We also assume that the liquidity functions $\alpha(.)$ and $\alpha_0(.)$ have a specific form given by

$$\alpha(t) = \alpha \times (T-t) + \beta, \quad \forall t \in [0,T] \tag{4.19}$$

$$\alpha_0(t) = \alpha_0 \times (T-t) + \beta_0, \quad \forall t \in [0,T] \tag{4.20}$$

where α, β, α_0 and β_0 are strictly positive constants. The liquidity functions are thus decreasing with time. This assumption relies on the fact that the market becomes more liquid as we get closer to the delivery time and it is less costly to trade when the market is liquid.

To simulate production forecasts we assume the following dynamics:

$$d\bar{X}_t = \bar{\sigma} d\bar{B}_t \tag{4.21}$$

$$d\check{X}_t^0 = \sigma^0 dB_t^0, \tag{4.22}$$

$$d\check{X}_t^i = \sigma^X dB_t^i, \qquad i \in \{1, \dots, N\} \tag{4.23}$$

where $\bar{\sigma}$, σ^0 and σ^X are constants and $(\bar{B}_t)_{t\in[0,T]}$, $(B_t^i)_{t\in[0,T]}$, $i \in \{0, \dots, N\}$ are independent Brownian motions, also independent from $(W_t)_{t\in[0,T]}$.

In this illustration, we choose the same parameters for the dynamics of the common and the individual production forecasts, as well as the forecast of the major agent. The common volatility is calibrated to wind energy forecasts in Germany over January 2015 during the last quotation hour, by using the classical volatility estimator

$$\bar{\sigma} = \sigma^0 = \sigma^X = \widehat{\sigma} = \frac{\sqrt{\Delta t}}{n' - 1} \sum_{i=1}^{n'} Y_i^2 \tag{4.24}$$

with Δt the time step between two observations, $Y_i = X_{t_i} - X_{t_{i-1}}$ the increment between two successive observations and n' the total number of observed increments. As the forecasts are updated every 15 minutes, there are three daily variations during the last hour of forecasts from the 3rd of January to the 31th of January. Thus, for each delivery hour we dispose of $n' = 87$ increments to estimate the volatility.

The model parameters are specified in Table 4.1.

Parameter	Value	Parameter	Value
S_0	40 €/MWh	a	1 €/MWh^2
σ^S	10 €/MWh.h$^{1/2}$	λ	100 €/MWh^2
$\overline{X}_0$	0 MWh	λ_0	100 €/MWh^2
$\bar{\sigma}$	73 MWh/h$^{1/2}$	α	0.3 €/h.MW^2
$\check{X}_0^i$	0 MWh	α_0	0.3 €/h.MW^2
σ^X, σ^0	73 MWh/h$^{1/2}$	β	0.1 €/MW^2
N	100	β_0	0.1 €/MW^2

Table 4.1: Parameters of the model

Equilibrium price and market impact In Figure 4.1, we plot the major agent production forecast and the common production forecast (respectively the orange and blue solid lines) together with the equilibrium position of the major agent and the aggregate position of the minor agents given by Proposition 30 and Corollary 35 (respectively the orange and blue dashed lines). For comparison, we also plot the aggregate position in the identical agent case (dotted green line). All trajectories have been computed with the same production forecasts, the same

fundamental price, initial values, volatilities and parameters, specified in Table 4.1, except for the price impact coefficients of major and minor player, which differ according to model specification. In the Stackelberg game we chose $a^0 = a = 0.5$ €/MWh^2, and in the homogeneous case we kept $a = 1$ €/MWh^2 .

Figure 4.1: Stackelberg game

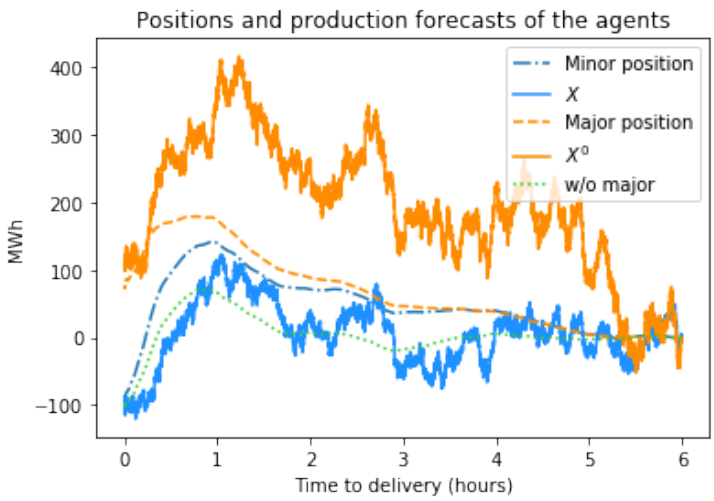

We observe that the strategy in the setting of identical agents and the strategy of the minor player in the Stackelberg setting converge to the same terminal value due to the terminal penalty. However, in the Stackelberg case, the minor agent position tends to follow the one of the major player during the first part of the trading period. In the case of identical agents, the fluctuations are not as strong since, contrary to the case when a major agent is present, the generic minor agent has no incentive to modify her trajectory to follow the leader. During the second half of the trading period, the minor agent position deviates further away from the one of the major agent to target the same terminal position as the mean field in the case of identical agents. We can argue that the strategy of the minor agent becomes more sensitive to the terminal constraint as we get closer to the delivery time: the weight of the terminal constraint in her strategy increases due to the decrease of the instantaneous trading cost.

Volatility and price-forecast correlation In this paragraph, we illustrate with simulations the effect of the presence of the major agent on the price characteristics such as the volatility and the correlation between the price and renewable infeed forecasts. The volatility was estimated from simulated price trajectories using a kernel-based non parametric estimator of the instantaneous volatility:

$$\widehat{\sigma}_t^2 = \frac{\sum_{i=1}^n K_h(t_{i-1} - t)\Delta \tilde{P}_{t_{i-1}}^2}{\sum_{i=1}^n K_h(t_{i-1} - t)(t_i - t_{i-1})}, \tag{4.25}$$

where $K(.)$ is the Epanechnikov kernel: $K(x) = \frac{3}{4}(1 - x^2)\mathbb{1}_{[-1,1]}(x)$ and $K_h(x) = \frac{1}{h}K(\frac{x}{h})$. The parameter h was taken equal to 0.08 hour ($\approx$ 5 minutes).

For a fixed scenario of production forecasts for the minor and major players, drawn in Figure

4.1, left graph, we estimated the volatility of the simulated market price for different values of the weights a_0 and a assigned, respectively, to the major player and the mean field of minor players in the price impact function. We studied three different combinations of weights to illustrate the impact of the minor players and the major player in the game: $a^0 = a = 0.5$ €/MWh^2, the impact of the major player and the minor players is the same; $a^0 = 0.9, a = 0.1$ €/MWh^2, the major player has a lot more impact than the minor players, and finally $a^0 = 0$ €/MWh^2, $a = 1$ €/MWh^2, equivalent to a market price without major player since she has no market impact in this case. These weights can be seen as the respective market shares held by the major agent and the minor players.

Figure 4.2, left graph, shows the estimated volatility trajectories for the three different cases of market shares of the major agent averaged over 1000 simulations. We note that the volatility of the market price depends on the strength of impact of the major player: the greater a^0, the higher the volatility. A possible explanation for this phenomenon is that stronger competition in the market (when the major agent is absent or has a small market share) reduces profit opportunities in the market and the agents therefore trade less actively.

For comparison, we also plot in Figure 4.2, right graph, the volatility estimated from empirical intraday electricity price data using the same estimator (4.25). This graph is taken from Féron et al. (2020). We see that the phenomenon of increasing volatility at the approach of the delivery date, clearly visible in the actual electricity markets, is well reproduced by our model.

Figure 4.2: Left: volatility of simulated prices for different market shares of the major agent. Right: volatility for different delivery hours, estimated empirically from EPEX spot intraday market data of January 2017 for the Germany delivery zone.

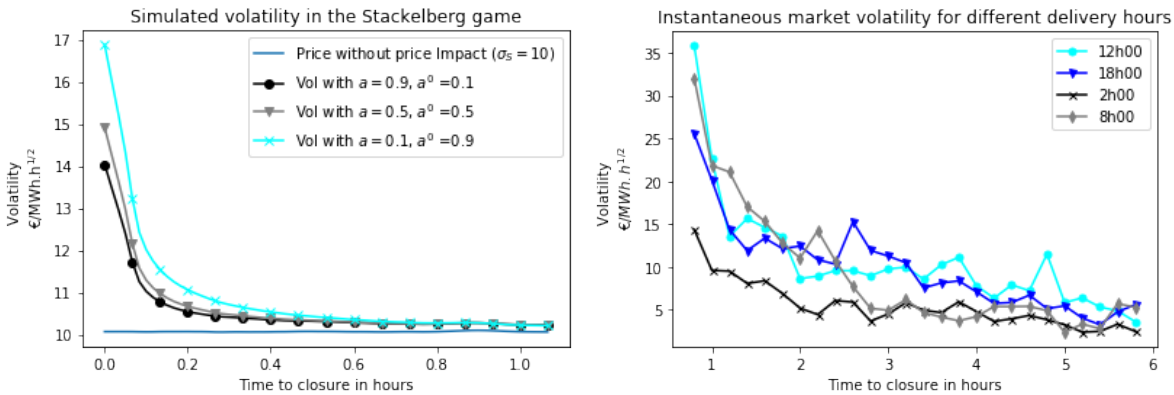

An important stylized feature of intraday market prices, observed empirically in Kiesel and Paraschiv (2017) and Féron et al. (2020) is the correlation between the price and the renewable production forecasts. In Figure 4.3, we plotted the correlation between the increments of the market price and the increments of the renewable production forecast of the major agent as function of time, in the market impact setting $a^0 = a = 0.5$ €/MWh^2; as well as the correlation between the price increments and the increments of the total aggregate forecast of both the major and minor players.

The correlation is computed over 15-minutes increments using the following estimator:

$$\widehat{\rho}_t = \frac{\sum_{k=1}^{N_{sim}} (\Delta Y_t^k - \overline{\Delta Y}_t)(\Delta P_t^k - \overline{\Delta P}_t)}{\sqrt{\sum_{k=1}^{N_{sim}} (\Delta Y_t^k - \overline{\Delta Y}_t)^2 \sum_{k=1}^{N_{sim}} (\Delta P_t^k - \overline{\Delta P}_t)^2}},$$

with N_{sim} the number of simulations (we considered $N_{sim} = 50000$) and where ΔY_t^k and $\Delta P_t^k = P_{t+dt}^{MF,k} - P_t^{MF,k}$ are the increments of, respectively, the forecast process and the market price.

Figure 4.3: Correlation between the price increments and the major player renewable production increments v.s the correlation between the price increments and the total renewable production increments

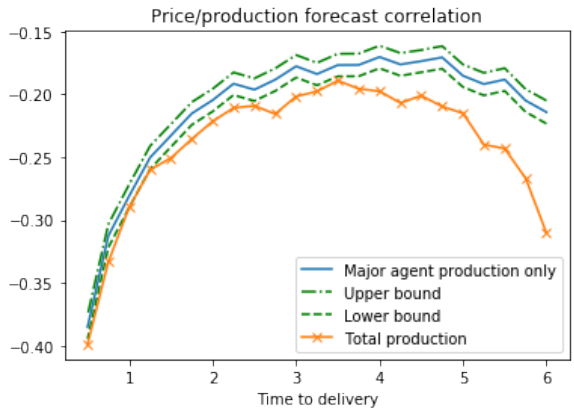

For the sake of clarity we only draw the Monte Carlo confidence interval for the case of the correlation between the major player production and the price considered on Figure 4.3. A similar confidence interval was obtained for the case of total production correlation. In Figure 4.3, we observe that the correlation between the production forecast increments of the major agent and the price is lower in absolute value than the correlation between the total production forecast increments and the price. However, the gap between the correlations diminishes as we approach the delivery time.

4.6 Conclusion

In this paper, we applied the theory of linear quadratic mean-field games with a major player to the analysis of intraday electricity markets. The linear quadratic setting, while allowing to obtain explicit formulas for the equilibrium price and the optimal strategies of the agents, requires one to impose a number of quite stringent assumptions on the price dynamics, such as linear market impact and quadratic hedging cost. While these assumptions have been used in a number of papers on optimal trading and order execution in financial markets, their validity in electricity markets remains to be studied. Contrary to equity and bond markets, the literature on market impact and microstructure of electricity markets is in the nascent state and we hope that our paper will motivate more in-depth analysis of this topic, and more detailed studies of order book data of electricity markets. Another important aspect of our study is the

theoretical demonstration of the correlation between the renewable production forecasts and market prices. As the renewable penetration and the participation of renewable producers in intraday market increases, these correlations will become more important and may erode the profits of the renewable producers, impeding the investment flows into this important domain. One may need therefore to develop alternative market structures facilitating the participation of renewable producers.

4.7 Appendix

4.7.1 Proofs of Lemma 29 and Proposition 30

Proof of Lemma 29. Assume first that $(\phi^*, \bar{\phi}^*, \phi^{0*})$ is a Stackelberg mean field equilibrium. Then, for every $\mathbb{F}^0$-adapted square integrable process ν,

$$J^{MF}(\phi^* + \int_0^{\cdot} \nu_s ds, \overline{\phi}^*, \phi^{0*}) \leq J^{MF}(\phi^*, \overline{\phi}^*, \phi^{0*}).$$

Developing the functionals we get,

$$\begin{aligned}
&\mathbb{E}\left[\frac{1}{2}\int_0^T \alpha(t)\nu_t^2 dt + \frac{\lambda}{2}\left(\int_0^T \nu_t dt\right)^2\right] \\
&+ \mathbb{E}\left[\int_0^T \nu_t \left\{\alpha(t)\dot{\phi}_t^* + S_t + a\bar{\phi}_t^* + a_0\phi_t^{0*}\right\} dt + \lambda(\phi_T^* - X_T)\int_0^T \nu_t dt\right] \geq 0,
\end{aligned}$$

and since ν is arbitrary, we see that this is equivalent to

$$\mathbb{E}\left[\int_0^T \nu_t \left\{\alpha(t)\dot{\phi}_t^* + S_t + a\bar{\phi}_t^* + a_0\phi_t^{0*}\right\} dt + \lambda(\phi_T^* - X_T)\int_0^T \nu_t dt\right] = 0.$$

Taking conditional expectations and using Fubini's theorem, we then get condition i. of the lemma.

Assume now that conditions i. and ii. of the lemma hold true, and let ϕ^* be given by Proposition 28 applied to the couple $(\bar{\phi}^*, \phi^{0*})$. Define $\tilde{\phi}_t^* := \mathbb{E}[\phi_t^* | \mathcal{F}_t^0]$. It remains to show that $\tilde{\phi}_t^* = \bar{\phi}_t^*$. Let Y^* be an $\mathbb{F}^0$-martingale satisfying $Y_T^* = \lambda(\bar{\phi}_T^* - \overline{X}_T)$. By integration by parts, condition i. of the lemma is equivalent to

$$\mathbb{E}\left[\int_0^T \nu_t \left\{\alpha(t)\dot{\bar{\phi}}_t^* + S_t a_0\phi_t^{0*} + a\bar{\phi}_t^* + Y_t^*\right\} dt\right] = 0,$$

and since ν is arbitrary,

$$\alpha(t)\dot{\bar{\phi}}_t^* + S_t + a_0\phi_t^{0*} + a\bar{\phi}_t^* + Y_t^* = 0,$$

for all t. On the other hand, by Proposition 28, taking the expectation with respect to $\mathbb{F}^0$, we get that there exists a $\mathbb{F}^0$-martingale $\widetilde{Y}$ with $\widetilde{Y}_T = \lambda(\tilde{\phi}_T^* - \overline{X}_T)$, and such that

$$\alpha(t)\dot{\tilde{\phi}}_t^* + S_t + a_0\phi_t^{0*} + a\bar{\phi}_t^* + \widetilde{Y}_t = 0.$$

Substracting this expression from the previous one, we get

$$\alpha(t)(\dot{\bar{\phi}}_t^* - \dot{\tilde{\phi}}_t^*) + Y_t^* - \widetilde{Y}_t = 0, \qquad Y_T^* - \widetilde{Y}_T = \lambda(\bar{\phi}_T^* - \tilde{\phi}_T^*)$$

Thus,

$$\bar{\phi}_t^* - \tilde{\phi}_t^* = \int_0^t \frac{\widetilde{Y}_s - Y_s^*}{\alpha(s)} ds$$

and therefore, using the terminal condition and the martingale property,

$$\widetilde{Y}_t - Y_t^* = \mathbb{E}[\widetilde{Y}_T - Y_T^* | \mathcal{F}_t^0] = \lambda \int_0^t \frac{\widetilde{Y}_s - Y_s^*}{\alpha(s)} ds + \lambda(\widetilde{Y}_t - Y_t^*) \int_t^T \frac{ds}{\alpha(s)}.$$

The unique solution of this linear equation is $\widetilde{Y}_t = Y_t^*$ for all t, and therefore $\tilde{\phi}_t^* = \bar{\phi}_t^*$ for all t.

Proof of Proposition 30. The optimization problem of the major agent consists in maximizing the objective function (4.3) under the constraint of Lemma 29, part i. Following the methodology of Bensoussan et al. (2016a), let us introduce the Lagrangian for this constrained optimization problem, which writes:

$$\begin{aligned} L(\phi^0, \bar{\phi}, \nu) = \mathbb{E}\left[\int_0^T \frac{\alpha_0(t)}{2} \dot{\phi}_t^{0\,2} + (S_t + a\bar{\phi}_t + a^0\phi_t^0)\dot{\phi}_t^0 dt + \frac{\lambda_0}{2}(\phi_T^0 - X_T^0)^2\right] \\ + \mathbb{E}\left[\int_0^T \nu_t \left\{\alpha(t)\dot{\bar{\phi}}_t + S_t + a^0\phi_t^0 + a\bar{\phi}_t\right\} dt + \lambda(\bar{\phi}_T - \overline{X}_T) \int_0^T \nu_t dt\right], \end{aligned}$$

where ν is a square integrable $\mathbb{F}^0$-adapted process. We claim that ϕ^0 is the solution of the problem (4.3) if and only if there exist ν and $\bar{\phi}$ such that $(\phi^0, \bar{\phi})$ maximizes the Lagrangian $L(\cdot, \cdot, \nu)$, and $\bar{\phi}$ satisfies the constraint of Lemma 29. Indeed, let $(\phi^0, \bar{\phi}, \nu)$ be such a triple and $(\phi^{0\prime}, \bar{\phi}')$ be another pair of strategies satisfying the constraint of Lemma 29. Then,

$$L(\phi^0, \bar{\phi}, \nu) \geq L(\phi^{0\prime}, \bar{\phi}', \nu),$$

and since both $\bar{\phi}$ and $\bar{\phi}'$ satisfy the constraint of Lemma 29, this implies that inequality (4.7) holds true.

We now turn to the problem of maximizing the Lagrangian. Let $N_t = \int_0^t \nu_s ds$. The first order condition for ϕ_0 writes: there exists a martingale M^0 such that

$$M_t^0 + \alpha_0(t)\dot{\phi}_t^0 + S_t + a\bar{\phi}_t - a^0 N_t = 0, \qquad M_T^0 = a^0 N_T + a^0\phi_T^0 + \lambda_0(\phi_T^0 - X_T^0).$$

The first order condition for $\bar{\phi}$ writes: there exists a martingale M such that

$$M_t - a\phi_t^0 + \alpha(t)\dot{N}_t - aN_t = 0, \qquad M_T = a\phi_T^0 + (a + \lambda)N_T.$$

Finally, the last condition is given by the constraint that $\bar{\phi}$ is optimal for the generic minor

agent. Hence conditioning (4.8) by the common noise, there exists a martingale $\overline{Y}$ such that

$$\overline{Y}_t + \alpha(t)\dot{\bar{\phi}}_t + S_t + a\bar{\phi}_t + a^0\phi_t^0 = 0, \qquad \overline{Y}_T = \lambda(\bar{\phi}_T - \overline{X}_T).$$

4.7.2 Proof of Theorem 31

From equations (4.8) in Proposition 28 and (4.10) in Proposition 30, we immediately deduce the expression of the characterizing differential equation of the equilibrium (4.11).

Let $\Phi(t)$ be the fundamental matrix solution of the equation $B(t)^{-1}A\,\Xi_t + \dot{\Xi}_t = 0$, that is, for every $C \in \mathbb{R}^3$, the solution with initial condition $\Xi_0 = C$ is given by $\Phi(t)C$. By variation of constants we have that the solution of (4.11) is given by:

$$\Xi_t = -\Phi(t)\int_0^t \Phi(s)^{-1}\begin{pmatrix}\alpha_0(s)^{-1}(M_s^0 + S_s)\\ \alpha(s)^{-1}M_s\\ \alpha(s)^{-1}(\overline{Y}_s + S_s)\end{pmatrix} ds.$$

Letting:

$$\mathcal{M}_s := \begin{pmatrix}M_s^0\\ M_s\\ \overline{Y}_s\end{pmatrix} \quad \text{and} \quad \widehat{\Xi}_t = \Xi_t - \Upsilon_t,$$

we obtain the simplified equation:

$$\widehat{\Xi}_t = -\Phi(t)\int_0^t \Phi(s)^{-1}B(s)^{-1}\mathcal{M}_s ds.$$

and finally, using (4.12) and the martingale property, the martingale components satisfy:

$$\mathcal{M}_t = -D\Phi(T)\int_0^t \Phi(s)^{-1}B(s)^{-1}\mathcal{M}_s ds - D\Pi_{t,T}\mathcal{M}_t + D\widetilde{\Upsilon}_t - \Lambda\mathcal{X}_t,$$

From this we deduce, on the one hand,

$$\mathcal{M}_0 = (I + D\Pi_{0,T})^{-1}(D\widetilde{\Upsilon}_0 - \Lambda\mathcal{X}_0),$$

and on the other hand,

$$(I + D\Pi_{t,T})d\mathcal{M}_t = Dd\widetilde{\Upsilon}_t - \Lambda d\mathcal{X}_t,$$

so that finally:

$$\begin{aligned}\Xi_t &= \Upsilon_t + \int_0^t d\Pi_{s,t}\cdot\mathcal{M}_s = \Upsilon_t - \Pi_{0,t}\mathcal{M}_0 - \int_0^t \Pi_{s,t}d\mathcal{M}_s\\ &= \Upsilon_t - \Pi_{0,t}(I + D\Pi_{0,T})^{-1}(D\widetilde{\Upsilon}_0 - \Lambda\mathcal{X}_0) - \int_0^t \Pi_{s,t}(I + D\Pi_{s,T})^{-1}(Dd\widetilde{\Upsilon}_s - \Lambda d\mathcal{X}_s).\end{aligned}$$

4.7.3 Proofs of Corollary 34 and Corollary 35

Proof of Corollary 34. The first part is a simplification of the proof of Theorem 31. Using the expressions of Y_t and $\widetilde{Y}_t$ in Theorem 31 and the martingale property of S_t , we can rewrite:

$$Y_t = -\Pi_{0,t}S_0 - \int_0^t \Pi_{s,t}dS_s, \quad \widetilde{Y}_t = -\Pi_{0,T}S_0 - \int_0^t \Pi_{s,T}dS_s.$$

Substituting these expressions in the equation (4.13) we obtain the result.

For the second part, we can rewrite:

$$\begin{aligned}\Xi_t &= \Upsilon_t - \Pi_{0,t}(I + D\Pi_{0,T})^{-1}(D\widetilde{\Upsilon}_0 - \Lambda\mathcal{X}_0) - \int_0^t \Pi_{s,t}(I + D\Pi_{s,T})^{-1}(Dd\widetilde{\Upsilon}_s - \Lambda d\mathcal{X}_s)\\ &= \Upsilon_t - \Pi_{0,t}(D^{-1} + \Pi_{0,T})^{-1}(\widetilde{\Upsilon}_0 - D^{-1}\Lambda\mathcal{X}_0) - \int_0^t \Pi_{s,t}(D^{-1} + \Pi_{s,T})^{-1}(d\widetilde{\Upsilon}_s - D^{-1}\Lambda d\mathcal{X}_s)\end{aligned}$$

and when $\lambda, \lambda_0 \longrightarrow \infty$, $D^{-1} \to 0$ and $D^{-1}\Lambda \to D_\infty$. The third part follows by direct substitution of $\lambda = \lambda_0 = 0$ into the general formula.

Proof of Corollary 35. Let $\check{\phi}^*_t = \phi^*_t - \bar{\phi}^*_t$, $\check{X}_t = X_t - \overline{X}_t$ and $\check{Y}_t := Y_t - \overline{Y}_t$. Then, from the explicit form of Y and $\overline{Y}$ in Proposition 28 it follows that $\check{Y}$ is an $\mathbb{F}$-martingale and satisfies

$$\check{Y}_T = -\lambda(\check{\phi}^*_T - \check{X}_T), \qquad \check{Y}_t = \alpha(t)\dot{\check{\phi}}^*_t.$$

Then,

$$\check{\phi}^*_t = \int_0^t \frac{\check{Y}_s}{\alpha(s)}ds, \tag{4.26}$$

and by the martingale property,

$$Y_t = -\lambda\mathbb{E}[\check{\phi}^*_T - \check{X}_T|\mathcal{F}_t] = -\lambda\int_0^t \frac{\check{Y}_s}{\alpha(s)}ds - \lambda\check{Y}_t\int_t^T \frac{ds}{\alpha(s)} + \lambda\check{X}_t.$$

Solving this linear equation for $\check{Y}$ then substituting into (4.26), we get the result.

4.7.4 Proof of Proposition 38

We need to show conditions i. and ii. of Definition 37. Condition i. is shown in the same way as in the case of homogeneous players (see proof of Proposition 2 in Féron et al. (2020)). We therefore focus on condition ii. Assume that all agents change their strategies to new ones $\phi^0, \ldots, \phi^N$, such that $\phi^1, \ldots, \phi^N$ are optimal responses to ϕ^0. Let $\bar{\phi}$ be the optimal "mean field" response to the major agent strategy ϕ^0.

Step 1. We first suppose that there exists a finite constant $A > 0$, independent of N, such that $\mathbb{E}\left[\int_0^T (\dot{\phi}^0_t)^2 dt\right] < A$.

By Proposition 1 in Féron et al. (2020), for some constants c and C which do not depend on N, and may change from line to line,

$$\mathbb{E}\left[\int_0^T (\bar{\phi}_t - \bar{\phi}_t^N)^2 dt\right] \leq \frac{C}{N^2}\mathbb{E}\left[\int_0^T (S_s + a^0\phi_s^0)^2 ds\right] + \frac{c}{N},$$

and by our assumption,

$$\begin{aligned}
&\mathbb{E}\left[\int_0^T (S_s + a^0\phi_s^0)^2 ds\right] \\
&= \mathbb{E}\left[\int_0^T S_s^2 ds\right] + 2a^0\mathbb{E}\left[\int_0^T S_s^2 ds\right]^{\frac{1}{2}} \mathbb{E}\left[\int_0^T (\phi_s^0)^2 ds\right]^{\frac{1}{2}} + \mathbb{E}\left[\int_0^T (a^0\phi_s^0)^2 ds\right] < C(1+A).
\end{aligned}$$

Thus using Cauchy-Schwartz inequality,

$$\begin{aligned}
&J^{N,0}(\phi^0, \phi^{-0}) - J^{N,0}(\phi^{0*}, \phi^{-0*}) \\
&= J^{N,0}(\phi^0, \phi^{-0}) - J^{MF}(\phi^0, \bar{\phi}) + J^{MF}(\phi^0, \bar{\phi}) - J^{MF}(\phi^{0*}, \bar{\phi}^*) + J^{MF}(\phi^{0^*}, \bar{\phi}^*) - J^{N,0}(\phi^{0*}, \phi^{-0*}) \\
&\leq J^{N,0}(\phi^0, \phi^{-0}) - J^{MF}(\phi^0, \bar{\phi}^0) + J^{MF}(\phi^{0^*}, \bar{\phi}^*) - J^{N,0}(\phi^{0*}, \phi^{-0*}) \\
&\leq a\left\{\mathbb{E}\left[\int_0^T (\dot{\phi}_t^0)^2 dt\right]^{\frac{1}{2}} \mathbb{E}\left[\int_0^T (\bar{\phi}_t - \bar{\phi}_t^N)^2 dt\right]^{\frac{1}{2}} + \mathbb{E}\left[\int_0^T (\dot{\phi}_t^{0*})^2 dt\right]^{\frac{1}{2}} \mathbb{E}\left[\int_0^T (\bar{\phi}_t^* - \bar{\phi}_t^{N*})^2 dt\right]^{\frac{1}{2}}\right\} \\
&= \mathcal{O}\left(N^{-\frac{1}{2}}\right)
\end{aligned}$$

Step 2. Let there exist a sufficiently large constant A, independent of N, such that

$$\mathbb{E}\left[\int_0^T \left(\dot{\phi}_t^0\right)^2 dt\right] > A$$

Letting $\bar{\alpha}_0 = \min_{0\leq t\leq T} \alpha_0(t)$, and $\bar{\alpha} = \min_{0\leq t\leq T} \alpha(t)$, we have, by definition,

$$J^{N,0}(\phi^0, \phi^{-0}) \leq -\mathbb{E}\left[\int_0^T \frac{\bar{\alpha}_0}{2}(\dot{\phi}_t^0)^2 + \dot{\phi}_t^0(S_t + a\bar{\phi}_t^N + a_0\phi_t^0)dt + \frac{\lambda_0}{2}(\phi_T^0 - X_T)^2\right]$$

On the other hand, since ϕ^i for $i = 1, \ldots, N$ are optimal responses to ϕ^0, we get that

$$\mathbb{E}\left[\int_0^T \left\{\frac{\alpha(t)}{2}(\dot{\phi}_t^i)^2 + \dot{\phi}_t^i(S_t + a\bar{\phi}_t^N + a_0\phi_t^0)\right\} dt + \frac{\lambda}{2}(\phi_T^i - X_T^i)^2\right] \leq \frac{\lambda}{2}\mathbb{E}[(X_T^i)^2] \leq C,$$

where $C < \infty$ is defined by $C := \max_i \frac{\lambda}{2}\mathbb{E}[(X_T^i)^2]$. Summing up the above inequality over $i = 1, \ldots, N$, dividing by N and using Jensen's inequality, we get

$$\mathbb{E}\left[\int_0^T \left\{\frac{\bar{\alpha}}{2}(\dot{\bar{\phi}}_t^N)^2 + \dot{\bar{\phi}}_t^N(S_t + a\bar{\phi}_t^N + a_0\phi_t^0)\right\} dt + \frac{\lambda}{2}(\bar{\phi}_T^N - \overline{X}_T^N)^2\right] \leq C.$$

Multiplying this inequality by $\frac{a}{a_0}$, adding it to the first one, and using integration by parts, we

finally get

$$
\begin{aligned}
J^{N,0}(\phi^0,\phi^{-0}) \leq C - \mathbb{E}\Bigg[\int_0^T \left\{ \frac{\bar{\alpha}_0}{2}(\dot{\phi}_t^0)^2 + \frac{a\bar{\alpha}}{2a_0}(\dot{\bar{\phi}}_t^N)^2 + \frac{S_t}{a_0}(a_0\dot{\phi}_t^0 + a\dot{\bar{\phi}}_t^N) \right\} dt \\
+ \frac{1}{2a_0}(a\bar{\phi}_T^N + a_0\phi_T^0)^2 + \frac{\lambda_0}{2}(\phi_T^0 - X_T)^2 + \frac{a\lambda}{2a_0}(\bar{\phi}_T^N - \overline{X}_T^N)^2 \Bigg] \\
\leq C - \mathbb{E}\left[\int_0^T \frac{\bar{\alpha}_0}{2}(\dot{\phi}_t^0)^2 dt\right] + \mathbb{E}\left[\int_0^T S_t^2 dt\right]^{\frac{1}{2}} \mathbb{E}\left[\int_0^T (\dot{\phi}_t^0)^2 dt\right]^{\frac{1}{2}} \\
- \mathbb{E}\left[\int_0^T \frac{a\bar{\alpha}}{2a_0}(\dot{\bar{\phi}}_t^N)^2 dt\right] + \frac{a}{a_0}\mathbb{E}\left[\int_0^T S_t^2 dt\right]^{\frac{1}{2}} \mathbb{E}\left[\int_0^T (\dot{\bar{\phi}}_t^N)^2 dt\right]^{\frac{1}{2}} \\
\leq C - \mathbb{E}\left[\int_0^T \frac{\bar{\alpha}_0}{2}(\dot{\phi}_t^0)^2 dt\right] + \mathbb{E}\left[\int_0^T S_t^2 dt\right]^{\frac{1}{2}} \mathbb{E}\left[\int_0^T (\dot{\phi}_t^0)^2 dt\right]^{\frac{1}{2}} + \frac{a}{\bar{\alpha}a_0}\mathbb{E}\left[\int_0^T S_t^2 dt\right]
\end{aligned}
$$

Thus, if

$$
\mathbb{E}\left[\int_0^T (\dot{\phi}_t^0)^2 dt\right] > A,
$$

for A sufficiently large (but not depending on N), then from the above estimate it follows that

$$
J^{N,0}(\phi^0,\phi^{-0}) \leq J^{N,0}(\phi^{0*},\phi^{-0*}),
$$

as well.

Chapter 5

Decision making with dynamic probabilistic forecasts

Abstract

We consider a sequential decision making process such as renewable energy trading or electrical production scheduling whose outcome depends on the future realization of a random factor, such as a meteorological variable. We assume that the decision maker disposes of a dynamically updated probabilistic forecast (predictive distribution) of the random factor. We propose several stochastic models for the evolution of the probabilistic forecast, and show how these models may be calibrated from ensemble forecasts, commonly provided by weather centers. We then show how these stochastic models can be used to determine optimal decision making strategies depending on the forecast updates. Applications to wind energy trading are given.

5.1 Introduction

Consider a sequential decision-making process, such as renewable energy trading or electrical production scheduling, whose outcome depends on the realization of a random factor, such as a meteorological variable. It is often the case, that at each point in time, the decision maker disposes of an imperfect probabilistic forecast of the random factor (such as, a confidence interval or a set of quantiles), and that this forecast is periodically, or continuously, updated. The goal of the decision maker is to optimally update her strategy according to the available information, to maximize a specific gain functional. To solve this problem in the framework of stochastic control, one needs to describe the dynamics of the predictive distribution with a stochastic model. Such a model determines the evolution of the predictive distribution and the relationship of the forecasts to the realization of the unknown random factor; in other words, the model describes the evolution of the forecast error as new information becomes available.

In the literature, stochastic decision update rules based on point forecasts have been proposed Tan and Tankov (2018); Collet et al. (2017); Aid et al. (2016); Skajaa et al. (2015), however probabilistic forecasts contain more dynamic information than point forecasts, as the expected forecast uncertainty can also vary dynamically. Figure 5.1 shows the evolution of the probabilistic forecast of power production of a wind plant in France as function of time, for a fixed production time.

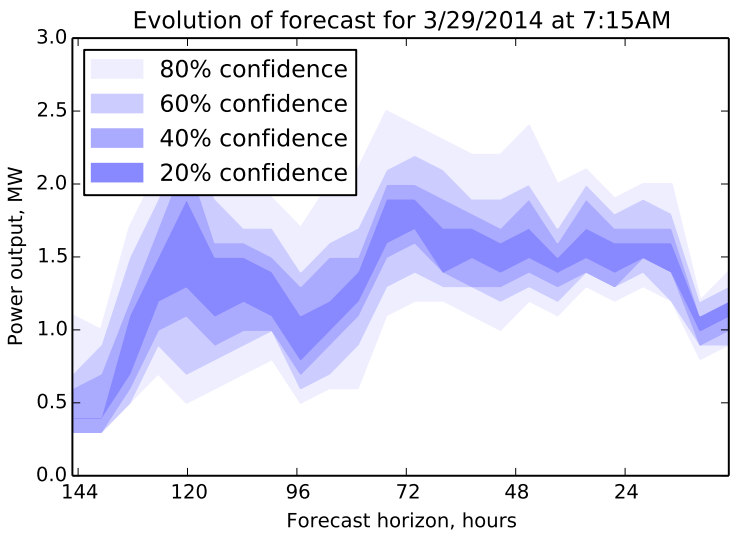

Figure 5.1: Evolution of the probabilistic forecast of power production of a wind plant as function of time, for a fixed production time.

It is clear that not only the average production varies with time, but also the width of the confidence interval changes: it does not always decrease with time and may not be fully correlated with the expected production level. This information reflects the varying forecast uncertainty and is not contained in the point forecast, but may be important for decision making. For example, a wind producer facing severe penalties in case of lack of production, or a network operator whose goal is to avoid shortages at all costs may need to purchase energy in the

intraday market to hedge the risk when forecast uncertainty increases, even if the predicted average production remains the same. This paper develops models of the dynamic evolution of probabilistic forecasts, allowing to take into account precisely this type of uncertainty in dynamic decision making.

In mathematical terms, let $(\Omega, \mathcal{F}, \mathbb{F}, \mathbb{P})$ be a filtered probability space and assume that $\mathbb{F}$ models the filtration of the decision maker. Fix a time horizon T, and let X be a real-valued $\mathcal{F}_T$-measurable random variable. We make a standing assumption that $\mathcal{F}_0$ is a trivial σ-field. Let μ_t denote the regular conditional distribution of X given $\mathcal{F}_t$. We call μ_t the probabilistic forecast of X at time t. See Gneiting et al. (2007) for the description of the mathematical framework of probabilistic forecasting and methods of forecast evaluation. The goal of this paper is to

a. Formulate the conditions that the dynamics of μ_t must satisfy and propose several tractable finite-dimensional models for this dynamics in the diffusion framework,

b. Show how these models may be calibrated with real meteorologic data, in the case where μ_t models the forecast of a meteorological variable.

c. Provide an example of using the methodology to solve stochastic control problems arising in the context of wind energy trading.

The full predictive distribution is an infinite-dimensional object, but the actual available information is always low-dimensional; for this reason we aim to summarize the dynamics of the full predictive distribution with a low number of factors, which are easy to interpret and estimate from the data (such as the conditional mean and variance of the predictive distribution). In addition, our objective of computing the optimal strategies using the tools of stochastic control precludes the use of high-dimensional specifications. More precisely, in this paper we consider parametric two-dimensional specifications where the predictive distribution is a function of two observable factors, say m_t and V_t. Here m_t represents the conditional expectation of X and V_t some measure of the error, such as the conditional variance. In our models, m_t and V_t have diffusion dynamics, and the predictive density μ_t corresponds to a distribution from some known class, such as Student t, normal inverse Gaussian, inverse Gaussian or log generalized hyperbolic, with parameters depending on m_t and V_t.

In practice, the forecast information received by the decision maker from a forecast provider may come, for example, in the form of a confidence interval around a point forecast, or in the form a set of quantiles of the predictive distribution. A particularly important case is that of ensemble forecasts. An ensemble forecast in meteorology is a set of several point forecasts aiming together to give an indication of the range of possible future states of the atmosphere. Members of the ensemble are obtained by running the forecasting model with perturbed initial conditions and / or parameters. An ensemble forecast is usually obtained with deterministic means, and therefore does not represent the best approximation of the predictive distribution of meteorological variables. In particular, ensemble forecasts are often uncalibrated (biased) and underdispersed compared to realizations Gneiting et al. (2005). However, techniques for statistical post-processing of ensemble forecast with the aim to improve calibration and sharpness

have been developed in the literature. Two such techniques are ensemble model output statistics (EMOS) Gneiting et al. (2005); Thorarinsdottir and Gneiting (2010) and Bayesian model averaging (BMA) Wilks (2002); Raftery et al. (2005). In Gneiting et al. (2005), the authors approximate the predictive density with a Gaussian distribution, whose parameters depend on the ensemble forecasts and are chosen to optimize calibration and sharpness of the resulting probabilistic forecast. In Wilks (2002); Raftery et al. (2005) the predictive density is reprensented by a mixture of normal distributions, whose weights are computed from the ensemble members. To account for positive random variables such as wind speed, EMOS with log-normal distributions has been used in Baran and Lerch (2015) and BMA with truncated normal components has been employed in Baran (2014). Other approaches to statistical post-processing of wind speed forecasts involve generalized extreme value distribution Lerch and Thorarinsdottir (2013) and weighted mixtures of log-normal and truncated normal distributions Baran and Lerch (2016).

In these papers, a single forecast horizon is fixed, and the calibration procedure uses a series of ensemble forecasts, obtained at different days of the training period for the fixed forecast horizon. At any given time, the calibrated method allows to compute the probabilistic forecast for this fixed horizon from the ensemble forecast, but no information about the evolution of the probabilistic forecast is available.

Our approach to calibrate the models presented in this paper is inspired by EMOS and also based on ensemble forecasts. However, we use more general predictive densities, potentially allowing for better calibration. More importantly, we do not fix a single forecast horizon, but model the dynamics of the predictive distribution for a given quantity at a given date, as time goes on and forecast horizon decreases. As a result, our calibrated model provides two types of information. First, as in statistical postprocessing methods, a predictive density in tractable form can be computed from an ensemble forecast. Secondly, the dynamics of this predictable distribution is given, in the form of a two-dimensional stochastic differential equation characterizing the evolution of the pair (m_t, V_t), the conditional mean of the predictive distribution and a measure of the error. This dynamics can be exploited in the decision making process, to make strategy updates based not only on the conditional mean of the variable of interest, but also on the evolution of our knowledge of the uncertainty around the mean.

Stochastic differential equations (SDE) have been used to model the dynamics of probabilistic forecasts by several authors, see e.g., Iversen et al. (2016); Bensoussan and Brouste (2016) in the context of wind speed, or Badosa et al. (2017) in the context of solar energy forecasting. In these approaches, the forecasted quantity (e.g., the wind speed) is modeled directly by a stochastic differential equation, from which the probabilistic forecasts at any horizons, as well as their dynamics, can be deduced. However, the predictive distributions are typically not in tractable form (e.g., in Iversen et al. (2016) they are approximated by Monte Carlo), and the dynamics of the forecasting error is hard-coded into the equation and cannot be calibrated independently from ensemble forecasts, in other words, the variance of the forecast is not stochastic. This makes it impossible to use information on forecast uncertainty in strategy updates.

Our approach provides a dynamic SDE-based model for forecast dynamics, tractable predictive distribution and possibility of model calibration with ensemble forecasts, in a sense taking

the best of both worlds to obtain a coherent and realistic model. Moreover, the results are exploited in a stochastic control problem to integrate the additional information provided by the probabilistic forecasts in the decision process.

Using probabilistic forecasts for decision making in wind energy trading and electricity scheduling has been studied e.g., in Pinson et al. (2007, 2013); Zugno et al. (2013a). These references suggest a static approach, where a probabilistic forecast of a quantity of interest is used to make the decision on e.g., the quantity of energy to sell in the day-ahead market. By contrast, our dynamic approach allows to continuously, or regularly, update the decision based on the evolution of the forecast and information about its uncertainty.

We illustrate our methodological contribution with an application to a wind power trading problem. In this problem, a wind power producer, who disposes of a dynamically updated probabilistic forecast of the wind speed, takes positions in the intraday electricity market to maximize the utility of terminal wealth. This setting gives rise to a three-dimensional stochastic control problem, which is solved using the dynamic programming principle. For the numerical solution we use the Least Squares Monte Carlo method implemented in the open-source library StOpt (see Gevret et al. (2018)) and based on the methods of Bouchard and Warin (2012) and Belomestny et al. (2010) generalizing the seminal approach of Longstaff and Schwartz (2001) and Tsitsiklis and Van Roy (1999). To assess the value of taking into account the dynamics of probabilistic forecasts, we compare the gains of an agent using our approach with the potential gains of another agent who uses only the point forecasts and show that our method leads to a 5% revenue increase in the simulation examples.

The paper is structured as follows. In Section 5.2 we describe several parametric models for the dynamics of probabilistic forecasts. In Section 5.3, we develop a procedure inspired by EMOS to calibrate the models of Section 5.2 for different lead times and show that our models have good prediction results and that EMOS increases accuracy of prediction compared with raw ensembles, as expressed with Continuous Ranked Probability Score. In Section 5.4, we present an application of our methodology to wind power trading.

5.2 Modeling probabilistic forecasts

As mentioned in the introduction, given a flow of information described by the filtration $(\mathcal{F}_t)_{0\leq t\leq T}$, a probabilistic forecast of an $\mathcal{F}_T$-measurable random variable $X \in \mathbb{R}^d$ is the conditional distribution μ_t of X given $\mathcal{F}_t$. A dynamic model for a probabilistic forecast is then a flow of probability measures $(\mu_t)_{0\leq t\leq T}$, which can be identified with a flow of conditional distributions of some $\mathcal{F}_T$-measurable random variable. This imposes strong constraints on the dynamics of μ_t, in particular, all moments of μ_t, when they exist, must be $(\mathcal{F}_t)$-martingales. A d-dimensional Markov specification of forecast dynamics is a Markov process $(X_t)_{0\leq t\leq T} \in \mathbb{R}^d$ such that, at every $t \in [0, T]$, $\mu_t = \mu(t, X_t)$, where $\mu : [0, T] \to \mathbb{R}^d \to \mathcal{P}(\mathbb{R})$ is a deterministic mapping, where $\mathcal{P}(\mathbb{R})$ is the set of probability measures on $\mathbb{R}$.

In this section, we develop several two-dimensional Markov specifications for forecast dy-

namics, which correspond to well-known tractable predictive distributions.

5.2.1 Forecast of a real-valued quantity

In this section we propose two tractable models for the dynamics of probabilistic forecast of a real-valued quantity, such as the temperature. The models are based on the time-changed Brownian motion. In the first paragraph, the predictive distribution at all times is the Student t distribution (with power law tails), and in the second paragraph, the predictive distribution is the normal inverse Gaussian distribution (with exponentially decaying tails).

Student t predictive distribution Let ρ be a positive deterministic function, continuous on $(0,\infty)$, with $\int_0^t \rho^2(t)dt = +\infty$ for all $t > 0$ (this function is singular at zero), let W and W' be independent standard Brownian motions, whose filtration will be denoted by $\mathbb{F}$, let $b > 0$ and consider the following pair of stochastic differential equations, defined for $t \in [0,T)$:

$$\frac{dV_t}{V_t} = -\rho^2(T-t)dt + b\rho(T-t)dW_t \tag{5.1}$$

$$dm_t = \sqrt{V_t}\rho(T-t)dW'_t, \tag{5.2}$$

Proposition 40. *The equation (5.1–5.2) admits a strong solution (m,V) on $[0,T)$. The limit $m_T = \lim_{t\to T} m_t$ exists in the almost sure sense, and for every $t \in [0,T)$, the conditional distribution of m_T given $\mathcal{F}_t$ is the Student t distribution with 2ν degrees of freedom, where $\nu = 1 + 2/b^2$:*

$$\begin{aligned}\frac{d}{dx}\mathbb{P}[m_T \in dx|\mathcal{F}_t] &= \frac{d}{dx}\mathbb{P}[m_T \in dx|m_t, V_t] \\ &= \frac{\Gamma(\nu+\frac{1}{2})}{\Gamma(\nu)}\frac{b}{2\sqrt{\pi V_t}}\left\{1 + \frac{(x-m_t)^2 b^2}{4V_t}\right\}^{-\nu-\frac{1}{2}}.\end{aligned}$$

In addition,

$$m_t = \mathbb{E}[m_T|\mathcal{F}_t] \quad \text{and} \quad V_t = \mathrm{Var}[m_T|\mathcal{F}_t].$$

Remark 41. *In this model, and in the other models of this section, the predictive distribution is parameterized by two (stochastic) variable parameters, m_t and V_t which typically determine the location and scale of the distribution and may change as time passes and the forecast horizon draws near, and one fixed parameter b (typically, the shape parameter), which remains constant throughout the lifetime of the forecast for a fixed date. In addition, the deterministic time-varying parameter ρ does not affect the predictive distribution, but affects the dynamics of the variables m_t and V_t.*

Remark 42. *The pair $(V_t, m_t)_{0\le t<T}$ can alternatively be written as $V_t = \widetilde{V}_{\theta_t}$ and $m_t = \widetilde{m}_{\theta_t}$ with*

$$\theta_t = \int_0^t \rho^2(T-s)ds, \quad 0 \le t < T,$$

and

$$\widetilde{V}_t = V_0 e^{-(1+\frac{b^2}{2})t + b\widetilde{W}_t}$$
$$\widetilde{m}_t = m_0 + \int_0^t \sqrt{\widetilde{V}_s} d\widetilde{W}'_s$$

on $[0,\infty)$*, with* $(\widetilde{W}, \widetilde{W}')$ *a standard* 2*-dimensional Brownian motion.*

Proof. Fixing $t < T$, as in the above remark, for $s \geq t$, we can write

$$V_s = V_t \exp\left(-\int_t^s (1+b^2/2)\rho^2(T-u)du + \int_t^s b\rho(T-u)dW_u\right) = V_t \overline{V}^{(-1-b^2/2,b)}_{\int_t^s \rho(T-u)^2 du}$$

where $\overline{V}_t^{(\mu,b)} = e^{\mu t + bB'_t}$ for a different Brownian motion B'. In addition,

$$\int_t^T \rho^2(T-s)V_s dt = V_t \int_t^T \rho^2(T-s)\overline{V}^{(-1-b^2/2,b)}_{\int_t^s \rho(T-u)^2 du} ds = V_t \int_0^\infty \overline{V}_s^{(-1-b^2/2,b)} ds$$
$$= \frac{4V_t}{b^2}\int_0^\infty \overline{V}_s^{-(4/b^2-2,2)} ds.$$

From (Dufresne, 1990, Proposition 4.4.4),

$$\int_0^\infty \overline{V}_t^{(-4/b^2-2,2)} dt \stackrel{d}{=} (2\gamma_\nu)^{-1}, \quad \nu = 2/b^2 + 1,$$

where γ_ν denotes a gamma random variable with parameter ν.

Then, for $t < s < T$,

$$m_s = m_t + B\left(\int_t^s V_r \rho^2(T-r)dr\right),$$

for a different Brownian motion B. Therefore,

$$m_T := \lim_{s\to T} m_s = m_t + B\left(\int_t^T V_r\rho^2(T-r)dr\right) = B\left(\frac{2V_t}{b^2\gamma_\nu}\right)$$

and finally

$$\frac{d}{dx}\mathbb{P}[m_T - m_t \in dx|\mathcal{F}_t] = \mathbb{E}\left[\frac{b\sqrt{\gamma_\nu}}{2\sqrt{\pi V_t}} e^{-\frac{x^2b^2\gamma_\nu}{4V_t}}\right]$$
$$= \frac{b}{2\Gamma(\nu)\sqrt{\pi V_t}}\int_0^\infty e^{-z-\frac{x^2b^2z}{4V_t}} z^{\nu-\frac{1}{2}} dz$$
$$= \frac{\Gamma(\nu+\frac{1}{2})}{\Gamma(\nu)}\frac{b}{2\sqrt{\pi V_t}}\left\{1+\frac{x^2b^2}{4V_t}\right\}^{-\nu-\frac{1}{2}},$$

which means that conditionnally on $\mathcal{F}_t$, $m_T - m_t$ follows the centered Student t distribution with $2\nu = 2 + 4/b^2$ degrees of freedom. The expressions for mean and variance are obtained from standard formulas for the Student t distribution. □

Figure 5.2 illustrates the dynamics of the predictive distribution in the model (5.1–5.2).

We see that the confidence interval has a nontrivial behavior, it shrinks around the middle of the graph as the conditional variance goes down before increasing in size when the conditional variance goes up and shrinking to zero again at the very end.

Figure 5.2: Confidence intervals for the Student predictive density, and the trajectories of the predictive distribution mean m_t and variance V_t.

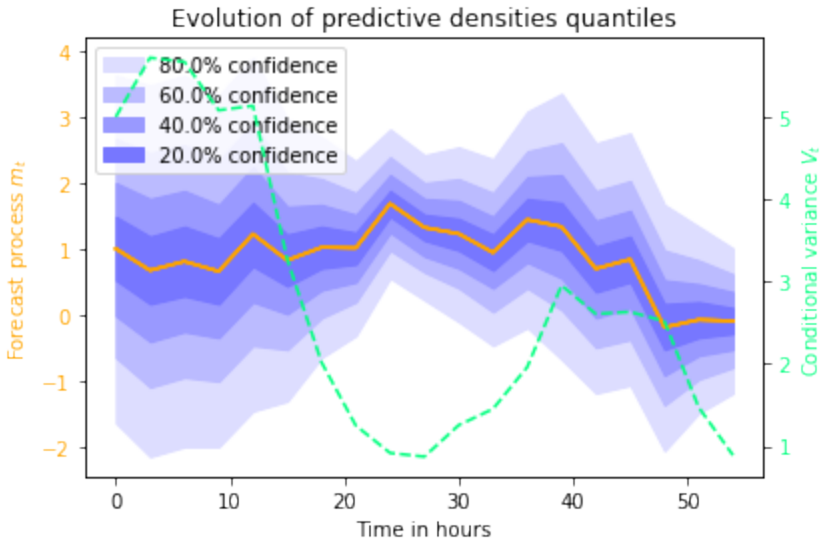

Normal inverse Gaussian predictive distribution Using the notation of the previous paragraph, consider the following pair of stochastic differential equations.

$$dm_t = \sqrt{V_t}\rho(T-t)dW_t \tag{5.3}$$

$$dV_t = -V_t\rho^2(T-t)dt + \sqrt{V_t}b\rho(T-t)dW_t'. \tag{5.4}$$

Here V is a time-changed square-root process which hits zero in finite time almost surely. We assume that this process remains at zero after the first hitting time. A sample evolution of (m, V) and the dynamics of the associated predictive distribution is shown in Figure 5.3. As in Remark 42, we can express $m_t = \widetilde{m}_{\theta_t}$ and $V_t = \widetilde{V}_{\theta_t}$, where the processes $\widetilde{m}$ and $\widetilde{V}$ have time-homogeneous dynamics (with different Brownian motions).

$$d\widetilde{m}_t = \sqrt{\widetilde{V}_t}dW_t \tag{5.5}$$

$$d\widetilde{V}_t = -\widetilde{V}_t dt + \sqrt{\widetilde{V}_t}bdW_t', \tag{5.6}$$

Proposition 43. *The equation (5.3–5.4) admits a strong solution (m, V) on $[0, T)$. The limit $m_T = \lim_{t\to T} m_t$ exists in the almost sure sense, and for every $t \in [0, T)$, the conditional distribution of m_T given $\mathcal{F}_t$ is the symmetric normal inverse Gaussian distribution on $\mathbb{R}$ with*

density

$$p(x) = \frac{\frac{V_t}{b^2} K_1\left(\frac{1}{b}\sqrt{(V_t/b)^2 + (x - m_t)^2}\right)}{\pi\sqrt{(V_t/b)^2 + (x - m_t)^2}} e^{V_t/b^2} \tag{5.7}$$

where K is the modified Bessel function of the third kind. Moreover, the conditional mean and variance of m_T are given by

$$m_t = \mathbb{E}[m_T|\mathcal{F}_t] \quad \textit{and} \quad \mathrm{Var}\,[m_T|\mathcal{F}_t] = V_t.$$

Proof. For the existence of the strong solution to (5.5)–(5.6), see (Jeanblanc et al., 2009, Section 6.3.1) From Remark 42 it follows that for $s \geq \theta_t$,

$$\tilde{m}_s - m_t = \int_{\theta_t}^{s} \sqrt{\widetilde{V}_u} dW_u = \widetilde{W}_{\int_{\theta_t}^{s} \widetilde{V}_u du}$$

for a different Brownian motion $\widetilde{W}$. In particular

$$m_T = \widetilde{m}_\infty = m_t + \widetilde{W}_{\int_{\theta_t}^{\infty} \widetilde{V}_s ds}.$$

$\tilde{V}$ is a square root process with zero long-term mean. The Laplace transform of the integrated square root process is known (Jeanblanc et al., 2009, Proposition 6.3.4.1):

$$\mathbb{E}\left[\exp\left(-u\int_{\theta_t}^{s} \widetilde{V}_u du\right)\Big|\widetilde{V}_{\theta_t}\right] = \exp\left(-\frac{2\widetilde{V}_{\theta_t} u}{1 + \gamma \coth\frac{\gamma(s-\theta_t)}{2}}\right),$$

where $\gamma = \sqrt{1 + 2ub^2}$. Integrating up to infinity, we then find:

$$\mathbb{E}\left[\exp\left(-u\int_{\theta_t}^{\infty} \widetilde{V}_s ds\right)\Big|\widetilde{V}_{\theta_t}\right] = \exp\left(-\frac{2\widetilde{V}_{\theta_t} u}{1 + \sqrt{1 + 2ub^2}}\right).$$

This allows us to compute the Fourier transform of the conditional distribution of m_T:

$$\begin{aligned}
\mathbb{E}[e^{iu(m_T - m_t)}|\mathcal{F}_t] &= \mathbb{E}\left[\exp\left(iu\widetilde{W}_{\int_{\theta_t}^{\infty} \widetilde{V}_s ds}\right)|\widetilde{V}_{\theta_t}\right] \\
&= \mathbb{E}\left[\exp\left(-\frac{u^2}{2}\int_{\theta_t}^{\infty} \widetilde{V}_s ds\right)\Big|\widetilde{V}_{\theta_t}\right] \\
&= \exp\left(-\frac{\widetilde{V}_{\theta_t} u^2}{1 + \sqrt{1 + u^2 b^2}}\right) \\
&= \exp\left(-\frac{V_t(\sqrt{1 + u^2 b^2} - 1)}{b^2}\right)
\end{aligned}$$

The characteristic function of the normal inverse Gaussian law with parameters μ, α, β, δ (Barndorff-Nielsen, 1997) is given by

$$e^{i\mu u + \delta(\sqrt{\alpha^2 - \beta^2} - \sqrt{\alpha^2 - (\beta + iu)^2})}.$$

Hence, $m_T - m_t$ conditionnally on $\mathcal{F}_t$ follows the normal inverse Gaussian law with parameters

$$\mu = 0, \quad \delta = \frac{V_t}{b}, \quad \beta = 0, \quad \alpha = \frac{1}{b}.$$

The expressions of the conditional moments may be easily obtained from the characteristic function. □

Figure 5.3: Confidence intervals for the NIG predictive density, and the trajectories of the predictive distribution mean m_t and variance V_t.

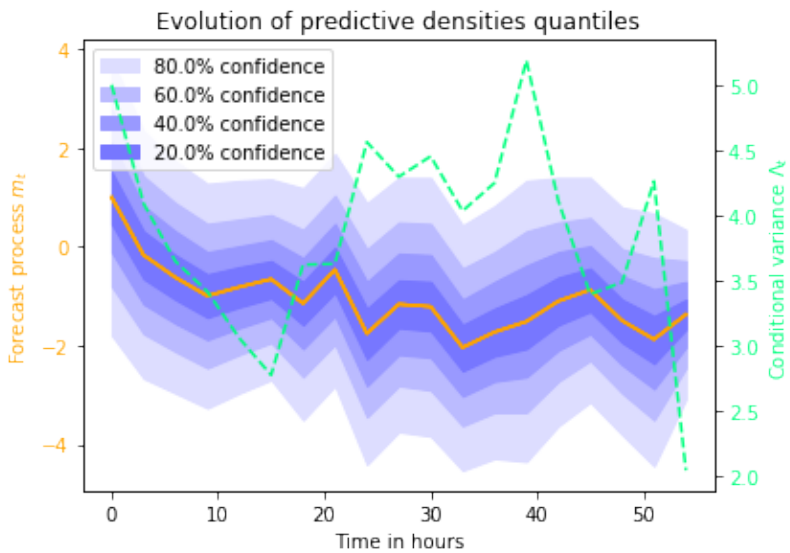

5.2.2 Forecast of a positive quantity

In this section we propose two models for the probabilistic forecast of a positive quantity such as the wind speed. The models are obtained from the ones of the previous section, replacing the Brownian motion with the martingale geometric Brownian motion.

Log-generalized hyperbolic predictive distribution Using the notation of the preceding section, consider the following pair of stochastic differential equations.

$$\frac{dV_t}{V_t} = -\rho^2(T-t)dt + b\rho(T-t)dW_t \tag{5.8}$$

$$\frac{dm_t}{m_t} = \sqrt{V_t}\rho(T-t)dW'_t. \tag{5.9}$$

As in Remark 42, we can write $m_t = \widetilde{m}_{\theta_t}$ and $V_t = \widetilde{V}_{\theta_t}$, where the processes $\widetilde{m}$ and $\widetilde{V}$ have time-homogeneous dynamics (with different Brownian motions).

$$\frac{d\widetilde{m}_t}{\widetilde{m}_t} = \sqrt{\widetilde{V}_t} dW_t$$
$$\frac{d\widetilde{V}_t}{\widetilde{V}_t} = -dt + b dW'_t.$$

Proposition 44. *Let (m, V) be a solution of (5.8–5.9). Then the conditional distribution of $\log m_T$ given $\mathcal{F}_t$ is the generalized hyperbolic distribution with density*

$$p(x) = \frac{b e^{\frac{x-\mu}{2}}}{\Gamma(\nu)\sqrt{\pi V_t}} \left(\frac{V_t}{b\sqrt{4V_t + (x-\mu)^2 b^2}} \right)^{\nu + \frac{1}{2}} K_{\nu+\frac{1}{2}} \left(\sqrt{\frac{V_t}{b^2} + \frac{(x-\mu)^2}{4}} \right)$$

with $\mu = \log m_t$ and $\nu = 1 + \frac{2}{b^2}$. In addition,

$$\mathbb{E}[m_T | \mathcal{F}_t] = m_t.$$

Remark 45. *This distribution is a particular case of the generalized hyperbolic distribution, known as generalized hyperbolic skew Student t distribution (Aas and Haff, 2006). With this distribution, m_T does not admit a second moment.*

Proof. With the notation of the proof of Proposition 40, we now get

$$\log \frac{m_T}{m_t} = W\left(\frac{2V_t}{b^2 \gamma_\nu}\right) - \frac{V_t}{b^2 \gamma_\nu},$$

and therefore

$$\begin{aligned}
\frac{d}{dx}\mathbb{P}[\log(m_T/m_t) \in dx | \mathcal{F}_t] &= \frac{d}{dx}\mathbb{E}[\mathbb{P}[\log(m_T/m_t) \in dx | \gamma_\mu, V_t]] \\
&= \mathbb{E}\left[\frac{b\sqrt{\gamma_\mu}}{2\sqrt{\pi V_t}} e^{-\frac{(x - \frac{V_t}{b^2\gamma_\mu})^2 b^2 \gamma_\mu}{4V_t}} \right] \\
&= \frac{b e^{\frac{x}{2}}}{2\Gamma(\mu)\sqrt{\pi V_t}} \int_0^\infty e^{-z\left(1 + \frac{x^2 b^2}{4V_t}\right) - \frac{V_t}{4b^2 z}} z^{\mu - \frac{1}{2}} dz \\
&= \frac{b e^{\frac{x}{2}}}{\Gamma(\mu)\sqrt{\pi V_t}} \left(\frac{V_t}{b\sqrt{4V_t + x^2 b^2}} \right)^{\mu + \frac{1}{2}} K_{\mu+\frac{1}{2}} \left(\sqrt{\frac{V_t}{b^2} + \frac{x^2}{4}} \right).
\end{aligned}$$

□

Figure 5.4: Confidence intervals for the log-generalized hyperbolic predictive density and the trajectories of the processes m and V

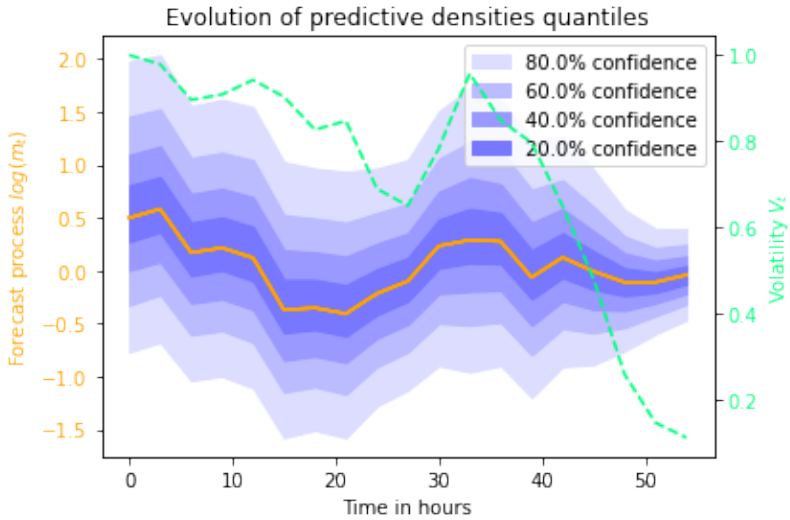

Log-normal inverse Gaussian predictable distribution Using the same notation as above, consider the following pair of SDEs.

$$\frac{dm_t}{m_t} = \sqrt{V_t}\rho(T-t)dW_t \tag{5.10}$$

$$dV_t = -V_t\rho^2(T-t)\left(1+\frac{b^2}{2}\right)dt + \sqrt{V_t}b\rho(T-t)dW_t'. \tag{5.11}$$

The equivalent time-changed representation takes the form

$$\frac{d\widetilde{m}_t}{\widetilde{m}_t} = \sqrt{\widetilde{V}_t}dW_t \tag{5.12}$$

$$d\widetilde{V}_t = -\widetilde{V}_t\left(1+\frac{b^2}{2}\right)dt + \sqrt{\widetilde{V}_t}bdW_t'. \tag{5.13}$$

The proof of the following proposition is very similar to that of Proposition 43 and will therefore be omitted.

Proposition 46. *Let (m, V) be a solution of (5.10–5.11). Then the conditional distribution of $\log m_T$ given $\mathcal{F}_t$ is the normal inverse Gaussian distribution on $\mathbb{R}$ with density*

$$p(x) = \frac{\alpha\delta K_1\left(\alpha\sqrt{\delta^2+(x-\mu)^2}\right)}{\pi\sqrt{\delta^2+(x-\mu)^2}}e^{\delta\gamma+\beta(x-\mu)} \tag{5.14}$$

where the parameters are given by

$$\mu = \log m_t, \quad \delta = \frac{V_t}{b}, \quad \beta = -\frac{1}{2}, \quad \alpha = \sqrt{(b^{-1}+\frac{b}{2})^2+\frac{1}{4}},$$

$\gamma = b^{-1}+\frac{b}{2}$, and K is the modified Bessel function of the third kind. Moreover, the first two

conditional moments of m_T *are given by*

$$m_t = \mathbb{E}[m_T|\mathcal{F}_t] \quad \textit{and} \quad \mathbb{E}[m_T^2|\mathcal{F}_t] = m_t^2 e^{V_t}.$$

Figure 5.5: Confidence intervals for the log normal inverse Gaussian predictive density and the trajectories of the processes m and V.

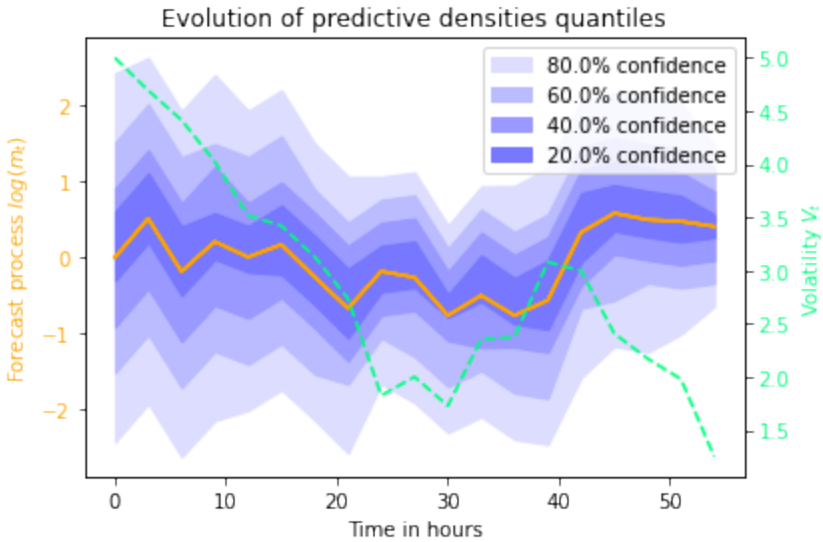

Note that for all processes presented in this section, the process V_t does not necessarily decrease with time (see Figures 5.2-5.5). This reflects the fact that uncertainty over the quantity to forecast can vary over time and does not always decrease as we approach the realization date. We also attract the reader's attention on the fact that while V_t does represent the forecast uncertainty, it does not coincide with the conditional variance of the predictive density for positive quantities – *e.g.*, for (m, V) solution of (5.10)–(5.11), the variance is given by $\mathrm{Var}[m_T|\mathcal{F}_t] = m_t^2(e^{V_t} - 1)$.

5.3 Fitting forecast models to data

In this section, we detail the procedure for calibrating our models for forecast dynamics from historical ensemble forecasts and the corresponding realizations. To illustrate the forecasting of a real-valued quantity, we shall use the normal inverse Gaussian model defined by the equations (5.3–5.4) and the predictive density (5.7), and apply it to ensemble forecasts of temperature. To illustrate the forecasting of a positive quantity, we shall use the model defined by the equations (5.10–5.11) and the predictive density (5.14), and apply it to ensemble forecasts of the wind speed.

5.3.1 Presentation of the dataset

The data is composed of meteorological ensemble forecasts from $K = 273$ different locations around Paris, France, plotted on the map in Figure 5.6, recorded over January 2015.

Figure 5.6: Locations of the meteorological records

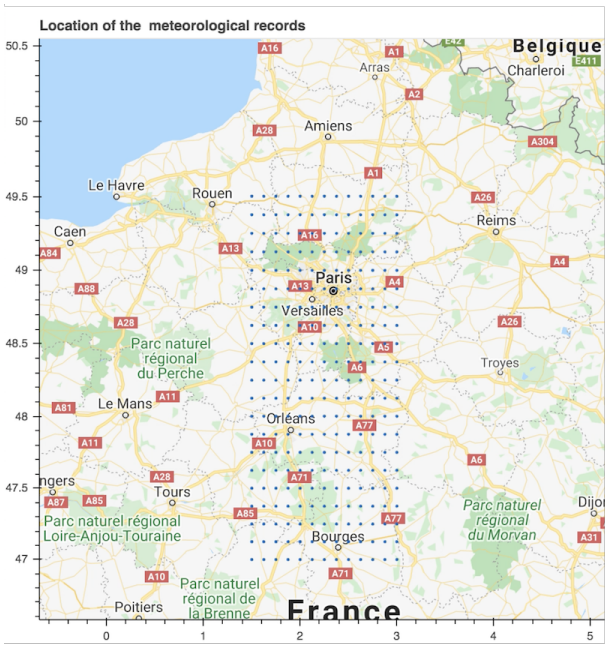

A new forecast ensemble becomes available at 12PM (noon) and at 12AM (midnight) on each day. Each forecast ensemble consists of 50 members, and each member provides a prediction for all meteorological variables for lead times from 1h to 48h, with a step of 3h. Since the forecasts are updated every 12 hours, in our study of forecasts dynamics, we use only the forecast horizons which are multiples of 12 hours, that is, $h \in \mathcal{H} = \{12, 24, 36, 48\}$.

As the locations are very close to each other we make the approximation that they form several ensemble forecasts of the same area. We will use all the ensemble forecasts to calibrate the model, making the approximation that, for each time horizon, the calibrated coefficient can be used to obtain the predictive densities in each of these locations.

Forecasts recorded on days from January 3 to January 21 constitute the training set and those recorded from January 22 to January 31 form the test set. Note that the first two days of January are not used for the calibration because we do not dispose of the full forecast data for them. Our training set is thus composed of $T = 38$ 12-hour periods.

In this application we are interested in two variables: the temperature at 2 meter height denoted by τ and the 10 meter wind speed denoted by w. The wind speed is not directly available in the data and we compute it from the two components w_x and w_y through the usual formula $w = \sqrt{w_x^2 + w_y^2}$ for each member of the forecast ensemble and for each realization.

Since the maximum lead time for our forecast is 48 hours and new forecast becomes available every 12 hours, to study the dynamics of the forecast of a given realization recorded at 12 AM or 12 PM, we dispose of 4 data points with lead times 48h, 36h, 24h and 12h. In Figure 5.7, we plot four forecast ensembles for the wind speed, recorded for a specific location in our dataset at four consecutive forecast update times (Jan 1st 12PM, Jan 2nd 12AM, Jan 2nd 12 PM and Jan 3rd 12AM). The forecasts for a fixed terminal time (Jan 3rd, 12 PM) are shown with the

vertical bar in the four graphs.

Figure 5.7: Forecast ensembles for the wind speed, recorded at four consecutive forecast update times (Jan 1st 12PM, Jan 2nd 12AM, Jan 2nd 12 PM and Jan 3rd 12AM). The forecasts for a fixed terminal time (Jan 3rd, 12 PM) are shown with the vertical bar

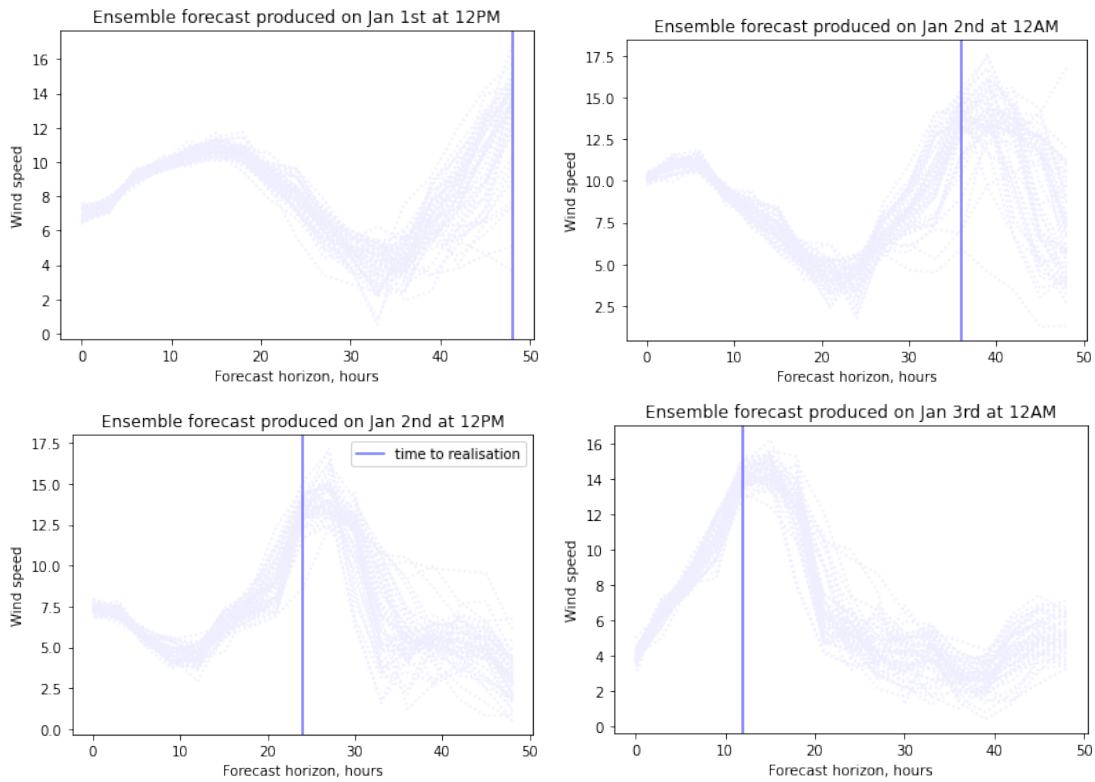

5.3.2 Model calibration

To calibrate our models for forecast dynamics from meteorological ensemble forecasts, we use an approach inspired by the EMOS methodology in Gneiting et al. (2005), to determine the conditional mean and variance of the predictive distribution from the ensemble forecasts. As explained in the introduction, the ensemble forecasts may be biased and underdispersed, so that the mean and variance of the predictive distribution are not necessarily equal to the mean and variance of the empirical distribution of the forecast members, although these quantities are certainly related to each other. Let x^m_{htk} denote the value of member m of the ensemble forecast of a given meteorological quantity (wind speed or temperature), recorded at time t, at location k, for the forecast horizon h, and by $\tilde{x}_{tk}$ the corresponding realization. We assume that the mean of the predictive distribution, denoted by m_{htk} is a linear function of the mean of ensemble members:

$$m_{htk} = m_{htk}(a^0_h, a^1_h) = a^0_h + \frac{a^1_h}{M}\sum_{m=1}^{M} x^m_{htk} := a^0_h + a^1_h \bar{x}_{htk}. \tag{5.15}$$

The coefficients a_h^0 and a_h^1 reflect the bias in the ensemble forecasts. In the case of unbiased forecasts we would have $a_h^0 = 0$ and $a_h^1 = 1$. Similarly, the variance of the predictive distribution depends on the spread of ensemble members, but the latter may not reflect the forecasting error entirely. Hence we assume that at each date t, each location k, and each lead time h, the variance of the predictive distribution is given by

$$\sigma_{htk}^2 = \sigma_{htk}^2(c_h, d_h) = c_h + d_h \frac{1}{M} \sum_{m=1}^{M} (x_{htk}^m - \bar{x}_{htk})^2 := c_h + d_h V_{htk}. \tag{5.16}$$

The coefficients in the expression for the variance can be interpreted as follows: c_h represents the part of the error that is not related to the spread of the ensemble members, whereas d_h represents the part of the conditional variance explained by the ensemble spread.

The full model specification for the temperature forecasts is thus given by equations (5.3–5.7) and (5.15–5.16), while the full specification for the wind forecasts is given by equations (5.10–5.14) and (5.15–5.16).

The full model is calibrated in a three-step procedure as detailed below.

Step 1 In the first step, we first calibrate, separately for each forecast horizon, the parameters a_h^0 and a_h^1 by linear regression:

$$(\widehat{a}_h^0, \widehat{a}_h^1) = \arg\min_{a^0, a^1} \sum_{t,k=1}^{T,K} (a^0 + a^1 \bar{x}_{htk} - \tilde{x}_{tk})^2.$$

Next, using the calibrated values $\widehat{a}_h^0$ and $\widehat{a}_h^1$, we calibrate c_h, d_h and b_h by maximum likelihood:

$$(\widehat{c}_h, \widehat{d}_h, \widehat{b}_h) = \arg\max_{c,d,b} \sum_{t,k=1}^{T,K} \log p(\tilde{x}_{tk}, m_{htk}(\widehat{a}_h^0, \widehat{a}_h^1), \sigma_{htk}^2(c,d), b),$$

where $p(x, m, \sigma^2, b)$ denotes the predictive density expressed in terms of the conditional mean m, the conditional variance σ^2 and the shape parameter b.

For the temperature, the predictive density (5.7) writes,

$$p^\tau(x, m, \sigma^2, b) = \frac{\frac{\sigma^2}{b^2} K_1\left(\frac{1}{b}\sqrt{(\sigma^2/b)^2 + (x-m)^2}\right)}{\pi\sqrt{(\sigma^2/b)^2 + (x-m)^2}} e^{\sigma^2/b^2},$$

and for the wind speed, the predictive density (5.14) writes,

$$p^w(x, m, \sigma^2, b) = \frac{\alpha V K_1\left(\frac{\alpha}{b}\sqrt{V^2 + b^2(x - \log m)^2}\right)}{\pi\sqrt{V^2 + b^2(x - \log m)^2}} e^{V(b^{-2} + \frac{1}{2}) - \frac{1}{2}(x - \log m)}$$

where

$$\alpha = \sqrt{(b^{-1} + \frac{b}{2})^2 + \frac{1}{4}} \quad \text{and} \quad V = \log\left(\frac{\sigma^2}{m^2} + 1\right).$$

Note that in this step, the shape parameter b_h is calibrated independently for each forecast horizon. A common value for all horizons will be fixed in the next step. The choice of the maximum likelihood procedure in this first step was motivated by the availability of predictive densities in explicit form. An alternative would be to use the Continuous Ranked Probability Score (CRPS) as in Gneiting et al. (2005), but for our models the CRPS is only available through heavy numerical computation, making the approach of Gneiting et al. (2005) difficult to implement. We also tested direct maximum likelihood estimation of the four parameters a_h^0, a_h^1, c_h, d_h and b_h, but the presented approach where linear regression is used for a_h^0, a_h^1 leads to better results.

Step 2 In the previous step, the shape parameter b was calibrated separately for each lead time. However, in our model, this parameter does not depend on the forecast lead time. Thus, once the parameters a_t^0, a_t^1, c_t, d_t have been estimated in the first step, we perform again an estimation of the parameter b by maximizing the likelihood including all forecast horizons:

$$\widehat{b} = \underset{b}{\operatorname{argmax}} \sum_{h \in \mathcal{H}} \sum_{t,k=1}^{T,K} \log p(\tilde{x}_{tk}, m_{htk}(\widehat{a}_h^0, \widehat{a}_h^1), \sigma_{htk}^2(\widehat{c}_h, \widehat{d}_h), b).$$

This formulation applies for both the temperature and the wind speed calibration. This procedure does not impact the goodness of fit in terms of first and second moment since the parameters of the mean and the variance are fixed. However, the shape of the distribution may change a bit with no major impact.

Step 3 Once we have estimated the 'static' properties of the model, that is, the parameters which appear in the predictive distribution (b and a_0, a_1, c and d for each time horizon), we need to estimate the 'dynamic parameter', that is, the function ρ, which describes how the forecast varies dynamically. While the static parameters are estimated by comparing the forecasts with their respective realizations, ρ can be estimated by comparing forecasts for the same quantity, obtained at different dates.

For the temperature model, from equation (5.4), we may write:

$$\mathbb{E}\left[\frac{V_{t+s}}{V_t}|\mathcal{F}_t\right] = \exp\left\{\left(-\int_t^{t+s} \rho(T-u)^2 du\right)\right\}.$$

Since the expectation in the right-hand side is deterministic, we can remove the conditioning and write:

$$\int_t^{t+s} \rho(T-u)^2 du = -\log\left(\mathbb{E}\left[\frac{V_{t+s}}{V_t}\right]\right) \tag{5.17}$$

Based on this identity, we suggest the following approach to calibrate ρ based on a given discrete set of forecast horizons $0 < h_1 < \cdots < h_H = T$, for which data are available. Here h_H corresponds to the longest available horizon (when the simulation starts) and h_1 corresponds to the shortest available horizon (last available forecast for a given realization). Assume that ρ is

constant on the intervals (h_i, h_{i+1}), $i = 1, \ldots, H-1$ and denote the value of ρ on the interval (h_i, h_{i+1}) by ρ_i. This assumption is without loss of generality since, in view of Equations (5.5–5.6) and (5.12–5.13), the law of $(m_{h_{i+1}}, V_{h_{i+1}})$ conditional on (m_{h_i}, V_{h_i}) depends on ρ only through the integral $\int_{h_i}^{h_{i+1}} \rho(T-u)^2 du$.

In view of (5.17), we propose to estimate the values $\rho_1, \ldots, \rho_{H-1}$ as follows:

$$\widehat{\rho}_i^2 = -\frac{1}{(h_{i+1}-h_i)} \log \frac{1}{TK} \sum_{t,k=1}^{T,K} \frac{\sigma^2_{h_{i+1}tk}(\widehat{c}_{h_{i+1}}, \widehat{d}_{h_{i+1}})}{\sigma^2_{h_i tk}(\widehat{c}_{h_i}, \widehat{d}_{h_i})}.$$

For the wind speed model, from equations (5.10-5.11), we may write:

$$\mathbb{E}\left[\log \frac{m_{t+h}}{m_t} | \mathcal{F}_t\right] = -\frac{1}{2} \int_t^{t+h} \rho^2(T-s) \mathbb{E}[V_s | \mathcal{F}_t] ds,$$

and

$$\mathbb{E}[V_s | \mathcal{F}_t] = V_t \exp\left(-\left(1+\frac{b^2}{2}\right) \int_t^s \rho^2(T-u) du\right),$$

so that

$$\mathbb{E}\left[\log \frac{m_{t+h}}{m_t} | \mathcal{F}_t\right] = -\frac{V_t}{2+b^2}\left(1 - \exp\left(-\left(1+\frac{b^2}{2}\right) \int_t^{t+h} \rho^2(T-u) du\right)\right).$$

Dividing both sides by V_t, we can remove the conditioning and rewrite this expression as follows:

$$\int_t^{t+h} \rho^2(T-u) du = -\frac{2}{2+b^2} \log\left(1 + (2+b^2)\mathbb{E}\left[\frac{1}{V_t} \log \frac{m_{t+h}}{m_t}\right]\right).$$

As before, assume without loss of generality that ρ is constant on every interval (h_i, h_{i+1}) for $i = 1, \ldots, H-1$, and denote its value on such interval by ρ_i. This suggests the following estimator for ρ_i:

$$\widehat{\rho}_i = -\frac{1}{h_{i+1}-h_i} \frac{2}{2+\widehat{b}^2} \log\left(1 + (2+\widehat{b}^2) \frac{1}{TK} \sum_{t,k=1}^{T,K} \frac{1}{\sigma^2_{h_i tk}(\widehat{c}_{h_i}, \widehat{d}_{h_i})} \log \frac{m_{t+h}}{m_t}\right).$$

5.3.3 Numerical illustrations

In this section we apply the methodology presented in Section 5.3.2 to ensemble forecasts for the wind speed and the temperature described in 5.3.1.

Estimated coefficients As explained in Section 5.3.2, we use an EMOS-inspired technique to improve the calibration of ensemble forecasts. To motivate this post-processing, we present in Figure 5.8 the Talagrand diagrams (rank histograms) for the ensemble forecasts of the log wind speed and the temperature, constructed using the test data. Talagrand diagram is a tool for checking the quality of calibration of ensemble forecasts and is a histogram of the ranks of observations within the corresponding forecast ensembles. In other words, for a given

forecast horizon h, we plot the histogram of $R(\tilde{x}_{tk}, (x^m_{htk})_{m=1\ldots M})^{T,K}_{t,k=1}$, where $R(\tilde{x}, (x^m)_{m=1\ldots M})$ is the normalized rank of the observation $\tilde{x}$ within the ensemble $(x^m)_{m=1\ldots M}$. For a perfectly calibrated ensemble forecast, the Talagrand diagram is within the confidence bounds of the uniform distribution. In the present case, histograms in Figure 5.8, for the lead time $12h00$, and Figure 5.9, for the lead time $24h00$, present a U-shaped profile, which is a clear indication of under-dispersion of our forecast ensembles. In addition, the asymmetric form of the diagram for the wind speed in Figure 5.8, and for the temperature in Figure 5.9, is an indication of the presence of a bias in the ensemble forecast. Histograms for lead times $36h00$ and $48h00$ are available in Appendix 5.5.1.

Figure 5.8: Talagrand diagrams for the wind speed and the temperature, lead time $12h00$

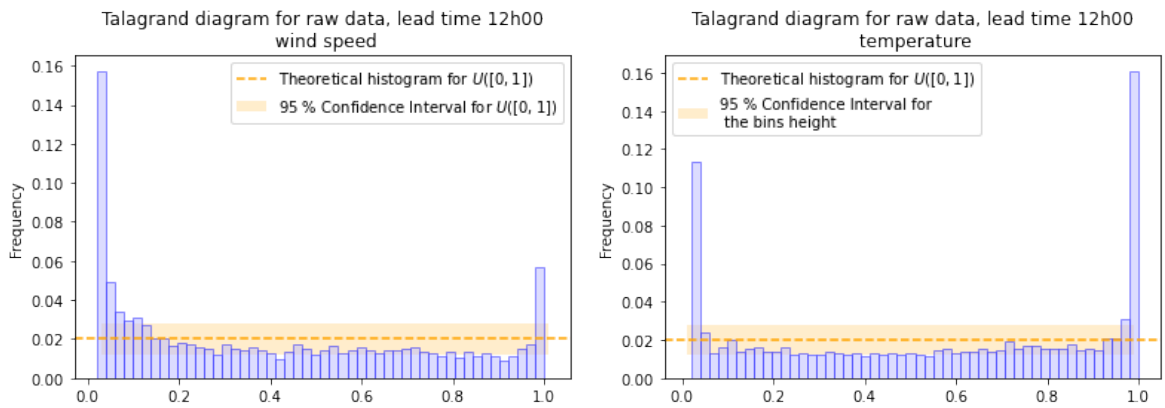

Figure 5.9: Talagrand diagrams for the wind speed and the temperature, lead time $24h00$

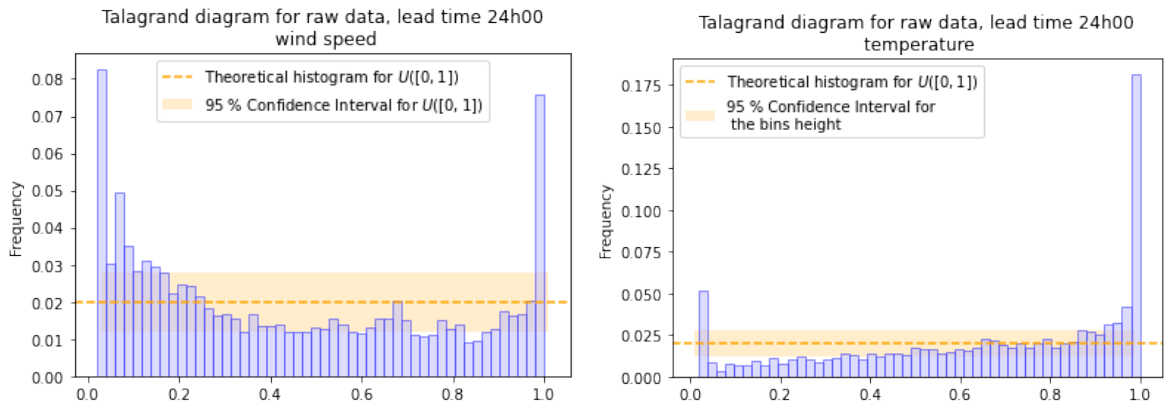

In view of the Talagrand diagrams discussed above, we apply our post-processig approach to obtain an unbiased and well-calibrated probabilistic forecast. Tables 5.1 and 5.2 show, respectively, the estimated coefficients for the wind speed and the temperature during the training period.

Lead time	a^0	a^1	c	d	b
12h	0.117	0.964	0.360	0.765	0.035
24h	-0.028	0.994	0.446	0.494	0.035
36h	-0.083	1.006	0.951	0.160	0.035
48h	-0.573	1.044	0.240	0.798	0.035

Lead time interval	$\rho(.)$
12h-24h	0.171
24h-36h	0.153
36h-48h	0.168

Table 5.1: Calibrated coefficients for the wind speed from the 1st to the 21th of January 2015

Lead time	a^0	a^1	c	d	b
12h	0.217	0.952	0.312	1.722	0.719
24h	0.103	1.006	0.416	0.916	0.719
36h	0.136	1.019	0.675	0.467	0.719
48h	0.202	1.015	0.306	0.913	0.719

Lead time interval	$\rho(.)$
12h-24h	0.160
24h-36h	0.163
36h-48h	0.180

Table 5.2: Calibrated coefficients for the temperature from the 1st to the 21th of January 2015

The spread coefficients c and d confirm that the raw ensemble forecasts are underdispersed: except for the $24h$ lead time wind speed forecasts, all provide a post process variance larger than the ensemble forecast spread. This is especially true for the $12h$ lead time temperature forecast which has an intercept of $c = 0.312$ and a slope of $d = 1.772$ which multiply almost by two the original variance.

We now proceed to analyse the diffusion coefficients b and ρ. The parameter b is quite low for the log wind speed distribution. This suggests that the log-wind speed distribution is closer to the Gaussian one, than the temperature distribution for which the coefficient b is higher. This indicates a heavier tailed model where extreme temperature values are more likely to happen than extreme wind spikes.

The values of the piecewise constant function $\rho(.)$ are very close to each other on the three time intervals considered for both the wind speed and the temperature. This feature will be useful when using this dynamics in the control problems in Section 5.4.

Goodness of fit In this paragraph, we check the goodness of fit of the estimated predictive distribution using the test dataset and provide some illustrations. We first present in Table 5.3 the mean square error computed using the test data before and after the pre-processing that is :

$$\mathrm{MSE}^{\mathrm{raw}}(h) = \sum_{t,k=1}^{T,K} (\bar{x}_{htk} - \tilde{x}_{tk})^2, \quad \mathrm{MSE}(h) = \sum_{t,k=1}^{T,K} (a_h^0 + a_h^1 \bar{x}_{htk} - \tilde{x}_{tk})^2.$$

Except for the wind speed at lead times $36h$ and $48h$, the correction of the bias brings the mean of the test data closer to the realization. Hence our estimation procedure doesn't overfit the training period and is robust when considering new data. However, we should mention that this is done over data for the same month and the same season. We may assume that the seasonality impacts the value of the coefficients and repeating the study over different months may provide additional insights.

Log wind speed		
Lead time	MSE Raw ensemble	MSE Model
12h	0.948	0.804
24h	0.955	0.931
36h	1.233	1.240
48h	1.633	1.795

Temperature		
Lead time	MSE Raw ensemble	MSE Model
12h	0.656	0.599
24h	0.851	0.778
36h	0.973	0.900
48h	1.539	1.424

Table 5.3: MSE for the log wind speed and the temperature over the period (22/01-31/01)

To evaluate the calibration of the predictive distribution we use the probability integral transform (PIT) histogram and to check both calibration and sharpness, we compute the continuous rank probability score (CRPS). The formal definition of the CRPS for a given realisation y and a predictive distribution with cumulative distribution function (CDF) F is given by:

$$\mathrm{CRPS}(F, y) = \int_{\mathbb{R}} (F(x) - \mathbb{1}_{\{y \leq x\}})^2 dx$$

For the normal inverve Gaussian and the log normal inverse Gaussian distributions, it is not possible to compute the analytical expression of the CDF. However, we can use the Plancherel formula and obtain an expression relying on the characteristic function ϕ of the normal inverse Gaussian predictive distribution:

$$\begin{aligned}\int_{\mathbb{R}} (F(x) - \mathbf{1}_{y \leq x})^2 dx &= \frac{1}{2\pi} \int_{\mathbb{R}} \frac{|\phi(u) - e^{iuy}|^2}{u^2} du \\ &= \frac{1}{2\pi} \int_{\mathbb{R}} \frac{\left| e^{i\mu u + \delta\left(\gamma - \sqrt{\alpha^2 - (\beta + iu)^2}\right)} - e^{iuy} \right|^2}{u^2} du,\end{aligned}$$

Following this formulation, for a predictive distribution for the location k, the lead time h and the date t, we denote the CDF F_{tkh} and the realisation y_{tkh}. The parameters $\alpha, \beta, \gamma, \delta$ and μ for temperature are given by,

$$\alpha = \frac{1}{b},\ \beta = 0,\ \gamma = \frac{1}{b},\ \delta = \frac{V_{htk}}{b},\ \mu = m_{htk},$$

and for the wind speed:

$$\alpha = \sqrt{\left(\frac{1}{b} + \frac{b}{2}\right)^2 + \frac{1}{4}},\ \beta = -\frac{1}{2},\ \gamma = \frac{1}{b} + \frac{b}{2},\ \delta = \frac{V_{htk}}{b},\ \mu = \log(m_{htk}).$$

We analyse the goodness of fit using the averaged CRPS over the period 22/01/15-31/01/15 for each lead time,

$$\mathrm{averageCRPS}_h(F, y) = \frac{1}{TK} \sum_{t,k=1}^{T,K} \int_{\mathbb{R}} \left(F_{tkh}(x) - \mathbb{1}_{\{y_{tkh} \leq x\}}\right)^2 dx$$

We compare it to the averaged CRPS obtained with the raw ensemble forecasts, that is:

$$\text{averageCRPS}_h(\text{Ensemble forecasts}, y_{tkh}) = \frac{1}{TK}\sum_{t,k=1}^{T,K}\int_{\mathbb{R}}\left(\widehat{F}^M_{tkh}(x) - \mathbb{1}_{\{y_{tkh}\leq x\}}\right)^2 dx,$$

where

$$\widehat{F}^M_{tkh}(x) = \frac{1}{M}\sum_{m=1}^{M}\mathbf{1}_{\{x^m_{tkh}\leq x\}}$$

and for each location, date and lead time, we may simplify the CRPS formula as follows:

$$\begin{aligned}\text{CRPS}(\widehat{F}^M_{tkh}, y_{tkh}) &= \int_{\mathbb{R}}\left(\widehat{F}^M_{tkh}(x) - \mathbb{1}_{\{y_{tkh}\leq x\}}\right)^2 dx \\ &= \frac{2}{M}\sum_{\ell=1}^{M}(x_{(\ell)} - y_{tkh})\left\{\mathbf{1}_{\{x_{(\ell)}>y_{tkh}\}} - \frac{\ell-\frac{1}{2}}{M}\right\},\end{aligned}$$

where $x_{(\ell)}, k = 1\ldots M$ is the order statistics of the sample $x^m_{tkh}, m = 1\ldots M$.

Table 5.4: CPRS for the wind speed ($\log m.s^{-1}$) and for the temperature (C°) over the test period 22/01/15-31/01/15

log wind speed				
Lead time	12h	24h	36h	48h
EMOS	0.039	0.056	0.062	0.128
Raw	0.080	0.095	0.100	0.102

temperature				
Lead time	12h	24h	36h	48h
EMOS	0.372	0.437	0.485	0.639
Raw	0.450	0.481	0.538	0.752

Table 5.4 compares the CRPS computed using the test period for raw ensemble forecast and for probabilistic forecasts obtained using our post-processing method. We observe a significant improvement for both wind speed and temperature forecasts, for all forecast horizons except the 48-hour forecast horizon for the wind speed.

As an independent illustration of the calibration of the post-processed forecasts we plot the probability integral transforms (PIT), which is the equivalent of Talagrand diagram in the context of probabilistic forecasts and consists in plotting the histogram of the predictive CDF evaluated at the realization point. If the predictive distribution is well calibrated then the histogram should be close to the uniform one.

The PIT of post-processed forecasts are shown in Figures 5.10 and 5.11 and may be compared to Talagrand diagrams in Figures 5.8–5.9. Here again, the improvement is considerable, although some deviations from the uniform distribution can still be observed. They may be explained by the fact that we use a parametric approach, which obviously cannot provide a perfect fit to the data, and our observations are not completely independent.

Figure 5.10: PIT histogram for the wind speed and the temperature, lead time 12h00

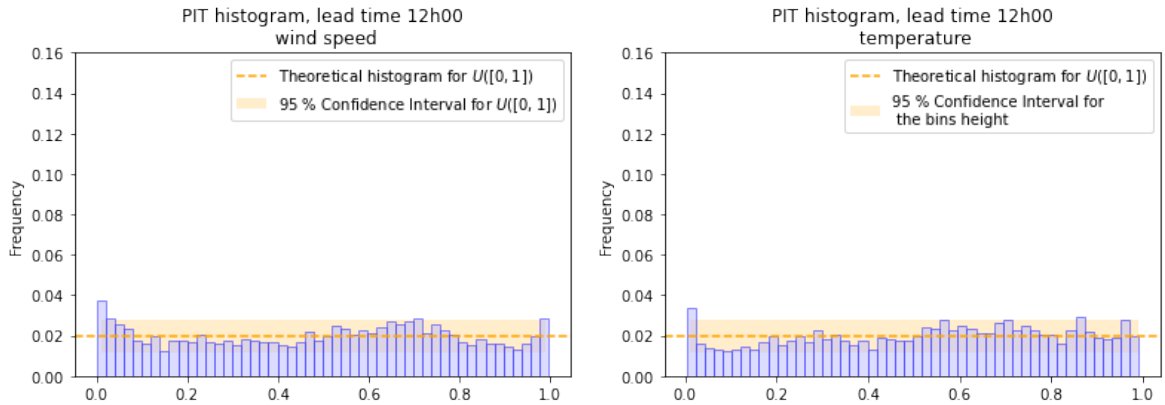

Figure 5.11: PIT histogram for the wind speed and the temperature, lead time 24h00

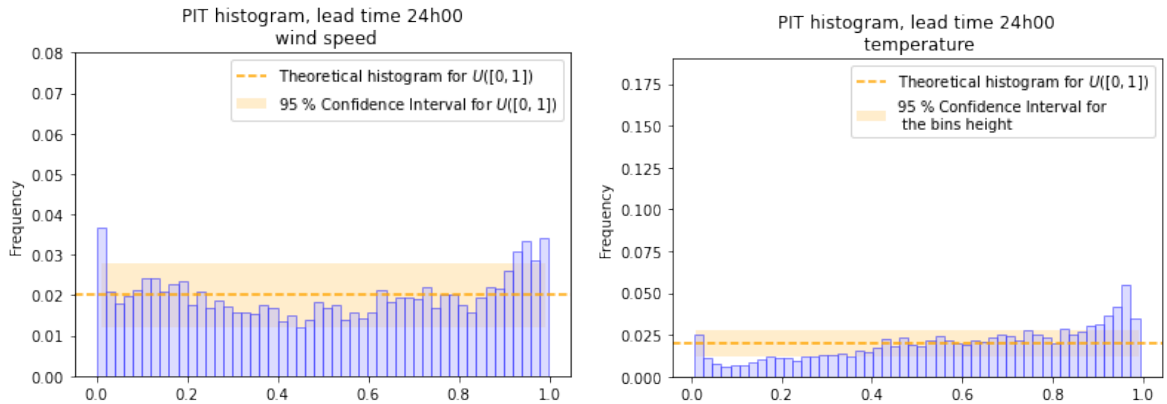

Behavior of the predictive density in test data To illustrate the shape of the predictive density obtained with our approach, we displayed in Figure 5.12, for a given realization date (22/01/15 at 12 a.m) and location, the predictive densities at each lead time as well as the realisation for the wind speed forecasts. We observe that the sharpness of the predictive distribution varies with the lead time. Interestingly it doesn't always improve as we approach the realisation time (e.g see the temperature predictive densities at lead times 12h and 24h). This is to be compared with the simulated predictive distributions in Figures 5.5 and 5.3.

At the same location we plotted the evolution of the point forecast (first moment of the predictive distribution) for a fixed lead time and the corresponding realization over the period 22/01/15-31/01/15 in Figure 5.13. We also show the confidence intervals around the point forecast. The realization always falls in the 90 % confidence interval and the width of the intervals varies throughout the simulation. The forecasts seem reasonably close to the realizations: this is especially true when the width of the confidence interval is small.

Figure 5.12: Example of predictive densities produced by the model

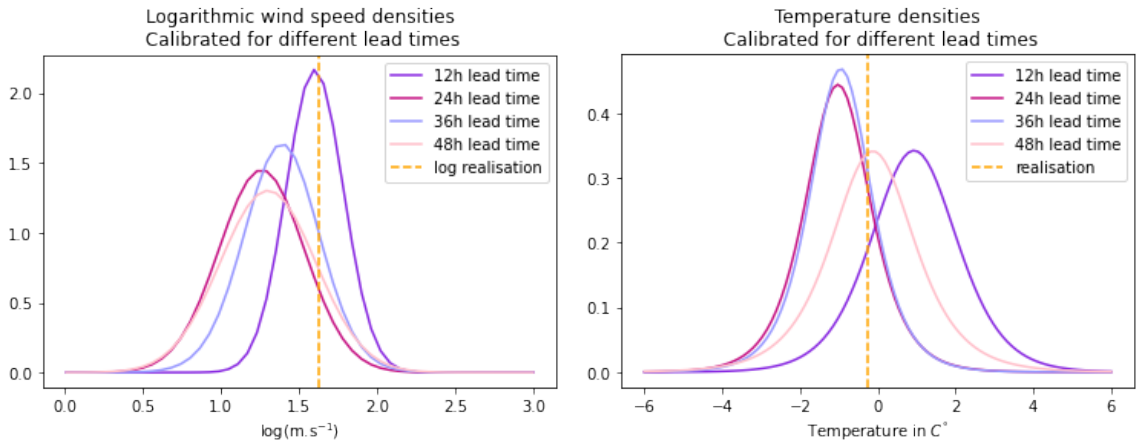

Table 5.6 shows the average width of the 90 % confidence interval, for the testing period: 22/01/2015-31/01/2015. For the first three lead times the width increases with lead time but the 36h lead time confidence intervals are on average sightly larger than the 48 h lead time ones. This confirms the intuition that the closer we get to the realisation date the less variations there are in forecasts updates (on average). On the other hand, finding a wider confidence interval average for the lead time $36h$ than for the lead time $48h$ is quite unexpected. We investigated the confidence interval width during the training period in Table 5.5 and found that the confidence interval width always decreases as we approach the realization time. This suggests that the model may not be fully consistent with the test data, perhaps owing to a possible non-stationarity of the forecasts.

Lead time	90% CI width Wind speed	90% CI width Temperature
12h	2.565	2.605
24h	2.869	2.923
36h	3.381	3.301
48h	3.628	3.585

Table 5.5: Average 90 % confidence interval for the period 01/01-21/01 for the wind speed and the temperature

Lead time	90% CI width Wind speed	90% CI width Temperature
12h	2.537	2.592
24h	2.801	2.747
36h	3.378	3.062
48h	3.254	2.956

Table 5.6: Average 90 % confidence interval for the period 22/01-31/01 for the wind speed and the temperature

Figure 5.13: Temperature and wind speed forecasts for the period 22/01/2015-31/01/2015

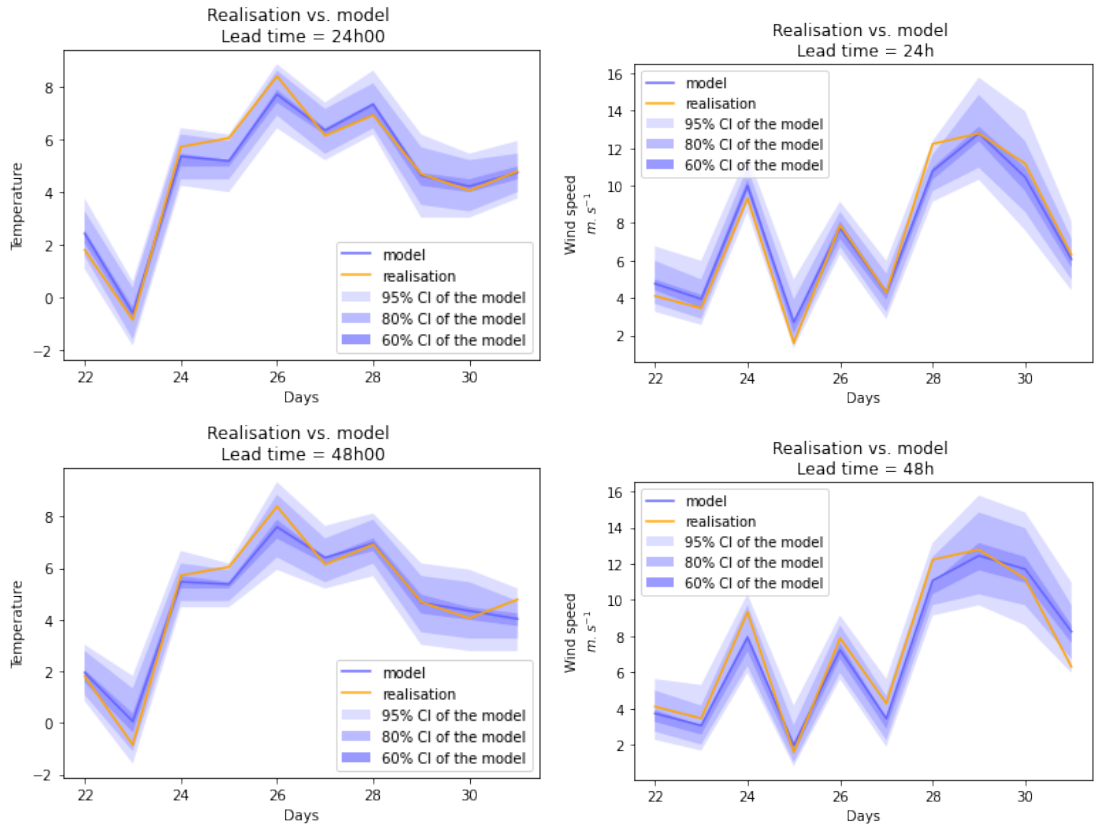

5.4 Application to wind power trading

In this section we present an application of our methodology based on the dynamic modeling of probabilistic forecasts to the problem of wind power trading in the intraday electricity market.

5.4.1 Description of the problem

Consider a wind power producer who aims to sell the output power in the intraday electricity market. To analyze the effect of market mechanisms we assume that there are no subsidies and no guaranteed purchase scheme. The intraday market opens every day at 3 p.m and allows continuous trading in all delivery hours of the next day. For a given delivery hour T, we consider the energy produced during a small time interval around this date T. We denote the average wind speed during this time interval by m_T, and the power curve of the wind turbine by f, that is, the rate of power production during this interval is given by $P_T = f(m_T)$. For the purpose of illustration we choose the stylized production function f defined by:

$$f(m) = \frac{(m - m_{\min})^+ - (m - m_{\max})^+}{m_{\max} - m_{\min}},$$

where $m_{\min}$ is the cut-in speed (at which the turbine starts to produce), and $m_{\max}$ is the rated speed (at which the turbine produces its maximum power), but the methodology applies without

modifications to any other production function.

This power can be sold at any time starting from the opening time of the intraday market up to 15 minutes before production. The fraction traded at the date t will have the price S_t, and we denote the total amount of power (for delivery at T) sold or bought up to date t by ϕ_t. Any power not sold in the intraday market prior to date T will be sold at date T at the balancing price denoted by S_T. In addition, balancing transactions are subject to imbalance penalty equal to a constant K times the volume of the transaction. Throughout this section, we assume the following dynamics for the price:

$$dS_t = \mu_S dt + \sigma_S dB_t, \quad \forall t \in [0, T] \tag{5.18}$$

where μ_S and σ_S are constants and $(B_t)_{t\in[0,T]}$ is a Brownian motion.

We make the assumption that the producer changes her position in the market only when a new forecast of the wind speed, and thus of the power production, becomes available. In other words, new trades are only triggered by new forecast information and not by price information which is available continuously. This is justified by the fact that most producers do not attempt to take advantage of potential price arbitrages but use the markets to compensate forecast errors. We denote by $t_0 < \cdots < t_{N-1}$ the discrete times at which the trades take place, $t_N = T$ being the time when the delivery starts. The profit of the producer is thus given by

$$\begin{aligned}\text{Profit} &= \underbrace{S_{t_0}\phi_{t_0} + \sum_{i=1}^{N-1} S_{t_i}(\phi_{t_i} - \phi_{t_{i-1}})}_{\text{Intraday market}} + \underbrace{S_T(f(m_T) - \phi_{t_{N-1}}) - K|f(m_T) - \phi_{t_{N-1}}|}_{\text{Imbalance payment}} \\ &= f(m_T)S_T - \sum_{i=0}^{N-1} \phi_{t_i}\Delta S_{t_i} - K|f(m_T) - \phi_{t_{N-1}}|,\end{aligned}$$

where $\Delta S_{t_i} = S_{t_{i+1}} - S_{t_i}$.

We assume that the producer aims to maximize the utility of profit at date T, that is, she solves the following control problem:

$$\max_{\phi:=(\phi_{t_i})_{i=1}^{N-1}} \mathbb{E}\left[u\left(f(m_T)S_T - \sum_{i=0}^{N-1} \phi_{t_i}\Delta S_{t_i} - K|f(m_T) - \phi_{t_{N-1}}|\right)\right], \tag{5.19}$$

where u is a utility function (concave, increasing and satisfying certain regularity conditions), and $\phi := (\phi_{t_i})_{i=1}^{N-1}$ belongs to a certain class of admissible strategies. In particular, the process ϕ must be adapted with respect to the filtration of the agent, generated by the history of the process S, the history of the forecast process m and a measure of forecast uncertainty if it is stochastic. For the numerical resolution we assume that the agent has an exponential CARA utility function given by:

$$u(x) = 1 - e^{-\alpha x}, \quad \alpha > 0.$$

To assess the importance of modeling the dynamics of forecast uncertainty, in the next section we perform the following numerical experiment.

- We consider two models, model A, which describes the dynamic evolution of forecast uncertainty, and model B, which does not include such a description. Model A, detailed in Section 5.2.2, uses the log-inverse Gaussian predictive distribution and the forecast evolution given by equations (5.10–5.11). Model B is a simplified version of model A, with a constant diffusion coefficient of the forecast process:

$$\frac{dm_t}{m_t} = \sigma_m dW_t, \tag{5.20}$$

  where σ_m is a constant such that $\sigma_m = \rho(T)\sqrt{V_0}$. Since empirical studies show a negative correlation between the market price and the wind production forecasts Kiesel and Paraschiv (2017); Féron et al. (2020), we assume $\langle W, B\rangle_t = \lambda t,\ \lambda < 0,\ \forall t \in [0, T]$.

- For each model, we compute the optimal feedback strategies $\phi^A_{t_i}(S, m, V)$ and $\phi^B_{t_i}(S, m)$, for $i = 0, \dots, N-1$ by solving the problem (5.19) using the Least Squares Monte Carlo algorithm.

- We then simulate the prices using model A, and compute the profit of the producer with the feedback strategies ϕ^A and ϕ^B. The difference between the two profit amounts allows to quantify the loss from using model B, that is, from not taking into account the dynamic evolution of the forecast uncertainty, when the data follow model A.

5.4.2 Numerical resolution

In the first part of this section we present the Least Squares Monte Carlo algorithm used to solve the control problem. Next we detail the parameter values and finally compare the profits obtained in the case of model A and model B.

Least Square Monte Carlo algorithm We consider the equivalent problem

$$\min_{\phi} \quad \mathbb{E}\left[\exp\left\{-\alpha\left(f(m_T)S_T - \sum_{i=0}^{N-1}\phi_{t_i}\Delta S_{t_i} - K|f(m_T) - \phi_{t_{N-1}}|\right)\right\}\right], \tag{5.21}$$

and define its value function at each time step $t_i,\ i = 0, \dots, N-1$,

$$v_{t_i}(X) = \min_{\phi_{t_i},\dots,\phi_{t_{N-1}}} \mathbb{E}^X_{t_i}\left[\exp\left\{-\alpha\left(f(m_T)S_T - \sum_{k=i}^{N-1}\phi_{t_k}\Delta S_{t_k} - K|f(m_T) - \phi_{t_{N-1}}|\right)\right\}\right]$$

where $X = (S, m, V)$ for Model A and $X = (S, m)$ for Model B.

Exploiting the exponential structure of the utility function, the dynamic programming prin-

ciple takes the following form.

$$v_{t_{N-1}}(X) = \min_{\phi_{t_{N-1}}} \quad \mathbb{E}^X_{t_{N-1}}\left[\exp\{\alpha(\phi_{t_{N-1}}\Delta S_{t_{N-1}} - f(m_{T_N})S_{T_N} + K|f(m_T) - \phi_{t_{N-1}}|)\}\right],$$
$$v_{t_i}(X) = \min_{\phi_{t_i}} \quad \mathbb{E}^X_{t_i}\left[\exp\{\alpha\phi_{t_i}\Delta S_{t_i}\}\right], \; i = 0, \dots, N-2.$$

For the numerical computation of the value functions, we use a regression approach based on adaptative local basis functions, described in Bouchard and Warin (2012) and implemented in the open source library StOpt (The STochastic OPTimization library, see Gevret et al. (2018) for a detailed documentation). We briefly describe the algorithm below and refer the reader to Bouchard and Warin (2012); Gevret et al. (2018) for further details.

At each time step t_i, $i = 1, \dots, N$, the state space is partitioned into Q cells, denoted by $D^{t_i}_q$, $q \in Q$, and on each cell, a linear local basis function ψ_q is defined. We denote by $\beta_q \in \mathbb{R}^{d+1}$ the coefficients of the function ψ_q, where $d = 3$ (resp. d=2) is the dimension of the problem.

Let $(X^j_{t_i})^{j=1,\dots M}_{i=1,\dots,N}$ be the Monte Carlo simulations of the discretized version of the processes $X := (S, m, V)$ for model A (resp $X := (S, m)$ for model B). We call these simulations the learning set. Let $\phi^\ell_{t_i}, \ell = 1, \dots, L$ be the discretized values of the control at time t_i. The algorithm for computing the optimal strategies consists in the following steps, performed backward in time, starting from $i = N-1$.

1. For each point $\phi^\ell_{t_i}$, we determine the vector $\widehat{\beta}_q(\phi^l_{t_i})$ as follows:

$$\widehat{\beta}^{N-1}_q(\phi^\ell_{t_{N-1}}) = \text{argmin}_{\beta_q \in \mathbb{R}^{d+1}} \sum_{x \in D^{t_{N-1}}_q} \left[\psi_q(\beta_q, x) - \exp\{\alpha(\phi_{t_{N-1}}\Delta S_{t_{N-1}} - f(m_{T_N})S_{T_N} + K|f(m_T) - \phi_{t_{N-1}}|)\}\right]^2,$$
$$\widehat{\beta}^i_q(\phi^\ell_{t_i}) = \text{argmin}_{\beta_q \in \mathbb{R}^{d+1}} \sum_{x \in D^{t_i}_q} \left[\psi_q(\beta_q, x) - \widehat{v}_{t_{i+1}} \times \exp\{\alpha\phi^\ell_{t_i}\Delta S_{t_i}\}\right]^2, \; i = 0, \dots, N-2$$

2. We then define the regression estimators for the conditional expectations:

$$\widehat{e}^M_i(x, \phi^\ell_{t_i}) = \sum_{q=1}^{Q} \psi_q(\beta_q(\phi^\ell_{t_i}), x)\mathbb{1}_{x \in D_q}, \; i = 0, \dots, N-1 \tag{5.22}$$

3. The optimal feedback strategy and the value function are estimated as follows.

$$\widehat{\phi}^*_{t_i}(x) = \text{argmin}_{\phi^\ell_{t_i}, \ell=1,\dots,L}\widehat{e}^M_i(x, \phi^\ell_{t_i}), \quad \widehat{v}_{t_i}(x) = \widehat{e}^M_i(x, \widehat{\phi}^*_{t_i}(x))$$

To compare the gains in model A and in model B we then simulate M' new trajectories of model A (the testing set) and compute the gains using the optimal feedback strategies derived above. In the next paragraph we detail the choice of parameters.

Setting of the numerical illustrations We use the empirical data and the parameters estimated in Section 5.3 for the simulations of the present illustration.

We consider the delivery hour 12 PM (noon), and assume that forecast updates become available and the trading takes place at 12 PM on the previous day (this corresponds to the day-ahead trade), at 6PM on the previous day, at midnight, at 6AM on the delivery date, and at 12PM on the delivery date. Letting $t = 0$ correspond to 12PM of the day preceding the delivery day, we then have: $N = 4$ and $(t_0, t_1, t_2, t_3, t_4) = (0H, 6H, 12H, 18H, 24H)$. The parameters of Model A are estimated as explained in Section 5.3.2, and for model B we fix m_0 to the same value as in Model A, and $\sigma_m = \rho\sqrt{V_0}$. The value of price volatility is calibrated as explained in Féron et al. (2020). The drift μ_S, is not fixed for now, we will make it vary in the next paragraph for our study. The absolute risk aversion coefficient α was chosen in an ad hoc manner, but in such a way that its numerical value is compatible with the average value of producer's revenues. The values of all parameters are summarized in Table 5.7.

Parameter	Value	Parameter	Value
S_0	40 €/MWh	V_0	0.032
σ^S	6 €/MWh.h$^{1/2}$	ρ	0.16
m_0	5.38 m/s	b	0.035
λ	-0.08	α	0.01 €$^{-1}$
$m_{\min}$	3.3 m/s	$m_{\max}$	25 m/s
K	10 €/MWh		

Table 5.7: Parameters of the model

For the estimation of the conditional expectations (the training set) we use $M = 200000$ MC trajectories, and the grid size $Q = 15\times15$ for the two-dimensional model and $Q = 15\times15\times15$ for the three-dimensional model. The actual grids are determined by the algorithm in an adaptive manner.

To evaluate the gains of the two strategies (the test set), we use $M' = 1000000$ trajectories. The control values (position of the agent) are also discretized on a grid, which depends on the setting of the problem. In the numerical experiments presented in the next section, we consider three different settings. When the price is martingale, we allow for positions between -1 and 1 with a step size of 0.01. When the price has a positive or negative trend, we allow for positions between -5 and 5 with a step size size 0.05. We chose these bounds in accordance with the observed shapes of the strategies during experiments, and by making a trade-off between accuracy and the computational cost.

Results After deriving the optimal strategies $(\phi^{j*}_{t_i})_{0\le i\le N-1}$ for model A and model B, we then simulate M' new trajectories under model A and compute the realized profit for the two producers:

$$f(m^j_T)S^j_T - \sum_{i=0}^{N-1} \phi^{j*}_{t_i}\Delta S^j_{t_i}.$$

This computation is repeated multiple times to evaluate the sensitivity of the average profit to the presence of price trend and the wind forecast uncertainty level. Figure 5.14 shows the range of wind speeds for different wind forecast uncertainty levels.

Figure 5.14: Wind speed range as function of the uncertainty parameter.

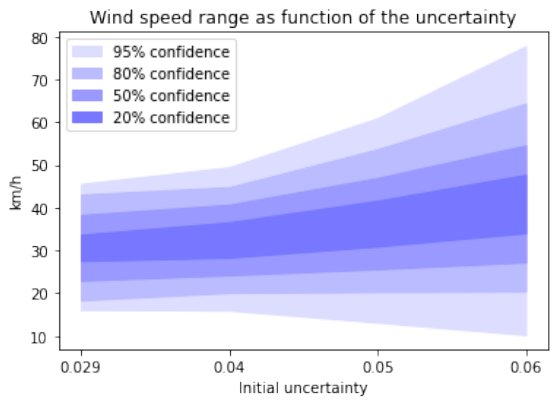

Figure 5.15: Realized profits for different values of the price trend

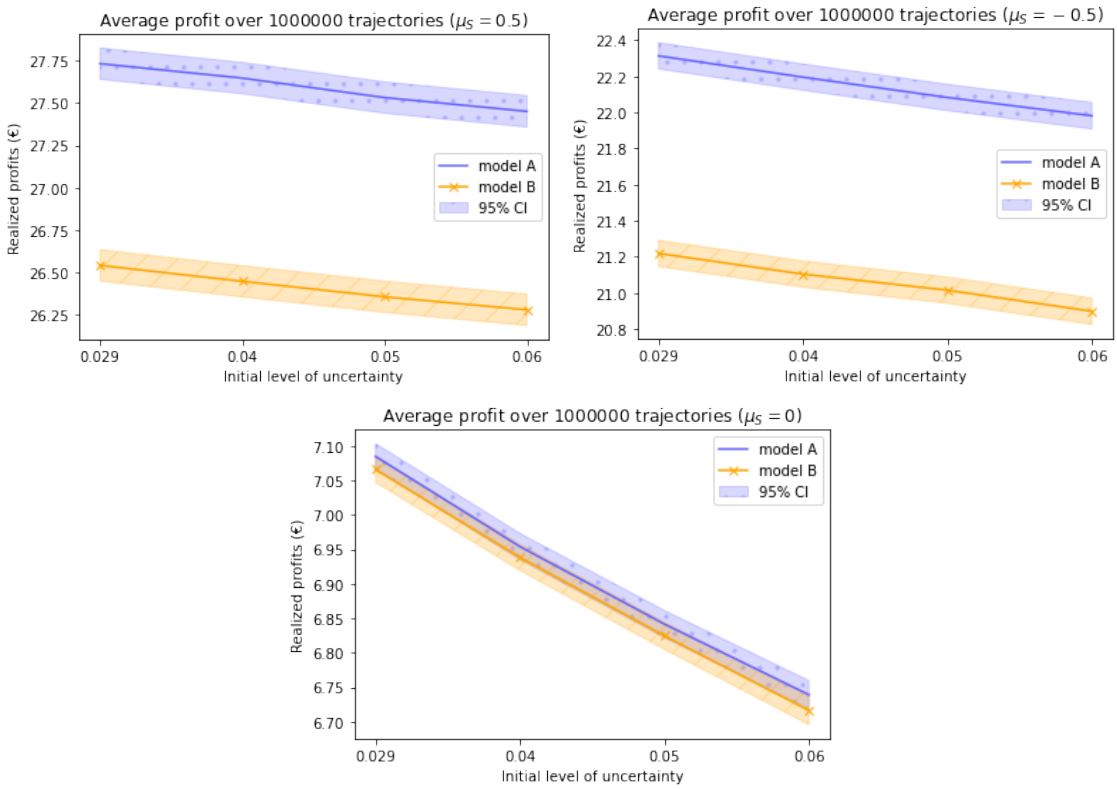

The simulated average profit for the two producers is shown in Figure 5.15. We see that in both models, the average profit decreases as function of the forecast uncertainty level. The presence of a price trend leads both for higher profits for the two models and a clear improvement of performance of model A compared to model B. In the case where the price is a martingale the result is not obvious since the confidence intervals of the average profits overlap to a large extent: this result needs to be further investigated.

The observations made from Figure 5.15 are confirmed in Table 5.8 where we show the relative profits computed as follows:

$$\frac{\text{Profit(Model A)} - \text{Profit(Model B)}}{\text{Profits(Model B)}} \times 100.$$

μ_S (€/MWh.h)	Relative profits in %
0	+0.27
0.5	**+4.49**
−0.5	**+5.15**

Table 5.8: Relative profits of Model A (taking into account stochastic forecast uncertainty) compared to Model B. Bold values are statistically significant at 95% level.

In the literature, estimations of the trend in the intraday electricity market price are scarce. A recent empirical study by Glas et al. (2020) reports the presence of a small trend composed of a constant part (0.0433 €/MWh^2) and a permanent price impact (0.0017 €/MWh^2). Hence taking into account the evolution of the forecast uncertainty as it is done in the model we propose seem to be adapted to the market reality and impacts the strategies in a way that increases significantly the profits.

5.5 Appendix

5.5.1 Talagrand diagrams and PIT histograms for lead times 36h and 48h

Figure 5.16: Talagrand diagrams for the wind speed and the temperature, lead time 36h00

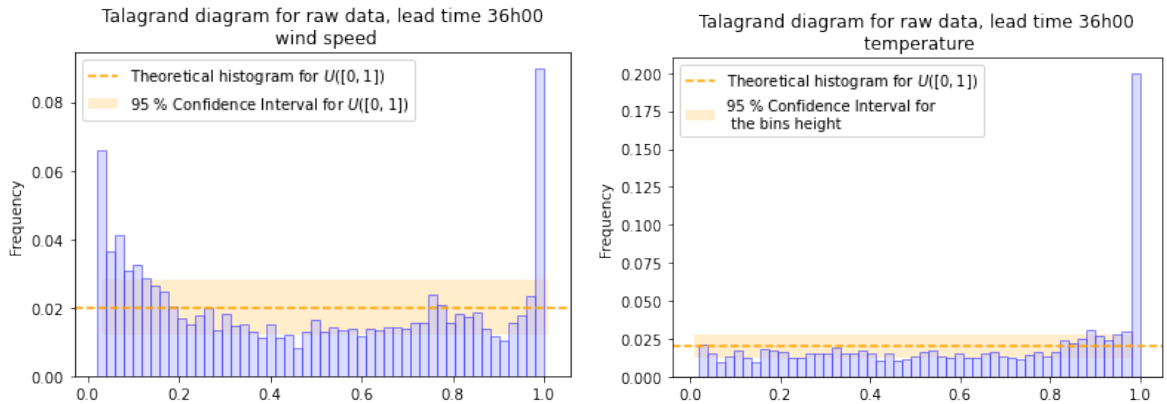

Figure 5.17: PIT histogram for the wind speed and the temperature, lead time 36h00

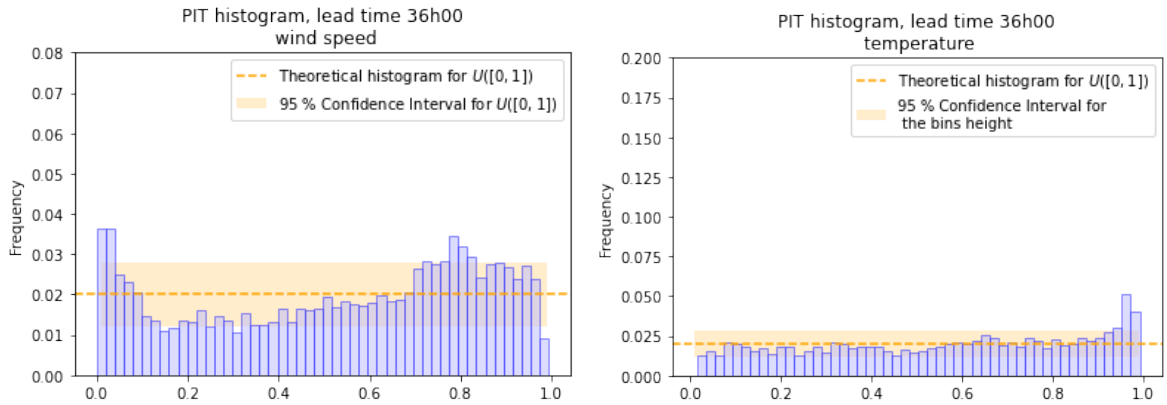

Figure 5.18: Talagrand diagrams for the wind speed and the temperature, lead time 48h00

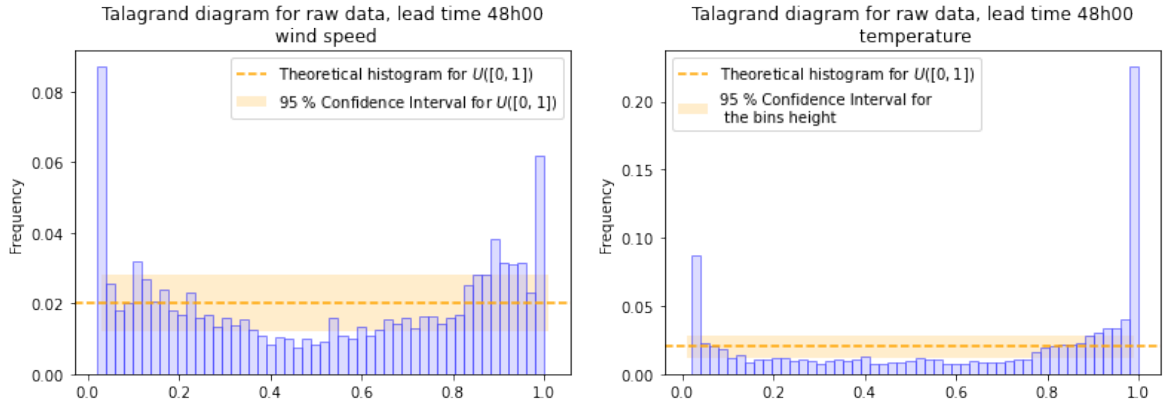

Figure 5.19: PIT histogram for the wind speed and the temperature, lead time 48h00

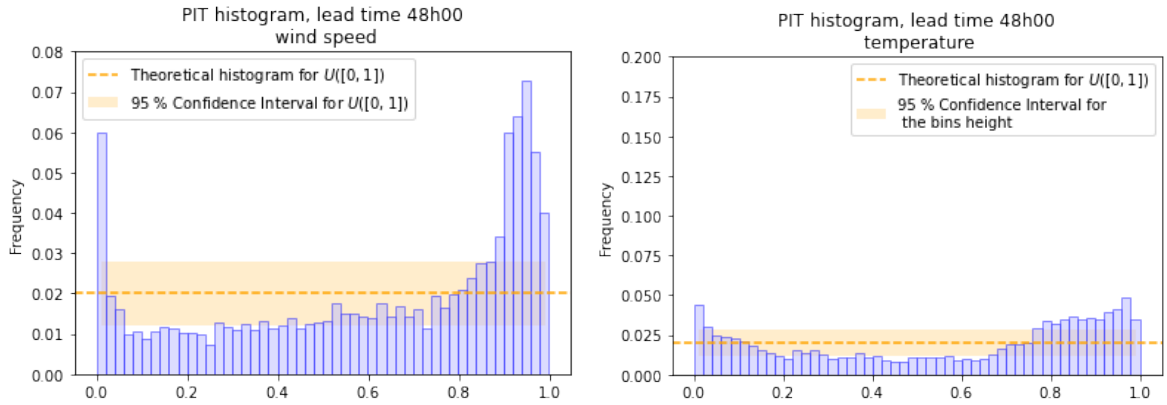

Chapter 6

Risk bounds for aggregated shallow neural networks using Gaussian priors

Abstract

Analyzing statistical properties of neural networks is a central topic in statistics and machine learning. However, most results in the literature focus on the properties of the neural network minimizing the training error. The goal of this paper is to consider aggregated neural networks using a Gaussian prior. The departure point of our approach is an arbitrary aggregate satisfying the PAC-Bayesian inequality. The main contribution is a precise nonasymptotic assessment of the estimation error appearing in the PAC-Bayes bound. Our analysis is sharp enough to lead to minimax rates of estimation over Sobolev smoothness classes.

6.1 Introduction

PAC-Bayesian learning has known a growing interest in the last decades. The seminal papers of McAllester (1999b, 2003) have opened the road to the development of a wide range of PAC-Bayesian bounds. These bounds have the advantage to be generic and apply to various learning problems. While McAllester developed empirical bounds depending on training data, Catoni (2007) enhanced this existing result in form of oracle inequalities 'in probability', thanks to deviation inequalities. Concomitantly, Leung and Barron (2006) initiated PAC-Bayesian bounds in expectation that have been further developed by Dalalyan and Tsybakov (2012a); Rigollet and Tsybakov (2012). We refer the reader to Guedj (2019) for a comprehensive review –and a very relevant historic on generalized bayesian learning.

Assume that we are interested in a class of estimators $\mathcal{F}_{\mathsf{W}} := \{f_{\boldsymbol{w}}, \boldsymbol{w} \in \mathsf{W}\}$, indexed by the measurable space $(\mathsf{W}, \mathcal{L})$. Let $\mathcal{P}_{\mathsf{W}}$ be a space of probability measures on $(\mathsf{W}, \mathcal{L})$, satisfying certain conditions. A probability measure from $\mathcal{P}_{\mathsf{W}}$ assigned to the class $\mathcal{F}_{\mathsf{W}}$ can also be seen as a mixing measure over the former class of estimators. This consideration drags us from the very general PAC-Bayesian bound framework to risk guaranties for aggregate estimators, arising in problems of model selection, and developed in varied settings. In the stochastic design, Yuditskii et al. (2005); Juditsky et al. (2008) develop a Mirror averaging procedure –based on the works of Nemirovsky and Yudin (1983); Nesterov (2009), consisting in an algorithmic scheme that allows for online learning with risk guaranties; which is particularly appropriate for the current increasing interest in reinforcement learning. The papers of Dalalyan and Tsybakov (2007, 2012b) focus on the development of a PAC-Bayesian risk bound for aggregate estimators with sparse priors in deterministic settings, relying on the well-known class of Exponentially weighted aggregate (EWA) estimators. In Dalalyan and Tsybakov (2012a), they show the PAC-Bayesian risk bound obtained in the previous deterministic design is also valid for stochastic online learning with sparse priors by a mirror averaging procedure for sparse priors. Not only these results are valid for the aforementioned different settings, but also, under certain assumptions, apply to a wide class of statistical problems such as regressions, classifications or density estimations; for which there exists an aggregate estimator $\widehat{f}_n$ such that the following inequality holds,

$$\mathbf{E}[\ell(\widehat{f}_n, f)] \leq C_{\mathsf{PB}} \inf_{p \in \mathcal{P}(\mathcal{F}_W)} \left\{ \int_{\mathsf{W}} \ell(f_{\boldsymbol{w}}, f)\, p(d\boldsymbol{w}) + \frac{\beta}{n} D_{\mathsf{KL}}(p||\pi) \right\}, \tag{6.1}$$

where C_{PB} is some universal constant, $\pi \in \mathcal{P}(\mathcal{F}_{\mathsf{W}})$ is a fixed prior distribution on W, $\beta > 0$ a temperature parameter and n is the sample size.

Their very generic use and their flexibility make PAC-Bayesian bounds appealing for the Machine learning community that suffers from a lack of theoretic results. In addition of this observation; Juditsky et al. (2008) prove that learning by mirror averaging, in the case of a discrete prior distribution with a finite class $\mathcal{F}_{\mathsf{W}}$, permits to obtain sharp oracle inequalities that would never be attained for the empirical risk minimizer. Yet, minimizing the empirical risk is the most standard way to calibrate and exploit a neural network. For this reason, investigate aggregate estimators in the framework of neural networks seems to be fully justified, as they

may be a serious and credible alternative method to obtain better theoretic guaranties. In recent works, the empirical PAC-Bayesian framework of McAllester (1999b) has been used to provide generalization bounds for Machine learning problems and neural network calibrations. These works relies on a new trend of the PAC-Bayesian theory – that makes another step aside from the classical Bayesian theory, granting specific data-driven priors such as in Rivasplata et al. (2018), or distribution dependent priors, as in Lever et al. (2013). Especially, in the work of Dziugaite and Roy (2017), they propose to optimize the generalization bound of a given deep neural network with a Gaussian distribution prior, whose parameters are set by performing a stochastic gradient descent over the training data. In the same trend, Neyshabur et al. (2017) also derive a generalization bound for a deep network with a specified architecture and ReLU activation functions. These works develop the notion of stochastic Neural networks and randomized classifiers; and show this kind of neural networks can achieve reasonable generalization bounds through extended numerical experiments for classification problems. They are data-driven and perform rigorous and numerous numerical tests that give credit to the PAC-Bayesian approach.

In the present work, we propose another way to use PAC-Bayesian theory in the framework of neural networks. Relying on the oracle bound (6.1), we search for a risk bound in the case where the parametrized set of functions $\mathcal{F}_{\mathsf{W}}$ is a set of neural networks with a single hidden layer –such that W represents the set of the possible weights of the neural network. We are interested in whether or not these aggregate estimators are able to improve existing risk bounds. In the framework of shallow neural networks, three questions arise naturally when considering a risk bound:

1. How does the weight initialization impact the tightness ?
2. How to choose the size of the hidden layer ?
3. What kind of risk guaranty these choices induce ?

As a consequence, rather than considering a fixed parametrized structure for the neural network, we question the obtained risk bound to derive practical tuning insights for the dependency on the prior distribution and the size of the hidden layer. To meet this goal and obtain a closed form inequality in place of (6.1), we derive a bound that depends on the approximation error and, on a remainder term. Explicit both quantities is the first step of our work. The remaining term can be tackled thanks to a distribution dependent prior and the structure of the neural networks. To bound the approximation error, we exploit existing results on approximation bounds for neural networks with sigmoid and ReLU activation functions.

Approximation bounds have been extensively studied in the basic setting of shallow networks. Most of the time, these studies concern sigmoid activation functions and literature focusing on ReLU/unbounded activation function type is scarcer. In the paper we will focus on Lipchtiz activation functions to cover most of the ones studied in theoritic works as well as used in practical problems. Sigmoid activation functions show good results in regression problems using shallow neural networks. For sigmoid activation functions we distinguish the random approach (Barron, 1993; Delyon et al., 1995; Maiorov and Meir, 2000; Maiorov, 2006)

from the deterministic and constructive approaches (Mhaskar and Micchelli, 1994; Petrushev, 1998; Burger and Neubauer, 2001; Cao et al., 2008; Costarelli and Spigler, 2013a,b). Concerning Piecewize activation functions, they have known an increase of interests these last years with the development of deep neural networks, as they allow for more computational efficiency. Though deep networks seem to be standard with ReLU, we wish to explore first what kind of results can be attained in the simplest case of shallow networks and compare it with result obtained with sigmoid activation functions. Literature on ReLU deep neural networks has known an important development these last years Yarotsky (2017, 2018); Yarotsky and Zhevnerchuk (2019); Gühring et al. (2020); Lu et al. (2020); Shen et al. (2019). While literature on shallow neural networks is scarcer, several relevant results exist Petrushev (1998); Bach (2017); Klusowski and Barron (2016b,a); Xu (2020); Siegel and Xu (2020). In section 4 we provide a more detailed review of the aforementioned works.

In this paper, we provide an optimized oracle PAC-Bayesian bound for aggregate shallow networks with Gaussian prior. Relying on approximation bounds for neural networks, we derive an explicit and tractable risk bound and allow for valuable insights on the structure a network should have, depending on the size of the training sample. The obtained bound outperforms existing bounds for shallow neural networks for the estimation of functions in the Sobolev space $W_2^r([0,1]^{D_0})$. More specifically, for sigmoid activation functions, we show aggregate estimators of neural networks reach the optimal minimax rate $O\big(n^{-\frac{2r}{2r+D_0}}\big)$ up to a $\log(n)$ factor. For ReLU activation functions, under some assumptions on r, we obtain a rate $O\big(n^{-\frac{2\bar{r}}{2\bar{r}+D_0+1}}\big)$ for $\bar{r} < r$, up to a $\log(n)$ factor, which improves on existing risk bounds for shallow networks but remains slightly less tight than some new results on deep neural networks attaining, under specific assumptions, the minimax rate (Schmidt-Hieber, 2020). The latter observation leaves room for future works and improvements on these deeper networks.

The paper is organized as follows. In Section 6.2, we define the generic PAC-Bayesian framework and specify it in the setting of shallow neural networks. Section 6.3 provides various examples of statistical problems where PAC-Bayesian bounds of type (6.1) apply. In Section 6.4, we derive an oracle bound for shallow neural networks with an optimized Gaussian prior. Section 6.5 is dedicated to the literature on approximation bounds and states results for both bounded (sigmoid type) and unbounded (ReLU type) activation functions. Finally, Section 6.6 contains the main result of this work: we propose a worst-case risk bound for a fully tuned shallow network; followed by a discussion on the related works and existing risk bounds. Section 6.7 concludes the paper and states the main outlooks.

6.2 Preliminaries and notations

In this section, we set the general framework of the PAC-Bayesian bound that will be the starting point of our work. We then precise it in the specific case of neural networks and provide some statistical examples for which this bound holds.

6.2.1 General framework

Let $(\mathcal{Z}, \mathscr{A})$ be a measurable space. We observe one realization of the random vector $\mathbf{Z}^n = (Z_1, \ldots, Z_n) \in \mathcal{Z}^n$ drawn from an unknown distribution $\mathcal{P}$ on $(\mathcal{Z}^n, \mathscr{A}^{\otimes n})$. Let $\mathcal{X} \subset \mathbb{R}^{D_0}$, $D_0 \geq 1$, be a Borel set and let μ be a σ-finite measure on $\big(\mathcal{X}, \mathscr{B}(\mathcal{X})\big)$ such that $M_2^2 = D_0^{-1} \int_{\mathcal{X}} \|\boldsymbol{x}\|_2^2 \, \mu(d\boldsymbol{x}) < +\infty$. In the sequel, we denote by $\mathbb{L}_q(\mu)$, $q \in [1, \infty)$, the set of all the functions $f : \mathcal{X} \to \mathbb{R}^{D_2}$ such that $\int_{\mathcal{X}} \|f(\boldsymbol{x})\|_2^q \, \mu(d\boldsymbol{x}) < \infty$.

We consider the problem of estimating a function $f_{\mathcal{P}} \in \mathbb{L}_2(\mu)$ characterized by the distribution $\mathcal{P}$. At this stage, one may think of $f_{\mathcal{P}}$ as the multidimensional regression function when $\mathcal{Z} = \mathcal{X} \times \mathbb{R}^{D_2}$, the Bayes classifier when $\mathcal{Z} = \mathcal{X} \times \{-1, 1\}$ or the density of observations when $\mathcal{Z} = \mathcal{X}$ (in the last two cases $D_2 = 1$). A common approach in statistics and statistical learning is to use a parametric set $\mathcal{F}_{\mathsf{W}} := \{f_{\boldsymbol{w}}, \boldsymbol{w} \in \mathsf{W}\} \subset \mathbb{L}_2(\mu)$, indexed by a measurable set $\mathsf{W} \subset \mathbb{R}^d$, for some $d \in \mathbb{N}$, for constructing an estimator of $f_{\mathcal{P}}$. Instances of this approach are the empirical risk minimizer, the Bayesian posterior mean, the exponentially weighted aggregate, etc. The quality of an estimator $\widehat{f}_n$ of $f_{\mathcal{P}}$ is measured by means of a loss function $\ell : \mathbb{L}_2(\mu) \times \mathbb{L}_2(\mu) \mapsto \mathbb{R}_+$; an estimator $\widehat{f}_n$ is good if its risk

$$\mathbf{E}_{\mathcal{P}}\big[\ell\big(\widehat{f}_n(\mathbf{Z}^n), f_{\mathcal{P}}\big)\big] = \int_{\mathcal{Z}^n} \ell\big(\widehat{f}_n(\mathbf{z}), f_{\mathcal{P}}\big) \, \mathcal{P}(d\mathbf{z})$$

is small. A widespread choice of the loss function, used throughout this paper except in Section 6.3, is the squared ℓ_2-norm

$$\ell(g, h) = \|g - h\|_{\mathbb{L}_2(\mu)}^2 = \int_{\mathcal{X}} \|g(\boldsymbol{x}) - h(\boldsymbol{x})\|_2^2 \, \mu(d\boldsymbol{x}), \quad \forall g, h \in \mathbb{L}_2(\mu).$$

6.2.2 PAC-Bayesian type upper bounds

Let us describe what we call the PAC-Bayes type upper bound, which is the central quantity we wish to optimize in the case where $\mathcal{F}_{\mathsf{W}}$ is the set of all neural networks with a given architecture. Let $\mathcal{P}_{\mathsf{W}}$ be the space of all probability measures on W and let

$$\mathcal{P}_1(\mathcal{F}_{\mathsf{W}}) = \bigg\{ p \in \mathcal{P}_{\mathsf{W}} : \int_{\mathsf{W}} \|f_{\boldsymbol{w}}(\boldsymbol{x})\|_2 \, p(d\boldsymbol{w}) < \infty, \quad \text{for all} \quad x \in \mathcal{X} \bigg\}.$$

Let $\pi \in \mathcal{P}_1(\mathcal{F}_{\mathsf{W}})$ be fixed; we will refer to π as prior distribution. We say that the estimator $\widehat{f}_n$ satisfies the PAC-Bayesian bound with prior π and temperature parameter β, if

$$\mathbf{E}_{\mathcal{P}}[\ell(\widehat{f}_n, f_{\mathcal{P}})] \leq C_{\mathsf{PB}} \inf_{p \in \mathcal{P}_1(\mathcal{F}_{\mathsf{W}})} \bigg\{ \int_{\mathsf{W}} \ell(f_{\boldsymbol{w}}, f_{\mathcal{P}}) \, p(d\boldsymbol{w}) + \frac{\beta}{n} D_{\mathsf{KL}}(p||\pi) \bigg\}, \tag{6.2}$$

for some universal constant C_{PB}. If $C_{\mathsf{PB}} = 1$, the bound is called exact or sharp. When the loss function is the squared $\mathbb{L}_2$-norm, the PAC-Bayesian bound reads as

$$\mathbf{E}_{\mathcal{P}}[\|\widehat{f}_n - f_{\mathcal{P}}\|_{\mathbb{L}_2}^2] \leq C_{\mathsf{PB}} \inf_{p \in \mathcal{P}_1(\mathcal{F}_{\mathsf{W}})} \bigg\{ \int_{\mathsf{W}} \|f_{\boldsymbol{w}} - f_{\mathcal{P}}\|_{\mathbb{L}_2(\mu)}^2 \, p(d\boldsymbol{w}) + \frac{\beta}{n} D_{\mathsf{KL}}(p||\pi) \bigg\}. \tag{6.3}$$

6.2.3 Shallow neural networks

In the rest of this section, we provide more details on the notations and assumptions that will stand when we estimate $f_{\mathcal{P}}$ by aggregation of neural networks. We consider the class of networks with a single hidden layer and denote by D_1 the number of units in this layer.

In order to merge weights and biases of a neural network, we note $\boldsymbol{x} = (1, x_1, \ldots, x_{D_0-1})^\top \in \mathcal{X}$. The set W of the weights of a neural network can be divided into the weights of the hidden layer, $\boldsymbol{w}_1$, and the weights of the output layer, $\boldsymbol{w}_2$ so that $\boldsymbol{w}_1 \in \mathbb{R}^{D_0 \times D_1}$ and $\boldsymbol{w}_2 \in \mathbb{R}^{D_1 \times D_2}$. Therefore, $\boldsymbol{w} = (\boldsymbol{w}_1, \boldsymbol{w}_2)^\top$ can be seen as an element of $\mathbb{R}^d$ with the overall dimension $d = D_0 D_1 + D_1 D_2$. The neural network parametrized by $\boldsymbol{w}$ has the form:

$$f_{\boldsymbol{w}}(\boldsymbol{x}) = \boldsymbol{w}_2^\top \bar{\sigma}(\boldsymbol{w}_1^\top \boldsymbol{x}) \in \mathbb{R}^{D_2}, \quad \forall \boldsymbol{x} \in \mathbb{R}^{D_0} \tag{6.4}$$

with

$$\bar{\sigma} : \boldsymbol{x} \in \mathbb{R}^{D_1} \mapsto \begin{bmatrix} \sigma(x_1) \\ \vdots \\ \sigma(x_{D_1}) \end{bmatrix} \in \mathbb{R}^{D_1},$$

where $\sigma : \mathbb{R} \to \mathbb{R}$ is an activation function.

In the next sections we will consider both the case of bounded and unbounded activation functions in order to cover most of the usual ones. We refer to the bounded case by means of the following assumption.

Assumption (σ-B). *The function σ is bounded by M_σ, i.e, $|\sigma(u)| \leq M_\sigma$ for all $u \in \mathbb{R}$.*

Let us stress that only some of our results require Assumption (σ-B) to be satisfied. However, all our results will require the Lipschitz assumption stated below, which is satisfied by sigmoid functions as well as piecewise continuous functions (including ReLU). Without loss of generality, we will assume that the Lipschitz constant is equal to one.

Assumption (σ-L). *For every pair of real numbers (u, u'), we have $|\sigma(u) - \sigma(u')| \leq |u - u'|$.*

6.2.4 Spherical Gaussian prior distribution

The prior distribution π defined in the PAC-Bayesian framework can be interpreted as the initial distribution of the weights in the neural network setting. Note that most of the time, in practice, the weights are initialized randomly from a Gaussian distribution $\mathcal{N}_d(0, \tau \mathbf{I}_d)$, $\tau > 0$. Building upon this remark, in the remaining part of the paper, we consider π to be a product of spherical Gaussians.

Recall that the weights of a neural network are split into two groups: the weights $\boldsymbol{w}_1$ of the hidden layer and the weights $\boldsymbol{w}_2$ of the output layer. To take into account their different roles we assume the distribution over $\boldsymbol{w}$, whether for the prior π or a distribution p involved in the right hand side of equation (6.3), is a product distribution that may have different parameters in function of the layers considered. We refer to π_1 and π_2 as the distribution of the hidden layer and the output layer respectively.

Assumption ($\mathcal{N}$). *The prior π satisfies* $\pi = \pi_1 \otimes \pi_2 = \mathcal{N}(0, \rho_1^2 \mathbf{I}_{D_0 D_1}) \otimes \mathcal{N}(0, \rho_2^2 \mathbf{I}_{D_1 D_2})$.

6.3 Examples of application

The bound (6.2) is stated in a very general fashion. From the literature in introduction, there are several frameworks and design assumptions under which it applies. In the next section, we state some examples of applications under which the result (6.2) holds.

Among the most common estimation problems; regression, density estimation and classification occur regularly. They may arise in the framework of deterministic (*e.g.*, signal and image processing) as well as in random design, that remains predominant in the machine learning literature. In the sequel, we provide four examples that cover these different frameworks and for which a version of (6.2) applies.

6.3.1 Fixed design regression

Regression with deterministic design and additive errors is often used in nonparametric modeling. In the case of Gaussian errors, it corresponds to the observations

$$\boldsymbol{Z}_i = f(\boldsymbol{x}_i) + \sigma \boldsymbol{\xi}_i, \qquad \boldsymbol{\xi}_i \overset{\text{iid}}{\sim} \mathcal{N}(0, \mathbf{I}_{D_2}), \quad i = 1, \dots, n,$$

where $\boldsymbol{x}_1, \dots, \boldsymbol{x}_n$ are given deterministic points and $\mathcal{Z} = \mathbb{R}^{D_2}$. In this case, the measure μ is the empirical uniform distribution: $\mu = \frac{1}{n} \sum_{i=1}^n \delta_{\boldsymbol{x}_i}$.

There are many results of type (6.2) in the literature for regression with fixed design. In particular, Leung and Barron (2006); Dalalyan and Tsybakov (2007, 2008); Dalalyan (2020) established a PAC-Bayesian bound for the exponentially weighted aggregate defined by $\widehat{f}_n(\mathbf{Z}, \boldsymbol{x}) = \int_{\mathsf{W}} f_{\boldsymbol{w}}(\boldsymbol{x})\, \widehat{\theta}_{n,\boldsymbol{w}}(\mathbf{Z})\, \pi(d\boldsymbol{w})$ with

$$\widehat{\theta}_{n,\boldsymbol{w}}(\mathbf{Z}) = \frac{\exp\Big\{-\frac{1}{\beta} \sum_{i=1}^n \|\boldsymbol{Z}_i - f_{\boldsymbol{w}}(\boldsymbol{x}_i))\|_2^2\Big\}}{\int_{\mathsf{W}} \exp\Big\{-\frac{1}{\beta} \sum_{i=1}^n \|\boldsymbol{Z}_i - f_{\widetilde{\boldsymbol{w}}}(\boldsymbol{x}_i))\|_2^2\Big\}\, \pi(d\widetilde{\boldsymbol{w}})}. \tag{6.5}$$

Note that $\boldsymbol{w} \mapsto \widehat{\theta}_{\boldsymbol{w}}$ is a probability density on (W, π), often referred to as posterior density. We refrain from stating here the precise conditions under which the PAC-Bayesian bound holds true, the interested reader is referred to the papers mentioned above.

6.3.2 Random design regression

In the remaining examples, we place ourselves in the setting of random design. In this case, the bound exposed in introduction is valid for sequential learning and relies on the mirror averaging (MA) estimator. It follows from (Juditsky et al., 2008; Dalalyan and Tsybakov, 2012a; Gerchinovitz, 2013) that this estimator satisfies (6.2) in various estimation problems. We cite below 3 examples without going too much into details. Interested reader is referred to (Juditsky

et al., 2008; Dalalyan and Tsybakov, 2012a) for more detailed and comprehensive account on the topic.

In this case, the estimation problem writes as in the previous example

$$\boldsymbol{Y}_i = f(\boldsymbol{X}_i) + \sigma\boldsymbol{\xi}_i, \qquad \boldsymbol{\xi}_i \sim \mathcal{N}(0, \mathbf{I}), \quad \boldsymbol{\xi}_i \perp\!\!\!\perp \boldsymbol{X}_i, \quad i = 1, \dots, n,$$

with $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, $\mathcal{X} \subset \mathbb{R}^{D_0}$, $\mathcal{Y} \subset \mathbb{R}^{D_2}$ and $\mathcal{P}$ is the distribution of n iid copies of $(\boldsymbol{X}_i, \boldsymbol{Y}_i)$ satisfying the regression equation of the last display. The natural choice of the measure μ here is the marginal distribution of $\boldsymbol{X}_i$ over $\mathcal{X}$.

We describe now the Mirror averaging procedure that allows to obtain the guaranties given by (6.2) for this problem. It is an online recursive algorithm which takes the form

$$\widehat{f}_n(\mathbf{Z}, x) = \int_{\mathsf{W}} f_{\boldsymbol{w}}(x)\widehat{\theta}^{\mathrm{MA}}_{n,\boldsymbol{w}}(\mathbf{Z})\pi(d\boldsymbol{w}) = \frac{1}{n+1}\sum_{m=0}^{n}\int_{\mathsf{W}} f_{\boldsymbol{w}}(x)\widehat{\theta}_{m,\boldsymbol{w}}(\mathbf{Z})\pi(d\boldsymbol{w}) \tag{6.6}$$

with $\widehat{\theta}_{0,\boldsymbol{w}} = 1$ and

$$\widehat{\theta}_{\boldsymbol{w}}(\mathbf{Z}) = \frac{1}{n+1}\sum_{m=0}^{n}\frac{\exp\Big\{-\frac{1}{\beta}\sum_{i=1}^{m} Q(\boldsymbol{Z}_i, f_{\boldsymbol{w}})\Big\}}{\int_{\mathsf{W}}\exp\Big\{-\frac{1}{\beta}\sum_{i=1}^{m} Q(\boldsymbol{Z}_i, f_{\widetilde{\boldsymbol{w}}})\Big\}\pi(d\widetilde{\boldsymbol{w}})}, \tag{6.7}$$

where $Q : \mathcal{Z} \times \mathbb{L}_2(\mu) \mapsto \mathbb{R}$ is a mapping satisfying some assumptions under which the minimizer of the loss $\ell : g \mapsto \ell(g, f)$ coincides with the minimizer of $g \mapsto \mathbf{E}_{\mathcal{P}}[Q(\mathbf{Z}, g)]$. In the case of regression, the Mirror averaging estimator can be evaluated with the ℓ_2-norm such that in (6.7), the function Q is given by $Q(\boldsymbol{Z}_i, f_{\boldsymbol{w}}) = \|\boldsymbol{Y}_i - f_{\boldsymbol{w}}(\boldsymbol{X}_i)\|_2^2$.

By its very general design, the estimator defined by (6.6)–(6.7) also applies for density estimation and classification as presented in the following paragraphs.

6.3.3 Density estimation

Consider the case where the elements of $\mathbf{Z}^n = (\boldsymbol{Z}_1, \dots, \boldsymbol{Z}_n) \in \mathcal{Z}^n$ are iid random variables drawn from a distribution having f as density with respect to a measure μ. We aim to estimate f and measure the risk of the mirror averaging estimator $\widehat{f}_n$ using the integrated risk:

$$\ell(\widehat{f}_n, f) = \|\widehat{f}_n - f\|^2_{\mathbb{L}_2(\mu)} = \int_{\mathcal{X}} \big(\widehat{f}_n(\boldsymbol{x}) - f(\boldsymbol{x})\big)^2 \mu(d\boldsymbol{x})$$

such that the mapping Q in (6.7) can by defined by $Q(\boldsymbol{x}, g) = \|g\|^2_{\mathbb{L}_2(\mu)} - 2g(\boldsymbol{x})$, and we can retrieve the MA estimator by the procedure (6.6)–(6.7).

6.3.4 Classification for Φ-risk

The measurable space $\mathcal{Z} = \mathcal{X} \times \mathcal{Y}$, gathers the state space $\mathcal{X} \subset \mathbb{R}^{D_0}$, and the labels $\mathcal{Y} = \{-1, +1\}$ such that $\mathbf{Z}^n = ((X_1, Y_1), \dots, (X_n, Y_n)) \in \mathcal{Z}^n$ are the iid observations from $\mathcal{P}$. The goal is to find a classifier $\widehat{f}_n$ such that for each $X_i \in \mathcal{X}$, it determines the label with a small

classification risk and such that it is close to the oracle classifier $f_\Phi \in \operatorname{argmin}_g R_\Phi[g]$ where the Φ-risk is defined as a twice differentiable convex function whose risk is given by

$$R_\Phi[\widehat{f}_n] = \int_{\mathcal{X}\times\{-1,+1\}} \Phi(-y\widehat{f}_n(x))P(dx,dy).$$

The loss is then given by the excess risk $\ell(\widehat{f}_n, f_\Phi) = R_\Phi[\widehat{f}_n] - R_\Phi[f_\Phi]$, and is such that the function Q can be defined by $Q(z,g) = \Phi(-yg(x))$. As in the previous examples, we can retrieve the MA estimator by the procedure (6.6)–(6.7).

The examples above confirm how wide is the area of applications of convex aggregate estimators verifying (6.2). In the next sections we work on the inequality (6.2) to find an appropriate way to take advantage of it in the framework of neural networks. There will be roughly speaking two main quantities we will focus on: the estimation bound and the approximation bound.

6.4 Risk bounds

In this section, we first derive a bound for the risk of the estimator $\widehat{f}_n$ when the prior has an arbitrary centered Gaussian distribution, and subsequently provide an oracle inequality for a carefully chosen Gaussian prior. Unless it is stated otherwise, we place ourselves in the general framework of Section 6.2.1.

Let $\bar{\boldsymbol{w}} \in \mathsf{W}$ be any parameter value. Using the triangle inequality, in conjunction with the fact that $\sqrt{(a+b)^2+c^2} \le a + \sqrt{b^2+c^2}$, one can infer from (6.3) that

$$\Big(C_{\mathsf{PB}}^{-1}\,\mathbf{E}_{\mathcal{P}}\big[\|\widehat{f}_n - f_{\mathcal{P}}\|^2_{\mathbb{L}_2(\mu)}\big]\Big)^{1/2} \le \|f_{\bar{\boldsymbol{w}}} - f_{\mathcal{P}}\|_{\mathbb{L}_2(\mu)} + \mathrm{Rem}_n(\bar{\boldsymbol{w}})^{1/2}, \tag{6.8}$$

with the remainder term given by

$$\mathrm{Rem}_n(\bar{\boldsymbol{w}}) \triangleq \inf_{p\in\mathcal{P}_1(\mathcal{F}_{\mathsf{W}})} \Big\{ \int_{\mathsf{W}} \|f_{\boldsymbol{w}} - f_{\bar{\boldsymbol{w}}}\|^2_{\mathbb{L}_2(\mu)}\, p(d\boldsymbol{w}) + \frac{\beta}{n} D_{\mathsf{KL}}(p\|\pi)\Big\}. \tag{6.9}$$

Taking $\bar{\boldsymbol{w}} \in \operatorname{argmin}_{\boldsymbol{w}\in\mathsf{W}} \|f_{\boldsymbol{w}} - f_{\mathcal{P}}\|_{\mathbb{L}_2(\mu)}$, the right hand side of inequality (6.8) can be decomposed into the approximation error $\inf_{\boldsymbol{w}\in\mathsf{W}} \|f_{\boldsymbol{w}} - f_{\mathcal{P}}\|_{\mathbb{L}_2(\mu)}$ and the remainder term $\mathrm{Rem}_n(\bar{\boldsymbol{w}})$ assessing the estimation error. The main goal of this paper is to analyze this estimation error. Our approach will consist in replacing the infinimum over all measure by the infinimum over suitably chosen Gaussian measures. This modification has the advantage of considerably simplifying the mathematical computations.

On the other hand, it is a well-known fact that for a fixed $\bar{\boldsymbol{w}}$, the infinimum in the expression (6.9) is attained at the distribution p^* given by thd Gibbs distribution (see, for example, McAllester (2003); Guedj (2019))

$$p^*(d\boldsymbol{w}) \propto \exp\Big\{-\frac{n}{\beta}\|f_{\boldsymbol{w}} - f_{\bar{\boldsymbol{w}}}\|^2_{\mathbb{L}_2(\mu)}\Big\}\pi(d\boldsymbol{w}).$$

Furthermore in this case,

$$\mathrm{Rem}_n(\bar{\boldsymbol{w}}) = -\frac{\beta}{n}\log\int_{\mathsf{W}}\exp\Big\{-\frac{n}{\beta}\|f_{\boldsymbol{w}} - f_{\bar{\boldsymbol{w}}}\|^2_{\mathbb{L}_2(\mu)}\Big\}\pi(d\boldsymbol{w}).$$

At this stage a question arises: *why don't we simply use this distribution p^* instead of minimizing over the Gaussian distributions?* The reason is that the expression of the last display, often referred to as the free energy, appears to be hard to simplify; therefore, answering the questions raised in the introduction in the framework of neural networks is not an easy task. In contrast with this, choosing as p a Gaussian distributions leads to a tractable expression that has a simple dependence on the relevant quantities such as the dimension, the sample size, etc.

6.4.1 Risk bound for an arbitrary centered Gaussian prior

In this section, we consider the prior chosen as in Assumption ($\mathcal{N}$). We intend to choose a "good" distribution p in $\mathcal{P}_1(\mathcal{F}_{\mathsf{W}})$ in the remainder term (6.9). The term "good" encodes the good properties in terms of tightness of the bound for the remainder term, but also in terms of tractability with respect to the relevant parameters.

Proposition 47. *Let Assumption (σ-L) and Assumption ($\mathcal{N}$) be satisfied. Recall that $d = D_0D_1 + D_1D_2$ is the number of weights of the neural network and n is the sample size.*

i) *If Assumption (σ-B) holds true, then*

$$\mathrm{Rem}_n(\bar{\boldsymbol{w}}) \le \frac{\beta}{2n}\left\{\frac{\|\bar{\boldsymbol{w}}_1\|^2_{\mathsf{F}}}{\rho_1^2} + \frac{\|\bar{\boldsymbol{w}}_2\|^2_{\mathsf{F}}}{\rho_2^2} + d\log\left(1 + \frac{2n(A_1\rho_1^2 + A_2\rho_2^2)}{d\beta}\right)\right\} \tag{6.10}$$

where $A_1 = D_0D_1M_2^2\|\bar{\boldsymbol{w}}_2\|^2_{\mathsf{F}}$ and $A_2 = D_1D_2\mu(\mathcal{X})M_\sigma^2$.

ii) *If the activation function is unbounded but vanishes at the origin, then*

$$\mathrm{Rem}_n(\bar{\boldsymbol{w}}) \le \frac{\beta}{2n}\left\{\frac{\|\bar{\boldsymbol{w}}_1\|^2_{\mathsf{F}}}{\rho_1^2} + \frac{\|\bar{\boldsymbol{w}}_2\|^2_{\mathsf{F}}}{\rho_2^2} + 2d\log\left(1 + \frac{n(A_1\rho_1^2 + A_2'\rho_2^2 + A_3'\rho_1^2\rho_2^2)}{d\beta}\right)\right\} \tag{6.11}$$

where $A_2' = \bar{M}_2^2\|\bar{\boldsymbol{w}}_1\|^2_{\mathsf{F}}D_2$ and $A_3' = M_2^2D_0D_2$.

The proof of the proposition is postponed to Section 6.8.2. It should be noted that the quantities A_1, A_2, A_2' and A_3' defined in the proposition are dimension dependent. They are, however, independent of the sample size n, the temperature parameter β and the variances ρ_1 and ρ_2 of the prior distribution. As already mentioned, the obtained bound is not the sharpest one, since restricting ourselves to Gaussian distributions p might have caused some loss in sharpness. However, this simplification allowed us to perform optimizations and obtain intelligible results in the following sections.

The above established upper bound on the remainder term offers an explicit dependence in the dimension, the parameters of the prior distribution and the sample size. Especially, the variance of the prior (ρ_1^2, ρ_2^2) tends to make the bound larger as it becomes smaller, which is

consistent with the intuition that in the Bayesian framework, too restrictive priors might lead to unwanted bias–as opposed to informative but reasonably general priors such as heavy tailed priors or large variance Gaussian priors. However, from the shape of the bound, one could make more precise choices for the variance of π than the previous qualitative considerations on 'large' or 'small' variances. In the next section, we show how the obtained bound can be used to select the parameters of the Gaussian prior.

6.4.2 Oracle inequality for optimized Gaussian prior

Combining (6.8) and Proposition 47, we get a risk bound that decreases when the number of observations n goes to infinity at a speed $n^{-1/2}$, when the dimension and other characteristics are fixed. This is the well known minimax rate for parametric estimation. However, the power of neural networks is in their attractive properties in nonparametric, high-dimensional setting. Thus, when the parameterization is not frozen beforehand but chosen by the practitioner, or number of observations is limited as compared to the dimensionality—which is the case most of the time in practice—it is important to take into account the other tuning parameters such as

(i) *The variances of the prior distribution* (ρ_1^2, ρ_2^2)

The optimization over (ρ_1^2, ρ_2^2) leads to values depending on the norm of $\bar{\boldsymbol{w}}$, the best value of the parameter in an approximation class, which we ignore. Nevertheless, in a minimax setting, we might know the order of magnitude of the relevant norms of $\bar{\boldsymbol{w}}$ and incorporate this knowledge in the choice of the prior.

(ii) *The dimension of the hidden layer* D_1

While in practice the size of the hidden layer is often fixed by cross validation, in a minimax setting one can exploit the flexibility of the size of this layer to tailor the estimator to the functional class the function $f_{\mathcal{P}}$ is supposed to belong to.

While the discussion on the choice of D_1 is the main topic of the subsequent sections, we present here an upper bound on the remainder term (6.9) obtained after a careful choice of the prior variances. The main idea is to choose ρ_1 and ρ_2, so that the worst-case value of the remainder,

$$\sup_{\bar{\boldsymbol{w}}:\|\bar{\boldsymbol{w}}_\ell\|_{\mathsf{F}}\leq B_\ell} \mathrm{Rem}_n(\bar{\boldsymbol{w}})$$

is minimized. Instead of solving this optimization problem, we will consider the one in which $\mathrm{Rem}_n(\bar{\boldsymbol{w}})$ is replaced by its upper bound furnished by Proposition 47. Furthermore, we will not look for the exact minimizer, but for a relatively simple expression of ρ_1 and ρ_2 that leads to a nearly smallest order of magnitude of the worst-case remainder term.

Corollary 48. *Let Assumptions (σ-L) and ($\mathcal{N}$) be satisfied with variances given by*

$$\rho_\ell^2 = \frac{B_\ell^2}{2D_{\ell-1}D_\ell}, \qquad \ell = 1, 2. \tag{6.12}$$

i) *If Assumption (σ-B) holds true, then*

$$\sup_{\bar{\boldsymbol{w}}:\|\bar{\boldsymbol{w}}_\ell\|_{\mathsf{F}}\leq B_\ell} \mathrm{Rem}_n(\bar{\boldsymbol{w}}) \leq \frac{\beta d}{n}\log\left(3+\frac{3nB_2^2(B_1^2M_2^2+\mu(\mathcal{X})M_\sigma^2)}{d\beta}\right).$$

ii) *If the activation function is unbounded but vanishes at the origin, then*

$$\sup_{\bar{\boldsymbol{w}}:\|\bar{\boldsymbol{w}}_\ell\|_{\mathsf{F}}\leq B_\ell} \mathrm{Rem}_n(\bar{\boldsymbol{w}}) \leq \frac{\beta d}{n}\log\left(3+\frac{3nB_1^2B_2^2(M_2^2+\bar{M}_2^2/D_1)}{d\beta}\right).$$

The proof, deferred to Section 6.8.3, consists in a simple substitution of the values of ρ_ℓ in the bounds obtained in Proposition 47. For pedagogical purposes, we have included in the proof the computation of the "optimal" variances which, after some simplification, lead to the proposal (6.12). Combining this corollary with inequality (6.8), we get the following risk bound.

Proposition 49. *Let $\widehat{f}_n$ be a method of aggregation of shallow neural networks $\mathcal{F}_{\mathsf{W}} = \{f_{\boldsymbol{w}}(\boldsymbol{x}) = \boldsymbol{w}_2^\top\bar{\sigma}(\boldsymbol{w}_1^\top\boldsymbol{x}) : \boldsymbol{w}_1 \in \mathbb{R}^{D_0\times D_1}; \boldsymbol{w}_2 \in \mathbb{R}^{D_1\times D_2}\}$, based on a prior distribution π, satisfying the* PAC-*Bayesian bound* (6.3). *Let Assumptions (σ-L) and ($\mathcal{N}$) be satisfied with variances given by $\rho_\ell^2 = \frac{B_\ell^2}{2D_{\ell-1}D_\ell}$ for $\ell = 1, 2$. Then*

$$\left(C_{\mathsf{PB}}^{-1}\mathbf{E}_{\mathcal{P}}\left[\|\widehat{f}_n - f_{\mathcal{P}}\|_{\mathbb{L}_2(\mu)}^2\right]\right)^{1/2} \leq \inf_{\substack{\|\boldsymbol{w}_1\|_{\mathsf{F}}\leq B_1\\ \|\boldsymbol{w}_2\|_{\mathsf{F}}\leq B_2}} \|f_{\bar{\boldsymbol{w}}} - f_{\mathcal{P}}\|_{\mathbb{L}_2(\mu)} + \left\{\frac{\beta d}{n}\log\left(3+\frac{nE}{d\beta}\right)\right\}^{1/2}, \tag{6.13}$$

where the constant E is defined by

$$E = \begin{cases} 3B_2^2(B_1^2M_2^2+\mu(\mathcal{X})M_\sigma^2), & \text{if } \sigma \text{ satisfies Assumption } (\sigma\text{-B}),\\ 3B_1^2B_2^2(M_2^2+\bar{M}_2^2/D_1), & \text{if } \sigma \text{ is unbounded but } \sigma(0)=0.\end{cases}$$

The important consequence of this proposition is the estimation error, upper bounded by the second term in (6.13), is of the order of $\sqrt{D_1/n}$ (we assume that the input and the output dimensions are fixed and neglect logarithmic factors). This is similar to many non-parametric estimation methods. For instance, if a regression function is estimated by a histogram with K bins, the estimation error is generally of the order $\sqrt{K/n}$. Thus, the number of units in the hidden layer of a neural network plays the same role as the number of bins in a histogram. This parameter D_1 has to be chosen carefully, in order to control the balance between the approximation error and the estimation error.

6.5 Approximation bounds

The goal of this section is to review the existing approximation bounds for shallow neural networks. In particular, we are interested in bounds that depend on the best possible way on the dimension of the hidden layer. We first present the results that are valid for sigmoid activation functions and then focus on the case of the ReLU activation function.

6.5.1 Bounds for sigmoidal activation functions

We provide here a selective overview on the literature dealing with single hidden layer neural networks with sigmoid activation functions. This literature covers different classes of functions to approximate. The quality of the approximation/tightness of the bound largely depends on the hypotheses made on the function f we want to approach. In what follows we gathered the most common classes encountered in the literature.

For the set of univariate locally α-Hölder continuous functions with $\alpha \in (0, 1]$, the constructive approach of (Cao et al., 2008) with fixed weights for the neural network uses discretized values of the target function, and achieves an approximation order of $O(D_1^{-\alpha})$ in ℓ_∞-norm. For $\alpha > 1$, Costarelli and Spigler (2013a) show that the order of the error of approximation is $O(D_1^{-1})$. They also extend the results to multivariate locally Hölder functions, and show that the rate is the same as in the univariate setting.

For other classes of functions, a common feature of the results is the requirement of the existence of some type of integral transform (*e.g.*, Fourier, Radon, wavelet) of the function $f_{\mathcal{P}}$. Each transform is tailored to a different specific "smoothness" class. An early example is the constructive approach from (Mhaskar and Micchelli, 1994) that focused on 2π-periodic functions from $\mathbb{L}_2\left([-\pi, \pi]^{D_0}\right)$ with absolutely convergent Fourier series. For such functions, the approximation error is shown to be of order $O(D_1^{-1/2})$. In the case of random design, the seminal paper (Barron, 1993) deals with functions f, the Fourier transform $\widehat{f}$ of which satisfies $\int_{\mathbb{R}^{D_0}} \|\boldsymbol{z}\|_2 |\widehat{f}(\boldsymbol{z})| \, d\boldsymbol{z} < \infty$, and establishes an upper bound of order $O(D_1^{-1/2})$ for the approximation error. The special characteristics of the target functions led to the definition of the Barron spectral space for which we will give a more precise definition afterwards.

Note that in papers mentioned in previous paragraph, the smoothness of the function and on the dimension of the input variable $\boldsymbol{X}$ do not appear in the risk bound. In contrast with this, for Sobolev spaces, Petrushev (1998) showed how the dimension of the input space and the order of the Sobolev space impact the approximation. He defined a very general class of shallow neural networks $\mathscr{R}_{D_1}$, known as ridge functions, such that $\mathscr{R}_{D_1} := \left\{ \sum_{d=1}^{D_1} \sigma_d(\boldsymbol{w}.x), \ x \in \Omega \right\}$ where Ω is a finite subset of vectors from the unit ball of $\mathbb{R}^{D_0}$ and $\sigma_d, \ d = 1, \ldots, D_1$ are univariate functions of $\mathbb{L}_2([-1, 1])$. This work proved that if we can approach a function from the Sobolev space $W_2^{r+(D_0-1)/2}([-1, 1]^{D_0})$ of smoothness r with a ridge function up to an error of order $O(D_1^{-r-(D_0-1)/2})$ in the univariate case $D_0 = 1$ then, this is also the case for multivariate functions. The advantage of this formulation is that a very large class of activation functions are considered. Later (Burger and Neubauer, 2001) showed that functions from the Sobolev space $W_2^r(\Omega)$, $\Omega \subset \mathbb{R}^{D_0}$, can be approximated to the order $O(D_1^{-1/2-\rho/D_0})$, when the activation function is ρ-Hölder continuous. Further, building on (Delyon et al., 1995), Maiorov and Meir (2000); Maiorov (2006) proved that the approximation error is of order $O(D_1^{-r/D_0})$ up to a $\log(D_1)$-factor.

In the present paper, we use the result of Maiorov and Meir (2000) and Maiorov (2006) to upper bound the approximation error in (6.13). To this end, consider a function $f \in \mathbb{L}_2(\mathbb{R}^{D_0}) \cap \mathbb{L}_1(\mathbb{R}^{D_0})$. We denote its Fourier transform $\widehat{f}$, such that $\widehat{f}(\boldsymbol{z}) = (2\pi)^{-D_0/2} \int_{\mathbb{R}^{D_0}} f(\boldsymbol{x}) \, e^{i\langle \boldsymbol{z}, \boldsymbol{x} \rangle} d\boldsymbol{x}$

and define the quantity $\mathsf{D}^{\boldsymbol{\alpha}} f := \mathscr{F}^{-1}\{\boldsymbol{z}^{\boldsymbol{\alpha}} \widehat{f}(\boldsymbol{z})\}$. The class of function considered in their paper is the unit ball of the Sobolev space that we define as

$$W_2^r([0,1]^{D_0}, \mu) = \left\{ f : \max_{0 \le \alpha \le r} \|\mathsf{D}^{\alpha} f\|_{\mathbb{L}^2(\mu)} \le 1 \right\}.$$

We now introduce additional conditions on the sigmoid functions before presenting the precise statement of the result. Let $\varphi \in \mathbb{L}_2(\mathbb{R}) \cap \mathbb{L}_1(\mathbb{R})$. Using φ, we can construct another function ψ satisfying

$$\int_0^{\infty} \frac{1}{a} \widehat{\varphi}(az) \overline{\widehat{\psi}}(az) \, da = 1, \ \forall z \tag{6.14}$$

where $\widehat{\varphi}$, $\widehat{\psi}$ are the Fourier transforms of φ and ψ. We define Φ^r as the set of all functions $\varphi \in \mathbb{L}_2(\mathbb{R}) \cap \mathbb{L}_1(\mathbb{R})$ such that there exists ψ satisfying (6.14) and $\forall \rho \in [0, r]$, $D^{\rho}\varphi \in \mathbb{L}_2(\mathbb{R})$, $D^{-\rho}\psi \in \mathbb{L}_1(\mathbb{R})$ and set the constant $B_{\varphi} = \max_{\rho \in [0,r]} \left\{ \|D^{\rho}\varphi\|_{\mathbb{L}_2(\mathbb{R})}, \|D^{-\rho}\psi\|_{\mathbb{L}_1(\mathbb{R})} \right\}$. We are now ready to recall the main theorem.

Theorem 50 (Maiorov (2006), Theorem 2.3)**.** *Let the measure μ be such that it admits a bounded density with respect to the Lebesgue measure and σ be any sigmoid function such that the function $\varphi(t) = \sigma(t+1) - \sigma(t)$ belongs to the class Φ^r. Then, for any function $f \in W_2^r([0,1]^{D_0}, \mu)$, there exists a neural network $f_{\boldsymbol{w}^*}$ defined as in (6.4) such that*

$$\|f - f_{\boldsymbol{w}^*}\|_{\mathbb{L}_2(\mu)} \le c_1 B_{\varphi} \frac{\log(D_1)}{D_1^{r/D_0}} \quad \textit{and} \quad |f_{\boldsymbol{w}^*}(\boldsymbol{x})| \le c_2 B_{\varphi} n^{(\frac{1}{2} - \frac{r}{D_0})_+}, \ \forall \boldsymbol{x} \in [0,1]^{D_0},$$

where c_1 and c_2 are constants depending only on the problem dimension D_0 and on the regularity parameter r.

Maiorov and Meir (2000) and Maiorov (2006) provide some examples of functions φ satisfying condition (6.14). However, they do not push further the analysis of the properties of the resulting function σ. To circumvent this shortcoming, we try here to have a closer look to the example obtained by using the function $\varphi(x) = \frac{1}{\sqrt{2}} e^{-x^2/2}$, which satisfies (6.14) with $\psi(x) = \frac{1}{\sqrt{2}}(1 - x^2)e^{-x^2/2}$. In this case, we can prove the following result.

Lemma 51. *The function $\sigma : \mathbb{R} \mapsto \mathbb{R}$ defined by*

$$\sigma(x) = \sum_{j=1}^{\infty} \varphi(x - j), \quad \forall x \in \mathbb{R}, \quad \textit{where} \quad \varphi(x) = \frac{1}{\sqrt{2}} e^{-x^2/2} \tag{6.15}$$

is a continuous function such that, $\sigma(x) \ge 0$, $\lim_{x \to -\infty} \sigma(x) = 0$ and $\sigma(x) \le \sqrt{\pi} + 1/\sqrt{2} \le 2.5$ for every $x \in \mathbb{R}$. Furthermore, σ is 1-Lipschitz continuous.

A proof of this result is postponed to Section 6.8.4. In Figure 6.1 we display the function σ defined in Lemma 51, as well as its limit behaviour when $|x| \to +\infty$.

Figure 6.1: Activation function satisfying the Maiorov condition with $\varphi(x) = \frac{1}{\sqrt{2}} e^{-x^2/2}$

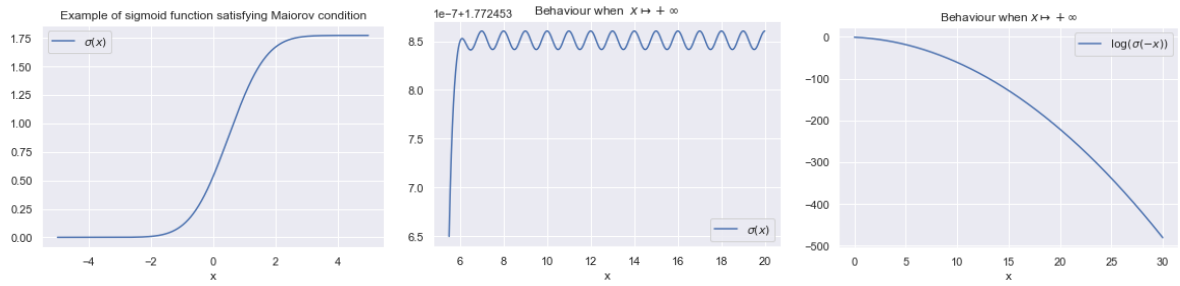

On the left hand side plot of Figure 6.1, the function shape looks very much like the standard sigmoid functions. In addition, the middle and the right hand side graphs show a limit behavior that is in line with the fact that $\sigma(x) \leq const \approx \sqrt{\pi}$ and $\lim_{x \to -\infty} \sigma(x) = 0$.

We can also consider the case where $\varphi(x) = \frac{x+1}{3} \mathbb{1}_{[-1,0]}(x) + \frac{1-x}{3} \mathbb{1}_{[0,1]}(x)$, for which we displayed, in Figure 6.2, the corresponding activation function σ. The function σ is derived using the same methodology as for the case of Lemma 51 (see also Section 6.8.4). We refer the reader to (Maiorov and Meir, 2000; Maiorov, 2006) for further examples.

Figure 6.2: Sigmoid function satisfying the Maiorov condition with $\varphi(x) = \frac{x+1}{3} \mathbb{1}_{[-1,0]}(x) + \frac{1-x}{3} \mathbb{1}_{[0,1]}(x)$

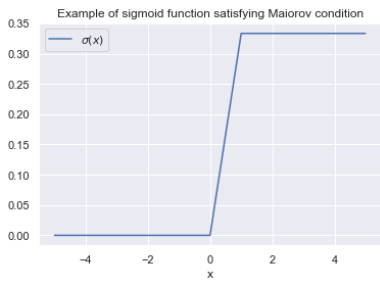

6.5.2 Bounds for the ReLU activation function

The literature on deep neural networks with ReLU activation has significantly grown these last years thanks to the computational benefits of considering piecewise linear activation functions. In this paper, we wish to derive risk bounds in the simplest possible case that are shallow neural networks. However, ReLU are increasingly—and almost exclusively in practice—used in deep networks. Hence, for the sake of completeness, in the first instance, we present important existing results for deeper neural networks.

Deep networks We should mention the work of Yarotsky (2017) that derives upper and lower bounds for deep neural networks. He shows a network of depth $O(\log(1/\varepsilon))$ and complexity $O\left(\epsilon^{D_0/r} \log(1/\varepsilon)\right)$ can approximate any functions of the unit ball of the Sobolev space $W^r_\infty([0,1]^{D_0})$ with an error measured in infinite norm smaller than $\varepsilon > 0$. Allowing for discontinuous weight assignment and introducing a "diagram phase" for ReLU approximation, Yarotsky

(2018) shows that the error rate proved in (Yarotsky, 2017) for continuous functions can be upgraded up to $O(\mathcal{W}^{-2/D_0})$, where $\mathcal{W}$ is the total number of neurons, which leads to a better quantification of the error. Gühring et al. (2020) generalized the results of Yarotsky (2017) to higher order Sobolev norm. Yarotsky and Zhevnerchuk (2019) showed results in (Yarotsky, 2018) could also apply to Sobolev space with regularity $r > 1$.

However, these results are expressed either as a function of an arbitrary error $\varepsilon > 0$, or, in terms of a total number of neurons $\mathcal{W}$. Recent works of Lu et al. (2020) and Shen et al. (2019) obtain a more general and quantified approximation bound as a function of the depth L and width N of the neural network, which is arguably more informative than the total number of neurons $\mathcal{W}$, especially for practical uses. In Lu et al. (2020) networks can approximate functions from the Sobolev space $W^r_\infty([0,1]^{D_0})$ with an error lower than $O(N^{-\frac{2r}{D_0}}L^{-\frac{2r}{D_0}})$, where $O(\log(L)L)$ is the depth order of the network and $O(\log(N)N)$ is its width. In Shen et al. (2019), an error bound of $O(N^{-\frac{2\alpha}{D_0}}L^{-\frac{2\alpha}{D_0}})$ for α-Hölder continuous functions is established.

Shallow networks Literature for shallow neural networks with specific ReLU activation function is scarcer. The seminal work of Petrushev (1998) already mentioned for sigmoid activation functions is also valid in the present case since it applies to all ridge functions. Bach (2017) obtains an approximation bound of order $O\big(\eta D_1^{-1/D_0}\big)$ for η-Lipschitz function . Results on Barron spectral space have been developed thanks to the seminal work of (Makovoz, 1996). Klusowski and Barron (2016a) show that we can reach an approximation bound of $O({D_1}^{-\frac{1}{2}-\frac{1}{D_0}})$ for functions from $\mathscr{B}^2([0,1]^{D_0})$, where $\mathscr{B}^s$ designates the Barron spectral space of order s defined as

$$\mathscr{B}^s(\Omega) := \left\{ f : \Omega \mapsto \mathbb{C} : \|f\|_{\mathscr{B}^s(\Omega)} := \inf_{f_e|\Omega = f} \int_{\mathbb{R}^d} (1 + \|\boldsymbol{z}\|_2)^s \, |\widehat{f_e}(\boldsymbol{z})| d\boldsymbol{z} < \infty \right\}, \tag{6.16}$$

where f_e is an $L^1(\mathbb{R}^d)$ extension of f, $\widehat{f_e}$ its Fourier transform, and $\Omega \subset \mathbb{R}^{D_0}$ is a bounded domain (see Siegel and Xu (2020)). Xu (2020) reaches the same bound with ReLU^k activation functions for target functions in $\mathscr{B}^{k+1}([0,1]^{D_0})$. Very recently, the paper of Siegel and Xu (2020), makes another step forward for the approximation guaranties of shallow neural networks. Their result is, to the best of our knowledge, the tightest upper bound for shallow neural networks with ReLU ^{k} activation functions. It shows that the approximation error for $\mathscr{B}^{k+1}$ is of order $O\big(D_1^{-(k+1)}\big)$ in $\mathbb{L}^2$ norm.

We choose to use the latter result (Siegel and Xu (2020)) in our work, in particular, we will use Theorem 4. This Theorem applies for any complex valued function of Barron spectral space and activation function ReLU^k defined as

$$\sigma^k(x) = \max(0, x)^k.$$

In our case, we are only interested in real-valued function approximation and the estimation results we have for unbounded activation functions only apply in the case it is Lipschtiz. Hence we state their theorem in this precise case, when $k = 1$.

Theorem 52 (Siegel and Xu (2020)). *Let $\Omega = [0,1]^{D_0}$ and $s \geq 1/2$. Suppose that $f : \Omega \to \mathbb{R}$ belongs to $\mathscr{B}^s(\Omega)$. Then, for $D_1 \geq 2$,*

$$\|f - f_{\boldsymbol{w}^*}\|_2 \leq C\|f\|_{\mathscr{B}^s} D_1^{-K} \log^m(D_1),$$

where K and m are given by

$$K = \begin{cases} 2 & \text{if } 2s \geq 3D_0 + 4 \\ \frac{1}{2} + \frac{2s-1}{2(D_0+1)} & \text{if } 2s < 3D_0 + 4 \end{cases} \qquad \text{and} \qquad m = \begin{cases} 0 & \text{if } 2s < 3D_0 + 4 \\ 1 & \text{if } 2s > 3D_0 + 4 \\ \frac{5}{2} & \text{if } 2s = 3D_0 + 4 \end{cases}$$

and C is a constant that depends on s and D_0 but not on D_1.

In the next part we use the Theorems 50 and 52 to obtain a tractable risk bound that allows an optimization with respect to D_1. Note however that in most of the papers, the value of the constants are not specified excepted for their dependency in D_1. A consequence of this is that we partially lose the information on the role of the input dimension in the risk. It follows that the curse of dimensionality might not be totally captured.

6.6 Worst-case risk bounds over smoothness classes

This section gathers the main results of this paper. We first derive a risk bound for networks with Sigmoid activation functions and then treat the case of ReLU. Finally, we provide some point of comparison with existing results.

6.6.1 Sigmoid activation functions

In this section we restrict ourselves to real valued functions (*i.e.*, $D_2 = 1$) belonging to the unit ball of the Sobolev space, $f_{\mathcal{P}} \in W_2^r(\mathcal{X}, \mu)$. Using claims of Proposition 49 and Theorem 50, we can express both the estimation error and the approximation error as a function of the number D_1 of neurons of the hidden layer for an activation function that satisfies the conditions of Theorem 50. We obtain the risk bound

$$C_{\mathsf{PB}}^{-1}\mathbf{E}_{\mathcal{P}}\left[\|\widehat{f}_n - f_{\mathcal{P}}\|^2_{\mathbb{L}_2(\mu)}\right] \leq 2c_1^2 B_\varphi^2 \frac{\log^2(D_1)}{D_1^{2r/D_0}} + \frac{4\beta D_1 D_0}{n} \log\left(3 + \frac{nE}{d\beta}\right), \tag{6.17}$$

where $c_1 B_\varphi$ is as defined in Theorem 50, and E is defined in Proposition 49. As mentioned in Section 6.4, the impact of the number of neurons on the hidden layer D_1, is opposite whether we consider the approximation error or the estimation one. This pertains to the well known estimation-approximation dichotomy leading to the bias-variance trade-off. Our goal is therefore to choose the size of the network guaranteeing the best trade-off in terms of the sample size.

If we neglect logarithmic factors and optimize the risk bound (6.17) with respect to D_1, we

see that the point of minimum should satisfy

$$D_1^{-2r/D_0} \asymp \frac{D_1}{n}.$$

With this value of D_1, we get the following worst-case risk bound.

Theorem 53 (Sigmoidal activation and Sobolev spaces)**.** *Let $W_2^r(\mathcal{X},\mu)$ be the unit ball of the Sobolev space of regularity $r>0$ and $\widehat{f}_n$ be an aggregate of neural networks satisfying the PAC-Bayesian risk bound (6.3). If the conditions of Theorem 50 hold and*

$$D_1 = \Big(\frac{\beta D_0}{n}\Big)^{-\frac{D_0}{2r+D_0}}$$

then the following risk bound is valid:

$$\sup_{\mathcal{P}:f_{\mathcal{P}}\in W_2^r(\mathcal{X},\mu)} \mathbf{E}_{\mathcal{P}}\big[\|\widehat{f}_n - f_{\mathcal{P}}\|^2_{\mathbb{L}_2(\mu)}\big] \le g(n)\Big(\frac{\beta D_0}{n}\Big)^{2r/(2r+D_0)}, \tag{6.18}$$

with the slowly varying function

$$g(n) = 2C_{\mathsf{PB}}\bigg(c_1^2 B_\varphi^2 \log^2(n/\beta) + 2\log\Big(3 + \frac{3nB_2^2(B_1^2M_2^2+\mu(\mathcal{X})M_\sigma^2)}{d\beta}\Big)\bigg).$$

This theorem follows from Corollary 48 and Theorem 50 by substituting D_1 by its expression in (6.17). Note that the rate, $n^{-2r/(2r+D_0)}$, is the classical minimax rate of estimation over classes of r-smooth functions. We further discuss this result and compare it to prior work in Section 6.6.3. In the next section, we derive the counterpart of the last risk bound in the case of ReLU activation function.

6.6.2 ReLU activation function

In this section, we propose a result for two classes of functions, the Barron spectral space and a specific Sobolev space. Let us first assume that $f_{\mathcal{P}} \in \mathscr{B}^s([0,1]^{D_0})$ and that the conditions of Theorem 52 are satisfied. In view of Proposition 49 and Theorem 52, we have the risk bound

$$\mathbf{E}_{\mathcal{P}}\big[\|\widehat{f}_n - f_{\mathcal{P}}\|^2_{\mathbb{L}_2(\mu)}\big] \le 2C_{\mathsf{PB}}C^2 D_1^{-2K(s,D_0)}\log^{2m}(D_1) + C_{\mathsf{PB}}\frac{4\beta D_1 D_0}{n}\log\Big(3+\frac{nE}{d\beta}\Big) \tag{6.19}$$

where C is as in Theorem 52, E is as in Proposition 49 and

$$K(s,D_0) = \begin{cases} 2 & \text{if } 2s \ge 3D_0+4 \\ \frac{2s+D_0}{2(D_0+1)} & \text{if } 2s < 3D_0+4 \end{cases} \quad \text{and} \quad m = \begin{cases} 0 & \text{if } 2s < 3D_0+4 \\ 1 & \text{if } 2s > 3D_0+4 \\ \frac{5}{2} & \text{if } 2s = 3D_0+4 \end{cases}.$$

The bias-variance balance equation advocates for choosing D_1 such that $D_1^{-2K} \asymp \beta D_1/n$, which leads to the following proposition.

Proposition 54 (ReLU activation and Barron spectral spaces)**.** *Let $\mathscr{B}^s([0,1]^{D_0})$ be the Barron spectral space and let $\widehat{f}_n$ be an aggregate of neural networks satisfying the PAC-Bayesian risk bound (6.3). If the conditions of Theorem 52 hold and, for $K = K(s, D_0)$,*

$$D_1 = \Big(\frac{\beta D_0}{n}\Big)^{-1/(2K+1)}$$

then the following risk bound is valid:

$$\sup_{\mathcal{P}:\|f_{\mathcal{P}}\|_{\mathscr{B}^s([0,1]^{D_0})}\le 1} \mathbf{E}_{\mathcal{P}}\big[\|\widehat{f}_n - f_{\mathcal{P}}\|^2_{\mathbb{L}_2(\mu)}\big] \le \bar{g}(n)\Big(\frac{\beta D_0}{n}\Big)^{2K/(2K+1)}, \tag{6.20}$$

with the slowly varying function

$$\bar{g}(n) = 2C_{\mathsf{PB}}C^2\log^{2m}(n/\beta) + 4C_{\mathsf{PB}}\log\left(3 + \frac{3nB_1^2B_2^2(M_2^2+\bar{M}_2^2)}{d\beta}\right)$$

and C is a constant that depends on s and D_0 but not on D_1.

The above result treats the case of Barron spectral space and, therefore, is not directly comparable with what we obtained in the case of a sigmoidal activation function. However, these two spaces are related. Indeed, according to (Xu, 2020, Lemma 2.5), the inclusion $W^{s+D_0/2+\varepsilon,2}(\mathcal{X}) \subset \mathscr{B}^s(\mathcal{X})$ holds true for arbitrarily small $\varepsilon > 0$. An equivalent way of writing this inclusion is $W^{r,2}(\mathcal{X}) \subset \mathscr{B}^{\bar{r}-D_0/2}(\mathcal{X})$ for every r, $\bar{r}$ such that $D_0/2 \le \bar{r} < r$. Depending on the order of the Barron spectral space and the dimension of the problem, this might require a significant level of smoothness for the function f we want to approximate. Keeping this constraint in mind, we proceed with the next proposition which is more easily comparable to Theorem 53.

Proposition 55 (ReLU activation and Sobolev space)**.** *Let $W_2^r([0,1]^{D_0})$ be the unit ball in the Sobolev space of regularity $r \in (\frac{D_0}{2}, 2D_0+2)$. If $\widehat{f}_n$ is an aggregate of neural networks with ReLU activation satisfying the PAC-Bayesian risk bound (6.3) then for every $\bar{r} < r$ there is a slowly varying function $g_{\bar{r}} : \mathbb{N} \to \mathbb{R}_+$ such that*

$$\sup_{\mathcal{P}:f_{\mathcal{P}}\in W_2^r([0,1]^{D_0})} \mathbf{E}_{\mathcal{P}}\big[\|\widehat{f}_n - f_{\mathcal{P}}\|^2_{\mathbb{L}_2(\mu)}\big] \le g_{\bar{r}}(n)n^{-\frac{2\bar{r}}{2\bar{r}+D_0+1}}. \tag{6.21}$$

A proof is available in appendix 6.8.5.

The present result is weaker than the one of Theorem 53 in three aspects. First, it has the constraint $r \in (\frac{D_0}{2}, 2D_0+2)$ limiting the order of smoothness of the Sobolev classes. Note that the constraint $r < 2D_0+2$ stems from the fact that we want $K(s, D_0)$ to take the value $(2s+D_0)/(2D_0+2)$. If $r \ge 2D_0+2$, the claim of the last proposition holds true if we replace $2\bar{r}/(2\bar{r}+D_0+1)$ by $4/5$. The second weakness of the result is that $\bar{r}$, present in the rate of convergence, is strictly smaller than the true smoothness r. Finally, the denominator in the exponent includes an additional term increasing the dimension D_0 by 1. This implies that the obtained rate of convergence is slightly slower than the minimax rate of estimation over Sobolev balls. This weakness is due to the absence, to our knowledge, of a sharp result on the

approximation properties of ReLU neural nets in Sobolev spaces.

We complete this section by comparing the obtained results with existing risk bounds for the empirical risk minimizer in the framework of both shallow and deep neural networks.

6.6.3 Related works for risk bound of neural networks

In the case of shallow neural networks with sigmoid activation functions, an early result from Barron (1994) shows that both the estimation error and approximation error are of order $O(n^{-1/2})$, up to a $\log n$ factor. McCaffrey and Gallant (1994) obtained the improved rate $O\left(n^{-2\bar{r}/(2\bar{r}+D_0+5)}\right)$, $\forall \bar{r} < r$, for specific cosine shape activation functions. To the best of our knowledge, for a single hidden layer, no better risk bound has been found yet for the empirical risk minimizer.

In the case of multiple hidden layers, still with sigmoid activation functions, Bauer and Kohler (2019) recently achieved a rate of $O(n^{-2r/(2r+D_0)})$ up to a cubic $\log n$ factor when using two hidden layers for functions satisfying a generalized hierarchical interaction model.

Our risk bound (6.18) is of order $O(n^{-2r/(2r+D_0)})$ up to $\log n$ factor. Interestingly, this matches with the nonparametric minimax rate of convergence (Stone, 1982; Tsybakov, 2008) and is better than the rates obtained for the empirical risk minimizer for networks with one hidden layer. This can be seen as an additional example of the fact that the aggregation might often outperform the selection. Roughly speaking, aggregation acts as an additional layer, so that the aggregated one-hidden-layer neural networks achieve the same rate of convergence as the two-hidden-layer neural networks minimizing the training error.

We switch now to neural networks with ReLU activation functions, for which the literature on risk bounds is scarce. In the case of neural networks with a single hidden layer, Bach (2017) established a risk bound of order $O(n^{-2/(D_0+3)})$, up to a logarithmic factor. Another important result, proved by Klusowski and Barron (2016b), concerns bounded ramp activation functions and the penalized empirical risk minimizer. In the special case, when D_0 is small compared to the sample size n, and functions belonging to the space $\mathscr{B}^2([0,1]^{D_0})$, they obtained a risk bound of order $O\left(n^{-(D_0+4)/(2D_0+6)}\right)$. This result can be directly compared to ours, in the particular case where $s = 2$. In this case, since

$$\frac{2K}{2K+1} = \frac{2s+D_0}{2s+2D_0+1} = \frac{D_0+4}{2D_0+5}$$

from Proposition 54, we infer a leading term in n of order $O\left(n^{-(D_0+4)/(2D_0+5)}\right)$, which tightens the result of Klusowski and Barron (2016b) by a factor $O(n^{-(D_0+4)/2(3+D_0)(2D_0+5)})$. For instance, if $D_0 = 3$ or $D_0 = 4$, we get the improvement factors $n^{-7/132}$ and $n^{-4/91}$, respectively. Note that this improvement becomes smaller when the dimension D_0 increases.

As mentioned earlier, these last years a special attention has been paid to deep neural networks and ReLU activation functions. For deep networks, Schmidt-Hieber (2020) established the counterpart of the risk bound of Bauer and Kohler (2019) in the case of ReLU activation function and the smoothness class of β-Hölder functions. In particular, the worst-case risk over

the Hölder class is shown to be of order $O(n^{-2\beta/(2\beta+D_0)})$, and an analogous result holds also for balls in Besov spaces (Suzuki, 2018). These results show that the minimax rate of estimation can be achieved by the ERM over deep ReLU networks, unlike the ERM over shallow networks. When compared with Proposition 55, this implies to some extent that deep neural networks minimizing the training error perform better by a factor $O(n^{-2\bar{r}/(2\bar{r}+D_0+1)(2\bar{r}+D_0)})$, in terms of the worst case risk, than the aggregate neural networks with a single hidden layer.

6.7 Conclusion and outlook

We have analyzed the estimation error of an aggregate of neural networks having one hidden layer and Lipschitz continuous activation function, under the condition that the aggregate satisfies the PAC-Bayesian inequality. We focused our attention to Gaussian priors and obtained risk bounds in which the dependence on all the involved parameters is explicit. All our bounds on the estimation error come with explicit constants.

We then combined our bounds on the estimation error with bounds on approximation error available in the literature. This allowed us to prove that aggregation of single-layer neural networks achieves the minimax risk over conventional smoothness classes. Note that these results provide risk bounds in which the constants are not explicit. This shortcoming is inherited from the bounds on the approximation error available in the literature. In our opinion, it would be highly relevant to refine the existing approximation bounds to make appear all the constants.

The results of the present work can be extended in different directions. First, it would be interesting to consider the problem of aggregation of deep neural networks in order to understand possible benefits of increasing the depth. Second, it might be relevant to analyze the case of a prior with heavier tails, such as the Laplace prior or the Student prior, with a hope to cover the case of high dimension $D_0 > n$ under some kind of sparsity assumption. Finally, another avenue of future research is to explore the computational benefits of considering aggregated neural networks in conjunction with the Langevin-type algorithms.

6.8 Appendix

As a preliminary remark let us note that, as a mixing measure, we expect the distribution p to aggregate the predictors $f_{\boldsymbol{w}}$ so that the resulting estimator is almost as good as the best predictors in $\mathcal{F}_{\mathsf{W}}$. A direct consequence of it is that "a good choice" of p should be centered in $\bar{\boldsymbol{w}}$. This is an heuristic way to choose the mean, and all along the appendix we will fix the distribution of p as

$$p = p_1 \otimes p_2 \sim \mathcal{N}(\bar{\boldsymbol{w}}_1, \tau_1^2 \mathbf{I}_{D_1 D_0}) \otimes \mathcal{N}(\bar{\boldsymbol{w}}_2, \tau_2^2 \mathbf{I}_{D_1 D_2}), \qquad \tau_1, \tau_2 > 0 \tag{6.22}$$

where $\bar{\boldsymbol{w}} \in \operatorname{argmin}_{\boldsymbol{w}\in\mathsf{W}} \|f_{\boldsymbol{w}} - f_{\mathcal{P}}\|_{\mathbb{L}_2(\mu)}$. The additional condition (6.22) is the starting point of our choice for p, it is now left to set values for the variance (τ_1^2, τ_2^2).

6.8.1 Some useful lemmas

In what follows, when appropriate, we will write $f_{\boldsymbol{w}_1,\boldsymbol{w}_2}$ instead of $f_{\boldsymbol{w}}$.

Lemma 56. *If the probability distribution p is such that $p(d\boldsymbol{w}) = p_1(d\boldsymbol{w}_1)p_2(d\boldsymbol{w}_2)$ with*

$$\int_{\mathsf{W}_2} \boldsymbol{w}_2 p_2(d\boldsymbol{w}_2) = \bar{\boldsymbol{w}}_2$$

then

$$\int_{\mathsf{W}} \|f_{\boldsymbol{w}} - f_{\bar{\boldsymbol{w}}}\|^2_{\mathbb{L}_2(\mu)}\, p(d\boldsymbol{w}) = \int_{\mathsf{W}} \|f_{\boldsymbol{w}} - f_{\boldsymbol{w}_1,\bar{\boldsymbol{w}}_2}\|^2_{\mathbb{L}_2(\mu)}\, p(d\boldsymbol{w}) + \int_{\mathsf{W}_1} \|f_{\boldsymbol{w}_1,\bar{\boldsymbol{w}}_2} - f_{\bar{\boldsymbol{w}}}\|^2_{\mathbb{L}_2}\, p_1(d\boldsymbol{w}_1).$$

Proof. Simple algebra yields

$$\begin{aligned}\int_{\mathsf{W}} \big(f_{\boldsymbol{w}} - f_{\bar{\boldsymbol{w}}}\big)^2(\boldsymbol{x})\, p(d\boldsymbol{w}) &= \int_{\mathsf{W}} (f_{\boldsymbol{w}} - f_{\boldsymbol{w}_1,\bar{\boldsymbol{w}}_2} + f_{\boldsymbol{w}_1,\bar{\boldsymbol{w}}_2} - f_{\bar{\boldsymbol{w}}})^2(\boldsymbol{x})\, p(d\boldsymbol{w}) \\ &= \int_{\mathsf{W}} (f_{\boldsymbol{w}} - f_{\boldsymbol{w}_1,\bar{\boldsymbol{w}}_2})^2(\boldsymbol{x})\, p(d\boldsymbol{w}) + \int_{\mathsf{W}} (f_{\boldsymbol{w}_1,\bar{\boldsymbol{w}}_2} - f_{\bar{\boldsymbol{w}}})^2(\boldsymbol{x})\, p(d\boldsymbol{w}) \\ &\quad + 2\underbrace{\int_{\mathsf{W}} (f_{\boldsymbol{w}_1,\boldsymbol{w}_2} - f_{\boldsymbol{w}_1,\bar{\boldsymbol{w}}_2})(\boldsymbol{x})(f_{\boldsymbol{w}_1,\bar{\boldsymbol{w}}_2} - f_{\bar{\boldsymbol{w}}_1,\bar{\boldsymbol{w}}_2})(\boldsymbol{x})\, p(d\boldsymbol{w})}_{\triangleq A}.\end{aligned}$$

To complete the proof it suffices to integrate the previous equality with respect to $\mu(d\boldsymbol{x})$ in virtue of Fubini-Tonelli theorem and to check that $A = 0$. The latter property follows from the fact that p is a product measure and, for all $\boldsymbol{w}_1 \in \mathsf{W}_1$,

$$\int_{\mathsf{W}_2} (f_{\boldsymbol{w}_1,\boldsymbol{w}_2} - f_{\boldsymbol{w}_1,\bar{\boldsymbol{w}}_2})(\boldsymbol{x})\, p_2(d\boldsymbol{w}_2) = \int_{\mathsf{W}_2} (\boldsymbol{w}_2 - \bar{\boldsymbol{w}}_2)^\top \bar{\sigma}(\boldsymbol{w}_1^\top \boldsymbol{x})\, p_2(d\boldsymbol{w}_2) = 0.$$

This yields the claim of the lemma. □

In this section and the next one, let us define the two quantities:

$$G_1(\boldsymbol{w}) = \|f_{\boldsymbol{w}} - f_{\boldsymbol{w}_1,\bar{\boldsymbol{w}}_2}\|^2_{\mathbb{L}_2(\mu)}, \qquad \text{and} \qquad G_2(\boldsymbol{w}_1) = \|f_{\boldsymbol{w}_1,\bar{\boldsymbol{w}}_2} - f_{\bar{\boldsymbol{w}}}\|^2_{\mathbb{L}_2(\mu)}.$$

Lemma 57. *If Assumptions (σ-L) and $M_2 < \infty$ are satisfied, and p is chosen as in (6.22), then*

$$\int_{\mathsf{W}_1} G_2(\boldsymbol{w}_1)\, p_1(d\boldsymbol{w}_1) \le D_0\big(M_2\tau_1\|\bar{\boldsymbol{w}}_2\|_{1,2}\big)^2 \le C_1 D_0 D_1 \tau_1^2 \tag{6.23}$$

with $C_1 = (M_2\|\bar{\boldsymbol{w}}_2\|_{\mathsf{F}})^2$.

Proof of Lemma 57. We first use the fact that σ is 1-Lipschitz. On the one hand, in conjunction

with the Fubini-Tonelli theorem, this yields

$$\begin{aligned}
\int_{\mathsf{W}_1} G_2(\boldsymbol{w}_1)\, p_1(d\boldsymbol{w}_1) &= \int_{\mathcal{X}} \int_{\mathsf{W}} \left\| \bar{\boldsymbol{w}}_2^\top \left\{ \bar{\sigma}(\boldsymbol{w}_1^\top \boldsymbol{x}) - \bar{\sigma}(\bar{\boldsymbol{w}}_1^\top \boldsymbol{x}) \right\} \right\|_2^2 p(d\boldsymbol{w}) \mu(d\boldsymbol{x}) \\
&\leq \int_{\mathcal{X}} \int_{\mathsf{W}} \sum_{j=1}^{D_2} \left(\sum_{i=1}^{D_1} |\bar{\boldsymbol{w}}_{2,ij}| |(\boldsymbol{w}_1 - \bar{\boldsymbol{w}}_1)_i^\top \boldsymbol{x}| \right)^2 p(d\boldsymbol{w}) \mu(d\boldsymbol{x}) \\
&\leq \int_{\mathcal{X}} \sum_{j=1}^{D_2} \left(\sum_{i=1}^{D_1} |\bar{\boldsymbol{w}}_{2,ij}| \left\{ \int_{\mathsf{W}} |(\boldsymbol{w}_1 - \bar{\boldsymbol{w}}_1)_i^\top \boldsymbol{x}|^2 p(d\boldsymbol{w}) \right\}^{1/2} \right)^2 \mu(d\boldsymbol{x}) \\
&= D_0 M_2^2 \tau_1^2 \sum_{j=1}^{D_2} \left(\sum_{i=1}^{D_1} |\bar{\boldsymbol{w}}_{2,ij}| \right)^2 \leq D_0 M_2^2 \tau_1^2 D_1 \|\bar{\boldsymbol{w}}_2\|_{\mathsf{F}}^2
\end{aligned}$$

and the claim of the lemma follows. □

In view of Lemma 56 and Lemma 57, we have

$$\int_{\mathsf{W}} \|f_{\boldsymbol{w}} - f_{\bar{\boldsymbol{w}}}\|_{\mathbb{L}_2(\mu)}^2 \, p(d\boldsymbol{w}) = \int_{\mathsf{W}} G_1(\boldsymbol{w})\, p(d\boldsymbol{w}) + \int_{\mathsf{W}} G_2(\boldsymbol{w}_1)\, p_1(d\boldsymbol{w}_1). \tag{6.24}$$

and

$$\int_{\mathsf{W}} G_2(\boldsymbol{w}_1)\, p(d\boldsymbol{w}) \leq D_1 (M_2 \|\bar{\boldsymbol{w}}_2\|_{\mathsf{F}})^2 \tau_1^2.$$

We now state two distinct lemmas to bound the quantity $\int_{\mathsf{W}} G_1(\boldsymbol{w})\, p(d\boldsymbol{w})$. Lemma 58 account for bounded activation functions whereas Lemma 59 focuses on unbounded ones.

Lemma 58. *Under Assumption (σ-B) and $M_2 < \infty$, if p is given by (6.22), we have*

$$\int_{\mathsf{W}} G_1(\boldsymbol{w})\, p(d\boldsymbol{w}) \leq (M_\sigma \tau_2)^2 \mu(\mathcal{X}) D_1 D_2.$$

Proof of Lemma 58. Using Fubini-Tonelli theorem, we get

$$\int_{\mathsf{W}} G_1(\boldsymbol{w})\, p(d\boldsymbol{w}) = \int_{\mathcal{X}} \int_{\mathsf{W}_1} \underbrace{\int_{\mathsf{W}_2} \left\| (\boldsymbol{w}_2 - \bar{\boldsymbol{w}}_2)^\top \bar{\sigma}(\boldsymbol{w}_1^\top \boldsymbol{x}) \right\|_{\mathbb{L}_2(\mu)}^2 p_2(d\boldsymbol{w}_2)}_{:= I(\boldsymbol{x}, \boldsymbol{w}_1)} \, p_1(d\boldsymbol{w}_1)\, \mu(d\boldsymbol{x}).$$

For the inner integral, simple algebra yields

$$\begin{aligned}
I(\boldsymbol{x}, \boldsymbol{w}_1) &= \int_{\mathsf{W}} \bar{\sigma}(\boldsymbol{w}_1^\top \boldsymbol{x})^\top (\boldsymbol{w}_2 - \bar{\boldsymbol{w}}_2)(\boldsymbol{w}_2 - \bar{\boldsymbol{w}}_2)^\top \bar{\sigma}(\boldsymbol{w}_1^\top \boldsymbol{x})\, p(d\boldsymbol{w}) && (6.25) \\
&= \bar{\sigma}(\boldsymbol{w}_1^\top \boldsymbol{x})^\top \left\{ \int_{\mathsf{W}_2} (\boldsymbol{w}_2 - \bar{\boldsymbol{w}}_2)(\boldsymbol{w}_2 - \bar{\boldsymbol{w}}_2)^\top p_2(d\boldsymbol{w}_2) \right\} \bar{\sigma}(\boldsymbol{w}_1^\top \boldsymbol{x}) && (6.26) \\
&= \tau_2^2 D_2 \|\bar{\sigma}(\boldsymbol{w}_1^\top \boldsymbol{x})\|_2^2. && (6.27)
\end{aligned}$$

Therefore,

$$\int_{\mathcal{X}} \int_{\mathsf{W}_1} I(\boldsymbol{x}, \boldsymbol{w}_1)\, p_1(d\boldsymbol{w}_1)\, \mu(d\boldsymbol{x}) \leq M_\sigma^2 \mu(\mathcal{X}) \tau_2^2 D_1 D_2.$$

This completes the proof of the lemma. □

Lemma 59. *Let $\bar{M}_2 = \| \int_{\mathcal{X}} \boldsymbol{x}\boldsymbol{x}^\top \mu(d\boldsymbol{x}) \|_{\mathsf{sp}}$ be the spectral norm of the "covariance" matrix of the design. Under Assumption (σ-L), if p is given by (6.22) and $\sigma(0) = 0$, we have*

$$\int_{\mathsf{W}} G_1(\boldsymbol{w})\, p(d\boldsymbol{w}) \leq M_2^2 D_0 D_1 D_2 \tau_1^2 \tau_2^2 + \bar{M}_2^2 D_2 \|\bar{\boldsymbol{w}}_1\|_{\mathsf{F}}^2 \tau_2^2. \tag{6.28}$$

Proof of Lemma 59. Using (6.27), we get

$$\begin{aligned}
\int_{\mathsf{W}} G_1(\boldsymbol{w})\, p(d\boldsymbol{w}) &= \tau_2^2 D_2 \int_{\mathcal{X}} \int_{\mathsf{W}_1} \|\bar{\sigma}(\boldsymbol{w}_1^\top \boldsymbol{x})\|_2^2\, p_1(d\boldsymbol{w}_1)\, \mu(d\boldsymbol{x}) \\
&\leq \tau_2^2 D_2 \int_{\mathcal{X}} \int_{\mathsf{W}_1} \|\boldsymbol{w}_1^\top \boldsymbol{x}\|_2^2\, p_1(d\boldsymbol{w}_1)\, \mu(d\boldsymbol{x}) \\
&= \tau_2^2 D_2 \int_{\mathcal{X}} \int_{\mathsf{W}_1} \|(\boldsymbol{w}_1 - \bar{\boldsymbol{w}}_1)^\top \boldsymbol{x}\|_2^2\, p_1(d\boldsymbol{w}_1)\, \mu(d\boldsymbol{x}) + \tau_2^2 D_2 \int_{\mathcal{X}} \|(\bar{\boldsymbol{w}}_1)^\top \boldsymbol{x}\|_2^2\, \mu(d\boldsymbol{x}) \\
&= M_2^2 D_0 D_1 D_2 \tau_1^2 \tau_2^2 + \bar{M}_2^2 D_2 \|\bar{\boldsymbol{w}}_1\|_{\mathsf{F}}^2 \tau_2^2.
\end{aligned}$$

This completes the proof of the lemma. □

Lemma 60. *Under Assumption (σ-L) and $M_2 < \infty$, if p is given by (6.22), then*

$$\int_{\mathsf{W}} \|f_{\boldsymbol{w}} - f_{\bar{\boldsymbol{w}}}\|_{\mathbb{L}_2(\mu)}^2\, p(d\boldsymbol{w}) \leq M_2^2 \|\bar{\boldsymbol{w}}_2\|_{\mathsf{F}}^2 D_0 D_1 \tau_1^2 + \bar{M}_2^2 D_2 \|\bar{\boldsymbol{w}}_1\|_{\mathsf{F}}^2 \tau_2^2 + M_2^2 D_0 D_1 D_2 \tau_1^2 \tau_2^2. \tag{6.29}$$

If, in addition, Assumption (σ-B) is satisfied, then

$$\int_{\mathsf{W}} \|f_{\boldsymbol{w}} - f_{\bar{\boldsymbol{w}}}\|_{\mathbb{L}_2(\mu)}^2\, p(d\boldsymbol{w}) \leq M_2^2 \|\bar{\boldsymbol{w}}_2\|_{\mathsf{F}}^2 D_0 D_1 \tau_1^2 + \mu(\mathcal{X}) M_\sigma^2 D_1 D_2 \tau_2^2. \tag{6.30}$$

Proof. In Lemma 56 we have checked that

$$\int_{\mathsf{W}} \|f_{\boldsymbol{w}} - f_{\bar{\boldsymbol{w}}}\|_{\mathbb{L}_2(\mu)}^2\, p(d\boldsymbol{w}) = \int_{\mathsf{W}} G_1(\boldsymbol{w})\, p(d\boldsymbol{w}) + \int_{\mathsf{W}} G_2(\boldsymbol{w}_1)\, p_1(d\boldsymbol{w}_1).$$

Lemma 58 and Lemma 57 take care of both integrals in the right hand side of the equality for bounded activation functions and we directly get (6.30). Similarly, Lemma 59 and Lemma 57 can be applied for unbounded activation functions, leading to (6.29). □

6.8.2 Proof of Proposition 47

Recall that the goal is to find an upper bound for the remainder term

$$\mathrm{Rem}_n(\bar{\boldsymbol{w}}) \triangleq \inf_{p\in\mathcal{P}_1(\mathcal{F}_{\mathsf{W}})} \left\{ \int_{\mathsf{W}} \|f_{\boldsymbol{w}} - f_{\bar{\boldsymbol{w}}}\|^2_{\mathbb{L}_2(\mu)}\, p(d\boldsymbol{w}) + \frac{\beta}{n} D_{\mathsf{KL}}(p||\pi) \right\}. \tag{6.31}$$

We start this proof by considering the case where Assumptions (σ-L), (σ-B) and ($\mathcal{N}$) are satisfied. We choose as p the product of two spherical Gaussian distributions with variances τ_1^2 and τ_2^2, as specified in (6.22). In this case, the Kullback-Leibler divergence $D_{\mathsf{KL}}(p||\pi)$ is given by

$$D_{\mathsf{KL}}(p||\pi) = \frac{1}{2}\sum_{\ell=1}^{2} \left\{ \frac{\|\bar{\boldsymbol{w}}_\ell\|^2_{\mathsf{F}}}{\rho_\ell^2} + D_{\ell-1}D_\ell \left[\left(\frac{\tau_\ell}{\rho_\ell}\right)^2 - 1 - \log\left(\frac{\tau_\ell^2}{\rho_\ell^2}\right) \right] \right\}. \tag{6.32}$$

It is now left to find good values for τ_1^2 and τ_2^2. Combining with the result (6.30) of Lemma 60, we get the inequality

$$\mathrm{Rem}_n(\bar{\boldsymbol{w}}) \le \frac{\beta\|\bar{\boldsymbol{w}}_1\|^2_{\mathsf{F}}}{2n\rho_1^2} + \frac{\beta\|\bar{\boldsymbol{w}}_2\|^2_{\mathsf{F}}}{2n\rho_2^2} + \frac{\beta}{2n}\sum_{\ell=1}^{2} D_{\ell-1}D_\ell \left\{ C_\ell \left(\frac{\tau_\ell}{\rho_\ell}\right)^2 - 1 - \log\left(\frac{\tau_\ell^2}{\rho_\ell^2}\right) \right\} \tag{6.33}$$

where

$$C_1 = \frac{2nM_2^2\|\bar{\boldsymbol{w}}_2\|^2_{\mathsf{F}}\rho_1^2}{\beta} + 1, \quad C_2 = \frac{2n\mu(\mathcal{X})M_\sigma^2\rho_2^2}{\beta} + 1.$$

One can easily check that the minimum of the function $u \mapsto Cu - 1 - \log u$ is attained at $u_{\min} = 1/C$ and the value at this point is $\log C$. This implies that

$$\begin{aligned}
\mathrm{Rem}_n(\bar{\boldsymbol{w}}) &\le \frac{\beta\|\bar{\boldsymbol{w}}_1\|^2_{\mathsf{F}}}{2n\rho_1^2} + \frac{\beta\|\bar{\boldsymbol{w}}_2\|^2_{\mathsf{F}}}{2n\rho_2^2} + \frac{\beta}{2n}\sum_{\ell=1}^{2} D_{\ell-1}D_\ell \log C_\ell && (6.34)\\
&\overset{(1)}{\le} \frac{\beta\|\bar{\boldsymbol{w}}_1\|^2_{\mathsf{F}}}{2n\rho_1^2} + \frac{\beta\|\bar{\boldsymbol{w}}_2\|^2_{\mathsf{F}}}{2n\rho_2^2} + \frac{\beta d}{2n}\log\left(\frac{D_0D_1C_1 + D_1D_2C_2}{d}\right) && (6.35)\\
&\le \frac{\beta\|\bar{\boldsymbol{w}}_1\|^2_{\mathsf{F}}}{2n\rho_1^2} + \frac{\beta\|\bar{\boldsymbol{w}}_2\|^2_{\mathsf{F}}}{2n\rho_2^2} + \frac{\beta d}{2n}\log\left(1 + \frac{2n(D_0D_1M_2^2\|\bar{\boldsymbol{w}}_2\|^2_{\mathsf{F}}\rho_1^2 + D_1D_2\mu(\mathcal{X})M_\sigma^2\rho_2^2)}{\beta d}\right), && (6.36)
\end{aligned}$$

where in (1) we have used the concavity of the function $u \mapsto \log u$. This completes the proof of the first claim of Proposition 47.

In the case where Assumption (σ-B) is not fulfilled, but instead $\sigma(0) = 0$, we repeat the same scheme of proof as above by using (6.29) instead of (6.30). This leads to

$$\begin{aligned}
\mathrm{Rem}_n(\bar{\boldsymbol{w}}) \le \frac{\beta d}{2n}&\left\{ C_1' \left(\frac{\tau_1}{\rho_1}\right)^2 + C_2' \left(\frac{\tau_2}{\rho_2}\right)^2 + C_3' \left(\frac{\tau_1}{\rho_1}\right)^2\left(\frac{\tau_2}{\rho_2}\right)^2 - 2 - \log\left(\frac{\tau_1^2\tau_2^2}{\rho_1^2\rho_2^2}\right) \right\} && (6.37)\\
&+ \frac{\beta\|\bar{\boldsymbol{w}}_1\|^2_{\mathsf{F}}}{2n\rho_1^2} + \frac{\beta\|\bar{\boldsymbol{w}}_2\|^2_{\mathsf{F}}}{2n\rho_2^2}. && (6.38)
\end{aligned}$$

where

$$C_1' = \frac{2nD_0M_2^2\|\bar{\boldsymbol{w}}_2\|_{\mathsf{F}}^2\rho_1^2}{\beta(D_0+D_2)} + 1, \quad C_2' = \frac{2n\bar{M}_2^2\|\bar{\boldsymbol{w}}_1\|_{\mathsf{F}}^2 D_2\rho_2^2}{\beta(D_0+D_2)D_1} + 1, \quad C_3' = \frac{2nM_2^2\rho_1^2\rho_2^2 D_0 D_2}{\beta(D_0+D_2)}.$$

We choose τ_1 and τ_2 so that

$$\left(\frac{\tau_1}{\rho_1}\right)^2 = \frac{1}{C_1' + C_3'(\tau_2/\rho_2)^2}, \qquad \left(\frac{\tau_2}{\rho_2}\right)^2 = 1/C_2'.$$

With this choice of τ_1 and τ_2 in (6.38) and simple algebra, we get

$$\mathrm{Rem}_n(\bar{\boldsymbol{w}}) \leq \frac{\beta d}{2n}\log\left(C_1'C_2' + C_3'\right) + \frac{\beta\|\bar{\boldsymbol{w}}_1\|_{\mathsf{F}}^2}{2n\rho_1^2} + \frac{\beta\|\bar{\boldsymbol{w}}_2\|_{\mathsf{F}}^2}{2n\rho_2^2}. \tag{6.39}$$

To complete the proof, we use the following inequalities

$$\begin{aligned}\ln\big(C_1'C_2' + C_3'\big) &\leq \log\left(C_1'(C_2' + C_3')\right)\\ &= \log C_1' + \log\big(C_2' + C_3'\big)\\ &\leq 2\ln\big((C_1' + C_2' + C_3')/2\big),\end{aligned}$$

where the first inequality follows from the fact that $C_2' \geq 1$ whereas the last inequality is a consequence of the concavity of the logarithm.

Remark 61. *The distribution p is centered on the oracle choice $\bar{\boldsymbol{w}}$ for the weights of the neural network and we observe that the optimized variances (τ_1^2, τ_2^2) in the proof of Proposition 47 are of the form $\tau_l^2 = \rho_\ell^2/(1 + c_\ell n\rho_\ell^2)$, $\ell = 1, 2$, for some positive constants c_1, c_2. These values of τ_ℓ arbitrate between the prior beliefs and the information brought by data. Indeed,* (1) *when no training data is available the uncertainty around $\bar{\boldsymbol{w}}$ corresponds to the prior uncertainty (ρ_1^2, ρ_2^2),* (2) *when the amount of observations n is unlimited and goes to infinity the uncertainty around the oracle value converges to* 0 *and p becomes close to the Dirac mass in $\bar{\boldsymbol{w}}$.*

6.8.3 Proof of Corollary 48

In view of (6.34), we have

$$\mathrm{Rem}_n(\bar{\boldsymbol{w}}) \leq \frac{\beta\|\bar{\boldsymbol{w}}_1\|_{\mathsf{F}}^2}{2n\rho_1^2} + \frac{\beta\|\bar{\boldsymbol{w}}_2\|_{\mathsf{F}}^2}{2n\rho_2^2} + \frac{\beta}{2n}\sum_{\ell=1}^{2} D_{\ell-1}D_\ell\log\big(1 + F_\ell\rho_\ell^2\big)$$

with

$$F_1 = \frac{2nM_2^2\|\bar{\boldsymbol{w}}_2\|_{\mathsf{F}}^2}{\beta} \qquad \text{and} \qquad F_2 = \frac{2n\mu(\mathcal{X})M_\sigma^2}{\beta}.$$

Taking the maximum over all $\bar{\boldsymbol{w}}$ such that the Frobenius norms of $\bar{\boldsymbol{w}}_1$ and $\bar{\boldsymbol{w}}_2$ are bounded by B_1 and B_2, we get

$$\sup_{\|\bar{\boldsymbol{w}}_1\|_{\mathsf{F}} \le B_1} \sup_{\|\bar{\boldsymbol{w}}_2\|_{\mathsf{F}} \le B_2} \mathrm{Rem}_n(\bar{\boldsymbol{w}}) \le \frac{\beta B_1^2}{2n\rho_1^2} + \frac{\beta B_2^2}{2n\rho_2^2} + \frac{\beta}{2n} \sum_{\ell=1}^{2} D_{\ell-1} D_\ell \log\big(1 + \bar{F}_\ell \rho_\ell^2\big) \tag{6.40}$$

with

$$\bar{F}_1 = \frac{2nM_2^2 B_2^2}{\beta} \qquad \text{and} \qquad \bar{F}_2 = \frac{2n\mu(\mathcal{X}) M_\sigma^2}{\beta}.$$

The first order necessary condition for optimizing the right hand side with respect to ρ_1^2 and ρ_2^2 reads as

$$-\frac{B_\ell^2}{\rho_\ell^4} + \frac{D_{\ell-1} D_\ell \bar{F}_\ell}{1 + \bar{F}_\ell \rho_\ell^2} = 0 \quad \Longleftrightarrow \quad \rho_\ell^4 - \frac{B_\ell^2}{D_{\ell-1} D_\ell} \rho_\ell^2 - \frac{B_\ell^2}{D_{\ell-1} D_\ell \bar{F}_\ell} = 0. \tag{6.41}$$

This second-order equation has only one positive root given by

$$\begin{aligned} \rho_\ell^2 &= \frac{B_\ell^2}{2D_{\ell-1} D_\ell} + \left(\frac{B_\ell^4}{4D_{\ell-1}^2 D_\ell^2} + \frac{B_\ell^2}{D_{\ell-1} D_\ell \bar{F}_\ell} \right)^{1/2} \\ &= \frac{B_\ell^2}{2D_{\ell-1} D_\ell} \left\{ 1 + \left(1 + \frac{4D_{\ell-1} D_\ell}{B_\ell^2 \bar{F}_\ell} \right)^{1/2} \right\}. \end{aligned}$$

We simplify computations by choosing

$$\rho_\ell^2 = \frac{B_\ell^2}{2D_{\ell-1} D_\ell}.$$

Replacing these values of ρ_ℓ^2 in (6.40), we get

$$\sup_{\|\bar{\boldsymbol{w}}_\ell\|_{\mathsf{F}} \le B_\ell} \mathrm{Rem}_n(\bar{\boldsymbol{w}}) \le \frac{\beta}{n} \sum_{\ell=1}^{2} D_{\ell-1} D_\ell \left\{ 1 + \frac{1}{2} \log \left(1 + \frac{B_\ell^2 \bar{F}_\ell}{2D_{\ell-1} D_\ell} \right) \right\} \tag{6.42}$$

$$\le \frac{\beta d}{n} \sum_{\ell=1}^{2} \left\{ 1 + \frac{1}{2} \log \left(1 + \frac{B_1^2 \bar{F}_1 + B_2^2 \bar{F}_2}{2d} \right) \right\}, \tag{6.43}$$

where the last inequality follows from the concavity of the logarithm. Replacing $\bar{F}_1$ and $\bar{F}_2$ with their respective expressions, we get the inequality

$$\begin{aligned} \sup_{\|\bar{\boldsymbol{w}}_\ell\|_{\mathsf{F}} \le B_\ell} \mathrm{Rem}_n(\bar{\boldsymbol{w}}) &\le \frac{\beta d}{n} \left(1 + \frac{1}{2} \log \left(1 + \frac{nB_2^2 (B_1^2 M_2^2 + \mu(\mathcal{X}) M_\sigma^2)}{d\beta} \right) \right) \\ &\le \frac{\beta d}{n} \log \left(3 + \frac{3nB_2^2 (B_1^2 M_2^2 + \mu(\mathcal{X}) M_\sigma^2)}{d\beta} \right), \end{aligned}$$

which coincides with the first claim of the corollary.

The second claim of the proposition is obtained by replacing ρ_ℓ's by their respective expres-

sions in the second claim of Proposition 47.

6.8.4 Proof of Lemma 51

Since

$$\sigma(x) = \sum_{j=1}^{\infty} \varphi(x-j). \tag{6.44}$$

we have

$$\begin{aligned}
\sigma(x+1) - \sigma(x) &= \sum_{j=1}^{\infty} \varphi(x+1-j) - \sum_{j=1}^{\infty} \varphi(x-j) \\
&= \sum_{j=0}^{\infty} \varphi(x-j) - \sum_{j=1}^{\infty} \varphi(x-j) \\
&= \varphi(x).
\end{aligned}$$

Now, recall that we use the function $\varphi(x) = \frac{1}{\sqrt{2}} e^{-x^2/2}$. It is clear, that the series

$$\sum_{j=1}^{\infty} |\varphi'(j-x)|$$

converges uniformly on any bounded interval. This implies that σ is differentiable and

$$\sigma'(x) = \sum_{j=1}^{\infty} \frac{(j-x)}{\sqrt{2}} e^{-(x-j)^2/2}.$$

Let us denote by $[x]$ the integer part of x and by $\{x\} = x - [x]$ the fractional part of x. Recall also that the function $u \mapsto e^{-u^2/2}$ is increasing on $(-\infty, 0]$ and decreasing on $[0, +\infty)$. Therefore, we have

$$\begin{aligned}
\sigma(x) &= \frac{1}{\sqrt{2}} \sum_{j=1}^{[x]} e^{-(x-j)^2/2} + \frac{1}{\sqrt{2}} \sum_{j=[x]+1}^{\infty} e^{-(x-j)^2/2} \\
&\leq \frac{1}{\sqrt{2}} \sum_{j=1}^{[x]-1} \int_{x-j-1}^{x-j} e^{-u^2/2}\, du + \frac{e^{-\{x\}^2/2} + e^{-(1-\{x\})^2/2}}{\sqrt{2}} + \frac{1}{\sqrt{2}} \sum_{j=[x]+2}^{\infty} \int_{j-x-1}^{j-x} e^{-u^2/2}\, du \\
&\leq \frac{1}{\sqrt{2}} \int_{\{x\}}^{x-1} e^{-u^2/2}\, du + \frac{e^{-\{x\}^2/2} + e^{-(1-\{x\})^2/2}}{\sqrt{2}} + \frac{1}{\sqrt{2}} \int_{-\infty}^{\{x\}-1} e^{-u^2/2}\, du \\
&\leq \frac{1}{\sqrt{2}} \int_{-\infty}^{+\infty} e^{-u^2/2}\, du + \frac{1}{\sqrt{2}} = \sqrt{\pi} + \frac{1}{\sqrt{2}}.
\end{aligned}$$

For $x > 0$, using similar arguments and the fact that the function $u \mapsto ue^{-u^2/2}$ is decreasing on $[1, \infty)$, we get

$$\begin{aligned}
\sqrt{2}\sigma'(x) &= -\sum_{j=1}^{[x]}(x-j)e^{-(x-j)^2/2} + \sum_{j=[x]+1}^{\infty}(j-x)e^{-(j-x)^2/2} \\
&\leq (1-\{x\})e^{-(1-\{x\})^2/2} + \sum_{j=[x]+2}^{\infty}\int_{j-x-1}^{j-x} ue^{-u^2/2}\,du \\
&\leq (1-\{x\})e^{-(1-\{x\})^2/2} + \int_{1-\{x\}}^{\infty} ue^{-u^2/2}\,du \\
&= (2-\{x\})e^{-(1-\{x\})^2/2} \leq \sqrt{2}.
\end{aligned}$$

In the same way, one can check that $\sqrt{2}\sigma'(x) \geq -\sqrt{2}$ for every $x > 0$. Therefore, $|\sigma'(x)| \leq 1$ for every positive x. On the other hand, for $x \leq 0$, we have $\sigma'(x) \geq 0$ and

$$\begin{aligned}
\sigma'(x) \leq \sigma'(0) &= \frac{1}{\sqrt{2}}\sum_{j=1}^{\infty} je^{-j^2/2} \\
&\leq \frac{1}{\sqrt{2}}\left(e^{-1/2} + \sum_{j=2}^{\infty}\int_{j-1}^{j} ue^{-u^2/2}\,du\right) \\
&= \frac{1}{\sqrt{2}}\left(e^{-1/2} + e^{-1/2}\right) \leq 1.
\end{aligned}$$

This completes the proof of the fact that σ is 1-Lipschitz.

6.8.5 Proof of Proposition 55

Assume $r \in (\frac{D_0}{2}, 2D_0 + 2)$ and $\bar{r} \in [D_0/2, r)$, then, $W^{r,2}(\mathcal{X}) \subset \mathscr{B}^{\bar{r}-D_0/2}(\mathcal{X})$ (Xu, 2020, Lemma 2.5). Since $\frac{2K}{2K+1} = \frac{2s+D_0}{2s+2D_0+1}$, substituting s by $\bar{r} - D_0/2$, we obtain:

$$\frac{2K}{2K+1} = \frac{2\bar{r}}{2\bar{r}+D_0+1}.$$

Substituting the terms in Proposition 54, this yields the result with

$$\bar{g}_{\bar{r}}(n) = 2C_{\mathsf{PB}}C^2 + 4C_{\mathsf{PB}}\log\left(3 + \frac{3nB_1^2B_2^2(M_2^2+\bar{M}_2^2)}{d\beta}\right).$$

Bibliography

K. Aas and I. H. Haff. The generalized hyperbolic skew student'st-distribution. Journal of financial econometrics, 4(2):275–309, 2006.

R. Aid, P. Gruet, and H. Pham. An optimal trading problem in intraday electricity markets. Mathematics and Financial Economics, 10(1):49–85, 2016.

R. Aid, A. Cosso, and H. Pham. Equilibrium price in intraday electricity markets. arXiv preprint arXiv:2010.09285, 2020.

C. Alasseur, I. B. Taher, and A. Matoussi. An extended mean field game for storage in smart grids. Journal of Optimization Theory and Applications, 184(2):644–670, 2020.

R. Almgren and N. Chriss. Value under liquidation. Risk, 12(12):61–63, 1999.

R. Almgren and N. Chriss. Optimal execution of portfolio transactions. Journal of Risk, 3:5–40, 2001.

P. Alquier and K. Lounici. PAC-Bayesian bounds for sparse regression estimation with exponential weights. Electronic Journal of Statistics, 5:127 – 145, 2011. doi: 10.1214/11-EJS601.

F. Bach. Breaking the curse of dimensionality with convex neural networks. The Journal of Machine Learning Research, 18(1):629–681, 2017.

A. Bachouch, C. Huré, N. Langrené, and H. Pham. Deep neural networks algorithms for stochastic control problems on finite horizon: numerical applications. Methodology and Computing in Applied Probability, pages 1–36, 2021.

E. Bacry, A. Iuga, M. Lasnier, and C.-A. Lehalle. Market impacts and the life cycle of investors orders. Market Microstructure and Liquidity, 1(02):1550009, 2015.

J. Badosa, E. Gobet, M. Grangereau, and D. Kim. Day-ahead probabilistic forecast of solar irradiance: a stochastic differential equation approach. In Forecasting and Risk Management for Renewable Energy, pages 73–93. Springer, 2017.

P. Bank, H. M. Soner, and M. Voß. Hedging with temporary price impact. Mathematics and financial economics, 11(2):215–239, 2017.

P. Bank, I. Ekren, and J. Muhle-Karbe. Liquidity in competitive dealer markets. Mathematical Finance, 31, 04 2021. doi: 10.1111/mafi.12305.

S. Baran. Probabilistic wind speed forecasting using bayesian model averaging with truncated normal components. Computational Statistics & Data Analysis, 75:227–238, 2014.

S. Baran and S. Lerch. Log-normal distribution based ensemble model output statistics models for probabilistic wind-speed forecasting. Quarterly Journal of the Royal Meteorological Society, 141(691):2289–2299, 2015.

S. Baran and S. Lerch. Mixture emos model for calibrating ensemble forecasts of wind speed. Environmetrics, 27(2):116–130, 2016.

O. E. Barndorff-Nielsen. Processes of normal inverse gaussian type. Finance and stochastics, 2 (1):41–68, 1997.

A. R. Barron. Universal approximation bounds for superpositions of a sigmoidal function. IEEE Transactions on Information theory, 39(3):930–945, 1993.

A. R. Barron. Approximation and estimation bounds for artificial neural networks. Machine learning, 14(1):115–133, 1994.

B. Bauer and M. Kohler. On deep learning as a remedy for the curse of dimensionality in nonparametric regression. Annals of Statistics, 47(4):2261–2285, 2019.

S. Becker, P. Cheridito, and A. Jentzen. Deep optimal stopping. Journal of Machine Learning Research, 20:74, 2019.

S. Becker, P. Cheridito, A. Jentzen, and T. Welti. Solving high-dimensional optimal stopping problems using deep learning. European Journal of Applied Mathematics, 32(3):470–514, 2021.

D. Belomestny, A. Kolodko, and J. Schoenmakers. Regression methods for stochastic control problems and their convergence analysis. SIAM Journal on Control and Optimization, 48(5): 3562–3588, 2010.

A. Bensoussan and A. Brouste. Cox–ingersoll–ross model for wind speed modeling and forecasting. Wind Energy, 19(7):1355–1365, 2016.

A. Bensoussan and S. C. P. Yam. Mean field approach to stochastic control with partial information. arXiv preprint arXiv:1909.10287v3, 2019.

A. Bensoussan, M. Chau, and S. Yam. Mean field games with a dominating player. Applied Mathematics & Optimization, 74(1):91–128, 2016a.

A. Bensoussan, K. Sung, S. C. P. Yam, and S.-P. Yung. Linear-quadratic mean field games. Journal of Optimization Theory and Applications, 169(2):496–529, 2016b.

N. Bershova and D. Rakhlin. The non-linear market impact of large trades: evidence from buy-side order flow. Quantitative finance, 13(11):1759–1778, 2013.

C. Bishop. Neural networks and their applications. Review of Scientific Instruments, 65:1803–1832, 1994.

B. Bouchard and N. Touzi. Discrete-time approximation and monte-carlo simulation of backward stochastic differential equations. Stochastic Processes and their applications, 111(2):175–206, 2004.

B. Bouchard and X. Warin. Monte-carlo valuation of american options: facts and new algorithms to improve existing methods. In Numerical methods in finance, pages 215–255. Springer, 2012.

B. Bouchard, M. Fukasawa, M. Herdegen, and J. Muhle-Karbe. Equilibrium returns with transaction costs. Finance and Stochastics, 22(3):569–601, 2018.

J.-P. Bouchaud. Price impact. Encyclopedia of quantitative finance, 2010.

F. Bucci, I. Mastromatteo, Z. Eisler, F. Lillo, J.-P. Bouchaud, and C.-A. Lehalle. Co-impact: Crowding effects in institutional trading activity. Quantitative Finance, 20(2):193–205, 2020.

M. Burger and A. Neubauer. Error bounds for approximation with neural networks. Journal of Approximation Theory, 112(2):235–250, 2001.

F. Cao, T. Xie, and Z. Xu. The estimate for approximation error of neural networks: a constructive approach. Neurocomputing, 71(4-6):626–630, 2008.

P. Cardaliaguet, F. Delarue, J.-M. Lasry, and P.-L. Lions. The Master Equation and the Convergence Problem in Mean Field Games:(AMS-201), volume 201. Princeton University Press, 2019.

P. Cardaliaguet, M. Cirant, and A. Porretta. Remarks on Nash equilibria in mean field game models with a major player. Proceedings of the American Mathematical Society, 148(10): 4241–4255, 2020.

R. Carmona and F. Delarue. Probabilistic Theory of Mean Field Games with Applications I. Springer, 2018.

R. Carmona, X. Zhu, et al. A probabilistic approach to mean field games with major and minor players. The Annals of Applied Probability, 26(3):1535–1580, 2016.

P. Casgrain and S. Jaimungal. Mean field games with partial information for algorithmic trading. arXiv preprint arXiv:1803.04094, 2018.

P. Casgrain and S. Jaimungal. Mean-field games with differing beliefs for algorithmic trading. Mathematical Finance, 30(3):995–1034, 2020.

O. Catoni. Pac-bayesian supervised classification: The thermodynamics of statistical learning. Lecture Notes-Monograph Series, 56:i–163, 2007. ISSN 07492170.

J. Collet, O. Féron, and P. Tankov. Optimal management of a wind power plant with storage capacity. In Forecasting and Risk Management for Renewable Energy, pages 229–246. Springer, 2017.

D. Costarelli and R. Spigler. Approximation results for neural network operators activated by sigmoidal functions. Neural Networks, 44:101–106, 2013a.

D. Costarelli and R. Spigler. Multivariate neural network operators with sigmoidal activation functions. Neural Networks, 48:72–77, 2013b.

A. S. Dalalyan. Exponential weights in multivariate regression and a low-rankness favoring prior. Ann. Inst. Henri Poincaré Probab. Stat., 56(2):1465–1483, 2020.

A. S. Dalalyan and A. B. Tsybakov. Aggregation by exponential weighting and sharp oracle inequalities. In International Conference on Computational Learning Theory, pages 97–111. Springer, 2007.

A. S. Dalalyan and A. B. Tsybakov. Aggregation by exponential weighting, sharp pac-bayesian bounds and sparsity. Machine Learning, 72(1-2):39–61, 2008.

A. S. Dalalyan and A. B. Tsybakov. Mirror averaging with sparsity priors. Bernoulli, 18(3): 914–944, 2012a.

A. S. Dalalyan and A. B. Tsybakov. Sparse regression learning by aggregation and langevin monte-carlo. Journal of Computer and System Sciences, 78(5):1423–1443, 2012b.

S. Delikaraoglou, A. Papakonstantinou, C. Ordoudis, and P. Pinson. Price-maker wind power producer participating in a joint day-ahead and real-time market. In 12th International Conference on the European Energy Market (EEM), pages 1–5. IEEE, 2015.

B. Delyon, A. Juditsky, and A. Benveniste. Accuracy analysis for wavelet approximations. IEEE Transactions on Neural Networks, 6(2):332–348, 1995. doi: 10.1109/72.363469.

J. Donier, J. Bonart, I. Mastromatteo, and J.-P. Bouchaud. A fully consistent, minimal model for non-linear market impact. Quantitative finance, 15(7):1109–1121, 2015.

D. Dufresne. The distribution of a perpetuity, with applications to risk theory and pension funding. Scandinavian Actuarial Journal, 1990(1):39–79, 1990.

G. K. Dziugaite and D. M. Roy. Computing nonvacuous generalization bounds for deep (stochastic) neural networks with many more parameters than training data. UAI, 2017.

D. Evangelista and Y. Thamsten. Finite population games of optimal execution. arXiv preprint arXiv:2004.00790, 2020.

S. Fécamp, J. Mikael, and X. Warin. Risk management with machine-learning-based algorithms. arXiv preprint arXiv:1902.05287, 2019.

O. Féron, P. Tankov, and L. Tinsi. Price formation and optimal trading in intraday electricity markets with a major player. Risks, 8(4):133, 2020.

G. Fu and U. Horst. Mean-field leader-follower games with terminal state constraint. SIAM Journal on Control and Optimization, 58(4):2078–2113, 2020a.

G. Fu and U. Horst. Mean-field leader-follower games with terminal state constraint. SIAM Journal on Control and Optimization, 58(4):2078–2113, 2020b.

G. Fu, U. Horst, and X. Xia. Portfolio liquidation games with self-exciting order flow. arXiv preprint arXiv:2011.05589, 2020.

G. Fu, P. Graewe, U. Horst, and A. Popier. A mean field game of optimal portfolio liquidation. Mathematics of Operations Research, 2021.

M. Fujii and A. Takahashi. A mean field game approach to equilibrium pricing with market clearing condition. CARF Working Paper CARF-F-473, 2020.

E. Garnier and R. Madlener. Balancing forecast errors in continuous-trade intraday markets. Energy Systems, 6(3):361–388, 2015.

J. Gatheral. No-dynamic-arbitrage and market impact. Quantitative finance, 10(7):749–759, 2010.

S. Gerchinovitz. Sparsity regret bounds for individual sequences in online linear regression. J. Mach. Learn. Res., 14:729–769, 2013. ISSN 1532-4435.

M. Germain, H. Pham, and X. Warin. Deep backward multistep schemes for nonlinear pdes and approximation error analysis. arXiv preprint arXiv:2006.01496, 2020.

H. Gevret, N. Langren'e, J. Lelong, X. Warin, and A. Maheshwari. Stochastic optimization library in c++. hal-01361291v1, 2018.

S. Glas, R. Kiesel, S. Kolkmann, M. Kremer, N. G. von Luckner, L. Ostmeier, K. Urban, and C. Weber. Intraday renewable electricity trading: Advanced modeling and numerical optimal control. Journal of Mathematics in Industry, 10(1):3, 2020.

T. Gneiting, A. E. Raftery, A. H. Westveld III, and T. Goldman. Calibrated probabilistic forecasting using ensemble model output statistics and minimum crps estimation. Monthly Weather Review, 133(5):1098–1118, 2005.

T. Gneiting, F. Balabdaoui, and A. E. Raftery. Probabilistic forecasts, calibration and sharpness. Journal of the Royal Statistical Society: Series B (Statistical Methodology), 69(2):243–268, 2007.

I. Goodfellow, Y. Bengio, and A. Courville. Deep Learning. MIT Press, 2016.

B. Guedj. A primer on pac-bayesian learning. arXiv preprint arXiv:1901.05353, 2019.

I. Gühring, G. Kutyniok, and P. Petersen. Error bounds for approximations with deep relu neural networks in w s, p norms. Analysis and Applications, 18(05):803–859, 2020.

K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. Neural networks, 2(5):359–366, 1989.

M. Huang. Large-population lqg games involving a major player: the nash certainty equivalence principle. SIAM Journal on Control and Optimization, 48(5):3318–3353, 2010.

M. Huang, R. P. Malhamé, and P. E. Caines. Large population stochastic dynamic games: closed-loop McKean-Vlasov systems and the nash certainty equivalence principle. Communications in Information & Systems, 6(3):221–252, 2006.

G. Huberman and W. Stanzl. Price manipulation and quasi-arbitrage. Econometrica, 72(4): 1247–1275, 2004.

C. Huré, H. Pham, and X. Warin. Deep backward schemes for high-dimensional nonlinear pdes. Mathematics of Computation, 89(324):1547–1579, 2020.

C. Huré, H. Pham, A. Bachouch, and N. Langrené. Deep neural networks algorithms for stochastic control problems on finite horizon: convergence analysis. SIAM Journal on Numerical Analysis, 59(1):525–557, 2021.

E. B. Iversen, J. M. Morales, J. K. Møller, and H. Madsen. Short-term probabilistic forecasting of wind speed using stochastic differential equations. International Journal of Forecasting, 32 (3):981–990, 2016.

E. Jaeck and D. Lautier. Volatility in electricity derivative markets: The Samuelson effect revisited. Energy Economics, 59:300–313, 2016.

M. Jeanblanc, M. Yor, and M. Chesney. Mathematical methods for financial markets. Springer Science & Business Media, 2009.

T. Jónsson, P. Pinson, and H. Madsen. On the market impact of wind energy forecasts. Energy Economics, 32(2):313–320, 2010.

A. Juditsky, P. Rigollet, and A. B. Tsybakov. Learning by mirror averaging. The Annals of Statistics, 36(5):2183 – 2206, 2008. doi: 10.1214/07-AOS546.

F. Karanfil and Y. Li. The role of continuous intraday electricity markets: the integration of large-share wind power generation in Denmark. The Energy Journal, 38(2), 2017.

R. Kiesel and F. Paraschiv. Econometric analysis of 15-minute intraday electricity prices. Energy Economics, 64:77–90, 2017.

J. M. Klusowski and A. Barron. Uniform approximation by neural networks activated by first and second order ridge splines. arXiv preprint arXiv:1607.07819v1, 2016a.

J. M. Klusowski and A. R. Barron. Risk bounds for high-dimensional ridge function combinations including neural networks. arXiv preprint arXiv:1607.01434, 2016b.

D. Kristensen. Nonparametric filtering of the realized spot volatility: A kernel-based approach. Econometric Theory, 26(1):60–93, 2010.

D. Lacker. On the convergence of closed-loop Nash equilibria to the mean field game limit. Annals of Applied Probability, 30(4):1693–1761, 08 2020.

J.-M. Lasry and P.-L. Lions. Mean field games. Japanese journal of mathematics, 2(1):229–260, 2007.

J.-M. Lasry and P.-L. Lions. Mean-field games with a major player. Comptes Rendus Mathematique, 356(8):886–890, 2018.

G. Lecué. Optimal rates of aggregation in classification under low noise assumption. Bernoulli, 13(4):1000 – 1022, 2007. doi: 10.3150/07-BEJ6044.

S. Lerch and T. L. Thorarinsdottir. Comparison of non-homogeneous regression models for probabilistic wind speed forecasting. Tellus A: Dynamic Meteorology and Oceanography, 65 (1):21206, 2013.

G. Leung and A. R. Barron. Information theory and mixing least-squares regressions. IEEE Trans. Inform. Theory, 52(8):3396–3410, 2006.

G. Leung and A. R. Barron. Information theory and mixing least-squares regressions. IEEE Transactions on Information Theory, 52(8):3396–3410, 2006. doi: 10.1109/TIT.2006.878172.

G. Lever, F. Laviolette, and J. Shawe-Taylor. Tighter pac-bayes bounds through distribution-dependent priors. Theor. Comput. Sci., 473:4–28, Feb. 2013. ISSN 0304-3975. doi: 10.1016/j.tcs.2012.10.013.

F. A. Longstaff and E. S. Schwartz. Valuing american options by simulation: a simple least-squares approach. The review of financial studies, 14(1):113–147, 2001.

J. Lu, Z. Shen, H. Yang, and S. Zhang. Deep network approximation for smooth functions. arXiv preprint arXiv:2001.03040, 2020.

V. Maiorov. Approximation by neural networks and learning theory. Journal of Complexity, 22 (1):102–117, 2006.

V. Maiorov and R. Meir. On the near optimality of the stochastic approximation of smooth functions by neural networks. Advances in Computational Mathematics, 13(1):79–103, 2000.

Y. Makovoz. Random approximants and neural networks. Journal of Approximation Theory, 85(1):98–109, 1996. ISSN 0021-9045.

D. A. McAllester. Pac-bayesian model averaging. In Proceedings of the twelfth annual conference on Computational learning theory, pages 164–170, 1999a.

D. A. McAllester. Some pac-bayesian theorems. Machine Learning, 37(3):355–363, 1999b.

D. A. McAllester. Pac-bayesian stochastic model selection. Machine Learning, 51(1):5–21, 2003.

D. F. McCaffrey and A. R. Gallant. Convergence rates for single hidden layer feedforward networks. Neural Networks, 7(1):147–158, 1994.

H. N. Mhaskar and C. A. Micchelli. Dimension-independent bounds on the degree of approximation by neural networks. IBM Journal of Research and Development, 38(3):277–284, 1994.

J. M. Morales, A. J. Conejo, and J. Pérez-Ruiz. Short-term trading for a wind power producer. IEEE Transactions on Power Systems, 25(1):554–564, 2010.

J. F. Nash. Equilibrium points in n-person games. Proceedings of the national academy of sciences, 36(1):48–49, 1950.

A. S. Nemirovsky and D. B. Yudin. Problem complexity and method efficiency in optimization. Chichester, Wiley, 1983.

Y. Nesterov. Primal-dual subgradient methods for convex problems. Mathematical programming, 120(1):221–259, 2009.

B. Neyshabur, S. Bhojanapalli, D. Mcallester, and N. Srebro. Exploring generalization in deep learning. In Advances in Neural Information Processing Systems, volume 30. Curran Associates, Inc., 2017.

M. Nourian and P. E. Caines. ϵ-nash mean field game theory for nonlinear stochastic dynamical systems with major and minor agents. SIAM Journal on Control and Optimization, 51(4): 3302–3331, 2013.

P. P. Petrushev. Approximation by ridge functions and neural networks. SIAM Journal on Mathematical Analysis, 30(1):155–189, 1998.

P. Pinson, C. Chevallier, and G. N. Kariniotakis. Trading wind generation from short-term probabilistic forecasts of wind power. IEEE Transactions on Power Systems, 22(3):1148–1156, 2007.

P. Pinson et al. Wind energy: Forecasting challenges for its operational management. Statistical Science, 28(4):564–585, 2013.

A. E. Raftery, T. Gneiting, F. Balabdaoui, and M. Polakowski. Using bayesian model averaging to calibrate forecast ensembles. Monthly weather review, 133(5):1155–1174, 2005.

P. Rigollet and A. B. Tsybakov. Sparse Estimation by Exponential Weighting. Statistical Science, 27(4):558 – 575, 2012. doi: 10.1214/12-STS393.

O. Rivasplata, E. Parrado-Hernandez, J. S. Shawe-Taylor, S. Sun, and C. Szepesvari. Pac-bayes bounds for stable algorithms with instance-dependent priors. In Advances in Neural Information Processing Systems, volume 31. Curran Associates, Inc., 2018.

P. A. Rowińska, A. Veraart, and P. Gruet. A multifactor approach to modelling the impact of wind energy on electricity spot prices. Available at SSRN 3110554, 2018.

J. Schmidt-Hieber. Nonparametric regression using deep neural networks with relu activation function. Annals of Statistics, 48(4):1875–1897, 2020.

Z. Shen, H. Yang, and S. Zhang. Deep network approximation characterized by number of neurons. arXiv preprint arXiv:1906.05497, 2019.

A. Shrivats, D. Firoozi, and S. Jaimungal. A mean-field game approach to equilibrium pricing, optimal generation, and trading in solar renewable energy certificate (srec) markets. arXiv preprint arXiv:2003.04938, 2020.

J. W. Siegel and J. Xu. High-order approximation rates for neural networks with relu ^{k} activation functions. arXiv preprint arXiv:2012.07205, 2020.

A. Skajaa, K. Edlund, and J. M. Morales. Intraday trading of wind energy. IEEE Transactions on power systems, 30(6):3181–3189, 2015.

C. J. Stone. Optimal global rates of convergence for nonparametric regression. The annals of statistics, pages 1040–1053, 1982.

T. Suzuki. Adaptivity of deep relu network for learning in besov and mixed smooth besov spaces: optimal rate and curse of dimensionality. arXiv preprint arXiv:1810.08033, 2018.

Z. Tan and P. Tankov. Optimal trading policies for wind energy producer. SIAM Journal on Financial Mathematics, 9(1):315–346, 2018.

T. L. Thorarinsdottir and T. Gneiting. Probabilistic forecasts of wind speed: Ensemble model output statistics by using heteroscedastic censored regression. Journal of the Royal Statistical Society: Series A (Statistics in Society), 173(2):371–388, 2010.

N. Touzi. Optimal stochastic control, stochastic target problems, and backward SDE, volume 29. Springer Science & Business Media, 2012.

J. N. Tsitsiklis and B. Van Roy. Optimal stopping of markov processes: Hilbert space theory, approximation algorithms, and an application to pricing high-dimensional financial derivatives. IEEE Transactions on Automatic Control, 44(10):1840–1851, 1999.

A. B. Tsybakov. Introduction to nonparametric estimation. Springer Science & Business Media, 2008.

L. G. Valiant. A theory of the learnable. Commun. ACM, 27(11):1134–1142, Nov. 1984. ISSN 0001-0782. doi: 10.1145/1968.1972.

M. Voß. A two-player price impact game. arXiv preprint arXiv:1911.05122, 2019.

D. S. Wilks. Smoothing forecast ensembles with fitted probability distributions. Quarterly Journal of the Royal Meteorological Society: A journal of the atmospheric sciences, applied meteorology and physical oceanography, 128(586):2821–2836, 2002.

J. Xu. Finite neuron method and convergence analysis. Communications in Computational Physics, 28(5):1707–1745, 2020. ISSN 1991-7120. doi: https://doi.org/10.4208/cicp.OA-2020-0191.

D. Yarotsky. Error bounds for approximations with deep relu networks. Neural Networks, 94: 103–114, 2017.

D. Yarotsky. Optimal approximation of continuous functions by very deep relu networks. In Conference on Learning Theory, pages 639–649. PMLR, 2018.

D. Yarotsky and A. Zhevnerchuk. The phase diagram of approximation rates for deep neural networks. arXiv preprint arXiv:1906.09477, 2019.

A. B. Yuditskii, A. V. Nazin, A. B. Tsybakov, and N. Vayatis. Recursive aggregation of estimators by mirror descent algorithm with averaging. Problemy Peredachi Informatsii, 41(4): 78–96, 2005.

M. Zugno, T. Jónsson, and P. Pinson. Trading wind energy on the basis of probabilistic forecasts both of wind generation and of market quantities. Wind Energy, 16(6):909–926, 2013a.

M. Zugno, J. M. Morales, P. Pinson, and H. Madsen. Pool strategy of a price-maker wind power producer. IEEE Transactions on Power Systems, 28(3):3440–3450, 2013b.

ECOLE
DOCTORALE
DE MATHEMATIQUES
HADAMARD

Titre: Modélisation et stratégies optimales sur les marchés court terme de l'énergie

Mots clés: Marchés infrajournaliers, énergie renouvelable, jeux à champ moyen, joueur majeur, prévision probabiliste, contrôle stochastique, réseaux de neurones, PAC-Bayes

Résumé: Cette thèse vise à fournir des outils théoriques pour soutenir le développement et la gestion des énergies renouvelables intermittentes sur les marchés court terme de l'électricité.

Dans la première partie, nous développons un modèle d'équilibre exploitable pour la formation des prix sur les marchés infrajournaliers de l'électricité. Pour cela, nous proposons un jeu non coopératif entre plusieurs producteurs interagissant sur le marché et faisant face à une production renouvelable intermittente. En utilisant la théorie des jeux et celle du contrôle stochastique, nous dérivons des stratégies optimales explicites pour ces producteurs ainsi qu'un prix d'équilibre en forme fermée pour différentes structures d'information et caractéristiques des joueurs. Notre modèle permet de reproduire et d'expliquer les principaux faits stylisés du marché intraday tels que la dépendance temporelle spécifique de la volatilité et la corrélation entre le prix et les prévisions de production renouvelable.

Dans la deuxième partie, nous étudions des prévisions probabilistes dynamiques sous la forme de processus de diffusion. Nous proposons plusieurs modèles d'équations différentielles stochastiques pour capturer l'évolution dynamique de l'incertitude associée à une prévision, nous dérivons les densités prédictives associées et nous calibrons le modèle sur des données météorologiques réelles. Nous l'appliquons ensuite au problème d'un producteur éolien recevant des mises à jour séquentielles des prévisions probabilistes de la vitesse du vent, utilisées pour prédire sa production, et prendre des décisions d'achat ou de vente sur le marché. Nous montrons dans quelle mesure cette méthode peut être avantageuse comparée à l'utilisation de prévisions ponctuelles dans les processus décisionnels.

Enfin, dans la dernière partie, nous proposons d'étudier les propriétés des réseaux de neurones peu profonds agrégés. Nous explorons le cadre PAC-Bayesien comme alternative à l'approche classique de minimisation du risque empirique. Nous nous concentrons sur les priors Gaussiens et dérivons des bornes de risque non asymptotiques pour les réseaux de neurones agrégés. Ces bornes atteignent des vitesses de convergence minimax pour l'estimation dans des espaces de Sobolev. Cette analyse fournit également une base théorique pour le réglage des paramètres et offre de nouvelles perspectives pour l'application des réseaux de neurones agrégés à des problèmes pratiques en grande dimension, de plus en plus présents dans les processus de décision liés à l'énergie et impliquant des moyens de production renouvelable ou du stockage.

Title: Modeling and optimal strategies in short-term energy markets

Keywords: Intraday electricity market, renewable energy, mean field games, major player, probabilistic forecasting, stochastic control, shallow neural networks, PAC-Bayes,

Abstract: This thesis focuses on providing theoretical tools to help in the development and management of intermittent renewable energy in short term electricity markets.
In the first part, we develop a tractable equilibrium model for price formation in intraday electricity markets. For this, we propose a non cooperative game between several producers interacting in the market and facing an intermittent renewable production. Using stochastic control and game theory, we derive explicit optimal strategies for these producers as well as a closed form equilibrium price for different information structures and player characteristics. Our model allows to reproduce and explain the main stylized features of the intraday market such as the specific time dependence of volatility and the correlation between the price and the renewable production forecasts.
In the second part, we study dynamic probabilistic forecasts in the diffusion framework. We propose several stochastic differential equation models to capture the dynamic evolution of the uncertainty associated to a forecast, derive the associated predictive densities and calibrate the model on real meteorological data. We then apply it to the problem of a wind energy producer receiving sequential updates of the probabilistic forecasts of the wind speed used to predict her production and make trading decisions in the market. We show to what extent this method can outperform the use of point forecasts in decision-making processes.
Finally, in the last part, we propose to study the properties of aggregated shallow neural networks. We explore the PAC-Bayesian framework as an alternative to the classical empirical risk minimization approach. We focus on Gaussian priors and derive non-asymptotic risk bounds for the aggregated neural networks. These bounds yield minimax rates of estimation over Sobolev smoothness classes. This analysis also provides a theoretical basis for tuning the parameters and offers new perspectives for applications of aggregated neural networks to practical high dimensional problems increasingly present in energy decision problems involving renewables or storage.

Institut Polytechnique de Paris
91120 Palaiseau, France