



HAL
open science

Some application of machine learning in quantitative finance: model calibration, volatility formation mechanism and news screening

Jianfei Zhang

► **To cite this version:**

Jianfei Zhang. Some application of machine learning in quantitative finance: model calibration, volatility formation mechanism and news screening. Computational Finance [q-fin.CP]. Institut Polytechnique de Paris, 2023. English. NNT : 2023IPPAX043 . tel-04528965

HAL Id: tel-04528965

<https://theses.hal.science/tel-04528965>

Submitted on 2 Apr 2024

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 : 2023IPPA043

Thèse de doctorat



Some applications of machine learning in quantitative finance: model calibration, volatility formation and news screening

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

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

Thèse présentée et soutenue à Palaiseau, le 27 June 2023, par

JIANFEI ZHANG

Composition du Jury :

| | |
|---|--------------------|
| Olivier Guéant Professeur, Université Paris 1 Panthéon-Sorbonne | Président |
| Ahmed Kebaier Professeur, Université d'Évry | Rapporteur |
| Jun Yu Professor, Singapore Management University | Rapporteur |
| Eduardo Abi Jaber Professeur Assistant, École polytechnique | Examineur |
| Stefano De Marco Professeur Associé, École polytechnique | Examineur |
| Sophie Laruelle Professeur Assistant, Université Paris-Est Créteil | Examinatrice |
| Sergio Pulido Professeur Associé, ENSIE | Examineur |
| Mathieu Rosenbaum Professeur, École polytechnique | Directeur de thèse |

Acknowledgement

Going through this journey has been invaluable for my expertise and personal development. I deeply appreciate all those who have contributed to the completion of this doctoral thesis. It would not be possible to reach this step without your support, guidance, and encouragement.

First and foremost, I would like to express my utmost gratitude to my supervisor, Mathieu Rosenbaum. Your dedication, patience, insightful comments, and scrupulosity about every detail are essential to the quality of this work, not speak of your remarkable expertise. I am truly fortunate to have had the opportunity to work with you. It is also a great pleasure to chat with you and with your son, and be impressed by his knowledge of the Paris transportation map.

I would also like to extend my sincere thanks to Ahmed Kabaier and Jun Yu for accepting to be the referees of my thesis, as well as Eduardo Abi Jaber, Stefano De Marco, Olivier Guéant, Sophie Laruelle, and Sergio Pulido. Your valuable feedback and suggestions have significantly contributed to the refinement and improvement of this thesis.

I feel privileged to work in the collaboration program between Exoduspoint Capital and Ecole polytechnique. The development of this program would not be realized without the commitment of Mathieu Rosenbaum and Nizar Touzi, and the generosity of Exoduspoint management team.

I would like to thank my colleagues inside the Exoduspoint Quantitative Strategy team: Marouane for your understanding and trust, Xiaofei for all the fruitful discussions and guidance, Qinkai for your technical help and I would definitely miss the countless afternoon cafe with you. My appreciation continues to the other members of the team, Alexandre, Quentin, Yuyang, Niem, Pierre-Jean, Jean-Sébastien, Mehdi, Nicolas, Johathan, Bruno, François, Benoît, Johan, Mukunda, Arthur and Thomas. Thanks all for your camaraderie and collaboration. Special thanks must go to Stéphane for your unconditional support.

Lastly, my deepest gratitude to my family for your love, belief, and encouragement.

Résumé

Nous commençons cette thèse au Chapitre I en présentant nos méthodologies basées sur l'apprentissage profond pour une calibration efficace du modèle *quadratic rough Heston*, qui est non-markovien et non-semimartingale. Une approximation multifactorielle du modèle est d'abord proposée. Deux réseaux neuronaux profonds sont ensuite entraînés sur des données simulées afin d'apprendre les fonctions de tarification des options SPX et VIX respectivement. Les résultats de la calibration jointe peuvent être obtenus instantanément. Des tests empiriques montrent que le modèle calibré peut très bien reproduire les surfaces de volatilité implicite du SPX et du VIX. Grâce à l'application de la différenciation adjointe automatique, ces deux réseaux neuronaux peuvent être utilisés pour calculer efficacement les quantités de couverture.

Dans le cadre de *quadratic rough Heston*, nous formulons ensuite un problème de *market making* optimale sur un panier composé de plusieurs dérivés de SPX, tels que des contrats à terme SPX et VIX, des options SPX et VIX. Le *market maker* maximise son profit en capturant les écarts entre l'achat et la vente et pénalise le risque d'inventaire du portefeuille, qui peut être expliqué principalement par la variation de SPX. Nous abordons la haute dimensionnalité du problème avec plusieurs approximations pertinentes. Des solutions asymptotiques à forme fermée peuvent ensuite être déduites.

Motivés par le fait largement établi que la volatilité est *rough* dans de nombreuses classes d'actifs et par les résultats prometteurs grâce à l'incorporation de l'effet *Zumbach*, nous nous intéressons à l'universalité des mécanismes de formation de la volatilité endogène au Chapitre III. Nous commençons par une approche quasi-nonparamétrique. Un réseau LSTM entraîné sur un ensemble de données regroupées couvrant des centaines d'actions liquides montre une performance de prévision de la volatilité supérieure aux autres dispositifs spécifiques à un actif ou à un secteur, ce qui suggère l'existence de l'universalité en question. Nous proposons ensuite une méthode de prévision paramétrique parcimonieuse combinant les principales caractéristiques des modèles *rough fractional stochastic volatility* et *quadratic rough Heston*. Avec des paramètres fixes, cette approche présente le même niveau de performance que le LSTM universel, confirmant à nouveau l'universalité d'un point de vue paramétrique.

Enfin, nous nous concentrons sur le lien entre les publications des nouvelles et les changements des conditions de liquidité intrajournalière dans le marché financier. En particulier, en détectant les sauts significatifs de la volatilité et des volumes, nous proposons une approche systématique pour distinguer les nouvelles impactantes des neutres présentées massivement. Un prédicteur du sentiment des nouvelles apprendit sur les échantillons sélectionnés s'avère beaucoup plus efficace que celui calibré sur l'ensemble des données brutes en termes de prédiction des mouvements de prix à court terme pour les actifs associés.

Abstract

We begin this thesis in Chapter I by introducing our deep learning-based methodologies for efficient calibration of the quadratic rough Heston model, which is non-Markovian and non-semimartingale. A multi-factor approximation of the model is first proposed. Two deep neural networks are then trained on simulated data to learn the resulting pricing functions of SPX and VIX options respectively. Given SPX/VIX smiles, joint calibration results can be instantaneously obtained. Empirical tests show that the calibrated model can reproduce both SPX and VIX implied volatility surfaces very well. Through the application of automatic adjoint differentiation, the neural networks can be used to compute efficiently related hedging quantities.

Under quadratic rough Heston, we then formulate an optimal market making problem on a basket composed of multiple derivatives of SPX, such as SPX and VIX futures, SPX and VIX options. The market maker maximizes its profit from spread capturing and penalizes the portfolio's inventory risk, which can be mostly explained by the variation of SPX. We tackle the high dimensionality of the problem with several relevant approximations. Closed-form asymptotic solutions can then be obtained.

Motivated by the widely established fact that volatility is rough across many asset classes and the promising results of the quadratic rough Heston model thanks to the incorporation of the Zumbach effect, we are interested in the universality of the endogenous volatility formation mechanism in Chapter III. We take first a quasi-nonparametric approach. An LSTM network trained on a pooled dataset covering hundreds of liquid stocks shows superior volatility forecasting performance than other asset-specific or sector-specific devices, suggesting the existence of the universality in question. We propose then a parsimonious parametric forecasting method combining the main features of the rough fractional stochastic volatility and quadratic rough Heston models. With fixed parameters, this approach presents the same level of performance as the universal LSTM, confirming again the universality from a parametric perspective.

Finally, we focus on the link between informative news releases and switches of intraday liquidity conditions. Particularly, by detecting significant jumps of volatility and trading volumes inside daily trading sessions, we propose a systematic approach to distinguish impactful news releases from the massively presented neutral ones. A news sentiment predictor fitted on the identified impactful data is shown to be much more effective than the one calibrated on the raw dataset in terms of short-term price movement prediction for the associated assets.

List of papers being part of this thesis

- M. Rosenbaum and J. Zhnag, *Deep calibration of the quadratic rough Heston model*, published in Risk Magazine, 2022.
- M. Rosenbaum and J. Zhang, *Multi-asset market making under quadratic rough Heston*, submitted, 2022.
- M. Rosenbaum and J. Zhang, *On the universality of the volatility formation process: when machine learning and rough volatility agree*, to appear in Frontiers of Mathematical Finance, 2023.
- J. Zhang, *Towards systematic intraday news screening: a liquidity-focused approach*, submitted, 2023.

Contents

| | | |
|--|--|-----------|
| Contents | | xi |
| I Introduction | | 1 |
| Motivations | | 1 |
| Outline | | 2 |
| 1 Chapter II: Deep calibration of the quadratic rough Heston model | | 4 |
| 1.1 The QRH model and its multi-factor approximation | | 4 |
| 1.2 Efficient calibration with deep learning | | 6 |
| 1.3 Hedging computation | | 8 |
| 2 Chapter III: Multi-asset market making under quadratic rough Heston | | 9 |
| 2.1 Problem description | | 9 |
| 2.2 Value function approximation | | 10 |
| 2.3 Numerical results | | 12 |
| 3 Chapter IV: On the universality of the volatility formation process: when machine learning and rough volatility agree | | 13 |
| 3.1 Realized volatility forecasting devices | | 14 |
| 3.2 Capturing universality with LSTM | | 14 |
| 3.3 Uncovering the universal volatility formation process | | 16 |
| 4 Chapter V: Towards systematic intraday news screening: a liquidity-focused approach | | 18 |
| 4.1 Liquidity mode fitting | | 18 |
| 4.2 News screening and sentiment learning | | 20 |
| I Introduction (en français) | | 23 |
| Motivations | | 23 |
| Schéma | | 25 |
| 1 Chapitre II: Calibration du modèle quadratique rough Heston avec l'apprentissage profond | | 26 |
| 1.1 Le modèle QRH et son approximation multifactorielle | | 27 |
| 1.2 Calibration efficace avec l'apprentissage profond | | 28 |
| 1.3 Calcul de couverture | | 30 |
| 2 Chapitre III: Market-making multi-actifs dans le cadre de la méthode quadra- tique de rough Heston | | 32 |

| | | |
|------------|---|-----------|
| 2.1 | Description du problème | 32 |
| 2.2 | Approximation de la fonction de valeur | 33 |
| 2.3 | Résultats numériques | 35 |
| 3 | Chapitre IV: Sur l'universalité du processus de formation de la volatilité : quand l'apprentissage automatique et la volatilité rugueuse s'accordent | 36 |
| 3.1 | Dispositifs de prévision de la volatilité réalisée | 37 |
| 3.2 | Capturer l'universalité avec LSTM | 38 |
| 3.3 | Découvrir le processus universel de formation de la volatilité | 39 |
| 4 | Chapitre V: Vers un filtrage systématique des nouvelles intrajournalières : une approche axée sur la liquidité | 40 |
| 4.1 | Calibration du mode de liquidité | 41 |
| 4.2 | Filtrage de l'actualité et apprentissage des sentiments | 43 |
| II | Deep calibration of the quadratic rough Heston model | 47 |
| 1 | Introduction | 48 |
| 2 | The quadratic rough Heston model and its multi-factor approximation | 49 |
| 3 | Model calibration with deep learning | 51 |
| 3.1 | Methodology | 51 |
| 3.2 | Pricing | 53 |
| 3.3 | Calibration | 54 |
| 4 | Toy examples for hedging | 58 |
| 4.1 | Hedging portfolio computation with neural networks | 58 |
| 4.2 | Numerical results | 61 |
| 5 | Conclusion | 63 |
| | Appendices | 65 |
| II.A | Kernel function approximation and simulation scheme | 65 |
| II.B | Network training with transfer learning | 66 |
| II.C | Pricing and calibration with neural networks | 67 |
| II.C.1 | On simulated data | 67 |
| II.C.2 | On market data | 69 |
| III | Multi-asset market making under quadratic rough Heston | 73 |
| 1 | Introduction | 74 |
| 2 | Description of the problem | 76 |
| 2.1 | Multi-factor approximation of the QRH model | 76 |
| 2.2 | Multi-asset market making | 77 |
| 3 | Value function approximation | 79 |
| 3.1 | The Hamilton-Jacobi-Bellman equations | 79 |
| 3.2 | Quadratic approximation | 81 |
| 4 | Numerical results | 83 |
| 4.1 | Daily hedging with SPX | 83 |
| 4.2 | Market making approximation | 84 |
| 4.3 | Example 3 | 89 |

| | | |
|-----------|--|------------|
| IV | On the universality of the volatility formation process: when machine learning and rough volatility agree | 91 |
| 1 | Introduction | 91 |
| 2 | Description of the forecasting devices | 94 |
| | 2.1 Parametric methods | 94 |
| | 2.2 Nonparametric forecasting with LSTM | 95 |
| 3 | Data and evaluation metrics | 96 |
| 4 | Capturing universality with LSTM | 97 |
| | 4.1 Parametric vs nonparametric | 97 |
| | 4.2 Network inspection | 100 |
| 5 | Uncovering the universal volatility formation process | 106 |
| 6 | Conclusion | 110 |
| V | Towards systematic intraday news screening: a liquidity-focused approach | 111 |
| 1 | Introduction | 112 |
| 2 | Data and preprocessing | 115 |
| 3 | Liquidity mode fitting | 117 |
| | 3.1 Jump model | 117 |
| | 3.2 Fitting results | 118 |
| 4 | News screening and learning | 120 |
| | 4.1 Methodology | 120 |
| | 4.2 Numerical results | 122 |
| 5 | Conclusion | 125 |
| | Appendices | 127 |
| | V.A Jump model fitting | 127 |
| | V.B Mutual information and naive Bayes classifier | 127 |
| | V.C Robustness tests | 129 |
| | Bibliography | 131 |

CHAPTER I

Introduction

Motivations

The widely accepted rough volatility paradigm can reproduce all the stylized facts of volatility time series and fit remarkably well the shape of implied volatility smiles and at-the-money skew curves. Particularly, the recently introduced quadratic rough Heston (QRH) model shows promising results on jointly modeling SPX/VIX smiles, where other continuous-time stochastic volatility models have been struggling for a long time. As the model is non-Markovian and non-semimartingale, it is intricate to calibrate it based on classical Monte Carlo simulations. Thus we are interested in:

Question 1. *Whether machine learning can provide alternative solutions for calibrating the QRH model?*

We propose a multi-factor approximated version of the original model. The approximation makes the model Markovian and can reproduce the “rough” dynamics at the time scales of interest. Deep neural networks are applied to learn the pricing maps of SPX/VIX options based on simulated data, from which joint calibration can be conducted efficiently. We show that the calibrated model can reproduce very well both SPX and VIX implied volatility surfaces. Moreover, the trained networks can be used to compute instantaneously related hedging quantities. This is very useful for risk control when practitioners hold SPX derivatives, for example in the context of market making activities, motivating us to consider the following question:

Question 2. *How to design and solve a multi-asset market making problem under the QRH model?*

We formulate an optimal market making problem on multiple SPX derivatives, including SPX futures, VIX futures, SPX and VIX options. To maximize the profit from spread capturing, the market maker decides whether to place limit orders in the queues of the best bid/ask limits, while controlling the portfolio’s holding risk. It leads to a high dimensional stochastic control problem. We surmount the curse of dimensionality with several approximations, whose accuracies and relevances are shown through numerical experiments. Closed-form asymptotic

solutions can then be obtained.

A key feature of the QRH model contributing to the implied volatility fitting is the inclusion of the so-called *Zumbach* effect, which describes essentially the feedback effect of price trends on volatility. This feature can also be helpful for realized volatility modeling. A lot of works have investigated the roles played by past market realizations on future volatility forecasting. With now the widely accepted rough volatility, we are particularly interested in the following question:

Question 3. *Is the volatility formation mechanism relating past returns and volatilities to future volatilities universal across assets?*

We tackle this question from both nonparametric and parametric perspectives. On the one hand, we train a long short-term memory (LSTM) network based on a pooled dataset covering hundreds of liquid stocks aiming to forecast the next daily volatility for all stocks. We show that this universal LSTM outperforms other asset-specific or sector-specific models. Its performance holds even across markets. On the other hand, a parsimonious parametric forecasting device, combining the main features of the rough fractional stochastic volatility (RFSV) and QRH models with fixed parameters, presents the same level of forecasting capability as the universal LSTM. We confirm thus the existence of the universality in question with these results.

In addition to the endogenous part induced by past market activities through this universal mechanism, market volatility is also contributed by exogenous events, a main source of which consists of news publications. Unexpected news are likely to be followed by significant market movements. To evaluate efficiently the impact of a news release on the price of the associated asset, institutional investors need to develop automatic news sentiment evaluation methods based on historical data. A large number of news articles are published every day and most of them are neutral, *i.e.* the views transferred are neither bullish nor bearish. In order to build a robust sentiment predictor, it is reasonable to fit it mostly on the news samples that have driven considerable market activities. For that, we inquire

Question 4. *How to efficiently filter impactful news releases out of the large raw news dataset?*

Naturally, bullish (bearish) news releases are likely to drive positive (negative) asset price changes. Yet selecting news samples solely based on post-release return information is not robust considering the low signal-to-noise ratio of return data. We believe that impactful exogenous signals will also induce significant changes in liquidity conditions, such as larger volatilities and more turnovers. Thus, we propose a approach that consists in detecting systematically the liquidity mode switches and taking the news releases around these switches as impactful ones.

Outline

In Chapter II we address Question 1 by building an efficient calibration procedure based on deep neural networks. We first transform the original QRH model into a multi-factor

approximated version. We generate large numbers of SPX/VIX option prices with randomly sampled model parameters. Two networks are then trained based on these samples to learn the pricing maps from model parameters to SPX/VIX implied volatility surfaces respectively. The joint calibration results are given by minimizing the distance between observed implied volatilities and the outputs of both networks. The accuracy of calibration is evaluated on simulated data and also on market data. In the setting of our model, only one Brownian motion is involved and European options can be hedged perfectly with only SPX. We show how to compute the hedging quantities with the trained networks in practice. For reasons of comparison, another hedging method based on differential machine learning is also introduced.

We give a practical solution to Question 2 in Chapter III. Given the promising results of the QRH model on joint modeling of SPX/VIX smiles, we consider a market making basket composed of numerous SPX derivatives. The assets considered are essentially large tick assets. The market maker tries to maximize its profit from spread capturing through sending limit orders at the best bid/ask limits, and controls its inventory risk, which can be mostly explained by the variation of SPX. This is formulated as a high-dimensional stochastic control problem. To solve it efficiently, we keep the stochastic volatility factors constant for each market making period. The HJB equation associated with the resulting optimization problem is proved to have a unique solution. Classical verification argument can be easily established. To get further practicability, we apply quadratic approximation to the Hamiltonian functions. A closed-form asymptotic solution can then be obtained. Numerical simulations show that the approximated solution performs closely to that given by the classical Euler scheme.

We turn to Question 3 in Chapter IV by evaluating the predictive power of past returns and volatilities on future volatilities. Several parametric model-based methods are investigated, including the heterogeneous autoregressive (HAR), the rough fractional stochastic volatility (RFSV) models, and a quasi-nonparametric method based on LSTM. The LSTM trained on the data of hundreds of liquid US stocks performs the best. Asset- or sector-specific training does not improve the results. Similar superior performances hold on assets not part of the training set, even those of a different market. These observations suggest the existence of a universal volatility formation mechanism. We then show that a parsimonious parametric forecasting device with five globally fixed parameters, incorporating the main features of the RFSV and QRH model can reproduce the same level of performance as the universal LSTM, indicating that the learned universality can be well described by the rough volatility paradigm boosted with the so-called Zumbach effect.

To tackle Question 4 in Chapter V, we propose a two-step approach. First, we measure continuously the intraday liquidity conditions of financial assets through four variables, *i.e.* volatility, trading volume, bid-ask spread, and book size. We categorize the liquidity states into two modes by using a nonparametric method called the jump model. Through numerical tests on S&P500 components, we find that the fitted two modes are distinct from each other in terms of volatility and trading volume. Second, we zoom in on the mode switches from that associated with lower volatility and trading volumes to the other one. We take all the news articles

published around these liquidity state jump moments as impactful ones. To illustrate the effect of our approach, we fit two naive Bayes classifiers (NBCs) on the raw and screened dataset respectively. The latter shows significantly better predictive power on short-term stock price movements, which indicates that our news screening procedure can diminish the impact of neutral news releases and thus make the capturing of sentiment-contributed features more effective.

We now give a rapid overview of the main results obtained in this thesis.

1 Chapter II: Deep calibration of the quadratic rough Heston model

The rough Heston model can fit very well the SPX implied volatility smiles [“el2018perfect”, n.d.; El Euch & Rosenbaum, 2019; El Euch et al., 2019]. However, one stylized fact of financial time series not reflected in this model is the feedback effect of past price trends on volatility, which is referred to as Zumbach effect. Dandapani et al. [2021] clarify the definition of Zumbach effect and introduce super-Heston rough volatility models encoding explicitly this effect. As a particular example of super-Heston rough volatility models, the QRH model introduced in Gatheral et al. [2020] shows its promising ability to fit jointly SPX/VIX implied volatility surfaces. In order to promote further the practicability of the QRH model, we aim at designing an efficient model calibration procedure in this chapter. Inspired by recent works about applications of deep learning in quantitative finance, see for example Hernandez [2016]; Horvath et al. [2019], we apply deep neural networks for this task. The effectiveness of the calibrated model for fitting jointly SPX and VIX smiles is illustrated through numerical experiments. Interestingly, under our model, the trained networks also allow us to hedge European options with instantaneous computation of hedging quantities.

1.1 The QRH model and its multi-factor approximation

In the QRH model, the dynamics of the price of an asset S (here the SPX) and its spot variance V under risk-neutral measure read

$$dS_t = S_t \sqrt{V_t} dW_t, \tag{I.1}$$

$$V_t = a(Z_t - b)^2 + c, \tag{I.2}$$

where W is a Brownian motion, a, b, c are all positive constants and Z_t is defined as

$$Z_t = \int_0^t \lambda \frac{(t-s)^{\alpha-1}}{\Gamma(\alpha)} (\theta_0(s) - Z_s) ds + \int_0^t \eta \frac{(t-s)^{\alpha-1}}{\Gamma(\alpha)} \sqrt{V_s} dW_s, \tag{I.3}$$

for $t \in [0, T]$. Here T is a positive time horizon, $\alpha \in (1/2, 1)$, $\lambda > 0$, $\eta > 0$ and $\theta_0(\cdot)$ is a deterministic function. El Euch & Rosenbaum [2019] show that the volatility trajectories under the rough Heston model have almost surely Hölder regularity $\alpha \sim \frac{1}{2} - \varepsilon$, for any $\varepsilon > 0$. This recalls the observations in Gatheral et al. [2018] that the dynamic of log-volatility is similar to that of a fractional Brownian motion with Hurst index of order 0.1. Compared to other

common stochastic volatility models, the QRH model involves only one Brownian motion. The process Z_t encodes the information of past returns since we have $\sqrt{V_t}dW_t = dS_t/S_t$. Then by (I.2), V_t incorporates naturally the *strong* Zumbach effect, as suggested in [Dandapani et al. \[2021\]](#). Similar to the rough Heston model, the fractional kernel $K(t) = \frac{t^{\alpha-1}}{\Gamma(\alpha)}$ in (I.3) enables us to generate rough volatility dynamics. However, it makes the QRH model non-Markovian and non-semimartingale, and thus brings difficulties on simulation and calibration. We apply the method introduced in [Abi Jaber & El Euch \[2019\]](#) to replace $K(t)$ by the following multi-factor approximated version

$$K^n(t) = \sum_{i=1}^n c_i^n e^{-\gamma_i^n t}, \quad n \geq 1,$$

where $(c_i^n)_{i=1, \dots, n}$ and $(\gamma_i^n)_{i=1, \dots, n}$ can be chosen as explicit functions of α . Thus, they are not free parameters to calibrate. Properly chosen parametrizations of the $2n$ parameters $(c_i^n, \gamma_i^n)_{i=1, \dots, n}$ in terms of α can make $K^n(t)$ converge to $K(t)$ in the L^2 sense as n goes to infinity and the multi-factor approximation models behave closely to their counterparts in the rough volatility paradigm [[Abi Jaber, 2019](#); [Abi Jaber & El Euch, 2019](#)]. Then, we can propose the following multi-factor approximation of the QRH model:

$$dS_t^n = S_t^n \sqrt{V_t^n} dW_t, \quad V_t^n = a(Z_t^n - b)^2 + c, \quad (\text{I.4})$$

$$Z_t^n = \sum_{i=1}^n c_i^n Z_t^{n,i}, \quad (\text{I.5})$$

$$dZ_t^{n,i} = (-\gamma_i^n Z_t^{n,i} - \lambda Z_t^n) dt + \eta \sqrt{V_t^n} dW_t, \quad Z_0^{n,i} = z_0^i, \quad (\text{I.6})$$

with $(z_0^i)_{i=1, \dots, n}$ some constants. In order to avoid predefining any specific formulation of $\theta(\cdot)$, we discard it in Equation (I.5) and we consider the starting values of $(z_0^i)_{i=1, \dots, n}$ as free parameters to calibrate from market data, which allows more model flexibilities. In fact, given a solution for (I.4) - (I.6), we can rewrite Equation (I.6) as

$$Z_t^{n,i} = z_0^i e^{-\gamma_i^n t} + \int_0^t e^{-\gamma_i^n(t-s)} (-\lambda Z_s^n ds + \eta \sqrt{V_s^n} dW_s). \quad (\text{I.7})$$

Then from (I.5) we can get

$$Z_t^n = g^n(t) + \int_0^t K^n(t-s) (b(Z_s^n) ds + \sigma(Z_s^n) dW_s), \quad (\text{I.8})$$

with $g^n(t) = \sum_{i=1}^n z_0^i c_i^n e^{-\gamma_i^n t}$, $b(Z_t^n) = -\lambda Z_t^n$ and $\sigma(Z_t^n) = \eta \sqrt{a(Z_t^n - b)^2 + c}$. By taking expectation on both sides of (I.8), we get

$$\mathbb{E}[Z_t^n] + \lambda \sum_{i=1}^n c_i^n \int_0^t e^{-\gamma_i^n(t-s)} \mathbb{E}[Z_s^n] ds = \sum_{i=1}^n z_0^i c_i^n e^{-\gamma_i^n t},$$

with which the dependence of $\mathbb{E}[Z_t^n]$ on $(z_0^i)_{i=1, \dots, n}$ can be explicitly obtained. Thus we can see that the $(z_0^i)_{i=1, \dots, n}$ allows us to encode initial ‘‘term-structure’’ of Z_t^n . Therefore, it can be understood as an analog of $\theta(t)$ for the variance process in the rough Heston model. Besides, the strong uniqueness of the solution of Equations (I.7)- (I.8) can be established by virtue of Proposition B.3 in [Abi Jaber & El Euch \[2019\]](#).

1.2 Efficient calibration with deep learning

- Methodology

We let $\boldsymbol{\omega} := (\lambda, \eta, a, b, c) \in \Omega \subset R^5$, $\mathbf{z}_0 := (z_0^1, \dots, z_0^n) \in \mathcal{Z} \subset R^n$. We set $n = 10$ in our tests, and we discard the label n in the following to simplify notations. We first generate large numbers of synthetic samples based on (I.4) - (I.6) and then train neural networks based on the simulated data.

For each random sampling $(\boldsymbol{\omega}, \mathbf{z}_0)$, we generate thousands of random paths of SPX and VIX with an explicit-implicit Euler scheme. Then Monte-Carlo prices are used to compute IVS_{SPX}^{MC} and IVS_{VIX}^{MC} , representing implied volatility surfaces (IVS) of SPX and VIX options respectively. Two neural networks, denoted by \mathcal{NN}_{SPX} and \mathcal{NN}_{VIX} , are trained on data pairs $\{(\boldsymbol{\omega}, \mathbf{z}_0), IVS_{SPX}\}$ and $\{(\boldsymbol{\omega}, \mathbf{z}_0), IVS_{VIX}\}$ respectively. Thus, \mathcal{NN}_{SPX} and \mathcal{NN}_{VIX} consist in learning the pricing functions of SPX and VIX options in the multi-factor approximated QRH model, as illustrated in Figure I.1.




Figure I.1: \mathcal{NN}_{SPX} and \mathcal{NN}_{VIX} approximate the pricing mappings from model parameters to SPX and VIX implied volatilities.

Given observed IVS_{SPX} and IVS_{VIX} , the joint calibration result is given as

$$\hat{\boldsymbol{\omega}}, \hat{\mathbf{z}}_0 = \operatorname{argmin}_{\boldsymbol{\omega} \in \Omega, \mathbf{z}_0 \in \mathcal{Z}} \|\mathcal{NN}_{SPX}(\boldsymbol{\omega}, \mathbf{z}_0) - IVS_{SPX}\|_2^2 + \|\mathcal{NN}_{VIX}(\boldsymbol{\omega}, \mathbf{z}_0) - IVS_{VIX}\|_2^2,$$

where $\|\cdot\|_2$ stands for L^2 norm. We apply L-BFGS-B, a gradient-based algorithm, for solving this optimization problem. The gradients needed are calculated directly with \mathcal{NN}_{SPX} and \mathcal{NN}_{VIX} via *automatic adjoint differentiation* (AAD).

- Numerical results

We perform model calibration on market data and then get IVS of SPX and VIX options with the calibrated parameters through Monte-Carlo simulations. Figure I.2 and I.3 give an example with the data of 19 May 2017. Our model fits very well the global shape of IVS of SPX and VIX options at the same time. Model's output fall essentially between the bid and ask quotes for VIX options. Moreover, excellent fits are obtained in terms of the at-the-money skew of SPX options.




Figure I.2: Implied volatilities of SPX options for 19 May 2017. Bid and ask of market volatilities are represented respectively by red and blue points. The green line is the output of the model with Monte-Carlo method.

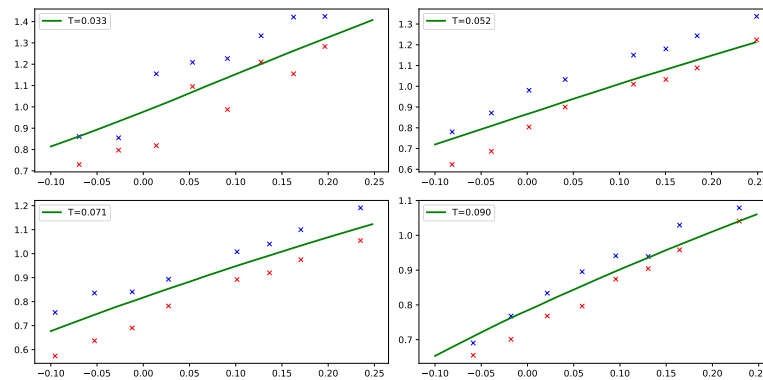


Figure I.3: Implied volatilities of VIX options for 19 May 2017. Bid and ask of market volatilities are represented respectively by red and blue points. The green line is the output of the model with Monte-Carlo method.

1.3 Hedging computation

- Methodology

In the model defined in (I.4)-(I.6), European options can be hedged with only SPX. We take vanilla SPX call as an illustrative example. Let $\mathbf{Z}_t := (Z_t^1, \dots, Z_t^{10})$ and $\mathbf{X}_t := (S_t, \mathbf{Z}_t)$. Given strike K , maturity T , and model parameters $\boldsymbol{\omega}$, the price of vanilla SPX call at time t is a function of \mathbf{X}_t . Let $P_t(\mathbf{X}_t; K, T, \boldsymbol{\omega})$ denote this quantity. We have

$$P_t(\mathbf{X}_t; K, T, \boldsymbol{\omega}) \simeq P^{BS}(S_t, K, T - t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)),$$

where $P^{BS}(S, K, T, \sigma)$ is the price of European call under Black-Scholes (BS) model and $\sigma_{\mathcal{N}\mathcal{N}}^{k, T}(\boldsymbol{\omega}, \mathbf{Z}_t)$ is the implied volatility corresponding to log-moneyness strike k and maturity T , calculated directly from $\mathcal{N}\mathcal{N}_{SPX}$ with $(\boldsymbol{\omega}, \mathbf{Z}_t)$ as input. We can have

$$\begin{aligned} dP^{BS}(S_t, K, T - t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)) &= \left(\frac{\partial P^{BS}}{\partial S_t} + \frac{\eta}{S_t} \sum_{i=1}^{10} \frac{\partial P^{BS}}{\partial Z_t^i} \right) (S_t, K, T - t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)) dS_t \\ &=: \delta_t dS_t, \end{aligned}$$

with

$$\frac{\partial P^{BS}(S_t, K, T - t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t))}{\partial S_t} \simeq \delta_{BS}(S_t, K, T - t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)) + \quad (\text{I.9})$$

$$v_{BS}(S_t, K, T - t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)) \frac{\partial \sigma_{\mathcal{N}\mathcal{N}}^{k, T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)}{\partial k} \Big|_{k=\log(K/S_t)} \frac{\partial \log(K/S_t)}{\partial S_t},$$

$$\frac{\partial P^{BS}(S_t, K, T - t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t))}{\partial Z_t^i} \simeq v_{BS}(S_t, K, T - t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)) \frac{\partial \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)}{\partial Z_t^i}, \quad (\text{I.10})$$

where $\delta_{BS}(S, K, T, \sigma)$ and $v_{BS}(S, K, T, \sigma)$ stand for the *Delta* and *Vega* respectively under BS model. The quantity $\frac{\partial \sigma_{\mathcal{N}\mathcal{N}}^{k, T}}{\partial Z_t^i}$ corresponds actually to the derivative of the outputs of $\mathcal{N}\mathcal{N}_{SPX}$ with respect to its inputs, which can be obtained instantaneously with built-in AAD. $\frac{\partial \sigma_{\mathcal{N}\mathcal{N}}^{k, T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)}{\partial k}$ in (I.9) can be approximated by the finite difference $\frac{(\sigma_{\mathcal{N}\mathcal{N}}^{k+\delta_k, T-t} - \sigma_{\mathcal{N}\mathcal{N}}^{k-\delta_k, T-t})}{2\delta_k}(\boldsymbol{\omega}, \mathbf{Z}_t)$.

- Numerical results

We perform discrete hedging with time step Δt . The Profit and Loss (P&L) of hedging at $t \in (0, T]$ is given by

$$\mathcal{I}_t = \mathcal{I}_t^\delta - \mathcal{I}_t^P,$$

where

$$\begin{aligned} \mathcal{I}_t^\delta &= \sum_{k=0}^{\lfloor t/\Delta t \rfloor - 1} \delta_{t_k} \Delta S_{t_{k+1}} + \delta_{t_{\lfloor t/\Delta t \rfloor}} (S_t - S_{\lfloor t/\Delta t \rfloor \Delta t}), \\ \mathcal{I}_t^P &= P_t - P_0, \end{aligned}$$

with $t_k = k\Delta t$, $\Delta S_{t_k} := S_{t_k} - S_{t_{k-1}}$, and P_t is the price of the SPX option to hedge at time t . As we can see, \mathcal{J}^δ stands for the P&L coming from holding the underlying, and \mathcal{J}^P reflects the price evolution of the option.

Figure I.4 shows the results of two examples. We remark that the hedging portfolio computed with $\mathcal{N}\mathcal{N}_{SPX}$ can follow more closely the market price of options than the simple BS approach.




Figure I.4: Hedging P&L on two SPX calls with maturity 16 June 2017, strike 2425 (left), and maturity 16 February 2018, strike 2750 (right).

2 Chapter III: Multi-asset market making under quadratic rough Heston

The QRH model shows promising results on jointly modeling the SPX derivatives, including VIX futures, SPX and VIX options. As shown above, deep learning can provide efficient model calibration solutions and can also help computing hedging quantities. This motivates us to consider a multi-asset market making problem on the derivatives of SPX.

2.1 Problem description

Given the Markovian nature of the model defined in (I.4)-(I.6), the price of many derivatives of SPX at time t , *e.g.* VIX futures, vanilla SPX and VIX options, can be obtained as a function of $X_t := (S_t, Z_t^1, \dots, Z_t^n)$. Let $P(t, X)$ denote this quantity. Under regularity conditions, we can write

$$dP(t, X) = \delta_t dS_t,$$

where

$$\delta_t = \frac{\partial P(t, X)}{\partial S} + \frac{\eta}{S} \sum_{i=1}^n \frac{\partial P(t, X)}{\partial Z^i}.$$

In the previous section, we show how to estimate efficiently δ_t of SPX options with neural networks. The same approach can be extended to VIX futures and other vanilla options.

We consider a market making problem on a multi-asset basket over the period $[0, T]$. For each asset $j \in \{1, \dots, d\}$, at each point in time, the market maker decides whether to make a market at the limits P_t^j plus/minus one-half tick size, denoted by D^j . Let $l_t^{j,b}, l_t^{j,a} \in \{0, 1\}$ represent the decisions concerning respectively the bid and ask side. The value 1 means the participation of the market maker with a limit order of constant size m^j , and 0 stands for the case where the order is placed far from the queue corresponding to the best limits so that the probability of execution is negligible. Let $(N_t^{j,b})_{t \in [0, T]}$ and $(N_t^{j,a})_{t \in [0, T]}$ denote respectively the processes modeling the number of transactions at the bid and ask side. Then the dynamics of the inventory process $(q_t^j)_{t \in [0, T]}$ of asset j is given by

$$dq_t^j = m^j dN_t^{j,b} - m^j dN_t^{j,a}.$$

We denote by $(\lambda_t^{j,b})_{t \in [0, T]}$ and $(\lambda_t^{j,a})_{t \in [0, T]}$ the intensity processes of $(N_t^{j,b})_{t \in [0, T]}$ and $(N_t^{j,a})_{t \in [0, T]}$ respectively, which satisfy

$$\lambda_t^{j,b} = l_t^{j,b} \Lambda^{j,b} \mathbb{1}_{\{q_t^j + m^j \leq Q^j\}}, \quad \lambda_t^{j,a} = l_t^{j,a} \Lambda^{j,a} \mathbb{1}_{\{q_t^j - m^j \geq -Q^j\}},$$

where Q^j stands for the maximum inventory of asset j that the market maker is willing to hold, which is a multiple of m^j without loss of generality. The constants $\Lambda^{j,b}$ and $\Lambda^{j,a}$ represent the execution probabilities at the bid and ask sides respectively. In real life, the chances of being executed for a limit order depend on its position in the queue, which is out of the scope of current work. Let $(\Pi_t)_{t \in [0, T]}$ be the process representing the mark-to-market value of the market maker's portfolio. Its dynamics can be written as

$$d\Pi_t = \sum_{j=1}^d \frac{D^j}{2} m^j (dN_t^{j,b} + dN_t^{j,a}) + q_t^j \delta_t^j dS_t.$$

As for the objective function, we consider a similar one to those used in [Cartea et al. \[2014, 2017, 2018\]](#), given by

$$\sup_{\substack{l^{1,b}, \dots, l^{d,b} \\ l^{1,a}, \dots, l^{d,a}}} \mathbb{E}[\Pi_T - \sum_{j=1}^d \frac{\kappa_j}{2} \int_0^T \sigma_t^2 (q_t^j \delta_t^j)^2 dt - \frac{\kappa}{2} \int_0^T \sigma_t^2 (\sum_{j=1}^d q_t^j \delta_t^j)^2 dt], \quad (\text{I.11})$$

with $\sigma_t = S_t \sqrt{V_t}$, $(\kappa_j)_{j=1, \dots, d}$, and κ nonnegative constants. It is easy to remark that the value function associated with (I.11) depends on $(2 + n + d)$ variables, *i.e.* $t, S, Z := (Z^1, \dots, Z^n)$ and $q := (q^1, \dots, q^d)$. In order to bypass this curse of dimensionality, we propose some approximations in the following.

2.2 Value function approximation

- The Hamilton-Jacobi-Bellman equation

We first suggest the following approximation to reduce the dimensionality coming from the multidimensional nature of Z :

Approximation I.1. We consider $(Z_t^i)_{t \in [0, T], i=1, \dots, n}$ and $(\delta_t^j)_{t \in [0, T], j=1, \dots, d}$ by their initial value, i.e.

$$\begin{aligned} Z_t^i &= Z_0^i \quad i = 1, \dots, n, \\ \delta_t^j &= \delta_0^j =: \delta^j \quad j = 1, \dots, d, \end{aligned}$$

for $t \in [0, T]$. Then we have $V_t = V_0$ for $t \in [0, T]$ by (I.2). We replace the dynamic of $(S_t)_{t \in [0, T]}$ by

$$dS_t = \mu dt + \sigma dW_t, \quad t \in [0, T],$$

where $\sigma := S_0 \sqrt{V_0}$ and we allow a constant drift μ .

This approximation is relevant, considering that the time horizon of the market making is usually less than several hours. The results in Figure I.4 and further tests on daily hedging indicate that it is reasonable to consider δ constant at the daily scale. And in practice, the market maker can always reset the algorithm with updated parameters in the case of significant market movements. We recall that this approximation is considered at the beginning of each market making period and the updates of δ across periods are ruled by the QRH model. Now the associated value function can be defined by

$$v : (t, q) \in [0, T] \times \mathcal{Q} \mapsto v(t, q),$$

where

$$\begin{aligned} v(t, q) = & \sup_{\left(\begin{smallmatrix} l_s^{1,b}, \dots, l_s^{d,b} \\ l_s^{1,a}, \dots, l_s^{d,a} \end{smallmatrix} \right)_{s \in [t, T]}} \mathbb{E}_{t, q} \left[\int_t^T \sum_{j=1}^d \left(\sum_{k=a, b} \frac{D^j}{2} m^j l_s^{j,k} \Lambda^{j,k} \mathbb{1}_{\{q_s^j + \phi^k m^j \in \mathcal{Q}^j\}} + q_s^j \delta^j \mu \right) ds \right. \\ & \left. - \sum_{j=1}^d \frac{\kappa_j}{2} \sigma^2 \int_t^T (q_s^j \delta^j)^2 ds - \frac{\kappa}{2} \sigma^2 \int_t^T \left(\sum_{j=1}^d q_s^j \delta^j \right)^2 ds \right], \end{aligned} \quad (\text{I.12})$$

with $\mathcal{Q}^j = \{-Q^j, -Q^j + m^j, \dots, Q^j\}$, $\mathcal{Q} = \prod_{j=1}^d \mathcal{Q}^j$ and

$$\phi^k = \begin{cases} +1 & \text{if } k = b, \\ -1 & \text{if } k = a. \end{cases}$$

The HJB equation associated with (I.12) can be easily established as

$$\begin{aligned} 0 = & -\frac{\partial v}{\partial t} - \mu \sum_{j=1}^d q^j \delta^j + \sum_{j=1}^d \frac{\kappa_j}{2} \sigma^2 (q^j \delta^j)^2 + \frac{\kappa}{2} \sigma^2 \left(\sum_{j=1}^d q^j \delta^j \right)^2 \\ & - \sum_{j=1}^d \sum_{k=a, b} \mathbb{1}_{\{q^j + \phi^k m^j \in \mathcal{Q}^j\}} m^j H^{j,k} \left(\frac{v(t, q) - v(t, q + \phi^k m^j e^j)}{m^j} \right) \end{aligned} \quad (\text{I.13})$$

with terminal condition

$$v(T, q) = 0, \quad (\text{I.14})$$

where $H^{j,k}(p) = \Lambda^{j,k} \mathbb{1}_{\{p \leq \frac{D^j}{2}\}} (\frac{D^j}{2} - p)$, and $\{e^j\}_{j=1}^d$ is the canonical basis of \mathbb{R}^d . Following the same methodologies as Guéant [2017], the existence and uniqueness of the solution of (I.13) can be proved.

- Quadratic approximation

Inspired by Bergault et al. [2021], we approximate the Hamiltonian functions involved in (I.13) with quadratic ones. Then the resulting value function can be written as a quadratic function of q . More precisely, we replace $H^{j,k}(p)$ with the following approximation:

$$\hat{H}^{j,k}(p) = \Lambda^{j,k} \left(\frac{\alpha^j}{2} p^2 - p + \frac{D^j}{2} \right)$$

where $\alpha^j > 0$. Then in the case without maximal inventory constraints for individual assets, *i.e.* $Q^j = +\infty$, $j \in \{1, \dots, d\}$, we obtain the following asymptotic formula for the approximated value function \hat{v} :

$$\hat{v}(0, q) \stackrel{T \rightarrow +\infty}{\sim} -q^T A q - q^T B,$$

where

$$\begin{aligned} A &= \frac{\sigma}{2} \sqrt{\kappa} \Gamma, \\ B &= -D_+^{-\frac{1}{2}} \hat{A}^+ D_+^{\frac{1}{2}} \delta \mu - D_+^{-\frac{1}{2}} \hat{A} \hat{A}^+ D_+^{-\frac{1}{2}} (V_- + \frac{\sigma}{2} \sqrt{\kappa} D_- \mathcal{D}(\Gamma)), \\ D_+ &= \text{diag}((\Lambda^{1,b} + \Lambda^{1,a}) \alpha^1 z^1, \dots, (\Lambda^{d,b} + \Lambda^{d,a}) \alpha^d z^d), \\ D_- &= \text{diag}((\Lambda^{1,b} - \Lambda^{1,a}) \alpha^1 z^1, \dots, (\Lambda^{d,b} - \Lambda^{d,a}) \alpha^d z^d), \\ V_- &= ((-\Lambda^{1,b} + \Lambda^{1,a}) z^1, \dots, (-\Lambda^{d,b} + \Lambda^{d,a}) z^d), \\ \Gamma &= D_+^{-\frac{1}{2}} (D_+^{\frac{1}{2}} \Sigma D_+^{\frac{1}{2}})^{\frac{1}{2}} D_+^{-\frac{1}{2}}, \quad \hat{A} = \sqrt{\kappa} V (D_+^{\frac{1}{2}} \Sigma D_+^{\frac{1}{2}})^{\frac{1}{2}}, \end{aligned}$$

\hat{A}^+ is the Moore-Penrose inverse of \hat{A} and \mathcal{D} is the linear operator mapping a matrix onto the vector of its diagonal coefficients. As suggested in Bergault et al. [2021], the market maker can take actions following the asymptotic value function. In our case, for any $(j, k) \in \{1, \dots, d\} \times \{a, b\}$, the following market making decisions can be obtained:

$$\tilde{l}^{j,k*}(q) = \mathbb{1}_{\{q^j + \phi^k m^j \in \mathcal{Q}^j, 2\phi^k (e^j)^T A q + m^j (e^j)^T A e^j + \phi^k (e^j)^T B \leq \frac{D^j}{2}\}}. \quad (\text{I.15})$$

2.3 Numerical results

We conduct simulations on synthetic data with different risk aversion preferences κ . We take $(\kappa_j)_{j=1, \dots, d} = \frac{\kappa}{2}$ for the sake of simplicity. Given certain κ , we simulate 5000 paths and the market maker takes actions by following (I.15) for each path. The mean and standard deviation of all resulted Π_T are evaluated. Figure I.5 shows the results with a basket of six assets. Our method, named as *multi-asset*, performs clearly better than the *uni-asset* approach, where we perform the risk control on each individual asset without considering the risk of portfolio level, *i.e.* $\kappa = 0$. This exemplifies the effect of risk mutualization among the considered assets.




Figure I.5: Mean performance and risk of the uni-asset and multi-asset market making strategies with varying risk penalization parameter κ .

3 Chapter IV: On the universality of the volatility formation process: when machine learning and rough volatility agree

Under the RFSV model, future volatility can be predicted based on past volatilities through an explicit formula, which depends essentially on only one parameter, namely the Hurst index [Gatheral et al., 2018]. Given the widely observed fact that the Hurst index is around 0.1 for a wide range of financial assets [Bennedsen et al., 2021; Bolko et al., 2022; Fukasawa et al., 2022; Gatheral et al., 2018; Wang et al., 2023; Wu et al., 2022], a quasi-universal forecasting formula of realized volatility emerges. This motivates us to consider the existence of a universal endogenous volatility formation mechanism across assets, linking past volatilities and returns to future volatilities. We aim to capture this formation process by fitting highly performant volatility forecasting devices. We first apply a quasi-nonparametric approach with LSTM. We show that the universal LSTM trained on a pooled dataset made of hundreds of liquid stocks consistently outperforms other asset-specific parametric forecasting devices, including the simple autoregressive (AR), the heterogeneous autoregressive (HAR), and the RFSV methods. Re-training the network asset by asset or sector by sector does not improve performance. We then confirm the universality in question from a parametric perspective. A parsimonious parametric forecasting device incorporating the main features of the RFSV and QRH models with fixed parameters results in the same level of performance as the universal LSTM.

3.1 Realized volatility forecasting devices

Let $(\sigma_{t-p})_{p=1,2,\dots}$ be past volatilities, we evaluate the following parametric volatility forecasting methods as benchmarks:

- AR(p):

$$\hat{\sigma}_t = \alpha_0 + \sum_{j=1}^p \beta_j \sigma_{t-j},$$

- HAR:

$$\hat{\sigma}_t = \alpha_0 + \beta_1 \sigma_{t-1} + \beta_2 \frac{1}{5} \sum_{j=1}^5 \sigma_{t-j} + \beta_3 \frac{1}{22} \sum_{j=1}^{22} \sigma_{t-j},$$

- RFSV:

$$\hat{\sigma}_t = \nu \exp(\widehat{\log \sigma}_t), \tag{I.16}$$

where

$$\widehat{\log \sigma}_t = \frac{\cos(H\pi)}{\pi} \int_{-\infty}^{t-1} \frac{\log \sigma_s}{(t-s+1)(t-s)^{H+1/2}} ds,$$

$$\nu = \exp\left(\frac{\Gamma(3/2-H)}{2\Gamma(H+1/2)\Gamma(2-2H)} \nu^2\right).$$

The parameters $\alpha_0, \beta_j \in \mathbb{R}$ in the AR and HAR methods can be estimated via ordinary least squares. As for the RFSV method, ν is a positive constant and $H < \frac{1}{2}$ measuring the smoothness of volatility series. We follow the methodology introduced in [Gatheral et al. \[2018\]](#) to estimate H and ν . The above integral is truncated to some length of history and then approximated through a Riemann sum. We refer to [Wang et al. \[2022\]](#) for a recent discussion on alternative forecasting methods of the fBm. As for the nonparametric approach based on LSTM, we train the network with samples containing the most recent L histories, as illustrated in Figure I.6.

3.2 Capturing universality with LSTM

The LSTM network is trained on the pooled dataset of over 800 liquid stocks of the US market, while the other parametric forecasting devices are fitted stock by stock. Figure I.7 summarizes the out-of-sample MSE relative to the HAR model. We remark:

- the simple AR(22) model, which imposes no structural constraint on the autoregressive coefficients, underperforms the HAR model. This is understandable as it is more likely to overfit,
- the RFSV outperforms the HAR model in terms of prediction error. It is remarkable as it involves only two free parameters, *i.e.* ν and H . It benefits from the special parametrization of the weights of past volatilities depending on H ,




Figure I.6: Structure of an LSTM cell (left) and simplified computational graph of the network based on LSTM (right).

- the nonparametric models, $LSTM_{var}^{us}$ and $LSTM_{ret}^{us}$ outperform the other parametric models, especially when we incorporate past returns data. This indicates that the potential universal volatility formation mechanism across assets, relating past volatilities and returns to future volatilities, allows us to calibrate a universal model based on all assets, where the risk of overfitting is reduced due to enriched realized scenarios.




Figure I.7: Boxplot showing each model's out-of-sample MSE relative to the HAR model for the stocks of the US market. $LSTM_{var}^{us}$ is trained based on past realized variance, and $LSTM_{ret}^{us}$ uses also past returns.

We fine-tune the universal LSTM on each individual stock to incorporate any potential asset-dependent features. As shown in Figure I.8, there is no universal reduction of forecasting

error after the parameters of the network are “adapted” locally to each stock. Moreover, the network trained based on the stocks of the US market can be directly applied on the stocks of the European market, with superior forecasting performance obtained as shown in Figure I.9. This suggests that the universality in question holds even across markets.




Figure I.8: Empirical distribution of out-of-sample MSE of the fine-tuned asset-specific LSTMs relative to the universal LSTM. The left corresponds to the case where all the parameters of the network can be updated. The right is where we only fine-tune the parameters of the output layer.

3.3 Uncovering the universal volatility formation process

We propose here a parsimonious parametric forecasting device to obtain similar level of performance to the universal LSTM. The main idea is to complete the RFSV method with an element encoding the feedback effect of past returns inspired by the QRH model. More precisely, we use the following forecast:

$$(1 - \lambda)\hat{\sigma}^{RFSV} + \lambda\hat{\sigma}^{QRH} \quad (\text{I.17})$$

with $\lambda \in (0, 1)$, $\hat{\sigma}^{RFSV}$ is the RFSV forecast given by (I.16), $\hat{\sigma}^{QRH}$ is given by

$$(\hat{\sigma}_t^{QRH})^2 = a(Z_{t-1} - b)^2 + c,$$

$$Z_t = \sum_{i=1}^n c_i^d Z_{i,t},$$

$$Z_{i,t} = e^{-\gamma_i^d} Z_{i,t-1} + r_t, \quad Z_{i,0} = 0, \quad i = 1, \dots, n,$$

where $a, b, c > 0$ and $(c_i^d, \gamma_i^d)_{i=1, \dots, n}$ are given by the same multi-factor approximation of the rough kernel function $K(t) := \frac{t^{H-1/2}}{\Gamma(H+1/2)}$ as in Rosenbaum & Zhang [2022a]. We take the following parameters for (I.17):

$$H = 0.059, v = 1.03, a = 0.065, b = 0.74, c = 0.55.$$





Figure I.9: Boxplot showing models' out-of-sample performances relative to the HAR model for the stocks of the European market. Note that LSTM^{EU} and LSTM^{US} are trained on stocks of the European and US markets respectively.

They are in fact the overall calibration results when we fit respectively the RFSV and QRH forecasting devices on all the stocks of the US market. As shown in Figure I.10, the parsimonious forecasting device with the above fixed parameters can obtain similar level performance to the LSTM network with $\lambda \simeq 0.1$. From a parametric perspective, this confirms again the relevant universality, showing that the main features of this universal volatility formation process can be well described by the rough volatility paradigm boosted with the feedback effect of price trends.



(a) US market

(b) European market

Figure I.10: Out-of-sample performance of the forecast $(1 - \lambda)\hat{\sigma}^{RFSV} + \lambda\hat{\sigma}^{QRH}$ relative to the universal LSTM forecast.

4 Chapter V: Towards systematic intraday news screening: a liquidity-focused approach

As financial markets become increasingly competitive, practitioners need to develop automatic news sentiment evaluation tools to anticipate the impact of news releases on financial asset prices. The presence of many neutral news articles, that do not imply directional views on the associated asset value, will result in biased model learning if one fits a model on the whole unscreened news dataset. We propose a two-step screening approach to pick out impactful news samples, that are thought to have influenced considerably market activities after being released and thus contain sentiment-charged features in the texts. First, we monitor the asset intraday liquidity conditions indicated by several variables, including volatility σ , trading volume V , bid-ask spread ϕ , and book size B . The jump model introduced in [Bemporad et al. \[2018\]](#) is used to categorize the liquidity conditions into two distinct modes. Second, we link the detected liquidity mode switches, from that associated with lower volatility and turnover to the other one with impactful news releases. In other words, all the news articles published around these liquidity changes are taken as impactful. To illustrate the effect of our method, we fit two naive Bayes classifiers (NBCs) respectively on the original news dataset and the screened one trying to evaluate news sentiments. Despite the much reduced size of the screened training set, out-of-sample tests show that the latter performs significantly better when being applied to predict short-term asset price movements, suggesting that the proposed approach can exclude effectively irrelevant news samples from the training set and lead to more effective news sentiment learning. Our numerical tests are conducted on the data of S&P 500 components to avoid results biased by small-cap names.

4.1 Liquidity mode fitting

In this part, we apply the jump model to output mode sequences indicating the intraday liquidity conditions of the assets of interest. For each stock-day pair (S, d) with $S = 1, \dots, N$ and $d = 1, \dots, D$, after removing the data of the first and last 15 minutes of the daily core trading period, we measure $x_t^{S,d} := (\phi_t^{S,d}, V_t^{S,d}, \sigma_t^{S,d}, B_t^{S,d})$ for the t -th 5-minute bin, where $t = 1, \dots, T$ and $T = 72$. We apply some preprocessing procedures on x to remove the well-known intraday seasonalities and regulate the observations of different stocks to have the same scale. Given the preprocessed $X^{S,d} := (x_t^{S,d})_{t=1, \dots, T}$, we aim at fitting a mode sequence $M^{S,d} := (m_t^{S,d})_{t=1, \dots, T}$ describing the dynamics of liquidity conditions. For ease of interpretation, we allow two liquidity modes only. One represents the ‘‘calm’’ state at which the market stays for most of time and the other describes the regime driven by some exogenous events. Thus we have $m_t \in \{1, 2\}$. Every day, we conduct liquidity mode fitting with the following particular objective function:

$$\min_{\Theta^{S,d}, M^{S,d}} J(X^{S,d}, \Theta^{S,d}, M^{S,d}) = \sum_{t=1}^T \|x_t^{S,d} - \theta_{m_t^{S,d}}^{S,d}\|_2^2 + \lambda \sum_{t=1}^{T-1} \mathbb{1}_{m_t^{S,d} \neq m_{t+1}^{S,d}}, \quad (\text{I.18})$$

where $\Theta^{S,d} := (\theta_1^{S,d}, \theta_2^{S,d}) \in (\mathbb{R}^4)^2$, and λ is some positive constant. Therefore, the results of minimization are a sequence of liquidity modes associated with each time bin and two vectors of dimension four. The time bins showing similar liquidity characteristics are tagged with

the same mode, which is represented by θ_1 or θ_2 . When $\lambda > 0$, the mode switches between two adjacent time bins are penalized since they will bring extra loss. In this way, some mode persistence is obtained and only significant changes in liquidity conditions are detected. Note that when $\lambda = 0$, we get the same objective as the classical K-means clustering method, which simply minimizes within-cluster variances.

For each stock-day pair (S, d) , the mode represents of dimension eight are obtained based on $T = 72$ observations, which implies some potential risk of overfitting. To reduce this risk, we expand the task on the observation sequences of all stocks of the same day. More precisely, on day d we run the fitting on the pooled set of sequences defined by $X^d := \{(x_t^{S,d})_{t=1,\dots,T}\}_{S=1,\dots,N}$. All stocks are assumed to have the same mode represents, denoted by $\Theta^d := (\theta_1^d, \theta_2^d)$. For ease of notation, we ignore the superscript d and the actual fitting objective function reads then

$$\min_{\Theta, M} J(X, \Theta, M) = \sum_{S=1}^N \left(\sum_{t=1}^T \|x_t^S - \theta_{m_t^S}\|_2^2 + \lambda \sum_{t=1}^{T-1} \mathbb{1}_{m_t^S \neq m_{t+1}^S} \right), \quad (\text{I.19})$$

where $M := \{(m_t^S)_{t=1,\dots,T}\}_{S=1,\dots,N}$. Problem (I.19) can be solved with a similar approach to the expectation-maximization algorithm that alternates minimization with respect to Θ and M .




Figure I.11: Historical evolution of fitted mode represents θ_1 and θ_2 when $\lambda = 0.5$.

Denoting the fitted liquidity mode associated with a lower level of volatility as *Mode 1* and the other as *Mode 2*, we plot the dynamics of the fitted mode means from 2017-01-01 to 2019-12-31

in Figure I.11. Interestingly, the two modes are well distinct in terms of volatility level and trading volume. It is natural to consider that *Mode 2* is mainly driven by some exogenous events given its high volatility level and increased trading activities.

4.2 News screening and sentiment learning

- Methodology




Figure I.12: Illustrative scenarios when a news release is considered impactful.

We apply the liquidity mode fitting results introduced above to identify “true” impactful news releases. Let the set of news released during the considered period be denoted by \mathcal{D} . Given a piece of news $a \in \mathcal{D}$ concerning stock S , we consider it impactful if the liquidity condition of S switches from *Mode 1* to *Mode 2* shortly before or after the news publication. As the scenarios shown in Figure I.12, we take a as impactful since we observe an abrupt change of liquidity modes. We let the set of the selected impactful news be \mathcal{D}_λ , where λ is the mode switch penalization parameter defined in (I.19). Denoting the market detrended return of S over h minutes after the publication of a by r_a^h , we let $r^{h,k}$ and $r^{h,100-k}$ be the k -th and $(100-k)$ -th percentile of $\mathcal{R}^h := \{r_a^h | a \in \mathcal{D}\}$ respectively. We define the following set:

$$\mathcal{Z}_{h,k}^- := \{a | r_a^h \leq r^{h,k}, a \in \mathcal{D}\} \quad \text{and} \quad \mathcal{Z}_{h,k}^+ := \{a | r_a^h \geq r^{h,100-k}, a \in \mathcal{D}\}.$$

Thus, Z^- and Z^+ are sets of bearish and bullish news releases respectively according to the amplitude of post-publication asset returns, which is the criterion commonly used in works such as Chen [2021]; Ding et al. [2014, 2015]; Hu et al. [2018]. In our approach, the sets of “true” impactful news releases are given by

$$\mathcal{N}_{\lambda,h,k}^- := \mathcal{D}_\lambda \cap \mathcal{Z}_{h,k}^- \quad \text{and} \quad \mathcal{N}_{\lambda,h,k}^+ := \mathcal{D}_\lambda \cap \mathcal{Z}_{h,k}^+.$$

Therefore, the proposed method takes also the changes of liquidity conditions into account. Based on a set of labeled samples $\mathcal{E} = \mathcal{Z}^+ \cup \mathcal{Z}^-$ or $\mathcal{E} = \mathcal{N}^+ \cup \mathcal{N}^-$, then we fit a multi-variate Bernoulli NBC to capture sentiment-contributed words, see Chapter V for more details. For news release a , the fitted classifier gives the probabilities of a being positive and negative,

denoted by $P_{\mathcal{E}}^+(a)$ and $P_{\mathcal{E}}^-(a)$ respectively. We define thus the sentiment score $F \in [-1, 1]$ of a by

$$F(a)|_{\mathcal{E}} := P_{\mathcal{E}}^+(a) - P_{\mathcal{E}}^-(a).$$

We are particularly interested in the predictive power of F on short-term asset price movements. A significantly high (low) F is expected to be followed by a positive (negative) price drift of the associated stock.

- Numerical results

The numerical study is based on Bloomberg news headlines¹. We fit two NBCs on $\mathcal{Z}_{15,10}^+ \cup \mathcal{Z}_{15,10}^-$ and $\mathcal{N}_{0.5,15,10}^+ \cup \mathcal{N}_{0.5,15,10}^-$ respectively, which are then applied on all the intraday news samples released during 2020-01-01 ~ 2021-12-31 to output corresponding sentiment scores. Figure I.13 shows the average short-term post-publication price drifts of the news releases associated with distinct sentiment scores allocated by the two classifiers. We remark clearly that the news dataset screened by our approach leads to better predictive performance of the news classifier.




Figure I.13: Average post-publication stock return for the news releases with significantly positive and negative sentiment scores assigned by the NBCs fitted on different training sets. The left represents the case with news screening solely based on return information. The right corresponds to our approach. Only the samples whose scores fall in the top/bottom 10% range are shown for ease of comparison. The shadow parts represent the empirical standard deviation of the average return.

Note that with $\lambda = 0.5$, the size of $\mathcal{N}_{0.5,15,10}^+ \cup \mathcal{N}_{0.5,15,10}^-$ is only about one-tenth of $\mathcal{Z}_{15,10}^+ \cup \mathcal{Z}_{15,10}^-$. To verify that our selection method is not equivalent to filtering news releases by

¹<https://www.bloomberg.com/professional/product/event-driven-feeds/>

the associated post-publication stock returns, we repeat the same test on $\mathcal{Z}_{15,1}^+ \cup \mathcal{Z}_{15,1}^-$. As shown in the left subfigure of Figure I.14, with a similar count of training samples to the case of $\mathcal{N}_{0.5,15,10}^+ \cup \mathcal{N}_{0.5,15,10}^-$, the prediction performance now is largely degraded. These results indicate that our method presents a more efficient way to identify impactful news releases than the approach based solely on return information. We show also the results of Bloomberg news sentiment scores given by the company’s proprietary classification algorithm in Figure I.14. Even though the classification model tested here, *i.e.* NBC, is very simple, slightly better performance is remarked than that with Bloomberg scores. Further numerical tests show that the results of our approach are robust to the involved parameters, *i.e.* λ, h and k .




Figure I.14: Average post-publication stock return for the news releases with significantly positive and negative sentiment scores assigned by the NBC fitted on the news samples associated with extreme realized return falling in the top/bottom 1% quantile (left), and the case with the sentiment scores given by Bloomberg (right).

Introduction (en français)

Motivations

Le paradigme de la volatilité rugueuse, largement accepté, peut reproduire tous les faits stylisés des séries temporelles de volatilité et s'adapter remarquablement bien à la forme des courbes de volatilité implicite et d'asymétrie à la monnaie. En particulier, le modèle *quadratic rough Heston* (QRH) récemment introduit montre des résultats prometteurs sur la modélisation jointe des options de SPX/VIX, où d'autres modèles de volatilité stochastique à temps continu sont en difficulté depuis longtemps. Comme le modèle est non-markovien et non-sémimartingale, il est difficile de le calibrer en basant sur des simulations Monte Carlo classiques. Nous nous intéressons donc à :

Question 5. *L'apprentissage statistique peut-il fournir des solutions alternatives pour calibrer le modèle QRH ?*

Nous proposons une version approximative multifactorielle du modèle original. L'approximation rend le modèle markovien et peut reproduire la dynamique *rough* aux échelles de temps qui nous intéressent. Des réseaux de neurones profonds sont appliqués pour apprendre les fonctions de tarification des options SPX/VIX sur la base de données simulées, à partir desquelles une calibration jointe peut être effectuée de manière efficace. Nous montrons que le modèle calibré peut très bien reproduire les surfaces de volatilité implicite du SPX et du VIX. En outre, les réseaux formés peuvent être utilisés pour calculer les quantités de couverture instantanément. Ceci est très utile pour le contrôle des risques lorsque les praticiens détiennent des produits dérivés du SPX, par exemple dans le cadre d'activités de *market making*, ce qui nous incite à nous pencher sur la question suivante :

Question 6. *Comment concevoir et résoudre un problème de market making multi-actifs dans le cadre du modèle QRH ?*

Nous formulons un problème de market making optimale sur de multiples dérivés du SPX, y compris les contrats à terme SPX, les contrats à terme VIX, les options SPX et VIX. Pour maximiser le profit venant de la capture des spreads, le *market maker* décide de placer ou non des ordres limités dans les files d'attente des meilleures limites d'achat et de vente, tout

en contrôlant le risque de détention du portefeuille. Cela conduit à un problème de contrôle stochastique de haute dimension. Nous surmontons la malédiction de la dimensionnalité avec plusieurs approximations, dont la précision et la pertinence sont démontrées par des expériences numériques. Des solutions asymptotiques de forme fermée peuvent alors être obtenues.

Une caractéristique clé du modèle QRH contribuant à la modélisation de la volatilité implicite est l'inclusion de ce que l'on appelle l'effet Zumbach, qui décrit essentiellement l'effet de rétroaction de l'évolution des prix sur la volatilité. Cette caractéristique peut également être utile pour modéliser la volatilité réalisée. De nombreux travaux ont étudié les rôles joués par les réalisations passées du marché sur la prévision de la volatilité future. La volatilité rugueuse étant désormais largement acceptée, nous nous intéressons particulièrement à la question suivante :

Question 7. *Le mécanisme de formation de la volatilité reliant les rendements et les volatilités passés aux volatilités futures est-il universel pour tous les actifs ?*

Nous abordons cette question d'un point de vue à la fois non-paramétrique et paramétrique. D'une part, nous entraînons un réseau LSTM basé sur un ensemble de données regroupées couvrant des centaines d'actions liquides, dans le but de prévoir la prochaine volatilité quotidienne pour toutes les actions. Nous montrons que ce LSTM universel surpasse d'autres modèles spécifiques à un actif ou à un secteur. Cette performance se maintient même entre les différents marchés. D'autre part, un dispositif de prévision paramétrique parcimonieux, combinant les principales caractéristiques des modèles *rough fractional stochastic volatility* (RFSV) et QRH avec des paramètres fixes, présente le même niveau de capacité de prévision que le LSTM universel. Ces résultats confirment donc l'existence de l'universalité en question.

Outre la part endogène induite par les activités passées du marché par le biais de ce mécanisme universel, la volatilité du marché est également due à des événements exogènes, dont la principale source est constituée par les publications de nouvelles. Les nouvelles inattendues sont susceptibles d'être suivies de mouvements importants sur le marché. Pour évaluer efficacement l'impact d'une nouvelle sur le prix de l'actif associé, les investisseurs institutionnels doivent développer des méthodes d'évaluation automatique du sentiment d'actualité basées sur des données historiques. Un grand nombre d'articles de presse sont publiés chaque jour et la plupart d'entre eux sont neutres, c'est-à-dire que les points de vue transférés ne sont ni haussiers ni baissiers. Afin de construire un prédicteur de sentiment robuste, il est raisonnable de l'utiliser principalement sur les échantillons de nouvelles qui ont entraîné des activités considérables sur le marché. Pour cela, nous demandons

Question 8. *Comment filtrer efficacement les nouvelles à fort impact à partir d'un vaste ensemble de données brutes ?*

Naturellement, les nouvelles haussiers (baissiers) sont susceptibles d'entraîner des variations positives (négatives) du prix des actifs. Cependant, la sélection de nouvelles uniquement sur la base des informations sur les rendements après la publication n'est pas solide compte tenu du

faible rapport signal/bruit des rendements. Nous pensons que les signaux exogènes ayant un impact induiront également des changements significatifs dans les conditions de liquidité, tels que des volatilités plus importantes et un volume de transactions plus élevé. Nous proposons donc une approche qui consiste à détecter systématiquement les changements de mode de liquidité et à considérer les nouvelles publiées autour de ces changements comme ayant un impact.

Schéma

Dans le Chapitre II, nous répondons à la Question 5 en élaborant une procédure de calibration efficace basée sur des réseaux neuronaux profonds. Nous transformons d'abord le modèle QRH original en une version approximative multifactorielle. Nous générons un grand nombre de prix d'options SPX/VIX avec des paramètres de modèle échantillonnés de manière aléatoire. Deux réseaux sont ensuite entraînés basés sur ces échantillons pour apprendre les fonctions de tarification à partir des paramètres du modèle aux surfaces de volatilité implicite SPX/VIX respectivement. Les résultats de la calibration jointe sont donnés en minimisant la distance entre les volatilités implicites observées et les sorties des deux réseaux. La précision de la calibration est évaluée sur des données simulées et sur des données de marché. Dans le cadre de notre modèle, un seul mouvement brownien est impliqué, et les options européennes peuvent être couvertes parfaitement avec seulement le SPX. Nous montrons comment calculer les quantités de couverture avec les réseaux entraînés dans la pratique. Pour des raisons de comparaison, une autre méthode de couverture basée sur l'apprentissage automatique différentiel est également introduite.

Nous donnons une solution pratique à la Question 6 au Chapitre III. Étant donné les résultats prometteurs du modèle QRH sur la modélisation jointe des sourires de volatilité implicite de SPX/VIX, nous considérons un panier de market making composé de nombreux dérivés de SPX. Les actifs considérés sont essentiellement des actifs de grande tick. Le market maker essaie de maximiser son profit provenant de la capture des spreads en envoyant des ordres limités aux meilleures limites bid/ask, et contrôle son risque d'inventaire, qui peut être principalement expliqué par la variation du SPX. Ce problème est formulé comme un problème de contrôle stochastique à haute dimension. Pour le résoudre efficacement, nous maintenons les facteurs de volatilité stochastiques constants pour chaque période de market making. Il est prouvé que l'équation HJB associée au problème d'optimisation résultant a une solution unique. L'argument de vérification classique peut être facilement établi. Pour plus de praticité, nous appliquons l'approximation quadratique aux fonctions hamiltoniennes. Une solution asymptotique de forme fermée peut alors être obtenue. Les simulations numériques montrent que la solution approximée est très proche de celle donnée par le schéma d'Euler classique.

Nous abordons la Question 7 au Chapitre IV en évaluant le pouvoir prédictif des rendements et des volatilités passés sur les volatilités futures. Plusieurs méthodes paramétriques basées sur des modèles sont étudiées, notamment les modèles autorégressifs hétérogènes (HAR), les

modèles de volatilité stochastique fractionnelle rugueuse (RFSV) et une méthode quasi non-paramétrique basée sur la LSTM. La LSTM entraînée sur les données de centaines d’actions américaines liquides est la plus performante. L’entraînement spécifique à un actif ou à un secteur n’améliore pas les résultats. Des performances supérieures similaires sont obtenues pour des actifs ne faisant pas partie de l’ensemble d’entraînement, même ceux d’un marché différent. Ces observations suggèrent l’existence d’un mécanisme universel de formation de la volatilité. Nous montrons ensuite qu’un dispositif de prévision paramétrique parcimonieux avec cinq paramètres fixés globalement, incorporant les principales caractéristiques des modèles RFSV et QRH, peut reproduire le même niveau de performance que la LSTM universelle, ce qui indique que l’universalité apprise peut être bien décrite par le paradigme de la volatilité rugueuse renforcé par ce que l’on appelle l’effet Zumbach.

Pour répondre à la Question 8 du Chapitre V, nous proposons une approche en deux étapes. Tout d’abord, nous mesurons en continu les conditions de liquidité intrajournalière des actifs financiers à l’aide de quatre variables, à savoir la volatilité, le volume des transactions, l’écart entre les cours acheteur et vendeur et le volume disponible dans le carnet d’ordres limités. Nous classons les états de liquidité en deux modes en utilisant une méthode non-paramétrique appelée le modèle de saut. Grâce à des tests numériques sur les composants du S&P500, nous constatons que les deux modes trouvés sont distincts l’un de l’autre en termes de volatilité et de volume de transactions. Ensuite, nous nous intéressons de plus près aux changements de mode, de celui qui est associé à une volatilité et à des volumes de transactions plus faibles à l’autre mode. Nous considérons tous les articles de presse publiés autour de ces moments de saut d’état de liquidité comme ayant un impact. Pour illustrer l’effet de notre approche, nous avons testé deux classificateurs de Bayes naïfs (NBC) sur l’ensemble de données brutes et sur l’ensemble de données filtrées. Ce dernier montre un pouvoir prédictif nettement meilleur sur les mouvements de prix des actions à court terme, ce qui indique que notre procédure de filtrage des nouvelles peut diminuer l’impact des nouvelles neutres et donc rendre plus efficace la capture des caractéristiques contribuant au sentiment.

Nous donnons maintenant un aperçu rapide des principaux résultats obtenus dans cette thèse.

1 Chapitre II: Calibration du modèle quadratic rough Heston avec l’apprentissage profond

Le modèle *rough Heston* peut reproduire très bien la volatilité implicite de l’indice SPX [“el2018perfect”, n.d.; El Euch & Rosenbaum, 2019; El Euch et al., 2019]. Cependant, un fait stylisé des séries temporelles financières qui n’est pas pris en compte dans ce modèle est l’effet de rétroaction des tendances passées des prix sur la volatilité, appelé l’effet Zumbach. Dandapani et al. [2021] clarifier la définition de l’effet Zumbach et introduire des modèles de *super-Heston rough volatility* codant explicitement cet effet. En tant qu’exemple particulier des modèles de cette catégorie, le modèle QRH présenté dans Gatheral et al. [2020] montre sa capacité prometteuse à s’adapter conjointement aux surfaces de volatilité implicite du SPX/VIX. Afin de promouvoir davantage la praticabilité du modèle QRH, nous visons à

concevoir une procédure efficace de calibration du modèle dans ce chapitre. Inspirés par des travaux récents sur les applications de l'apprentissage profond en finance quantitative, voir par exemple [Hernandez \[2016\]](#); [Horvath et al. \[2019\]](#), nous appliquons des réseaux neuronaux profonds pour cette tâche. L'efficacité du modèle calibré pour reproduire conjointement les sourires SPX et VIX est illustrée par des expériences numériques. Il est intéressant de noter que, dans le cadre de notre modèle, les réseaux entraînés nous permettent également de couvrir les options européennes avec un calcul instantané des quantités de couverture.

1.1 Le modèle QRH et son approximation multifactorielle

Dans le modèle QRH, la dynamique du prix d'un actif S (ici le SPX) et de sa variance au comptant V sous une mesure neutre vis-à-vis du risque est la suivante

$$dS_t = S_t \sqrt{V_t} dW_t, \quad (\text{I.1})$$

$$V_t = a(Z_t - b)^2 + c, \quad (\text{I.2})$$

où W est un mouvement brownien, a, b, c sont des constantes positives et Z_t est défini comme

$$Z_t = \int_0^t \lambda \frac{(t-s)^{\alpha-1}}{\Gamma(\alpha)} (\theta_0(s) - Z_s) ds + \int_0^t \eta \frac{(t-s)^{\alpha-1}}{\Gamma(\alpha)} \sqrt{V_s} dW_s, \quad (\text{I.3})$$

pour $t \in [0, T]$. Ici, T est un horizon temporel positif, $\alpha \in (1/2, 1)$, $\lambda > 0$, $\eta > 0$ et $\theta_0(\cdot)$ est une fonction déterministe. [El Euch & Rosenbaum \[2019\]](#) montrent que les trajectoires de volatilité dans le cadre du modèle rough Heston ont presque sûrement une régularité Hölder $\alpha \sim \frac{1}{2} - \varepsilon$, pour tout $\varepsilon > 0$. Cela rappelle les observations de [Gatheral et al. \[2018\]](#) selon lesquelles la dynamique de la log-volatilité est similaire à celle d'un mouvement brownien fractionnaire avec un indice de Hurst d'ordre 0.1. Par rapport à d'autres modèles de volatilité stochastique courants, le modèle QRH n'implique qu'un seul mouvement brownien. Le processus Z_t encode l'information des rendements passés puisque nous avons $\sqrt{V_t} dW_t = dS_t / S_t$. Ensuite, par (I.2), V_t incorpore naturellement la version forte de l'effet Zumbach, comme suggéré par [Dandapani et al. \[2021\]](#). Comme dans le modèle rough Heston, le noyau fractionnaire $K(t) = \frac{t^{\alpha-1}}{\Gamma(\alpha)}$ dans (I.3) nous permet de générer une dynamique de volatilité rugueuse. Cependant, il rend le modèle QRH non-markovien et non-sémimartingale, ce qui entraîne des difficultés en matière de simulation et de calibration. Nous appliquons la méthode introduite dans [Abi Jaber & El Euch \[2019\]](#) pour remplacer $K(t)$ par la version approximée multifactorielle suivante

$$K^n(t) = \sum_{i=1}^n c_i^n e^{-\gamma_i^n t}, \quad n \geq 1,$$

où $(c_i^n)_{i=1, \dots, n}$ et $(\gamma_i^n)_{i=1, \dots, n}$ peuvent être choisis comme des fonctions explicites de α . Ce ne sont donc pas des paramètres libres à calibrer. Des paramétrages correctement choisis des $2n$ paramètres $(c_i^n, \gamma_i^n)_{i=1, \dots, n}$ en termes de α peuvent faire converger $K^n(t)$ vers $K(t)$ au sens L^2 lorsque n va à l'infini et les modèles d'approximation multifactorielle se comportent étroitement avec leurs homologues dans le paradigme de la volatilité rugueuse [[Abi Jaber, 2019](#); [Abi Jaber & El Euch, 2019](#)]. Nous pouvons donc proposer l'approximation multifactorielle suivante du modèle QRH:

$$dS_t^n = S_t^n \sqrt{V_t^n} dW_t, \quad V_t^n = a(Z_t^n - b)^2 + c, \quad (\text{I.4})$$

$$Z_t^n = \sum_{i=1}^n c_i^n Z_t^{n,i}, \quad (\text{I.5})$$

$$dZ_t^{n,i} = (-\gamma_i^n Z_t^{n,i} - \lambda Z_t^n) dt + \eta \sqrt{V_t^n} dW_t, \quad Z_0^{n,i} = z_0^i, \quad (\text{I.6})$$

avec $(z_0^i)_{i=1,\dots,n}$ quelques constantes. Afin d'éviter de prédéfinir une formulation spécifique de $\theta(\cdot)$, nous le discardons dans l'Équation (I.5) et nous considérons les valeurs initiales de $(z_0^i)_{i=1,\dots,n}$ comme des paramètres libres à calibrer à partir des données du marché, ce qui permet une plus grande flexibilité du modèle. En fait, étant donné une solution pour (I.4) - (I.6), nous pouvons réécrire l'Équation (I.6) comme suit

$$Z_t^{n,i} = z_0^i e^{-\gamma_i^n t} + \int_0^t e^{-\gamma_i^n(t-s)} (-\lambda Z_s^n ds + \eta \sqrt{V_s^n} dW_s). \quad (\text{I.7})$$

Puis, à partir de (I.5), nous obtenons

$$Z_t^n = g^n(t) + \int_0^t K^n(t-s) (b(Z_s^n) ds + \sigma(Z_s^n) dW_s), \quad (\text{I.8})$$

avec $g^n(t) = \sum_{i=1}^n z_0^i c_i^n e^{-\gamma_i^n t}$, $b(Z_t^n) = -\lambda Z_t^n$ et $\sigma(Z_t^n) = \eta \sqrt{a(Z_t^n - b)^2 + c}$. En prenant l'expectation sur les deux côtés de (I.8), nous obtenons

$$\mathbb{E}[Z_t^n] + \lambda \sum_{i=1}^n c_i^n \int_0^t e^{-\gamma_i^n(t-s)} \mathbb{E}[Z_s^n] ds = \sum_{i=1}^n z_0^i c_i^n e^{-\gamma_i^n t},$$

avec laquelle la dépendance de $\mathbb{E}[Z_t^n]$ sur $(z_0^i)_{i=1,\dots,n}$ peut être obtenue explicitement. Nous pouvons donc voir que $(z_0^i)_{i=1,\dots,n}$ nous permet d'encoder la "structure des termes" initiaux de Z_t^n . Par conséquent, il peut être compris comme un analogue de $\theta(t)$ pour le processus de variance dans le modèle rough Heston. En outre, l'unicité forte de la solution des Équations (I.7)- (I.8) peut être établie en vertu de la Proposition B.3 de [Abi Jaber & El Euch \[2019\]](#).

1.2 Calibration efficace avec l'apprentissage profond

- Méthodologie

Nous notons $\boldsymbol{\omega} := (\lambda, \eta, a, b, c) \in \Omega \subset R^5$, $\mathbf{z}_0 := (z_0^1, \dots, z_0^n) \in \mathcal{Z} \subset R^n$. Nous fixons $n = 10$ dans nos tests et nous supprimons l'étiquette n dans ce qui suit pour simplifier les notations. Nous commençons par générer un grand nombre d'échantillons synthétiques suivant (I.4) - (I.6), puis nous entraînons des réseaux de neurones sur la base des données simulées.

Pour chaque échantillon aléatoire $(\boldsymbol{\omega}, \mathbf{z}_0)$, nous générons des milliers de trajectoires aléatoires du SPX et du VIX à l'aide d'un schéma d'Euler explicite-implicite. Les prix de Monte-Carlo sont ensuite utilisés pour calculer IVS_{SPX}^{MC} et IVS_{VIX}^{MC} , qui représentent les surfaces de

volatilité implicite (IVS) des options SPX et VIX respectivement. Deux réseaux de neurones, désignés par \mathcal{NN}_{SPX} et \mathcal{NN}_{VIX} , sont entraînés sur les paires de données $\{(\boldsymbol{\omega}, \mathbf{z}_0), IVS_{SPX}\}$ et $\{(\boldsymbol{\omega}, \mathbf{z}_0), IVS_{VIX}\}$ respectivement. Ainsi, \mathcal{NN}_{SPX} et \mathcal{NN}_{VIX} consistent à apprendre les fonctions d'évaluation des options SPX et VIX dans le modèle QRH approximé multifactoriel, comme illustré à la Figure I.1.




Figure I.1: \mathcal{NN}_{SPX} et \mathcal{NN}_{VIX} approchent les fonctions entre les paramètres du modèle et les volatilités implicites du SPX et du VIX.

Compte tenu des observations IVS_{SPX} et IVS_{VIX} , le résultat de la calibration jointe est donné par la formule suivante

$$\hat{\boldsymbol{\omega}}, \hat{\mathbf{z}}_0 = \operatorname{argmin}_{\boldsymbol{\omega} \in \Omega, \mathbf{z}_0 \in \mathcal{Z}} \|\mathcal{NN}_{SPX}(\boldsymbol{\omega}, \mathbf{z}_0) - IVS_{SPX}\|_2^2 + \|\mathcal{NN}_{VIX}(\boldsymbol{\omega}, \mathbf{z}_0) - IVS_{VIX}\|_2^2,$$

où $\|\cdot\|_2$ représente la norme L^2 . Nous appliquons L-BFGS-B, un algorithme basé sur le gradient, pour résoudre ce problème d'optimisation. Les gradients nécessaires sont calculés directement avec \mathcal{NN}_{SPX} et \mathcal{NN}_{VIX} via la différenciation adjointe automatique (DAA).

- Résultats numériques

Nous effectuons la calibration du modèle sur les données du marché et obtenons ensuite les IVS des options SPX et VIX avec les paramètres calibrés par le biais de simulations de Monte-Carlo. Les Figures I.2 et I.3 donnent un exemple avec les données du 19 mai 2017. Notre modèle reproduit très bien la forme globale des IVS des options SPX et VIX en même temps. Les résultats du modèle se situent essentiellement entre les cours acheteur et vendeur des options VIX. En outre, d'excellents alignements sont obtenus en termes de asymétrie à la monnaie pour les options SPX.




Figure I.2: Volatilités implicites des options SPX pour le 19 mai 2017. Les volatilités du marché à l'achat et à la vente sont représentées respectivement par des points rouges et bleus. La ligne verte est la sortie du modèle avec la méthode de Monte-Carlo.

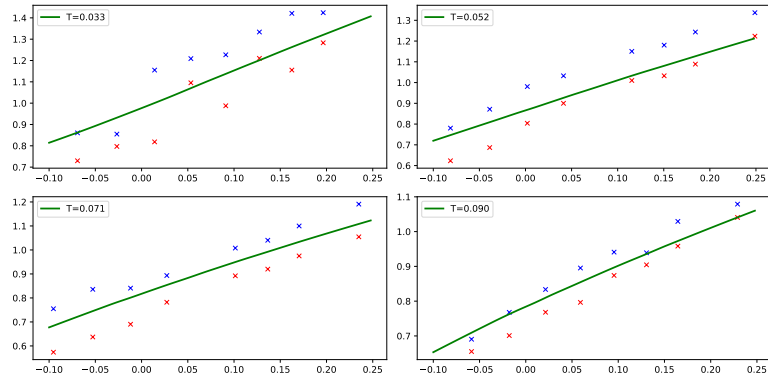


Figure I.3: Volatilités implicites des options VIX pour le 19 mai 2017. Les volatilités du marché à l'achat et à la vente sont représentées respectivement par des points rouges et bleus. La ligne verte est la sortie du modèle avec la méthode de Monte-Carlo.

1.3 Calcul de couverture

- Méthodologie

Dans le modèle défini en (I.4)-(I.6), les options européennes peuvent être couvertes avec seulement le SPX. Nous prenons l'option d'achat vanille SPX comme exemple illustratif. Soit $\mathbf{Z}_t := (Z_t^1, \dots, Z_t^{10})$ et $\mathbf{X}_t := (S_t, \mathbf{Z}_t)$. Étant donné le prix d'exercice K , l'échéance T et les paramètres du modèle $\boldsymbol{\omega}$, le prix de l'option d'achat SPX vanille à l'instant t est une fonction de \mathbf{X}_t . Soit $P_t(\mathbf{X}_t; K, T, \boldsymbol{\omega})$ cette quantité. Nous avons

$$P_t(\mathbf{X}_t; K, T, \boldsymbol{\omega}) \simeq P^{BS}(S_t, K, T - t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)),$$

où $P^{BS}(S, K, T, \sigma)$ est le prix du call européen selon le modèle Black-Scholes (BS) et $\sigma_{\mathcal{N}\mathcal{N}}^{k, T}(\boldsymbol{\omega}, \mathbf{Z}_t)$ est la volatilité implicite correspondant au prix d'exercice de log-moneyness k , et à l'échéance

T , calculée directement à partir de $\mathcal{N}\mathcal{N}_{SPX}$ avec $(\boldsymbol{\omega}, \mathbf{Z}_t)$ comme entrée. Nous pouvons avoir

$$\begin{aligned} dP^{BS}(S_t, K, T-t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)) &= \left(\frac{\partial P^{BS}}{\partial S_t} + \frac{\eta}{S_t} \sum_{i=1}^{10} \frac{\partial P^{BS}}{\partial Z_t^i} \right) (S_t, K, T-t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)) dS_t \\ &=: \delta_t dS_t, \end{aligned}$$

avec

$$\begin{aligned} \frac{\partial P^{BS}(S_t, K, T-t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t))}{\partial S_t} &\simeq \delta_{BS}(S_t, K, T-t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)) + \\ &v_{BS}(S_t, K, T-t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)) \frac{\partial \sigma_{\mathcal{N}\mathcal{N}}^{k, T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)}{\partial k} \Big|_{k=\log(K/S_t)} \frac{\partial \log(K/S_t)}{\partial S_t}, \end{aligned} \quad (\text{I.9})$$

$$\frac{\partial P^{BS}(S_t, K, T-t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t))}{\partial Z_t^i} \simeq v_{BS}(S_t, K, T-t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)) \frac{\partial \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)}{\partial Z_t^i}, \quad (\text{I.10})$$

où $\delta_{BS}(S, K, T, \sigma)$ et $v_{BS}(S, K, T, \sigma)$ sont le *Delta* et *Vega* respectivement dans le cadre du modèle BS. La quantité $\frac{\partial \sigma_{\mathcal{N}\mathcal{N}}^{k, T-t}}{\partial Z_t^i}$ correspond en fait à la dérivée des sorties de $\mathcal{N}\mathcal{N}_{SPX}$ par rapport à ses entrées, qui peut être obtenue instantanément avec la DAA intégrée. Le terme $\frac{\partial \sigma_{\mathcal{N}\mathcal{N}}^{k, T-t}(\boldsymbol{\omega}, \mathbf{Z}_t)}{\partial k}$ dans (I.9) peut être approximée par la différence finie $\frac{(\sigma_{\mathcal{N}\mathcal{N}}^{k+\delta_k, T-t} - \sigma_{\mathcal{N}\mathcal{N}}^{k-\delta_k, T-t})}{2\delta_k}(\boldsymbol{\omega}, \mathbf{Z}_t)$.

- Résultats numériques

Nous effectuons une couverture discrète avec un pas de temps Δt . Le profit et la perte (P&L) de la couverture à $t \in (0, T]$ est donné par

$$\mathcal{I}_t = \mathcal{I}_t^\delta - \mathcal{I}_t^P,$$

où

$$\begin{aligned} \mathcal{I}_t^\delta &= \sum_{k=0}^{\lfloor t/\Delta t \rfloor - 1} \delta_{t_k} \Delta S_{t_{k+1}} + \delta_{t_{\lfloor t/\Delta t \rfloor}} (S_t - S_{\lfloor t/\Delta t \rfloor \Delta t}), \\ \mathcal{I}_t^P &= P_t - P_0, \end{aligned}$$

avec $t_k = k\Delta t$, $\Delta S_{t_k} := S_{t_k} - S_{t_{k-1}}$, et P_t est le prix de l'option SPX à couvrir au temps t . Comme nous pouvons le voir, \mathcal{I}_t^δ représente le P&L provenant de la détention du sous-jacent, et \mathcal{I}_t^P reflète l'évolution du prix de l'option.

La Figure I.4 montre les résultats de deux exemples. Nous remarquons que le portefeuille de couverture calculé avec $\mathcal{N}\mathcal{N}_{SPX}$ peut suivre de plus près le prix du marché des options que l'approche BS simple.




Figure I.4: P&L de couverture de deux options d'achat SPX à échéance 16 juin 2017, strike 2425 (à gauche), et à échéance 16 février 2018, strike 2750 (à droite).

2 Chapitre III: Market-making multi-actifs dans le cadre de la méthode quadratique de rough Heston

Le modèle QRH montre des résultats prometteurs pour la modélisation jointe des dérivés du SPX, y compris les contrats à terme VIX, les options SPX et VIX. Comme nous l'avons montré ci-dessus, l'apprentissage profond peut fournir des solutions efficaces de calibration du modèle et peut également aider à calculer les quantités de couverture. Ceci nous motive à considérer un problème de market making multi-actifs sur les dérivés du SPX.

2.1 Description du problème

Étant donné la nature markovienne du modèle défini dans (I.4) - (I.6), le prix de nombreux dérivés du SPX à l'instant t , par exemple les contrats à terme VIX, les options vanilles de SPX et de VIX, peut être obtenu comme une fonction de $X_t := (S_t, Z_t^1, \dots, Z_t^n)$. Soit $P(t, X)$ pour désigner cette quantité. Sous des conditions de régularité, nous pouvons écrire

$$dP(t, X) = \delta_t dS_t,$$

où

$$\delta_t = \frac{\partial P(t, X)}{\partial S} + \frac{\eta}{S} \sum_{i=1}^n \frac{\partial P(t, X)}{\partial Z^i}.$$

Dans la section précédente, nous avons montré comment estimer efficacement le δ_t des options SPX à l'aide de réseaux de neurones. La même approche peut être étendue aux contrats à terme VIX et à d'autres options vanille de SPX.

Nous considérons un problème de market making sur un panier multi-actifs sur la période $[0, T]$. Pour chaque actif $j \in \{1, \dots, d\}$, à chaque instant, le market maker décide de faire ou non un marché aux limites P_t^j plus/moins une demi-taille de tick, notée par D^j . $l_t^{j,b}, l_t^{j,a} \in \{0, 1\}$ représentent les décisions concernant respectivement le côté acheteur et le côté vendeur. La valeur 1 signifie la participation du market maker avec un ordre limité de taille constante

m^j , et 0 représente le cas où l'ordre est placé loin de la file d'attente correspondant aux meilleures limites, de sorte que la probabilité d'exécution est négligeable. Laissons $(N_t^{j,b})_{t \in [0, T]}$ et $(N_t^{j,a})_{t \in [0, T]}$ désigner respectivement les processus modélisant le nombre de transactions du côté de l'offre et de la demande. La dynamique du processus d'inventaire $(q_t^j)_{t \in [0, T]}$ de l'actif j est alors donnée par

$$dq_t^j = m^j dN_t^{j,b} - m^j dN_t^{j,a}.$$

Nous désignons par $(\lambda_t^{j,b})_{t \in [0, T]}$ et $(\lambda_t^{j,a})_{t \in [0, T]}$ les processus d'intensité de $(N_t^{j,b})_{t \in [0, T]}$ et $(N_t^{j,a})_{t \in [0, T]}$ respectivement, qui satisfont à

$$\lambda_t^{j,b} = l_t^{j,b} \Lambda^{j,b} \mathbb{1}_{\{q_t^j + m^j \leq Q^j\}}, \quad \lambda_t^{j,a} = l_t^{j,a} \Lambda^{j,a} \mathbb{1}_{\{q_t^j - m^j \geq -Q^j\}},$$

où Q^j représente le stock maximal de l'actif j que le market maker est disposé à détenir, qui est un multiple de m^j sans perte de généralité. Les constantes $\Lambda^{j,b}$ et $\Lambda^{j,a}$ représentent les probabilités d'exécution du côté de l'offre et de la demande respectivement. Dans la réalité, les chances d'exécution d'un ordre à cours limité dépendent de sa position dans la file d'attente, ce qui n'entre pas dans le cadre du présent travail. Soit $(\Pi_t)_{t \in [0, T]}$ le processus représentant la valeur de marché du portefeuille du market maker. Sa dynamique peut s'écrire comme suit

$$d\Pi_t = \sum_{j=1}^d \frac{D^j}{2} m^j (dN_t^{j,b} + dN_t^{j,a}) + q_t^j \delta_t^j dS_t.$$

En ce qui concerne la fonction objective, nous considérons une fonction similaire à celles utilisées dans [Cartea et al. \[2014, 2017, 2018\]](#), donnée par

$$\sup_{\substack{l^{1,b}, \dots, l^{d,b} \\ l^{1,a}, \dots, l^{d,a}}} \mathbb{E}[\Pi_T - \sum_{j=1}^d \frac{\kappa_j}{2} \int_0^T \sigma_t^2 (q_t^j \delta_t^j)^2 dt - \frac{\kappa}{2} \int_0^T \sigma_t^2 (\sum_{j=1}^d q_t^j \delta_t^j)^2 dt], \quad (\text{I.11})$$

avec $\sigma_t = S_t \sqrt{V_t}$, $(\kappa_j)_{j=1, \dots, d}$ et κ constantes non négatives. Il est facile de remarquer que la fonction de valeur associée à (I.11) dépend de $(2 + n + d)$ variables, c'est-à-dire t , S , $Z := (Z^1, \dots, Z^n)$ et $q := (q^1, \dots, q^d)$. Afin de contourner cette malédiction de la dimensionnalité, nous proposons quelques approximations dans ce qui suit.

2.2 Approximation de la fonction de valeur

- L'équation de Hamilton-Jacobi-Bellman

Nous proposons d'abord l'approximation suivante pour réduire la dimensionnalité due à la nature multidimensionnelle de Z :

Approximation I.1. *Nous considérons $(Z_t^i)_{t \in [0, T], i=1, \dots, n}$ et $(\delta_t^j)_{t \in [0, T], j=1, \dots, d}$ par leur valeurs initiaux, c'est-à-dire*

$$\begin{aligned} Z_t^i &= Z_0^i \quad i = 1, \dots, n, \\ \delta_t^j &= \delta_0^j =: \delta^j \quad j = 1, \dots, d, \end{aligned}$$

pour $t \in [0, T]$. Nous avons donc $V_t = V_0$ pour $t \in [0, T]$ par (I.2). Nous remplaçons la dynamique de $(S_t)_{t \in [0, T]}$ par

$$dS_t = \mu dt + \sigma dW_t, \quad t \in [0, T],$$

où $\sigma := S_0 \sqrt{V_0}$ et une tendance constante μ est autorisée.

Cette approximation est pertinente si nous considérons que l'horizon temporel de market making est généralement inférieur à plusieurs heures. Les résultats de la Figure I.4 et d'autres tests sur la couverture quotidienne indiquent qu'il est raisonnable de considérer que δ est constant à l'échelle quotidienne. En pratique, le market maker peut toujours réinitialiser l'algorithme avec des paramètres mis à jour en cas de mouvements importants du marché. Nous rappelons que cette approximation est considérée au début de chaque période de market making et que les mises à jour de δ entre les périodes sont régies par le modèle QRH. La fonction de valeur associée peut maintenant être définie comme suit

$$v : (t, q) \in [0, T] \times \mathcal{Q} \mapsto v(t, q),$$

où

$$v(t, q) = \sup_{\left(\begin{smallmatrix} l_s^{1,b}, \dots, l_s^{d,b} \\ l_s^{1,a}, \dots, l_s^{d,a} \end{smallmatrix} \right)_{s \in [t, T]}} \mathbb{E}_{t, q} \left[\int_t^T \sum_{j=1}^d \left(\sum_{k=a, b} \frac{D^j}{2} m^j l_s^{j, k} \Lambda^{j, k} \mathbb{1}_{\{q_s^j + \phi^k m^j \in \mathcal{Q}^j\}} + q_s^j \delta^j \mu \right) ds \right. \\ \left. - \sum_{j=1}^d \frac{\kappa_j}{2} \sigma^2 \int_t^T (q_s^j \delta^j)^2 ds - \frac{\kappa}{2} \sigma^2 \int_t^T \left(\sum_{j=1}^d q_s^j \delta^j \right)^2 ds \right], \quad (\text{I.12})$$

avec $\mathcal{Q}^j = \{-Q^j, -Q^j + m^j, \dots, Q^j\}$, $\mathcal{Q} = \prod_{j=1}^d \mathcal{Q}^j$ et

$$\phi^k = \begin{cases} +1 & \text{if } k = b, \\ -1 & \text{if } k = a. \end{cases}$$

L'équation de HJB associée à (I.12) peut être facilement établie comme suit

$$0 = -\frac{\partial v}{\partial t} - \mu \sum_{j=1}^d q^j \delta^j + \sum_{j=1}^d \frac{\kappa_j}{2} \sigma^2 (q^j \delta^j)^2 + \frac{\kappa}{2} \sigma^2 \left(\sum_{j=1}^d q^j \delta^j \right)^2 \\ - \sum_{j=1}^d \sum_{k=a, b} \mathbb{1}_{\{q^j + \phi^k z^j \in \mathcal{Q}^j\}} m^j H^{j, k} \left(\frac{v(t, q) - v(t, q + \phi^k m^j e^j)}{m^j} \right) \quad (\text{I.13})$$

avec la condition terminale

$$v(T, q) = 0, \quad (\text{I.14})$$

où $H^{j, k}(p) = \Lambda^{j, k} \mathbb{1}_{\{p \leq \frac{D^j}{2}\}} \left(\frac{D^j}{2} - p \right)$, et $\{e^j\}_{j=1}^d$ est la base canonique de \mathbb{R}^d . En suivant les mêmes méthodologie que Guéant [2017], l'existence et l'unicité de la solution de (I.13) peuvent être prouvées.

- Approximation quadratique

Inspirés par [Bergault et al. \[2021\]](#), nous approximations les fonctions hamiltoniennes impliquées dans (I.13) par des fonctions quadratiques. La fonction de valeur résultante peut alors être écrite comme une fonction quadratique de q . Plus précisément, nous remplaçons $H^{j,k}(p)$ par l'approximation suivante :

$$\hat{H}^{j,k}(p) = \Lambda^{j,k} \left(\frac{\alpha^j}{2} p^2 - p + \frac{D^j}{2} \right)$$

où $\alpha^j > 0$. Ensuite, dans le cas où il n'y a pas de contraintes d'inventaire maximal pour les actifs individuels, c'est-à-dire $Q^j = +\infty$, $j \in \{1, \dots, d\}$, nous obtenons la formule asymptotique suivante pour la fonction de valeur approchée \hat{v} :

$$\hat{v}(0, q) \stackrel{T \rightarrow +\infty}{\sim} -q^T A q - q^T B,$$

où

$$\begin{aligned} A &= \frac{\sigma}{2} \sqrt{\kappa} \Gamma, \\ B &= -D_+^{-\frac{1}{2}} \hat{A}^+ D_+^{\frac{1}{2}} \delta \mu - D_+^{-\frac{1}{2}} \hat{A} \hat{A}^+ D_+^{-\frac{1}{2}} \left(V_- + \frac{\sigma}{2} \sqrt{\kappa} D_- \mathcal{D}(\Gamma) \right), \\ D_+ &= \text{diag}((\Lambda^{1,b} + \Lambda^{1,a}) \alpha^1 z^1, \dots, (\Lambda^{d,b} + \Lambda^{d,a}) \alpha^d z^d), \\ D_- &= \text{diag}((\Lambda^{1,b} - \Lambda^{1,a}) \alpha^1 z^1, \dots, (\Lambda^{d,b} - \Lambda^{d,a}) \alpha^d z^d), \\ V_- &= ((-\Lambda^{1,b} + \Lambda^{1,a}) z^1, \dots, (-\Lambda^{d,b} + \Lambda^{d,a}) z^d), \\ \Gamma &= D_+^{-\frac{1}{2}} (D_+^{\frac{1}{2}} \Sigma D_+^{\frac{1}{2}})^{\frac{1}{2}} D_+^{-\frac{1}{2}}, \quad \hat{A} = \sqrt{\kappa V} (D_+^{\frac{1}{2}} \Sigma D_+^{\frac{1}{2}})^{\frac{1}{2}}, \end{aligned}$$

\hat{A}^+ est l'inverse de Moore-Penrose de \hat{A} et \mathcal{D} est l'opérateur linéaire qui fait correspondre une matrice au vecteur de ses coefficients diagonaux. Comme le suggère dans [Bergault et al. \[2021\]](#), le market maker peut prendre des décisions en suivant la fonction de valeur asymptotique. Dans notre cas, pour tout $(j, k) \in \{1, \dots, d\} \times \{a, b\}$, les décisions de market making suivantes peuvent être obtenues:

$$\tilde{I}^{j,k*}(q) = \mathbb{1}_{\{q^j + \phi^k m^j \in \mathcal{Q}^j, 2\phi^k(e^j)^T A q + m^j(e^j)^T A e^j + \phi^k(e^j)^T B \leq \frac{D^j}{2}\}}. \quad (\text{I.15})$$

2.3 Résultats numériques

Nous effectuons des simulations sur des données synthétiques avec différentes préférences d'aversion au risque κ . Nous prenons $(\kappa_j)_{j=1, \dots, d} = \frac{\kappa}{2}$ pour des raisons de simplicité. Étant donné un certain κ , nous simulons 5000 trajectoires et le teneur de marché prend des mesures en suivant (I.15) pour chaque trajectoire. La moyenne et l'écart-type de tous les Π_T sont évalués. La Figure I.5 montre les résultats avec un panier de six actifs. Notre méthode, appelée multi-actifs, est nettement plus performante que l'approche uni-actifs, dans laquelle nous effectuons le contrôle du risque sur chaque actif individuel sans tenir compte du risque au niveau du portefeuille, c'est-à-dire $\kappa = 0$. Cela illustre l'effet de la mutualisation du risque entre les actifs considérés.




Figure I.5: Performance moyenne et risque des stratégies de tenue de marché uni-actifs et multi-actifs avec différents paramètres de pénalisation du risque κ .

3 Chapitre IV: Sur l'universalité du processus de formation de la volatilité : quand l'apprentissage automatique et la volatilité rugueuse s'accordent

Dans le cadre du modèle RFSV, la volatilité future peut être prédite sur la base des volatilités passées au moyen d'une formule explicite, qui dépend essentiellement d'un seul paramètre, à savoir l'indice de Hurst [Gatheral et al., 2018]. Étant donné le fait largement observé que l'indice de Hurst se situe autour de 0.1 pour un large éventail d'actifs financiers [Bennedsen et al., 2021; Bolko et al., 2022; Gatheral et al., 2018; Wu et al., 2022], une formule de prévision quasi-universelle de la volatilité réalisée émerge. Cela nous incite à envisager l'existence d'un mécanisme endogène universel de formation de la volatilité entre les actifs, reliant les volatilités et les rendements passés aux volatilités futures. Nous cherchons à saisir ce processus de formation en adaptant des dispositifs de prévision de la volatilité très performants. Nous appliquons tout d'abord une approche quasi-non paramétrique avec LSTM. Nous montrons que la LSTM universelle entraînée sur un ensemble de données regroupant des centaines d'actions liquides surpasse systématiquement les autres dispositifs de prévision paramétriques spécifiques aux actifs, y compris les méthodes autorégressives simples (AR), autorégressives hétérogènes (HAR) et RFSV. Le réentraînement du réseau actif par actif ou secteur par secteur n'améliore pas les performances. Nous confirmons ensuite l'universalité en question d'un point de vue paramétrique. Un dispositif de prévision paramétrique parcimonieux incorporant les principales caractéristiques des modèles RFSV et QRH avec des paramètres fixes permet d'obtenir le même niveau de performance que la LSTM universelle.

3.1 Dispositifs de prévision de la volatilité réalisée

Soit $(\sigma_{t-p})_{p=1,2,\dots}$ les volatilités passées, nous évaluons les méthodes paramétriques de prévision de la volatilité suivantes en tant que références :

- AR(p):

$$\hat{\sigma}_t = \alpha_0 + \sum_{j=1}^p \beta_j \sigma_{t-j},$$

- HAR:

$$\hat{\sigma}_t = \alpha_0 + \beta_1 \sigma_{t-1} + \beta_2 \frac{1}{5} \sum_{j=1}^5 \sigma_{t-j} + \beta_3 \frac{1}{22} \sum_{j=1}^{22} \sigma_{t-j},$$

- RFSV:

$$\hat{\sigma}_t = v \exp(\widehat{\log \sigma}_t), \tag{I.16}$$

où

$$\widehat{\log \sigma}_t = \frac{\cos(H\pi)}{\pi} \int_{-\infty}^{t-1} \frac{\log \sigma_s}{(t-s+1)(t-s)^{H+1/2}} ds,$$

$$v = \exp\left(\frac{\Gamma(3/2-H)}{2\Gamma(H+1/2)\Gamma(2-2H)} v^2\right).$$

Les paramètres $\alpha_0, \beta_j \in \mathbb{R}$ dans les méthodes AR et HAR peuvent être estimés par les moindres carrés ordinaires. Comme pour la méthode RFSV, v est une constante positive et $H < \frac{1}{2}$ mesurant le caractère lisse des séries de volatilité. Nous suivons la méthodologie introduite dans [Gatheral et al. \[2018\]](#) pour estimer H et v . Comme pour l'approche non-paramétrique basée sur la LSTM, nous entraînons le réseau avec des échantillons contenant les L historiques les plus récents, comme illustré à la Figure I.6.




Figure I.6: Structure d'une cellule LSTM (à gauche) et graphique simplifié du réseau basé sur la LSTM (à droite).

3.2 Capturer l'universalité avec LSTM

Le réseau LSTM est entraîné sur l'ensemble des données de plus de 800 actions liquides du marché américain, tandis que les autres dispositifs de prévision paramétriques sont calibrés action par action. La Figure I.7 résume l'erreur quadratique moyenne (EQM) hors échantillon par rapport au modèle HAR. Nous remarquons :

- le modèle AR(22), qui n'impose aucune contrainte structurelle sur les coefficients autorégressifs, est moins performant que le modèle HAR. Ceci est compréhensible car il est plus susceptible d'être *overfitted*,
- le RFSV surpasse le modèle HAR en termes d'erreur de prédiction. Il est remarquable car il ne fait intervenir que deux paramètres libres, à savoir ν et H . Il bénéficie d'une paramétrisation spéciale des poids des volatilités passées en fonction de H ,
- les modèles non-paramétriques, $LSTM_{var}^{us}$ et $LSTM_{ret}^{us}$ surpassent les autres modèles paramétriques, en particulier lorsque nous incorporons des données sur les rendements passés. Cela indique que le mécanisme universel potentiel de formation de la volatilité entre les actifs, reliant les volatilités et les rendements passés aux volatilités futures, nous permet de calibrer un modèle universel basé sur tous les actifs, où le risque de *overfitting* est réduit en raison de l'enrichissement des scénarios réalisés.




Figure I.7: Diagramme en boîte montrant l'EQM hors échantillon de chaque modèle par rapport au modèle HAR pour les actions du marché américain. $LSTM_{var}^{us}$ est formé sur la base de la variance réalisée dans le passé, et $LSTM_{ret}^{us}$ utilise également les rendements passés.

Nous affinons le LSTM universel sur chaque action individuelle afin d'incorporer toute caractéristique potentielle dépendant de l'actif. Comme le montre la Figure I.8, il n'y a pas

de réduction universelle de l'erreur de prévision après que les paramètres du réseau ont été "adaptés" localement à chaque action. En outre, le réseau formé sur la base des actions du marché américain peut être directement appliqué aux actions du marché européen, avec des performances de prévision supérieures obtenues, comme le montre la Figure I.9. Cela suggère que l'universalité en question se vérifie même entre les marchés.




Figure I.8: Distribution empirique de l'EQM hors échantillon des LSTM spécifiques à l'actif finement réglées par rapport à la LSTM universelle. La partie gauche correspond au cas où tous les paramètres du réseau peuvent être mis à jour. La droite correspond au cas où seuls les paramètres de la couche de sortie peuvent être mis à jour.

3.3 Découvrir le processus universel de formation de la volatilité

Nous proposons ici un dispositif de prévision paramétrique parcimonieux permettant d'obtenir un niveau de performance similaire à celui de la LSTM universelle. L'idée principale est de compléter la méthode RFSV par un élément codant l'effet de rétroaction des rendements passés inspiré du modèle QRH. Plus précisément, nous utilisons la prévision suivante :

$$(1 - \lambda)\hat{\sigma}^{RFSV} + \lambda\hat{\sigma}^{QRH} \tag{I.17}$$

avec $\lambda \in (0, 1)$, $\hat{\sigma}^{RFSV}$ est la prévision donné par la méthode RFSV en suivant (I.16), $\hat{\sigma}^{QRH}$ est donné par

$$\begin{aligned} (\hat{\sigma}_t^{QRH})^2 &= a(Z_{t-1} - b)^2 + c, \\ Z_t &= \sum_{i=1}^n c_i^d Z_{i,t}, \\ Z_{i,t} &= e^{-\gamma_i^d} Z_{i,t-1} + r_t, \quad Z_{i,0} = 0, \quad i = 1, \dots, n, \end{aligned}$$

où $a, b, c > 0$ et $(c_i^d, \gamma_i^d)_{i=1, \dots, n}$ sont donné par la même approximation multifactorielle de la fonction noyau rugueuse $K(t) := \frac{t^{H-1/2}}{\Gamma(H+1/2)}$ comme dans Rosenbaum & Zhang [2022a]. Nous




Figure I.9: Diagramme en boîte montrant les performances hors échantillon des modèles par rapport au modèle HAR pour les actions du marché européen. Il convient de noter que les modèles $LSTM^{eu}$ et $LSTM^{us}$ sont entraînés respectivement sur les actions des marchés européen et américain.

prenons les paramètres suivants pour (I.17):

$$H = 0.059, v = 1.03, a = 0.065, b = 0.74, c = 0.55.$$

Il s'agit en fait des résultats de calibration globaux lorsque nous entraînons respectivement les dispositifs de prévision RFSV et QRH sur tous les titres du marché américain. Comme le montre la Figure I.10, le dispositif de prévision parcimonieux avec les paramètres fixes ci-dessus peut obtenir un niveau de performance similaire à celui du réseau LSTM avec $\lambda \approx 0.1$. D'un point de vue paramétrique, cela confirme à nouveau l'universalité pertinente, montrant que les principales caractéristiques de ce processus universel de formation de la volatilité peuvent être bien décrites par le paradigme de la volatilité approximative renforcé par l'effet de rétroaction des tendances des prix.

4 Chapitre V: Vers un filtrage systématique des nouvelles intrajournalières : une approche axée sur la liquidité

Les marchés financiers devenant de plus en plus compétitifs, les praticiens ont besoin de développer des outils d'évaluation automatique du sentiment d'actualité afin d'anticiper l'impact des communiqués de presse sur les prix des actifs financiers. La présence de nombreux articles neutres, qui n'impliquent pas d'opinion directionnelle sur la valeur de l'actif associé, entraînera un apprentissage de modèle biaisé si l'on forme un modèle sur l'ensemble des données d'actualité non filtrées. Nous proposons une approche de filtrage en deux étapes




Figure I.10: Performance relative hors échantillon de la prévision $(1 - \lambda)\hat{\sigma}^{RFSV} + \lambda\hat{\sigma}^{QRH}$ par rapport à la prévision universelle LSTM.

pour sélectionner les échantillons d'actualités ayant un impact, dont on pense qu'ils ont considérablement influencé les activités du marché après leur publication et qui contiennent donc des éléments chargés de sentiment dans les textes. Premièrement, nous surveillons les conditions de liquidité intrajournalière de l'actif indiquées par plusieurs variables, notamment la volatilité σ , le volume de transactions V , l'écart entre les cours acheteur et vendeur ϕ et le volume disponible dans le carnet d'ordres limité B . Le modèle de saut introduit dans [Bemporad et al. \[2018\]](#) est utilisé pour classer les conditions de liquidité en deux modes distincts. Deuxièmement, nous établissons un lien entre les changements de mode de liquidité détectés, de celui associé à une volatilité et à une activité de transactions plus faibles à l'autre associé à des communiqués de presse percutants. En d'autres termes, tous les articles de presse publiés autour de ces changements de liquidité sont considérés comme ayant un impact. Pour illustrer l'effet de notre méthode, nous avons adapté deux classificateurs de Bayes naïfs (NBC) respectivement sur l'ensemble de données d'actualités original et sur l'ensemble de données filtré, en essayant d'évaluer les sentiments à l'égard de l'actualité. Malgré la taille beaucoup plus réduite de l'ensemble d'entraînement filtré, les tests hors échantillon montrent que ce dernier est nettement plus performant lorsqu'il est appliqué pour prédire les mouvements à court terme du prix des actifs, ce qui suggère que l'approche proposée peut exclure de l'ensemble d'entraînement les échantillons d'actualités non pertinents et conduire à un apprentissage plus efficace des sentiments à l'égard de l'actualité. Nos tests numériques sont effectués sur les données des composants du S&P 500 afin d'éviter que les résultats ne soient biaisés par les titres à faible capitalisation.

4.1 Calibration du mode de liquidité

Dans cette partie, nous appliquons le modèle de saut pour produire des séquences de modes indiquant les conditions de liquidité intrajournalière des actifs qui nous intéressent. Pour chaque paire actif-jour (S, d) avec $S = 1, \dots, N$ et $d = 1, \dots, D$, après avoir supprimé les données des 15 premières et dernières minutes de la période ouverte du marché, nous mesurons $x_t^{S,d} := (\phi_t^{S,d}, V_t^{S,d}, \sigma_t^{S,d}, B_t^{S,d})$ pour la t -ième tranche de 5 minutes, où $t = 1, \dots, T$ et $T = 72$.

Nous appliquons certaines procédures de prétraitement à x afin d'éliminer les saisonnalités intrajournalières bien connues et de faire en sorte que les observations des différents stocks aient la même échelle. Compte tenu de $X^{S,d} := (x_t^{S,d})_{t=1,\dots,T}$, nous cherchons une séquence de modes $M^{S,d} := (m_t^{S,d})_{t=1,\dots,T}$ décrivant la dynamique des conditions de liquidité. Pour faciliter l'interprétation, nous n'autorisons que deux modes de liquidité. L'un représente l'état "calme" dans lequel le marché reste pendant la plupart du temps et l'autre décrit le régime induit par certains événements exogènes. Nous avons donc $m_t \in \{1,2\}$. Chaque jour, nous effectuons une calibration des modes de liquidité avec la fonction objective particulière suivante :

$$\min_{\Theta^{S,d}, M^{S,d}} J(X^{S,d}, \Theta^{S,d}, M^{S,d}) = \sum_{t=1}^T \|x_t^{S,d} - \theta_{m_t^{S,d}}^{S,d}\|_2^2 + \lambda \sum_{t=1}^{T-1} \mathbb{1}_{m_t^{S,d} \neq m_{t+1}^{S,d}}, \quad (\text{I.18})$$

où $\Theta^{S,d} := (\theta_1^{S,d}, \theta_2^{S,d}) \in (\mathbb{R}^4)^2$, et λ est une constante positive. Par conséquent, les résultats de la minimisation sont une séquence de modes de liquidité associés à chaque intervalle de temps et deux vecteurs de dimension quatre. Les intervalles de temps présentant des caractéristiques de liquidité similaires sont associés au même mode, représenté par θ_1 ou θ_2 . Lorsque $\lambda > 0$, les changements de mode entre deux intervalles de temps adjacents sont pénalisés car ils entraînent des pertes supplémentaires. De cette manière, une certaine persistance du mode est obtenue et seuls les changements significatifs dans les conditions de liquidité sont détectés. Il convient de noter que lorsque $\lambda = 0$, nous obtenons le même objectif que la méthode classique de regroupement par K-moyennes, qui minimise simplement les variances à l'intérieur des clusters.

Pour chaque paire (S, d) , les représentants de mode de dimension huit sont basés sur $T = 72$ observations, ce qui implique un risque potentiel de overfitting. Pour réduire ce risque, nous étendons la tâche aux séquences d'observation de toutes les actions du même jour. Plus précisément, pour le jour d , nous exécutons la calibration sur l'ensemble des séquences définies par $X^d := \{(x_t^{S,d})_{t=1,\dots,T}\}_{S=1,\dots,N}$. Tous les stocks sont supposés avoir le même mode représente, dénoté par $\Theta^d := (\theta_1^d, \theta_2^d)$. Pour faciliter la notation, nous ne tenons pas compte de l'exposant d et la nouvelle fonction objective est la suivante :

$$\min_{\Theta, M} J(X, \Theta, M) = \sum_{S=1}^N \left(\sum_{t=1}^T \|x_t^S - \theta_{m_t^S}\|_2^2 + \lambda \sum_{t=1}^{T-1} \mathbb{1}_{m_t^S \neq m_{t+1}^S} \right), \quad (\text{I.19})$$

où $M := \{(m_t^S)_{t=1,\dots,T}\}_{S=1,\dots,N}$. Le Problème (I.19) peut être résolu par une approche similaire à l'algorithme *expectation-maximization* qui alterne la minimisation par rapport à Θ et M .

En désignant le mode de liquidité associé à un niveau de volatilité plus faible comme le *Mode 1* et l'autre comme le *Mode 2*, nous donnons la dynamique des représentants de mode de 2017-01-01 à 2019-12-31 dans la Figure I.11. Il est intéressant de noter que les deux modes sont bien distincts en termes de niveau de volatilité et de volume de transactions. Il est naturel de considérer que le *Mode 2* est principalement alimenté par certains événements exogènes étant donné son niveau de volatilité élevé et ses activités de transaction accrues.




Figure I.11: Évolution historique des représentants de mode θ_1 et θ_2 lorsque $\lambda = 0.5$.

4.2 Filtrage de l'actualité et apprentissage des sentiments

- Méthodologie




Figure I.12: Scénarios illustrant les cas où une nouvelle est considéré comme ayant un impact.

Nous appliquons les résultats de la calibration des modes de liquidité présentés ci-dessus pour identifier les "véritables" nouvelles impactantes. L'ensemble des nouvelles publiées au cours de la période considérée est désigné par \mathcal{D} . Étant donné une nouvelle $a \in \mathcal{D}$ concernant l'action S , nous considérons qu'elle a un impact si la liquidité de S passe du *Mode 1* au *Mode 2* peu de

temps avant ou après la publication de la nouvelle. Dans les scénarios présentés à la Figure I.12, nous considérons que la nouvelle a a un impact puisque nous observons un changement brutal des modes de liquidité. L'ensemble des nouvelles impactantes sélectionnées est \mathcal{D}_λ , où λ est le paramètre de pénalisation du changement de mode défini dans (I.19). En désignant par r_a^h le rendement du marché de S sur h minutes après la publication de a , nous laissons $r^{h,k}$ et $r^{h,100-k}$ être respectivement le k -ième et $(100-k)$ -ième centile de $\mathcal{R}^h := \{r_a^h | a \in \mathcal{D}\}$. Nous définissons l'ensemble suivant :

$$\mathcal{Z}_{h,k}^- := \{a | r_a^h \leq r^{h,k}, a \in \mathcal{D}\} \quad \text{et} \quad \mathcal{Z}_{h,k}^+ := \{a | r_a^h \geq r^{h,100-k}, a \in \mathcal{D}\}.$$

Ainsi, Z^- et Z^+ sont des ensembles de nouvelles baissiers et haussiers respectivement en fonction de l'amplitude des rendements des actifs après les publications associées, ce qui est le critère couramment utilisé dans des travaux tels que [Chen \[2021\]](#); [Ding et al. \[2014, 2015\]](#); [Hu et al. \[2018\]](#). Dans notre approche, les ensembles de "vraies" nouvelles ayant un impact sont donnés par

$$\mathcal{N}_{\lambda,h,k}^- := \mathcal{D}_\lambda \cap \mathcal{Z}_{h,k}^- \quad \text{et} \quad \mathcal{N}_{\lambda,h,k}^+ := \mathcal{D}_\lambda \cap \mathcal{Z}_{h,k}^+.$$

Par conséquent, la méthode proposée prend également en compte les changements des conditions de liquidité. Sur la base d'un ensemble d'échantillons étiquetés $\mathcal{E} = \mathcal{Z}^+ \cup \mathcal{Z}^-$ ou $\mathcal{E} = \mathcal{N}^+ \cup \mathcal{N}^-$, nous formons un NBC de Bernoulli à plusieurs variables pour capturer les mots contribuant au sentiment, voir le Chapitre V pour plus de détails. Pour la nouvelle a , le classificateur entraîné donne les probabilités que a soit positif et négatif, dénotées respectivement par $P_{\mathcal{E}}^+(a)$ et $P_{\mathcal{E}}^-(a)$. Nous définissons ainsi le score de sentiment $F \in [-1, 1]$ de a par

$$F(a)|_{\mathcal{E}} := P_{\mathcal{E}}^+(a) - P_{\mathcal{E}}^-(a).$$

Nous nous intéressons particulièrement au pouvoir prédictif de F sur les mouvements de prix des actifs à court terme. Un F significativement élevé (faible) devrait être suivi d'une tendance positive (négative) du prix de l'action associée.

- Résultats numériques

L'étude numérique est basée sur les titres de l'actualité de Bloomberg². Nous entraînonns deux NBCs sur $\mathcal{Z}_{15,10}^+ \cup \mathcal{Z}_{15,10}^-$ et $\mathcal{N}_{0.5,15,10}^+ \cup \mathcal{N}_{0.5,15,10}^-$ respectivement, qui sont ensuite appliqués sur tous les échantillons de nouvelles intrajournalières publiées au cours de 2020-01-01 ~ 2021-12-31 pour produire les scores de sentiment correspondants. La Figure I.13 montre les mouvements moyennes à court terme des prix après les publications des nouvelles associés à des scores de sentiment distincts attribués par les deux classificateurs. Nous remarquons clairement que l'ensemble de données de nouvelles sélectionnées par notre approche conduit à une meilleure performance prédictive du classificateur de nouvelles.

Notons qu'avec $\lambda = 0.5$, la taille de $\mathcal{N}_{0.5,15,10}^+ \cup \mathcal{N}_{0.5,15,10}^-$ n'est que d'environ un dixième de $\mathcal{Z}_{15,10}^+ \cup \mathcal{Z}_{15,10}^-$. Pour vérifier que notre méthode de sélection n'équivaut pas à filtrer les

²<https://www.bloomberg.com/professional/product/event-driven-feeds/>




Figure I.13: Rendement moyen des actions après les publications des nouvelles dont les notes de sentiment significativement positives et négatives ont été attribuées par les NBCs calibrés sur différents ensembles de formation. La partie gauche représente le cas où le filtrage des nouvelles est uniquement basé sur les informations relatives au rendement. La droite correspond à notre approche. Seuls les échantillons dont les scores se situent dans la fourchette supérieure/inférieure de 10% sont présentés pour faciliter la comparaison. Les parties ombrées représentent l'écart-type empirique du rendement moyen.

nouvelles en fonction des rendements boursiers associés, nous répétons le même test sur $\mathcal{Z}_{15,1}^+ \cup \mathcal{Z}_{15,1}^-$. Comme le montre la sous-figure de gauche de la Figure I.14, avec un nombre d'échantillons d'entraînement similaire au cas de $\mathcal{N}_{0.5,15,10}^+ \cup \mathcal{N}_{0.5,15,10}^-$, la performance de prédiction est maintenant largement dégradée. Ces résultats indiquent que notre méthode présente un moyen plus efficace d'identifier les nouvelles impactantes que l'approche basée uniquement sur les informations de rendement. Nous présentons également les résultats des scores de sentiment des nouvelles de Bloomberg donnés par l'algorithme de classification propre à l'entreprise dans la Figure I.14. Bien que le modèle de classification testé ici, c'est-à-dire NBC, soit très simple, les performances sont légèrement meilleures que celles obtenues avec les scores de Bloomberg. D'autres tests numériques montrent que les résultats de notre approche sont robustes aux paramètres impliqués, c'est-à-dire λ , h et k .




Figure I.14: Rendement moyen des actions après les publications des nouvelles dont les notes de sentiment significativement positives et négatives ont été attribuées par le NBC entraîné sur les échantillons associés à un rendement réalisé extrême tombant dans le quantile supérieur/inférieur de 1% (à gauche), et le cas avec les notes de sentiment attribuées par Bloomberg (à droite).

Deep calibration of the quadratic rough Heston model

Abstract

The quadratic rough Heston model provides a natural way to encode Zumbach effect in the rough volatility paradigm. We apply multi-factor approximation and use deep learning methods to build an efficient calibration procedure for this model. We show that the model is able to reproduce very well both SPX and VIX implied volatilities. We typically obtain VIX option prices within the bid-ask spread and an excellent fit of the SPX at-the-money skew. Moreover, we also explain how to use the trained neural networks for hedging with instantaneous computation of hedging quantities.

Keywords— Quadratic rough Heston, multi-factor approximation, SPX smile, VIX smile, deep learning, joint calibration, hedging

1 Introduction

The rough volatility paradigm introduced in Gatheral et al. [2018] is now widely accepted, both by practitioners and academics. On the macroscopic side, rough volatility models can fit with remarkable accuracy the shape of implied volatility smiles and at-the-money skew curves. They also reproduce amazingly well stylized facts of realized volatilities, see for example Bayer et al. [2016]; Bennedsen et al. [2021]; El Euch et al. [2019]; Gatheral et al. [2018]; Livieri et al. [2018]. On the microstructural side, Dandapani et al. [2021]; El Euch et al. [2018]; Jaisson & Rosenbaum [2016] show that the rough Heston model introduced and developed in “el2018perfect” [n.d.]; El Euch & Rosenbaum [2019] naturally emerges from agents behaviors at the microstructural scale.

Nevertheless, one stylized fact of financial time series that is not reflected in the rough Heston model is the feedback effect of past price trends on future volatility, which is discussed by Zumbach in Zumbach [2010]. Super-Heston rough volatility models introduced in Dandapani et al. [2021] fill this gap by considering quadratic Hawkes processes from microstructural level, and showing that the Zumbach effect remains explicit in the limiting models. As a particular example of super-Heston rough volatility models, the authors in Gatheral et al. [2020] propose the quadratic rough Heston model, and show its promising ability to calibrate jointly SPX smiles and VIX smiles, where other continuous-time models have been struggling for a long time [Guyon, 2020].

The VIX index is in fact by definition a derivative of the SPX index S , which can be represented as

$$\text{VIX}_t = \sqrt{-2\mathbb{E}[\log(S_{t+\Delta}/S_t)|\mathcal{F}_t]} \times 100, \quad (1)$$

where $\Delta = 30$ days and \mathbb{E} is the risk-neutral expectation. Consequently, VIX options are also derivatives of SPX. Finding a model which jointly calibrates the prices of SPX and VIX options is known to be very challenging, especially for short maturities. As indicated in Guyon [2020], “the very negative skew of short-term SPX options, which in continuous models implies a very large volatility of volatility, seems inconsistent with the comparatively low levels of VIX implied volatilities”. Through numerical examples, Gatheral et al. [2020] show the relevance of the quadratic rough Heston model in terms of pricing simultaneously SPX and VIX options. In this chapter, in the spirit of Abi Jaber & El Euch [2019], we propose a multi-factor approximated version of this model and an associated efficient calibration procedure.

Under the rough Heston model, the characteristic function of the log-price has semi-closed form formula, and thus fast numerical pricing methods can be designed, see El Euch & Rosenbaum [2019]. However, pricing in the quadratic rough Heston model is more intricate. The multi-factor approximation method for the rough kernel function developed in Abi Jaber & El Euch [2019] makes rough volatility models Markovian in high dimension. Thus Monte-Carlo simulations become more feasible in practice. Still, in our case, model calibration remains a difficult task. Inspired by recent works about applications of deep learning in quantitative finance, see for example Bayer et al. [2019]; Hernandez [2016]; Horvath et al. [2021], we use deep neural networks to speed up model calibration. The effectiveness of the calibrated model for fitting jointly SPX and VIX smiles is illustrated through numerical experiments. Interestingly, under our model, the trained networks also allow us to hedge options with instantaneous computation of hedging quantities.

This chapter is organized as follows. In Section 2, we give the definition of our model and introduce the approximation method. In Section 3, we develop the model calibration with deep neural networks. Validity of the methods is tested both on simulated data and market data. Finally in Section 4, we show how to perform hedging in the model with neural networks through some toy examples.

2 The quadratic rough Heston model and its multi-factor approximation

The quadratic rough Heston model, proposed in [Gatheral et al. \[2020\]](#), for the price of an asset S (here the SPX) and its spot variance V under risk-neutral measure is

$$dS_t = S_t \sqrt{V_t} dW_t, \quad V_t = a(Z_t - b)^2 + c,^3 \quad (2)$$

where W is a Brownian motion, a, b, c are all positive constants and Z_t is defined as

$$Z_t = \int_0^t \lambda \frac{(t-s)^{\alpha-1}}{\Gamma(\alpha)} (\theta_0(s) - Z_s) ds + \int_0^t \eta \frac{(t-s)^{\alpha-1}}{\Gamma(\alpha)} \sqrt{V_s} dW_s, \quad (3)$$

for $t \in [0, T]$. Here T is a positive time horizon, $\alpha \in (1/2, 1)$, $\lambda > 0$, $\eta > 0$ and $\theta_0(\cdot)$ is a deterministic function. Z_t is driven by the returns through $\sqrt{V_t} dW_t = dS_t/S_t$. Then the square in V_t can be understood as a natural way to encode the so called *strong Zumbach effect*, which means that the conditional law of future volatility depends not only on path volatility trajectory but also on past returns. Note that in this case we have a pure-feedback model as S_t and V_t are driven by the same Brownian motion, see [Dandapani et al. \[2021\]](#) for more details on the derivation of this type of models. We will see in Section 4 that this setting enables us to hedge perfectly European options with SPX only. We recall the parameter interpretation given in [Gatheral et al. \[2020\]](#):

- a stands for the strength of the feedback effect on volatility.
- b encodes the asymmetry of the feedback effect. It reflects the empirical fact that negative price returns can lead to volatility spikes, while it is less pronounced for positive returns.
- c is the base level of variance, independent from past prices information.

It is shown in [El Euch & Rosenbaum \[2019\]](#) that under the rough Heston model the volatility trajectories have almost surely Hölder regularity $\alpha - 1/2 - \varepsilon$, for any $\varepsilon > 0$. This actually recalls the observation in [Gatheral et al. \[2018\]](#) that the dynamic of log-volatility is similar to that of a fractional Brownian motion with Hurst parameter of order 0.1. Similarly the fractional kernel $K(t) = \frac{t^{\alpha-1}}{\Gamma(\alpha)}$ in (3) enables us to generate rough volatility dynamics, which is highly desirable as explained in the introduction. However, it makes the quadratic rough Heston model non-Markovian and non-semimartingale, and thus difficult to simulate efficiently. In this chapter, we apply the multi-factor approximation proposed in [Abi Jaber & El Euch \[2019\]](#) to do so. The key idea is to write the fractional kernel $K(t)$ as the Laplace transform of a positive measure μ

$$K(t) = \int_0^\infty e^{-\gamma t} \mu(d\gamma), \quad \mu(d\gamma) = \frac{\gamma^{-\alpha}}{\Gamma(\alpha)\Gamma(1-\alpha)} d\gamma.$$

Then we approximate μ by a finite sum of Dirac measures $\mu^n = \sum_{i=1}^n c_i^n \delta_{\gamma_i^n}$ with positive weights $(c_i^n)_{i=1, \dots, n}$ and discount coefficients $(\gamma_i^n)_{i=1, \dots, n}$, with $n \geq 1$. This gives us the approximated kernel function

$$K^n(t) = \sum_{i=1}^n c_i^n e^{-\gamma_i^n t}, \quad n \geq 1.$$

³To ensure the martingale property of S_t , we can in fact use $V_t = a\phi(Z_t - b) + c$, where ϕ is defined as

$$\phi(x) = \begin{cases} x^2 & \text{if } x < x^* \\ (x^*)^2 & \text{otherwise} \end{cases}$$

with x^* sufficiently large. In this chapter, for ease of notation we keep writing the square function.

II. Deep calibration of the quadratic rough Heston model

A well-chosen parametrization of the $2n$ parameters $(c_i^n, \gamma_i^n)_{i=1, \dots, n}$ in terms of α can make $K^n(t)$ converge to $K(t)$ in the L^2 sense as n goes to infinity, and the multi-factor approximation models behave closely to their counterparts in the rough volatility paradigm, see [Abi Jaber \[2019\]](#); [Abi Jaber & El Euch \[2019\]](#) for more details. We recall in Appendix II.A the parametrization method proposed in [Abi Jaber \[2019\]](#). Then given the time horizon T and n , $(c_i^n)_{i=1, \dots, n}$ and $(\gamma_i^n)_{i=1, \dots, n}$ are just deterministic functions of α , and therefore not free parameters to calibrate. We can give the following multi-factor approximation of the quadratic rough Heston model:

$$dS_t^n = S_t^n \sqrt{V_t^n} dW_t, \quad V_t^n = a(Z_t^n - b)^2 + c, \quad (4)$$

$$Z_t^n = \sum_{i=1}^n c_i^n Z_t^{n,i}, \quad (5)$$

$$dZ_t^{n,i} = (-\gamma_i^n Z_t^{n,i} - \lambda Z_t^n) dt + \eta \sqrt{V_t^n} dW_t, \quad Z_0^{n,i} = z_0^i, \quad (6)$$

with $(z_0^i)_{i=1, \dots, n}$ some constants. Contrary to the case of the rough Heston model, $\theta(t)$ cannot be easily written as a functional of the forward variance curve in the quadratic rough Heston model. Then instead of making the factors $(Z_t^{n,i})_{i=1, \dots, n}$ starting from 0 and taking $Z_t^n = \lambda \sum_{i=1}^n c_i^n \int_0^t e^{-\gamma_i^n(t-s)} \theta(s) ds + \sum_{i=1}^n c_i^n Z_t^{n,i}$, as the authors do in [Abi Jaber \[2019\]](#); [Abi Jaber & El Euch \[2019\]](#), here we discard $\theta(\cdot)$ in Equation (5) and consider the starting values of factors $(z_0^i)_{i=1, \dots, n}$ as free parameters to calibrate from market data. This setting allows Z_0 and V_0 to adapt to market conditions and also encodes various possibilities for the “term-structure” of Z_t^n . To see this, given a solution for (4)-(6), (6) can be rewritten as

$$Z_t^{n,i} = z_0^i e^{-\gamma_i^n t} + \int_0^t e^{-\gamma_i^n(t-s)} (-\lambda Z_s^n ds + \eta \sqrt{V_s^n} dW_s). \quad (7)$$

Then from (5) we can get

$$Z_t^n = g^n(t) + \int_0^t K^n(t-s) (b(Z_s^n) ds + \sigma(Z_s^n) dW_s), \quad (8)$$

with $g^n(t) = \sum_{i=1}^n z_0^i c_i^n e^{-\gamma_i^n t}$, $b(Z_t^n) = -\lambda Z_t^n$ and $\sigma(Z_t^n) = \eta \sqrt{a(Z_t^n - b)^2 + c}$. By taking expectation on both sides of (8), we get

$$\mathbb{E}[Z_t^n] + \lambda \sum_{i=1}^n c_i^n \int_0^t e^{-\gamma_i^n(t-s)} \mathbb{E}[Z_s^n] ds = \sum_{i=1}^n z_0^i c_i^n e^{-\gamma_i^n t}.$$

Thus we can see that the $(z_0^i)_{i=1, \dots, n}$ allows us to encode initial “term-structure” of Z_t^n . Therefore, it can be understood as an analogy of $\theta(t)$ for the variance process in the rough Heston model. Besides, we will see in Section 4 that this setting allows us to hedge options perfectly with only SPX.

By virtue of Proposition B.3 in [Abi Jaber & El Euch \[2019\]](#), for given n , Equations (8) and equivalently (7) admit a unique strong solution, since $g^n(t)$ is Hölder continuous, $b(\cdot)$ and $\sigma(\cdot)$ have linear growth, and K^n is continuously differentiable admitting a resolvent of the first kind. We stress again the fact that Model (4-6) does not bring new parameters to calibrate compared to the quadratic rough Heston model defined in (2-3), with the idea of the correspondance between $\theta(t)$ and $(z_0^i)_{i=1, \dots, n}$. The fractional kernel in the rough volatility paradigm helps us to build the factors $(Z_t^{n,i})_{i=1, \dots, n}$. Factors with large discount coefficient γ_i^n can mimic roughness and account for short timescales, while ones with small γ_i^n capture information from longer timescales. The quantity Z_t^n aggregates these factors and therefore encodes the multi-timescales nature of volatility processes, which is discussed for example in [Fouque et al. \[2011\]](#); [Gatheral et al. \[2018\]](#).

Remark 1. *With multi-factor approximation, Model (4-6) is Markovian with a state vector of dimension $n+1$ given by $\mathbf{X}_t^n := (S_t^n, Z_t^{n,1}, \dots, Z_t^{n,n})$. Hence the price of SPX options at time t is fully determined by \mathbf{X}_t^n .*

As discussed in Appendix II.A, we choose $n = 10$ in our numerical experiments. To simplify notations, we discard the label n in the following, and let $\boldsymbol{\omega} := (\lambda, \eta, a, b, c) \in \Omega \subset \mathbb{R}^5$, $\mathbf{z}_0 := (z_0^1, \dots, z_0^{10}) \in \mathcal{Z} \subset \mathbb{R}^{10}$.

3 Model calibration with deep learning

The Markovian nature of Model (4-6) makes the calibration with Monte-Carlo simulations more feasible. However, besides parameters $\boldsymbol{\omega}$, the initial state of factors \mathbf{z}_0 is also supposed to be calibrated from market data. In this case pricing with Monte-Carlo is not well adapted to classical optimal parameter search methods for model calibration, as it leads to heavy computation and the results are not always satisfactory. To bypass this “curse of dimensionality”, we apply deep learning to speed up further model calibration. Deep learning has already achieved remarkable success with high-dimensional data like images and audio. Recently its potentials for model calibration in quantitative finance has been investigated for example in Bayer et al. [2019]; Hernandez [2016]; Horvath et al. [2021]. Two types of methods are proposed in the literature:

- From prices to model parameters (PtM) [Hernandez, 2016]: deep neural networks are trained to approximate the mapping from prices of some financial contracts, *e.g.* options, to model parameters. With this method, we can get directly the calibrated parameters from market data, without use of some numerical methods searching for optimal parameters.
- From model parameters to prices (MtP) [Bayer et al., 2019; Horvath et al., 2021]: in the first step, deep neural networks are trained to approximate the pricing function, that is the mapping from model parameters to prices of financial contracts. Then in the second step, traditional optimization algorithms can be used to find optimal parameters, to minimize the discrepancy between market data and model outputs.

It is hard to say that one method is always better than the other. PtM method is faster and avoids computational errors caused by optimization algorithms, while MtP method is more robust to the varying nature of options data (strike, maturity, ...). In the following, the two methods are applied with implied volatility surfaces (IVS), represented by certain points with respect to some predetermined strikes and maturities. During model calibration, all these points need to be built from market quotes for PtM method, while we could focus on some points of interest for MtP method, for example those near-the-money. For comparison, we will test both methods in the following with simulated data.

3.1 Methodology

The neural networks used in our tests are all multilayer perceptrons. They are trained with synthetic dataset generated from the model. Our methodology is mainly based on two steps: **data generation** and **model training**.

- Synthetic data generation

II. Deep calibration of the quadratic rough Heston model

The objective is to generate data samples $\{\omega^4, \mathbf{z}_0^5, IVS_{SPX}, IVS_{VIX}\}$, where IVS_{SPX} and IVS_{VIX} stand for implied volatility surface of SPX and VIX options respectively. We randomly sample ω and \mathbf{z}_0 with the following distribution:

$$\omega \in \mathcal{U}[0.5, 2.5] \times \mathcal{U}[1.0, 1.5] \times \mathcal{U}[0.1, 0.6] \times \mathcal{U}[0.01, 0.5] \times \mathcal{U}[0.0001, 0.03],$$

$$\mathbf{z}_0 \in \mathcal{U}[-0.5, 0.5]^{10}, \text{ where } \mathcal{U} \text{ denotes the uniform distribution.}$$

For each sampling from above distribution, we generate 50,000 random paths of SPX and VIX with the explicit-implicit Euler scheme (18). Then Monte-Carlo prices are used to compute IVS_{SPX}^{MC} and IVS_{VIX}^{MC} with respect to some predetermined log-moneyness strikes and maturities:

- log-moneyness strikes of SPX options: $k_{SPX} = \{-0.15, -0.12, -0.1, -0.08, -0.05, -0.04, -0.03, -0.02, -0.01, 0.0, 0.01, 0.02, 0.03, 0.04, 0.05\}$,
- log-moneyness strikes of VIX options: $k_{VIX} = \{-0.1, -0.05, -0.03, -0.01, 0.01, 0.03, 0.05, 0.07, 0.09, 0.11, 0.13, 0.15, 0.17, 0.19, 0.21\}$,
- maturities $T = \{0.03, 0.05, 0.07, 0.09\}$.

Then IVS_{SPX}^{MC} is represented by a vector of size $m = \#k_{SPX} \times \#T$, where $\#A$ is the cardinality of set A . We use flattened vectors instead of matrices as the former is more adapted to multilayer perceptrons, and analogously for IVS_{VIX}^{MC} . We generate in total 150,000 data pairs as training set, 20,000 data pairs as validation set, which is used for early stopping to avoid overfitting neural networks, and 10,000 pairs as test set for evaluating the performance of trained neural networks.

- Model training




Figure II.1: Scheme of PtM and MtP method.

We denote the neural network of PtM method by $\mathcal{NN}_{(SPX,VIX)}^{PtM} : \mathbf{R}_+^{2m} \mapsto \Omega \times \mathcal{Z}$. It takes IVS_{SPX} and IVS_{VIX} as input, and outputs estimation of parameters ω and \mathbf{z}_0 . As for MtP method, the network

⁴In our tests we do not calibrate α and fix it to be 0.51. In fact $c_i^n, \gamma_i^n, i = 1, \dots, n$ in K^n only depend on α , so making α constant fixes also $c_i^n, \gamma_i^n, i = 1, \dots, n$. Through experiments with market data, we find $\alpha = 0.51$ is a consistently relevant choice. Results in this chapter are not sensitive to the choice of α , and in practice we can generate few samples with other α and train neural networks with “transfer learning”. One example is given in Appendix II.B.

⁵We do not include S_0 since we look at prices of options with respect to log-moneyness strikes.

consists of two sub-networks, denoted with $\mathcal{N}\mathcal{N}_{SPX}^{MiP} : \Omega \times \mathcal{Z} \mapsto \mathbb{R}_+^m$ and $\mathcal{N}\mathcal{N}_{VIX}^{MiP} : \Omega \times \mathcal{Z} \mapsto \mathbb{R}_+^m$. They aim at approximating the mappings from model parameters to IVS_{SPX} and IVS_{VIX} respectively. The methodology is illustrated in Figure II.1. Table II.1 summarizes some key characteristics of these networks and the training process⁶. Note that for model training, each element of ω and \mathbf{z}_0 is standardized to be in $[-1, 1]$, every point of the IVS^{MC} is subtracted by the sample mean, and divided by the sample standard deviation across training set.

| | $\mathcal{N}\mathcal{N}_{(SPX,VIX)}^{PiM}$ | $\mathcal{N}\mathcal{N}_{SPX}^{MiP}$ | $\mathcal{N}\mathcal{N}_{VIX}^{MiP}$ |
|------------------|--|--------------------------------------|--------------------------------------|
| Input dimension | 120 | 15 | 15 |
| Output dimension | 15 | 60 | 60 |
| Hidden layers | 7 with 25 hidden nodes for each, followed by SiLU activation function, see Hendrycks & Gimpel [2016] | | |
| Training epochs | 150 epochs with early stopping if not improved on validation set for 5 epochs | | |
| Others | Adam optimizer, initial learning rate 0.001, reduced by a factor of 2 every 10 epochs, mini-batch size 128 | | |

Table II.1: Some key characteristics of the networks and the training process.

3.2 Pricing




Figure II.2: Average Monte-Carlo absolute errors of implied volatilities across test set, defined as the half 95% confidence interval of Monte-Carlo simulations.

In this part we first check the ability of neural networks to approximate the pricing function of the model, *i.e.* the mapping from model parameters to IVS_{SPX} and IVS_{VIX} . To see this, we compare the estimations \widehat{IVS}_{SPX} and \widehat{IVS}_{VIX} , given by $\mathcal{N}\mathcal{N}_{SPX}^{MiP}$ and $\mathcal{N}\mathcal{N}_{VIX}^{MiP}$, with the “true” counterparts given by Monte-Carlo method.

Figure II.2 presents the benchmark, given by the half 95% confidence interval of Monte-Carlo simulations for each points on IVS_{SPX} and IVS_{VIX} . We then apply $\mathcal{N}\mathcal{N}_{SPX}^{MiP}$ and $\mathcal{N}\mathcal{N}_{VIX}^{MiP}$ on test set and evaluate the average absolute errors $(\text{abs}(\widehat{IVS}_I - IVS_I^{MC}))_{I=SPX,VIX}$. The results are shown in Figure II.3. We can see that the estimations given by neural networks are close to Monte-Carlo references,

⁶We actually do not have the same number of parameters for the neural networks of the two methods. In fact it is found empirically that the depth of network plays a more important role than the number of parameters, see [Goodfellow et al. \[2016\]](#). Besides, results presented here are not sensitive to the width of hidden layers.

II. Deep calibration of the quadratic rough Heston model

with the majority of points of IVS falling in the 95% confidence interval of Monte-Carlo. We also tested average relative errors as an alternative metric, defined as $(\text{abs}(\overline{IVS}_I - IVS_I^{MC}) \oslash IVS_I^{MC})_{I=SPX,VIX}$, where \oslash means element-wise division between vectors. The results are given in Figure II.C.1 and Figure II.C.2 in Appendix, and are consistent with the observations in Figures II.2-II.3.




Figure II.3: Average absolute errors of implied volatilities across test set for $\mathcal{N}\mathcal{N}_{SPX}^{MtP}$ and $\mathcal{N}\mathcal{N}_{VIX}^{MtP}$. Errors are calculated with respect to Monte-Carlo counterparts.

At this stage, it is reasonable to conclude that $\mathcal{N}\mathcal{N}_{SPX}^{MtP}$ and $\mathcal{N}\mathcal{N}_{VIX}^{MtP}$ are able to learn the pricing functions from data of Monte-Carlo simulations. With these networks we can generate implied volatility surfaces for SPX and VIX for any model parameters.

3.3 Calibration

In this part, we use $\mathcal{N}\mathcal{N}_{(SPX,VIX)}^{PtM}$, $\mathcal{N}\mathcal{N}_{SPX}^{MtP}$ and $\mathcal{N}\mathcal{N}_{VIX}^{MtP}$ to perform model calibration. For PtM method, the output of network gives directly calibration results. For MtP method, the result is given as

$$\hat{\omega}, \hat{\mathbf{z}}_0 = \underset{\omega \in \Omega, \mathbf{z}_0 \in \mathcal{Z}}{\text{argmin}} \|\mathcal{N}\mathcal{N}_{SPX}^{MtP}(\omega, \mathbf{z}_0) - IVS_{SPX}\|_2^2 + \|\mathcal{N}\mathcal{N}_{VIX}^{MtP}(\omega, \mathbf{z}_0) - IVS_{VIX}\|_2^2. \quad (9)$$

where $\Omega = [0.5, 2.5] \times [1.0, 1.5] \times [0.1, 0.6] \times [0.01, 0.5] \times [0.0001, 0.03]$, $\mathcal{Z} = [-0.5, 0.5]^{10}$. We use the same weight for all points of IVS in (9). In practice we could consider different weights to adapt to the varying nature of market quotes in terms of liquidity, bid-ask spread, etc. We apply L-BFGS-B algorithm for this optimization problem. Note that it is a gradient-based algorithm, and the gradients needed are calculated directly with $\mathcal{N}\mathcal{N}_{SPX}^{MtP}$ and $\mathcal{N}\mathcal{N}_{VIX}^{MtP}$ via *automatic adjoint differentiation* (AAD), which is already implemented in popular deep learning frameworks like TensorFlow and PyTorch. One can refer to Section 4.1 for calculation principles.

Calibration on simulated data

We apply the two methods on test set generated by Monte-Carlo simulations, and we evaluate the calibration by Normalized Absolute Errors (NAE) of parameters and the reconstruction Root Mean Square Errors (RMSE) of IVS:

$$\text{NAE} = \frac{|\hat{\theta} - \theta|}{\theta_{up} - \theta_{low}},$$

$$\text{RMSE} = \sqrt{\frac{1}{\#k_I \times \#T} \|\mathcal{N}\mathcal{N}_I^{MtP}(\hat{\omega}, \hat{\mathbf{z}}_0) - IVS_I^{MC}\|_2^2}, \quad I \in \{SPX, VIX\},$$

where θ is one element of $\boldsymbol{\omega}$ or \mathbf{z}_0 , and $\theta_{up}, \theta_{low}$ stand for the upper bound and lower bound of the uniform distribution for sampling θ . Note that we could alternatively use Monte-Carlo to reconstruct IVS with the calibrated parameters instead of $\mathcal{N}\mathcal{N}_{SPX}^{MtP}$ and $\mathcal{N}\mathcal{N}_{VIX}^{MtP}$. However it would be much slower and we have seen in the above that the outputs given by $\mathcal{N}\mathcal{N}_{SPX}^{MtP}$ and $\mathcal{N}\mathcal{N}_{VIX}^{MtP}$ are very close to those of Monte-Carlo. Figure II.C.3 in Appendix shows the empirical cumulative distribution function of NAE for all the calibrated parameters. Figure II.4 gives the empirical distribution of RMSE of IVS with the calibrated parameters from the two methods. We can make the following remarks:

- From Figure II.C.3 in Appendix, we see that PtM method can usually get smaller discrepancy for calibrated parameters than MtP method. This is not surprising since the latter may end with locally optimal solutions when we use gradient-based optimization algorithms.
- MtP method performs better in terms of reconstruction error, which is expected since it is exactly what the algorithm tries to minimize.
- Our results are of course not as accurate as those reported in [Horvath et al. \[2021\]](#) for some other rough volatility models, especially in terms of calibration errors for model parameters. This is because our model is more complex and has more parameters. Consequently more data and more complicated network architecture are demanded for training. The algorithms are also more likely to end with locally optimal solutions.




Figure II.4: Empirical distribution of reconstruction RMSE of IVS.

Calibration on market data

We use MtP approach on market data since there are no “true” reference model parameter values in this case, and the objective is to minimize the discrepancy between model outputs and market observations. The arbitrage-free implied volatility interpolation method presented in [Gatheral & Jacquier \[2014\]](#) is used to generate IVS with the same log-moneyness strikes and maturities as before. Taking the data of 19 May 2017 tested in [Gatheral et al. \[2020\]](#) as example, we get the following calibration results:

$$\boldsymbol{\omega} = (2.5, 1.485, 0.401, 0.235, 0.001),$$

$$\mathbf{z}_0 = (-0.033, 0.015, -0.004, 0.017, 0.028, 0.098, 0.192, -0.076, 0.072, -0.062).$$

We then use Monte-Carlo to get whole IVS of SPX and VIX with these parameters. The slices corresponding to several existing maturities in the market are shown in Figure II.5 and Figure II.6.

Our model fits very well the globe shape of IVS of SPX and VIX options at the same time. Model’s outputs fall essentially between bid and ask quotes for VIX options. In addition, excellent fits are

II. Deep calibration of the quadratic rough Heston model




Figure II.5: Implied volatilities on SPX options for 19 May 2017. Bid and ask of market volatilities are represented respectively by red and blue points. Green line is the output of model with Monte-Carlo method.




Figure II.6: Implied volatilities on VIX options for 19 May 2017. Bid and ask of market volatilities are represented respectively by red and blue points. Green line is the output of model with Monte-Carlo method.

obtained in terms of at-the-money skew of SPX options ⁷. Note that we do not require the market quotes of SPX options and VIX options to have same maturities. Another example of joint calibration is given in Appendix II.C.2.

Figure II.C.6 in Appendix presents the historical dynamics of parameters from daily calibration on market data ⁸. It is interesting to remark that during the beginning of COVID-19 crisis, a, b, c all increased, which means stronger feedback effect, more distinct feedback asymmetry and larger base variance level. The quantity $Z_0 := \sum_{i=1}^{10} c_i z_0^i$ became negative. All these changes contributed to larger volatility.




Figure II.7: Historical reconstruction RMSE of IVS_{SPX} .




Figure II.8: Historical reconstruction RMSE of IVS_{VIX} .

⁷For quotes far from the money, we can remark a discrepancy between the market and the output with Monte-Carlo method. In fact, the mismatch between the interpolated implied volatility surface and the market on these points, and the deviation of Monte-Carlo means from neural networks' outputs can both induce this discrepancy.

⁸In this chapter we limited the parameters in some restricted intervals to illustrate the methodology with reasonable size of random sampling. The observation that the predetermined bounds are reached for some parameters indicates that these intervals cannot cover all market situations. Interested readers can choose wider ones or use unbounded distribution like Gaussian for random sampling to relax this issue, although it may demand more synthetic data for network training.

Figures II.7 and II.8 show historical RMSE for IVS_{SPX} and IVS_{VIX} . For most time periods, the model can fit very well the two IVS at the same time, with small RMSE. For certain dates with extremely large market moves, we observe some spikes of errors. In fact for these dates market liquidity is usually very concentrated on one type of contract (Call or Put) with specific strikes and maturities. In this case, market IVS is not smooth enough so that the model may fail to output satisfying fit. Alternative more robust IVS interpolation methods could be tested in practice to generate smoother IVS. We could also focus only on some points of interest on IVS by excluding other points from the calibration objective (9). Note that all the empirical results presented here are with $\alpha = 0.51$ and $n = 10$. Interested readers can try other α or bigger n to improve globally historical RMSE, while new synthetic data need to be generated for each new setting.

4 Toy examples for hedging

We have seen in the previous section that the proposed model can jointly fit IVS_{SPX} and IVS_{VIX} with small errors. Then it is important to know how to hedge options with this model. In this section, we give toy examples on synthetic data and market data as well to show how to use the neural networks to perform hedging for vanilla SPX calls. We will see that in our model perfect hedging for these products is possible with only SPX.

4.1 Hedging portfolio computation with neural networks

Let $\mathbf{Z}_t := (Z_t^1, \dots, Z_t^{10})$ and $\mathbf{X}_t := (S_t, \mathbf{Z}_t)$. As indicated in Remark (1), Model (4-6) being Markovian, given strike K , maturity T and model parameters $\boldsymbol{\omega}$, the price of vanilla SPX call at time t is then a function of \mathbf{X}_t . Let $P_t(\mathbf{X}_t; K, T, \boldsymbol{\omega})$ denote this quantity. With the dynamics of $(Z_t^i)_{i=1, \dots, 10}$ in Equation (6), we then have

$$dP_t(\mathbf{X}_t; K, T, \boldsymbol{\omega}) = \delta_t dS_t, \quad (10)$$

where

$$\delta_t = \frac{\partial P_t(\mathbf{X}_t; K, T, \boldsymbol{\omega})}{\partial S_t} + \frac{\eta}{S_t} \sum_{i=1}^{10} \frac{\partial P_t(\mathbf{X}_t; K, T, \boldsymbol{\omega})}{\partial Z_t^i}. \quad (11)$$

Note that the factors $(Z_t^i)_{i=1, \dots, 10}$ can be fully traced because they are assumed to be driven by the same Brownian motion as S , which is observable from market data. With the neural networks approximating the pricing function of the model, we will see that we can then obtain approximation of δ_t for any t . Of course continuous hedging is impossible in practice. Here we perform discrete hedging with time step Δt . The Profit and Loss (P&L) of hedging at $t \in (0, T]$ is given by

$$\mathcal{J}_t = \mathcal{J}_t^\delta - \mathcal{J}_t^P, \quad (12)$$

where

$$\begin{aligned} \mathcal{J}_t^\delta &= \sum_{k=0}^{\lfloor t/\Delta t \rfloor - 1} \hat{\delta}_{t_k} \Delta S_{t_{k+1}} + \hat{\delta}_{t_{\lfloor t/\Delta t \rfloor}} (S_t - S_{\lfloor t/\Delta t \rfloor \Delta t}), \\ \mathcal{J}_t^P &= P_t - P_0, \end{aligned}$$

with $t_k = k\Delta t$, $\Delta S_{t_k} := S_{t_k} - S_{t_{k-1}}$, $\hat{\delta}_{t_k}$ is the hedging ratio given by neural networks at k -th hedging time t_k , and P_t is the price of the SPX option to hedge at time t . As we can see, \mathcal{J}^δ stands for the P&L coming from holding the underlying and \mathcal{J}^P reflects the price evolution of the option. We show in the following that $\hat{\delta}_{t_k}$ can be given directly from $\mathcal{N} \mathcal{N}_{SPX}^{MIP}$ in our model.

Note that $\mathcal{N}\mathcal{N}_{SPX}^{MtP}$ behaves like a “global” pricer that is reusable under any model parameters ω . Given ω , we could actually train a finer network as a “local” pricer taking only \mathbf{z}_0 as input. Of course the same methodology as before can be used with fixed ω . Here we apply alternatively Differential Machine Learning as a fast method to obtain approximation of pricing function from simulated paths, under a given calibration of the model, see [Huge & Savine \[2020\]](#).

Method 1: with $\mathcal{N}\mathcal{N}_{SPX}^{MtP}$

With $\mathcal{N}\mathcal{N}_{SPX}^{MtP}$ outputting implied volatilities with respect to log-moneyness strikes, we have

$$P_t(\mathbf{X}_t; K, T, \omega) \simeq P^{BS}(S_t, K, T - t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\omega, \mathbf{Z}_t)),$$

where $P^{BS}(S, K, T, \sigma)$ is the price of European call under Black-Scholes model and $\sigma_{\mathcal{N}\mathcal{N}}^{k, T}(\omega, \mathbf{Z}_t)$ is the implied volatility corresponding to log-moneyness strike k and maturity T , calculated directly by $\mathcal{N}\mathcal{N}_{SPX}^{MtP}$ with (ω, \mathbf{Z}_t) as input. Then the partial derivatives in (11) are given by

$$\begin{aligned} \frac{\partial P_t(\mathbf{X}_t; K, T, \omega)}{\partial S_t} &\simeq \delta_{BS}(S_t, K, T - t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\omega, \mathbf{Z}_t)) + \\ &v_{BS}(S_t, K, T - t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\omega, \mathbf{Z}_t)) \frac{\partial \sigma_{\mathcal{N}\mathcal{N}}^{k, T-t}(\omega, \mathbf{Z}_t)}{\partial k} \Big|_{k=\log(K/S_t)} \frac{\partial \log(K/S_t)}{\partial S_t}, \end{aligned} \quad (13)$$

$$\frac{\partial P_t(\mathbf{X}_t; K, T, \omega)}{\partial Z_t^i} \simeq v_{BS}(S_t, K, T - t, \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\omega, \mathbf{Z}_t)) \frac{\partial \sigma_{\mathcal{N}\mathcal{N}}^{\log(K/S_t), T-t}(\omega, \mathbf{Z}_t)}{\partial Z_t^i}, \quad (14)$$

where $\delta_{BS}(S, K, T, \sigma)$ and $v(S, K, T, \sigma)$ stand for the *Delta* and *Vega* respectively under Black-Scholes model. The quantity $\frac{\partial \sigma_{\mathcal{N}\mathcal{N}}^{k, T}}{\partial Z_t^i}$ corresponds actually to the derivative of the outputs of $\mathcal{N}\mathcal{N}_{SPX}^{MtP}$ with respect to its inputs. Thus it can be obtained instantaneously with built-in AAD. $\frac{\partial \sigma_{\mathcal{N}\mathcal{N}}^{k, T-t}(\omega, \mathbf{Z}_t)}{\partial k}$ in (13) can be approximated by the finite difference $\frac{(\sigma_{\mathcal{N}\mathcal{N}}^{k+\delta_k, T-t} - \sigma_{\mathcal{N}\mathcal{N}}^{k-\delta_k, T-t})}{2\delta_k}(\omega, \mathbf{Z}_t)$. Note that some interpolation methods need to be applied for arbitrary pair $(K, T - t)$ since $\mathcal{N}\mathcal{N}_{SPX}^{MtP}$ has fixed log-moneyness strikes and maturities.

Method 2: with Differential Machine Learning

Given parameters ω , we can simulate a path of model state $(\mathbf{X}_t)_{0 \leq t \leq T}$ starting from the initial state \mathbf{X}_0 . The pathwise payoff $(S_T - K)_+$ is in fact an unbiased estimation of $P_0(\mathbf{X}_0; K, T, \omega)$. Under some regularity conditions, the pathwise derivative $\frac{\partial (S_T - K)_+}{\partial \mathbf{X}_0}$ is also an unbiased estimation of $\frac{\partial P_0(\mathbf{X}_0; K, T, \omega)}{\partial \mathbf{X}_0}$. We show in the following how to calculate this quantity with the simulation scheme proposed in (18). The basic idea of Differential Machine Learning [[Giles & Glasserman, 2006](#); [Huge & Savine, 2020](#)] is to concatenate pathwise payoff and pathwise derivatives as targets to train a neural network, denoted by $\mathcal{N}\mathcal{N}_{DML}$, to approximate the pricing mapping from \mathbf{X}_0 to P_0 under some fixed ω . Thus, the training samples are like $\{\mathbf{X}_0, ((S_T - K)_+, \frac{\partial (S_T - K)_+}{\partial \mathbf{X}_0})\}$, and the loss function for training is like

$$\mathcal{L} = \mathcal{L}_1(\mathcal{N}\mathcal{N}_{DML}(\mathbf{X}_0), (S_T - K)_+) + \mathcal{L}_2\left(\frac{\partial \mathcal{N}\mathcal{N}_{DML}(\mathbf{X}_0)}{\partial \mathbf{X}_0}, \frac{\partial (S_T - K)_+}{\partial \mathbf{X}_0}\right), \quad (15)$$

with \mathcal{L}_1 and \mathcal{L}_2 some suitably chosen loss functions. Similarly to the case with $\mathcal{N}\mathcal{N}_{SPX}^{MtP}$, $\frac{\partial \mathcal{N}\mathcal{N}_{DML}(\mathbf{X}_0)}{\partial \mathbf{X}_0}$ can be calculated efficiently with AAD. In this way, $\mathcal{N}\mathcal{N}_{DML}$ aims at learning both the pricing function

II. Deep calibration of the quadratic rough Heston model

and its derivatives during training. This can help the networks converge with few samples, see [Huge & Savine \[2020\]](#).

The quantity $\frac{\partial(\hat{S}_T - K)_+}{\partial \mathbf{X}_0}$ can be calculated with AAD, which is based on the chain rule of derivatives computation, see [Giles & Glasserman \[2006\]](#) for more details. In our case, with the Euler scheme in (18), let Δt the simulation step with $N\Delta t = T$. We have

$$\begin{aligned}\frac{\partial \hat{S}_{k+1}}{\partial X_0^j} &= (1 + \sqrt{\hat{V}_k}(W_{k+1} - W_k)) \frac{\partial \hat{S}_k}{\partial X_0^j} + \frac{a \hat{S}_k (W_{k+1} - W_k) (\sum_{i=1}^{10} c_i \hat{Z}_k^i - b)}{\sqrt{\hat{V}_k}} \sum_{i=1}^{10} c_i \frac{\partial \hat{Z}_k^i}{\partial X_0^j}, \\ \frac{\partial \hat{Z}_{k+1}^i}{\partial X_0^j} &= \frac{1}{1 + \gamma_i \Delta t} \left(\frac{\partial \hat{Z}_k^i}{\partial X_0^j} + (-\lambda \Delta t + \frac{\eta a (W_{k+1} - W_k) (\sum_{i=1}^{10} c_i \hat{Z}_k^i - b)}{\sqrt{\hat{V}_k}}) \sum_{i=1}^{10} c_i \frac{\partial \hat{Z}_k^i}{\partial X_0^j} \right),\end{aligned}$$

where X_0^j is the j -th element of \mathbf{X}_0 . Let $\hat{\mathbf{X}}_k := (\hat{S}_k, \hat{Z}_k^1, \dots, \hat{Z}_k^{10})$. This can be rewritten in matrix form $\Delta(k+1) = \mathbf{D}(k)\Delta(k)$, with $\Delta_{i,j}(k) = \frac{\partial \hat{X}_k^i}{\partial X_0^j}$, $i, j = 1, \dots, 11$, and

$$\mathbf{D}(k) = \begin{pmatrix} 1 + \sqrt{\hat{V}_k} \Delta W_{k+1}, & \hat{S}_k M_k^1 c_1, & \dots & \hat{S}_k M_k^1 c_{10} \\ 0, & \frac{1 + c_1 M_k^2}{1 + \gamma_1 \Delta t}, & \dots & \frac{c_{10} M_k^2}{1 + \gamma_{10} \Delta t} \\ \vdots & \vdots & \ddots & \vdots \\ 0 & \frac{c_1 M_k^2}{1 + \gamma_1 \Delta t}, & \dots & \frac{1 + c_{10} M_k^2}{1 + \gamma_{10} \Delta t} \end{pmatrix},$$

where

$$\begin{aligned}\Delta W_{k+1} &:= W_{k+1} - W_k, \\ M_k^1 &:= \frac{a \Delta W_{k+1} (\sum_{i=1}^{10} c_i \hat{Z}_k^i - b)}{\sqrt{\hat{V}_k}}, \\ M_k^2 &:= \eta M_k^1 - \lambda \Delta t.\end{aligned}$$

Then we have

$$\begin{aligned}\frac{\partial(\hat{S}_T - K)_+}{\partial \mathbf{X}_0} &= \Delta(N)^T \frac{\partial(\hat{S}_T - K)_+}{\partial \hat{\mathbf{X}}_N} \\ &= \Delta(0)^T \mathbf{D}(0)^T \dots \mathbf{D}(N-2)^T \mathbf{D}(N-1)^T \frac{\partial(\hat{S}_T - K)_+}{\partial \hat{\mathbf{X}}_N} \\ &= \Delta(0)^T \mathbf{V}(0),\end{aligned}\tag{16}$$

where $\Delta(0)$ is simply the identity matrix by definition, and $\mathbf{V}(0)$ can be calculated recursively:

$$\begin{aligned}\mathbf{V}(k) &= \mathbf{D}(k)^T \mathbf{V}(k+1), \quad k = 0, \dots, N-1, \\ \mathbf{V}(N) &= \frac{\partial(\hat{S}_T - K)_+}{\partial \hat{\mathbf{X}}_N} = \mathbf{1}(\hat{S}_T > K)(1, 0, \dots, 0)^T.\end{aligned}\tag{17}$$

Note that for each simulated path, $(\mathbf{D}(k))_{k=0, \dots, N-1}$ can be readily obtained, so the quantity $\frac{\partial(\hat{S}_T - K)_+}{\partial \mathbf{X}_0}$ can be efficiently computed following (16, 17).

Since we use the same trained network for hedging at any $t \in (0, T]$, to accommodate the varying time to maturity $(T - t)$, we consider multiple outputs corresponding to different maturities T_1, T_2, \dots, T_m for \mathcal{NN}_{DML} , with $T_1 < T < T_m$. The quantities $(\frac{\partial(\hat{S}_{T_i-K})_+}{\partial \mathbf{X}_0})_{i=1, \dots, m}$ can be computed following (16, 17). For training, we use the average of m derivatives in (15) for simplification, see Huge & Savine [2020] for more details on designs with multi-dimensional output. After training \mathcal{NN}_{DML} with respect to strike K and parameter $\boldsymbol{\omega}$, we have

$$P_t(\mathbf{X}_t; K, T, \boldsymbol{\omega}) \simeq \mathcal{NN}_{DML}^{T-t}(\mathbf{X}_t),$$

with $\mathcal{NN}_{DML}^{T-t}(\cdot)$ the output corresponding to maturity $(T - t)$. Thus we get

$$\begin{aligned} \frac{\partial P_t(\mathbf{X}_t; K, T, \boldsymbol{\omega})}{\partial S_t} &\simeq \frac{\partial \mathcal{NN}_{DML}^{T-t}(\mathbf{X}_t)}{\partial S_t}, \\ \frac{\partial P_t(\mathbf{X}_t; K, T, \boldsymbol{\omega})}{\partial Z_t^i} &\simeq \frac{\partial \mathcal{NN}_{DML}^{T-t}(\mathbf{X}_t)}{\partial Z_t^i}. \end{aligned}$$

Then the formula in (11) is used to obtain hedging ratio. As in the case with \mathcal{NN}_{SPX}^{MTP} , some interpolation methods are needed for arbitrary $(T - t)$. In our tests, we choose the following characteristics for \mathcal{NN}_{DML} and its training:

- 4 hidden layers, with 20 hidden nodes for each, SiLU as activation function,
- input dimension is 11, output dimension is 5 corresponding to maturities 0.02, 0.04, 0.06, 0.08, 0.1,
- mini-batch gradient descent with batch size 128, initial learning rate 0.001, divided by 2 every 5 epochs,
- sample uniformly \mathbf{X}_0 and simulate 50,000 paths, train \mathcal{NN}_{DML} with 20 epochs.

4.2 Numerical results

Hedging on synthetic data

Without loss of generality, we take the following parameters in the experiments:

- $\boldsymbol{\omega} = (1, 1.2, 0.35, 0.2, 0.0025)$
- $S_0 = 100, Z_0^i = 0, i = 1, \dots, 10$
- $K = 98, T = 0.08$

First we generate 50,000 paths of $(S_t)_{0 \leq t \leq T}$ by following the scheme 18 with time step $\Delta t = 0.0012$. The price P_0 is estimated by the average of pathwise payoffs. Then we evaluate \mathcal{J}_T for 5000 paths among them, with hedging time step $\Delta t \in \{0.0012, 0.0036\}$. From Figure II.9 we can see that both methods lead to hedging payoffs around 0. Hedging less frequently brings slightly larger variance of payoffs. We also remark that the method with \mathcal{NN}_{SPX}^{MTP} generates smaller variance than the other with \mathcal{NN}_{DML} . It is expected as the latter is trained with pathwise labels while the former is trained with “true” labels given by Monte-Carlo means, which have certainly smaller variance.

Hedging on market data

We perform daily hedging on two SPX monthly calls:

1. Maturity date 16 June 2017, strike 2425, hedged since 18 May 2017.

II. Deep calibration of the quadratic rough Heston model




Figure II.9: Hedging P&L with \mathcal{NN}_{DML} and \mathcal{NN}_{SPX}^{MTP} on simulated data. P_0 given by Monte-Carlo is 2.9.




Figure II.10: Hedging P&L on two SPX calls with maturity 16 June 2017, strike 2425 (left), and maturity 16 February 2018, strike 2750 (right).

2. Maturity date 16 February 2018, strike 2750, hedged since 18 January 2018.

On first day of hedging, we take market data for model calibration and get ω and \mathbf{z}_0 . \mathcal{NN}_{DML} is then trained on paths generated under ω . Then for each following day, we update the value of factors $(Z^i)_{i=1,\dots,10}$ with respect to the evolution of SPX, and we compute the hedging portfolio with \mathcal{NN}_{SPX}^{MTP} and \mathcal{NN}_{DML} as explained in the above. Besides, we also test with Black-Scholes model where the implied volatility of the first day is used to compute *Delta* as hedging ratio. Figure II.10 presents the evolution of \mathcal{J}^P and \mathcal{J}^δ for these two examples. Note that all quantities are divided by P_0 to be unitless. We see that \mathcal{NN}_{SPX}^{MTP} and \mathcal{NN}_{DML} can follow very well the market price of options, with smaller $|\mathcal{J}_T|$ than Black-Scholes approach. Of course, in practice we need to consider more elements like hedging cost, slippage, and to do more tests for systematic comparison, but this is out of the scope of our current work.

5 Conclusion

We have seen that the deep neural networks can be used in calibrating the quadratic rough Heston model with reliable results. The training of network demands indeed lots of simulated data, especially when the dimension of model parameters is high. However, it is done off-line only once and the network will be reusable in many situations. Under the particular setting of our model, we can also use the network for risk hedging. Certainly, we can still improve the results presented in the above, for example fixing finer grids of strikes and maturities, or using more factors for the approximation. We emphasize that the methodologies presented in our work are of course not limited to the model introduced here, and can be adapted to other models and other financial products.

Appendix

II.A Kernel function approximation and simulation scheme

Here we recall the geometric partition of $(c_i^n, \gamma_i^n)_{i=1, \dots, n}$ proposed in [Abi Jaber \[2019\]](#):

$$c_i^n = \int_{\eta_{i-1}^n}^{\eta_i^n} \mu(dx), \quad \gamma_i^n = \frac{1}{c_i^n} \int_{\eta_{i-1}^n}^{\eta_i^n} x \mu(dx), \quad i = 1, \dots, n,$$

with $\mu(dx) = \frac{x^{-\alpha}}{\Gamma(\alpha)\Gamma(1-\alpha)} dx$ and $\eta_i^n = x_n^{i-n/2}$. Hence we have

$$c_i^n = \frac{x_n^{(1-\alpha)(i-n/2)}(1-x_n^{\alpha-1})}{(1-\alpha)\Gamma(\alpha)\Gamma(1-\alpha)}, \quad \gamma_i^n = \frac{(1-\alpha)x_n^{i-1-n/2}(x_n^{2-\alpha}-1)}{(2-\alpha)(x_n^{1-\alpha}-1)}.$$

Given n , α and T , we can determine the “optimal” x_n as

$$x_n^*(\alpha, T) = \arg \min_{x_n^* > 1} \|K^n - K\|_{L^2(0, T)}.$$

We fix $T = 0.1$ as we are more interested in short maturities and we get the optimal x_n^* for different n and α as shown in Figure II.A.1. It is consistent with the analysis in [Abi Jaber \[2019\]](#) that given α , we need to increase x_n to mimic roughness with less factors. Given n , x_n^* does not change a lot with α , which indicates that in practice we can actually fix x_n^* independently of α .

Choosing a good n is a trade-off between simulation efficiency and good approximation of rough volatility models. In our test, we take $n = 10$ and $x_{10}^* = 3.92$. On one hand, we can see from Figure II.A.2 that the approximation of K^{10} to K is not far away from other K^n with larger n . On the other hand, it means a margin for improvement with larger n . Figure II.A.3 gives the resulting c_i^n and γ_i^n . We see that the wide range of γ contributes to the multi-timescales nature of volatility processes.

We discard the notation n and use the following modified explicit-implicit Euler scheme for simulating Model (4-6):

$$\begin{aligned} \hat{S}_{k+1} &= \hat{S}_k + \hat{S}_k \sqrt{\hat{V}_k} (W_{k+1} - W_k), & \hat{V}_k &= a(\hat{Z}_k - b)^2 + c, \\ \hat{Z}_{k+1}^i &= \frac{1}{1 + \gamma_i \Delta t} (\hat{Z}_k^i - \lambda \hat{Z}_k \Delta t + \eta \sqrt{\hat{V}_k} (W_{k+1} - W_k)), \\ \hat{Z}_k &= \sum_{i=1}^n c_i \hat{Z}_k^i, \end{aligned} \tag{18}$$

for a time step $\Delta t = T/N$, $k = 1, \dots, N$ and $(W_{k+1} - W_k) \sim \mathcal{N}(0, \Delta t)$. One could also use the explicit scheme for Z^i given by

$$\hat{Z}_{k+1}^i = (1 - \gamma_i \Delta t) \hat{Z}_k^i - \lambda \hat{Z}_k \Delta t + \eta \sqrt{\hat{V}_k} (W_{k+1} - W_k).$$

II. Deep calibration of the quadratic rough Heston model




Figure II.A.1: x_n^* for different n and α .

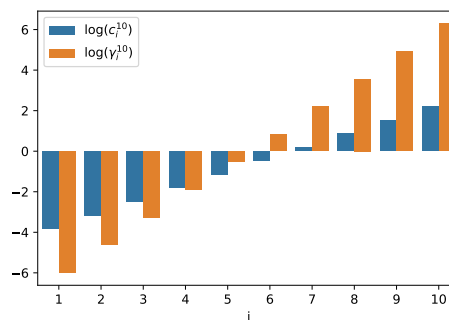



Figure II.A.2: $K(t)$ and $K^n(t)$ with different n Figure II.A.3: Logarithm of c_i^{10} and γ_i^{10} by taking $x_{10}^* = 3.92$.

However, with above x_{10}^* , we get $\gamma_{10} = 542.32$. One then need Δt to be necessarily small to ensure the scheme's stability. Instead we could use

$$\hat{Z}_{k+1}^i - \hat{Z}_k^i = -\gamma_i \Delta t \hat{Z}_{k+1}^i - \lambda \hat{Z}_k \Delta t + \eta \sqrt{\hat{V}_k} (W_{k+1} - W_k),$$

which leads to Scheme 18 and avoids this issue.

II.B Network training with transfer learning

When we switch to models with α not equal to 0.51, we can apply the idea of transfer learning to accelerate network training. More precisely, we use the parameters of the network corresponding to the case $\alpha = 0.51$ to initialise the network for cases with different α . Here we give an example on the training of $\mathcal{N}\mathcal{N}_{SPX}^{MTP}$ with $\alpha = 0.65$. With 10,000 training samples, we can see from Figure II.B.1 that




Figure II.B.1: Training loss evolution for $\alpha = 0.65$ with 10,000 samples.

transfer learning can help the training converge much faster to a lower loss than the one with random parameter initialization.

I.I.C Pricing and calibration with neural networks

I.I.C.1 On simulated data

Here we present the average relative pricing errors across test as an alternative evaluation metric. Figure II.C.1 stands for the benchmark given by Monte-Carlo and Figure II.C.2 gives the results with $\mathcal{N}\mathcal{N}_{SPX}^{MTP}$ and $\mathcal{N}\mathcal{N}_{VIX}^{MTP}$. Figure II.C.3 shows the accuracy of calibrated parameters by plotting the empirical CDFs of NAE.




Figure II.C.1: Average Monte-Carlo relative errors of implied volatilities across test set, defined as the normalized half 95% confidence interval of Monte-Carlo simulations.

II. Deep calibration of the quadratic rough Heston model




Figure II.C.2: Average relative errors of implied volatilities across test set for $\mathcal{N}\mathcal{N}_{SPX}^{MIP}$ and $\mathcal{N}\mathcal{N}_{VIX}^{MIP}$.

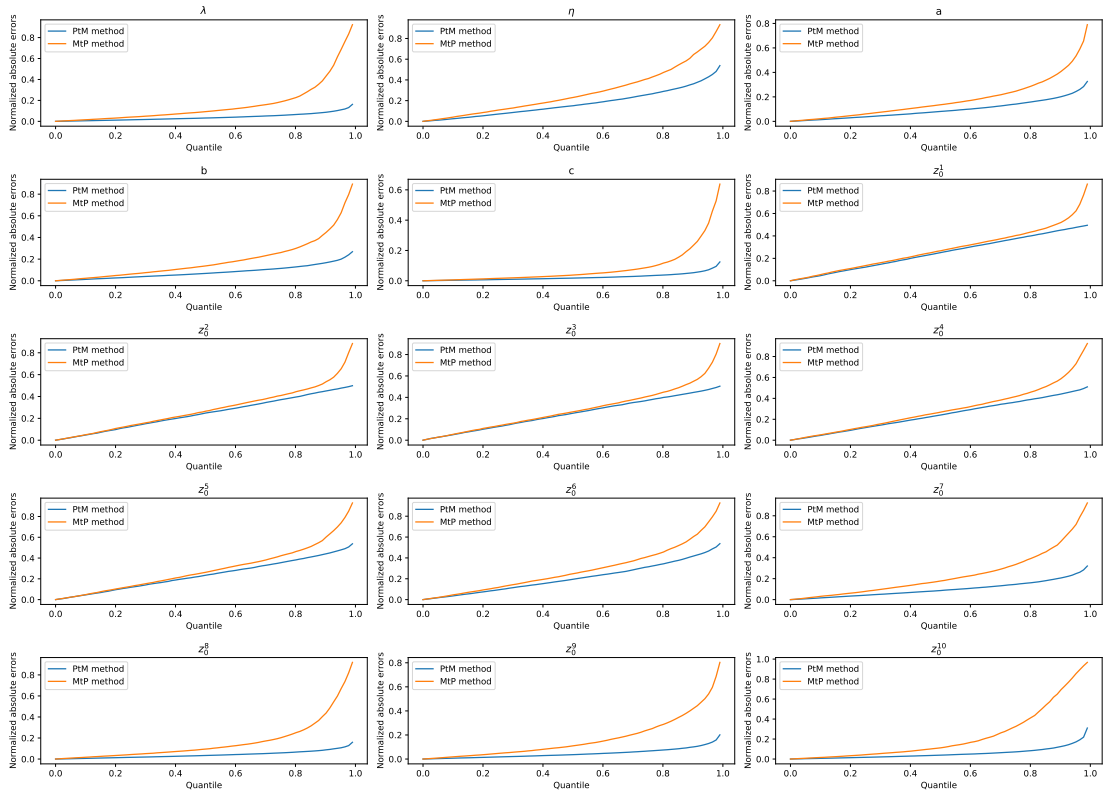


Figure II.C.3: Empirical CDF of NAE of calibrated parameters.

I.I.C.2 On market data

We give another IVS fit example on 10 September 2019. The parameters calibrated are:

$$\omega = (1.717, 1.5, 0.265, 0.246, 0.0001),$$

$$\mathbf{z}_0 = (-0.009, 0.015, 0.011, 0.036, 0.002, -0.011, -0.018, 0.074, 0.142, -0.171).$$

The Monte-Carlo results with these parameters are shown in Figure I.I.C.4 and Figure I.I.C.5. Figure I.I.C.6 presents the historical dynamics of parameters from daily calibration on market data.




Figure I.I.C.4: Implied volatilities fit on SPX options for 10 September 2019. Bid and ask of market volatilities are represented respectively by red and blue points. Green line is the output of model with Monte-Carlo method.

II. Deep calibration of the quadratic rough Heston model




Figure II.C.5: Implied volatilities fit on VIX options for 10 September 2019. Bid and ask of market volatilities are represented respectively by red and blue points. Green line is the output of model with Monte-Carlo method.




Figure II.C.6: Historical dynamics of calibrated parameters. The orange lines are 10-day moving average of parameters. Recall that $Z_0 := \sum_{i=1}^{10} c_i z_0^i$.

Multi-asset market making under quadratic rough Heston

Abstract

Given the promising results on joint modeling of SPX/VIX smiles of the recently introduced quadratic rough Heston model, we consider a multi-asset market making problem on the derivatives of SPX, *e.g.* SPX/VIX futures, SPX/VIX options. The market maker tries to maximize its profit from spread capturing while controlling the portfolio's inventory risk, which can be fully explained by the value change of SPX under the particular setting of the quadratic rough Heston model. The high dimensionality of the resulting optimization problem is relaxed by several approximations. An asymptotic closed-form solution can be obtained. The accuracy and relevance of the approximations are illustrated through numerical experiments.

Keywords— Multi-asset market making, quadratic rough Heston, quadratic approximation

1 Introduction

It is well known that the celebrated Black-Scholes model's constant volatility assumption is inconsistent with empirical observations of financial time series and uneven implied volatility surfaces. Tackling this issue by regarding the volatility as a continuous-time random process, stochastic volatility models are able to reproduce several stylized facts of historical data, such as the fat-tailed distribution of returns and the volatility clustering phenomenon. However, the shapes of the implied volatility surfaces generated by conventional stochastic volatility models, such as the Hull and White, Heston, and SABR models, usually differ substantially from those of empirical observations, see for instance Bayer et al. [2016]. The rough volatility paradigm brings new solutions, being able to achieve superior fits of implied volatility surfaces than the former models [Bayer et al., 2016; El Euch et al., 2019; Gatheral et al., 2018]. The recently introduced quadratic rough Heston (QRH) model shows its potential in calibrating jointly SPX and VIX smiles [Gatheral et al., 2020]. It models the price of an asset S (here the SPX) and its spot variance V under risk-neutral measure as

$$dS_t = S_t \sqrt{V_t} dW_t, \quad V_t = a(Z_t - b)^2 + c,$$

where W is a Brownian motion, a, b, c are all positive constants and Z_t is defined as

$$Z_t = \int_0^t \lambda \frac{(t-s)^{\alpha-1}}{\Gamma(\alpha)} (\theta_0(s) - Z_s) ds + \int_0^t \eta \frac{(t-s)^{\alpha-1}}{\Gamma(\alpha)} \sqrt{V_s} dW_s, \quad (1)$$

where $\alpha \in (1/2, 1)$, $\lambda > 0$, $\eta > 0$, and $\theta_0(\cdot)$ is a deterministic function. The fractional kernel $K(t) = \frac{t^{\alpha-1}}{\Gamma(\alpha)}$ enables us to generate rough volatility dynamics. It is also used in the rough Heston model, under which the volatility trajectories have almost surely Hölder regularity $\alpha - 1/2 - \varepsilon$, for any $\varepsilon > 0$ [El Euch & Rosenbaum, 2019]. This actually recalls the observation in Gatheral et al. [2018] that the dynamic of log-volatility is similar to that of a fractional Brownian motion with Hurst parameter of order 0.1. Importantly, the QRH model gives a natural way to encode the *strong Zumbach effect* [Dandapani et al., 2021], which means that the conditional law of future volatility depends on the past not only through past volatility trajectories but also through past returns.

Note that in this model, only one Brownian motion is involved to generate the stochastic nature of the model and the dynamic of the volatility can be fully explained by past returns. While this is quite different from conventional settings, where volatility is exposed to additional random factors, satisfactory results on market data are reported [Gatheral et al., 2020; Rosenbaum & Zhang, 2022a]. Guyon & Lekeufack [2022] also emphasize that price returns largely explain volatility. Rosenbaum & Zhang [2022b] present an extensive empirical study on the formation process of realized volatility, underlying the predictive power of price returns on future realized volatility. The QRH model gives opportunities to model consistently SPX derivatives, including SPX/VIX futures and SPX/VIX options. This motivates us to consider a multi-asset market making problem on the aforementioned derivatives. Today these assets are actively traded and most of the volume comes from electronic trading platforms based on limit order books, which justifies the relevance of developing automated market making algorithms.

The academic literature on optimal market making dates back to the 80s. Ho & Stoll [1981] formulate the challenge faced by a market maker, which consists in maximizing its profit from spread capturing with a particular focus on execution uncertainty and inventory risk, as a dynamic programming problem. Approximate solutions are given for specific functional forms of the order arrival rates. The approach is revived in the seminal work Avellaneda & Stoikov [2008] with tools of stochastic optimal control. Since then, extensive literature on optimal market making has been developed. Guéant et al. [2013]

show that under inventory constraints, the Hamilton-Jacobi-Bellman (HJB) equations associated with the control problem introduced in [Avellaneda & Stoikov \[2008\]](#) can be reduced to a linear system of ordinary differential equations (ODEs), for which the asymptotic of the closed-form solution is given when the considered time horizon tends to infinity. Various features, such as model misspecification, short-term alpha signals, and microstructural characteristics are investigated in [Cartea et al. \[2014, 2017, 2018\]](#). In these papers, a mean-variance type objective function is considered instead of the Von Neumann-Morgenstern expected utility as in [Avellaneda & Stoikov \[2008\]](#); [Guéant et al. \[2013\]](#). [Guéant \[2017\]](#) extend the results of [Guéant et al. \[2013\]](#) to more general order execution rates and the two aforementioned objective functions are reconciled. It is shown that the HJB equations associated with the two classes of objective functions can be sorted out into the same family of ODEs.

Multi-asset market making problems are usually exposed to the curse of dimensionality. The conventional grid methods typically used in the case of single-asset become inadmissible given the huge amount of computational resources required. [Bergault & Guéant \[2021\]](#) develop a method based on factor decomposition to reduce the dimension of the problem. [Guéant & Manziuk \[2019\]](#) use deep neural networks to approximate the value function and the optimal quoting strategy. Closed-form approximations are obtained in [Bergault et al. \[2021\]](#) replacing the Hamiltonian functions with quadratic ones. A similar idea is followed in [Baldacci, Derchu, & Manziuk \[2021\]](#) in the context of option market making. We refer to [Cartea et al. \[2015\]](#) and [Guéant \[2016\]](#) for a panorama on algorithmic market making.

All the models proposed in the above papers consider continuous control variables, the quotes' distances with respect to the reference price, which makes them more adapted to quote-driven markets such as corporate bond markets and FX markets. The case of the markets ruled by limit order books is addressed in [Guilbaud & Pham \[2013, 2015\]](#), where the market maker decides whether to submit limit orders at the best limits or to get immediate execution using market orders. More dedicated microstructure models reflecting some important empirical observations at the microstructural scale are chosen in [Fodra & Pham \[2015\]](#) and [Lu & Abergel \[2018\]](#).

In this chapter, we consider an agent continuously quoting for SPX derivatives including SPX/VIX futures and SPX/VIX options, via limit orders on both the bid and ask sides. The market maker tries to maximize its expected gain from spread capturing while controlling the inventory risk. In the option market making problems considered in [Baldacci, Bergault, & Guéant \[2021\]](#) and [Baldacci, Derchu, & Manziuk \[2021\]](#), the underlying follows a one-factor stochastic volatility model, and perfect Delta hedging is assumed for European options. The market maker then penalizes the option portfolio's total Vega. In the QRH model, the price dynamics of all the aforementioned derivatives can be explained by the value variation of SPX. Thus, in the context of market making, the principal inventory risk is summarized by the sensitivities of their price change with respect to that of SPX. Even if theoretically this risk can be hedged out perfectly using only SPX⁹ as shown in [Rosenbaum & Zhang \[2022a\]](#), continuous hedging is not realistic and the cost of frequent hedging is far from negligible. More particularly, as the typical inventory process of market making shows a mean reverting behavior around zero, frequent hedging results in unnecessary buys and sells of SPX. In our approach, instead of hedging out completely the net inventory risk of the held derivatives with SPX, the agent also market makes on SPX with limit orders, and some net inventory risk is authorized.

The market making basket selected here is essentially made of large tick assets, for which the effective spread is almost always equal to one tick. Thus the quoting strategy of the market maker considered in

⁹SPX, or formally S&P 500 index, is not tradable. When no ambiguity occurs, it refers to SPX futures without loss of generality. It can also be other liquid index-tracking products such as index funds and ETFs.

this chapter consists in whether to send limit orders on the two best limits given its portfolio inventory, *i.e.* whether to make a market on the best limits, instead of choosing the optimal quoting price as most of the models inspired from that of [Avellaneda & Stoikov \[2008\]](#). The resulting modeling framework is in the spirit of [Guilbaud & Pham \[2013\]](#), while the dimension of the problem is much higher. The high dimensionality comes from two aspects. First, we use here the multi-factor approximation of the QRH model, which is Markovian and suitable for the computation of hedging quantities [[Abi Jaber & El Euch, 2019](#); [Rosenbaum & Zhang, 2022a](#)]. However, this brings additional model states. Second, the portfolio of the market maker should be scaled to dozens of assets, including futures and options of various specifications.

To relax the first one, we assume the volatility of SPX and the price sensitivities of other assets against SPX to be constant during the time horizon of market making. At first sight, it seems to deviate from the important feature of the QRH model that volatility spikes can happen in the case of significant price trends. However, this assumption is relevant given the short time horizon of market making problems, which is usually less than several hours. In fact, we always follow the approximated QRH model for the computation of related quantities when switching from one market making period to another, *e.g.* at the beginning of each day, and we only take them as constant during the following period. As for the multi-asset nature of the problem, we apply the idea introduced in [Bergault et al. \[2021\]](#) of approximating the HJB equation with another one whose asymptotic closed-form solution can be deduced. Having closed-form approximate solutions makes the recalibration of the algorithm very efficient, and thus it is even possible to design an “online” market making algorithm, *i.e.* the algorithm is fed regularly with the latest parameter values inside each market making period.

This chapter is organized as follows. In Section 2, we recall first the multi-factor approximation of the QRH model and then describe the considered multi-asset market making problem. Several approximations are introduced in Section 3 to reduce the dimension of the problem. The asymptotic value function associated with the approximated problem is given in closed-form. Finally, in Section 4, we evaluate the relevance of these approximations through numerical experiments.

2 Description of the problem

2.1 Multi-factor approximation of the QRH model

Inspired by [Abi Jaber & El Euch \[2019\]](#), in [Rosenbaum & Zhang \[2022a\]](#) a multi-factor approximation of the QRH model is introduced to make the model Markovian, and an efficient calibration procedure based on deep learning is designed. Essentially, the multi-factor approximation consists in replacing $K(t) := \frac{t^{\alpha-1}}{\Gamma(\alpha)}$ in Equation (1) with the approximated kernel function $K^n(t) = \sum_{i=1}^n c_i e^{-\gamma_i t}$. Then we have

$$dZ_t^i = (-\gamma_i Z_t^i - \lambda \sum_{i=1}^n c_i Z_t^i) dt + \eta \sqrt{V_t} dW_t, \quad Z_0^i = z_0^i, \quad (2)$$

and $Z_t = \sum_{i=1}^n c_i Z_t^i$. The $2n$ parameters $(c_i, \gamma_i)_{i=1, \dots, n}$ are explicit functions of $\alpha \in (\frac{1}{2}, 1)$, instead of being free parameters to calibrate, see [Rosenbaum & Zhang \[2022a\]](#) and the references therein for more details. Now given the Markovian nature of the approximated version of the model, the price of many derivatives of SPX at time t , *e.g.* VIX futures, vanilla SPX and VIX options, can be obtained as a function of $X_t := (S_t, Z_t^1, \dots, Z_t^n)$. In fact, with $E_t[\cdot] := E[\cdot | \mathcal{F}_t]$ being the conditional expectation under the risk-neutral measure, we have the following representations:

- VIX future with expiration T : $\mathbb{E}_t[\text{VIX}_T]$ with $\text{VIX}_t := 100 \sqrt{-2\mathbb{E}_t[\log(S_{t+\Delta}/S_t)]}$, $\Delta = 30$ days.

- Vanilla SPX call or put with expiration T and strike K : $\mathbb{E}_t[(S_T - K)_+]$ or $\mathbb{E}_t[(K - S_T)_+]$.
- Vanilla VIX call or put with expiration T and strike K : $\mathbb{E}_t[(\text{VIX}_T - K)_+]$ or $\mathbb{E}_t[(K - \text{VIX}_T)_+]$.

Let $P(t, X)$ denote one of these quantities. Under usual regularity conditions¹⁰, we can write

$$dP(t, X) = \delta_t dS_t,$$

where

$$\delta_t = \frac{\partial P(t, X)}{\partial S} + \frac{\eta}{S} \sum_{i=1}^n \frac{\partial P(t, X)}{\partial Z^i}.$$

An example with vanilla SPX call is taken in [Rosenbaum & Zhang \[2022a\]](#) to show the computation of δ_t using neural networks. The same approach can be extended to VIX futures and other vanilla options¹¹.

In [Rosenbaum & Zhang \[2022a\]](#), the approximation of the pricing mapping $P(\cdot)$ is learned directly from samples generated through Monte-Carlo simulations. One can also obtain $P(\cdot)$ through solving the following partial differential equation (PDE):

$$\frac{\partial P}{\partial t} + \sum_{i=1}^n \frac{\partial P}{\partial Z^i} (-\gamma_i Z^i - \lambda \sum_i c_i Z^i) + \frac{1}{2} \sum_{i=1}^n \frac{\partial^2 P}{\partial S \partial Z^i} \eta SV + \frac{1}{2} \sum_{i,j} \frac{\partial^2 P}{\partial Z^i \partial Z^j} \eta^2 V = 0, \quad (3)$$

with terminal condition $P(T, X) = g(X)$, where $g(\cdot)$ represents the payoff function of the considered derivative. It is *a priori* intricate to solve Equation (3) with conventional methods, especially when n is large. One can apply the method with the help of deep learning introduced in [Sirignano & Spiliopoulos \[2018\]](#), which is reported to be effective on high-dimensional PDEs. For the tests presented in Section 4, we rely on the method introduced in [Rosenbaum & Zhang \[2022a\]](#) for the computation of δ_t .

2.2 Multi-asset market making

Our market making problem is considered over the period $[0, T]$, with T smaller than the expiration of any asset in our selected basket. For asset $j \in \{1, \dots, d\}$, at each point in time, the market maker decides whether to make a market at the limits P_t^j plus/minus one-half tick size. Let $l_t^{j,b}, l_t^{j,a} \in \{0, 1\}$ represent the decisions concerning respectively the bid and ask sides. The value 1 means the participation of the market maker with a limit order of constant size m^j , and 0 stands for the case where the order is not placed in the queue corresponding to the best limits so that the probability of execution is small. We denote the two point processes modeling the number of transactions at the bid and ask sides by $(N_t^{j,b})_{t \in [0, T]}$ and $(N_t^{j,a})_{t \in [0, T]}$ respectively. Then the dynamics of the inventory process $(q_t^j)_{t \in [0, T]}$ of asset j is given by

$$dq_t^j = m^j dN_t^{j,b} - m^j dN_t^{j,a}.$$

We denote by $(\lambda_t^{j,b})_{t \in [0, T]}$ and $(\lambda_t^{j,a})_{t \in [0, T]}$ the intensity processes of $(N_t^{j,b})_{t \in [0, T]}$ and $(N_t^{j,a})_{t \in [0, T]}$ respectively, which verify

$$\lambda_t^{j,b} = l_t^{j,b} \Lambda^{j,b} \mathbb{1}_{\{q_t^j + m^j \leq -Q^j\}}, \quad \lambda_t^{j,a} = l_t^{j,a} \Lambda^{j,a} \mathbb{1}_{\{q_t^j - m^j \geq -Q^j\}},$$

¹⁰In practice the price of any derivative concerned here can be approached by a regular approximator, such as the neural network used in [Rosenbaum & Zhang \[2022a\]](#).

¹¹For SPX we have naturally $\delta_t \equiv 1$.

III. Multi-asset market making under quadratic rough Heston

where Q^j stands for the maximum inventory of asset j that the market maker is willing to hold, which is a multiple of m^j without loss of generality. We do not differentiate the queue positions of the orders in this work, and denote the constant execution probability at the bid and ask sides by $\Lambda^{j,b}$ and $\Lambda^{j,a}$ respectively.

Denoting by D^j the tick size of asset j , the resulting dynamics of the cash process $(Y_t)_{t \in [0, T]}$ of the market maker is:

$$\begin{aligned} dY_t &= \sum_{j=1}^d -\left(P_t^j - \frac{D^j}{2}\right) m^j dN_t^{j,b} + \left(P_t^j + \frac{D^j}{2}\right) m^j dN_t^{j,a}, \\ &= \sum_{j=1}^d \frac{D^j}{2} m^j (dN_t^{j,b} + dN_t^{j,a}) - P_t^j dq_t^j. \end{aligned}$$

Let $(\Pi_t)_{t \in [0, T]}$ be the process representing the mark-to-market value of the market maker's portfolio, *i.e.* $\Pi_t := Y_t + \sum_{j=1}^d P_t^j q_t^j$. Its dynamics are given by

$$\begin{aligned} d\Pi_t &= dY_t + \sum_{j=1}^d P_t^j dq_t^j + \sum_{j=1}^d q_t^j dP_t^j, \\ &= \sum_{j=1}^d \frac{D^j}{2} m^j (dN_t^{j,b} + dN_t^{j,a}) + q_t^j \delta_t^j dS_t. \end{aligned}$$

The objective function of the market maker considered here is akin to those used in [Cartea et al. \[2017, 2014, 2018\]](#), where the expected terminal wealth, *i.e.* Π_T , is maximized while the inventory risk is penalized. More precisely, we are interested in the following optimal control problem:

$$\sup_{\substack{l^{1,b}, \dots, l^{d,b} \\ l^{1,a}, \dots, l^{d,a}}} \mathbb{E}[\Pi_T - \sum_{j=1}^d \frac{\kappa_j}{2} \int_0^T \sigma_t^2 (q_t^j \delta_t^j)^2 dt - \frac{\kappa}{2} \int_0^T \sigma_t^2 (\sum_{j=1}^d q_t^j \delta_t^j)^2 dt], \quad (4)$$

with $\sigma_t = S_t \sqrt{V_t}$, $(\kappa_j)_{j=1, \dots, d}$, and κ nonnegative constants. The two items with integration represent the quadratic variation of the inventory risk of individual assets and the portfolio respectively. The market maker can thus consider only the whole portfolio's net inventory risk ($\kappa_j = 0, j = 1, \dots, d$) or express also some specific views on the inventory risk of certain particular assets ($\kappa_j \neq 0, j \in \{1, \dots, d\}$). Of course, other objective functions can be thought of, such as the widely used constant absolute risk aversion utility function as in [Avellaneda & Stoikov \[2008\]](#). We refer to [Guéant \[2017\]](#) for a discussion on the equivalence between these two classes of objective functions. Now we can define the value function

$$u : (t, S, Z, q) \in [0, T] \times \mathbb{R}_+ \times \mathbb{R}^n \times \mathcal{Q} \mapsto u(t, S, Z, q)$$

associated with (4) as

$$\begin{aligned} u(t, S, Z, q) &= \sup_{\substack{(l_s^{1,b}, \dots, l_s^{d,b}) \\ (l_s^{1,a}, \dots, l_s^{d,a})}_{s \in [t, T]}} \mathbb{E}_{t, S, Z, q} \left[\int_t^T \sum_{j=1}^d \left(\sum_{k=a, b} \frac{D^j}{2} m^j l_s^{j,k} \Lambda^{j,k} \mathbb{1}_{\{q_s^j + \phi^k m^j \in \mathcal{Q}^j\}} \right) ds + q_s^j \delta_s^j dS_s \right. \\ &\quad \left. - \sum_{j=1}^d \frac{\kappa_j}{2} \int_t^T \sigma_s^2 (q_s^j \delta_s^j)^2 ds - \frac{\kappa}{2} \int_t^T \sigma_s^2 (\sum_{j=1}^d q_s^j \delta_s^j)^2 ds \right], \end{aligned}$$

where $Z := (Z^1, \dots, Z^n)$, $q := (q^1, \dots, q^d)$, $\mathcal{Q}^j = \{-Q^j, -Q^j + m^j, \dots, Q^j\}$, $\mathcal{Q} = \prod_{j=1}^d \mathcal{Q}^j$ and

$$\phi^k = \begin{cases} +1 & \text{if } k = b, \\ -1 & \text{if } k = a. \end{cases}$$

Therefore, u has $(2 + n + d)$ variables. We introduce in the following some approximations to bypass this curse of dimensionality.

3 Value function approximation

3.1 The Hamilton-Jacobi-Bellman equations

To reduce the dimensionality coming from the multidimensional nature of the QRH model, we suggest the following approximation:

Approximation III.1. We approximate $(Z_t^i)_{t \in [0, T], i=1, \dots, n}$ and $(\delta_t^j)_{t \in [0, T], j=1, \dots, d}$ by its initial value, i.e.

$$\begin{aligned} Z_t^i &= Z_0^i & i = 1, \dots, n, \\ \delta_t^j &= \delta_0^j =: \delta^j & j = 1, \dots, d, \end{aligned}$$

for $t \in [0, T]$. Then we have $V_t = V_0$ for $t \in [0, T]$ by (2). We replace the dynamic of $(S_t)_{t \in [0, T]}$ by

$$dS_t = \mu dt + \sigma dW_t, \quad t \in [0, T],$$

where $\sigma := S_0 \sqrt{V_0}$ and we allow a constant drift μ .

The above approximation is relevant when the time horizon of the problem is relatively short. [Rosenbaum & Zhang \[2022a\]](#) illustrate the effectiveness of daily hedging under the QRH model through numerical tests on simulated and market data. We give additional daily hedging results in the next section. All these examples indicate that it is reasonable to consider δ constant at the daily scale. In practice, the market maker can always reset the algorithm with updated parameters in the case of significant market movement. We recall that the computation of $(\delta_0^j)_{j=1, \dots, d}$ and V_0 at the beginning of the day is ruled by the QRH model, see [Rosenbaum & Zhang \[2022a\]](#) for more details. Now the value function satisfies

$$\forall (t, S, Z, q) \in [0, T] \times \mathbb{R}_+ \times \mathbb{R}^n \times \mathcal{Q}, u(t, S, Z, q) = v(t, q),$$

where

$$\begin{aligned} v(t, q) = & \sup_{\left(\begin{smallmatrix} I_s^{1,b}, \dots, I_s^{d,b} \\ I_s^{1,a}, \dots, I_s^{d,a} \end{smallmatrix} \right)_{s \in [t, T]}} \mathbb{E}_{t, q} \left[\int_t^T \sum_{j=1}^d \left(\sum_{k=a, b} \frac{D^j}{2} m^j I_s^{j, k} \Lambda^{j, k} \mathbb{1}_{\{q_s^j + \phi^k m^j \in \mathcal{Q}^j\}} + q_s^j \delta^j \mu \right) ds \right. \\ & \left. - \sum_{j=1}^d \frac{\kappa_j}{2} \sigma^2 \int_t^T (q_s^j \delta_s^j)^2 ds - \frac{\kappa}{2} \sigma^2 \int_t^T \left(\sum_{j=1}^d q_s^j \delta^j \right)^2 ds \right]. \end{aligned} \quad (5)$$

III. Multi-asset market making under quadratic rough Heston

Now the value function $u(\cdot)$ is transformed to one with $(1+d)$ variables. Following [Øksendal & Sulem \[2005\]](#), the HJB equation associated with (5) is given by a system of ODEs verifying

$$0 = -\frac{\partial v}{\partial t} - \mu \sum_{j=1}^d q^j \delta^j + \sum_{j=1}^d \frac{\kappa_j}{2} \sigma^2 (q^j \delta^j)^2 + \frac{\kappa}{2} \sigma^2 \left(\sum_{j=1}^d q^j \delta^j \right)^2 - \sum_{j=1}^d \sum_{k=a,b} \mathbb{1}_{\{q^j + \phi^k z^j \in \mathcal{Q}^j\}} \sup_{l^{j,k} \in \{0,1\}} l^{j,k} \Lambda^{j,k} \left(m^j \frac{D^j}{2} + v(t, q + \phi^k m^j e^j) - v(t, q) \right), \quad (6)$$

with terminal condition

$$v(T, q) = 0, \quad (7)$$

where $\{e^j\}_{j=1}^d$ is the canonical basis of \mathbb{R}^d . We can rewrite the terms involving the control variables as

$$\begin{aligned} & \mathbb{1}_{\{q^j + \phi^k m^j \in \mathcal{Q}^j\}} \sup_{l^{j,k} \in \{0,1\}} l^{j,k} \Lambda^{j,k} \left(m^j \frac{D^j}{2} + v(t, q + \phi^k m^j e^j) - v(t, q) \right) \\ & =: \mathbb{1}_{\{q^j + \phi^k z^j \in \mathcal{Q}^j\}} m^j H^{j,k} \left(\frac{v(t, q) - v(t, q + \phi^k m^j e^j)}{m^j} \right), \end{aligned}$$

with

$$H^{j,k}(p) = \Lambda^{j,k} \mathbb{1}_{\{p \leq \frac{D^j}{2}\}} \left(\frac{D^j}{2} - p \right). \quad (8)$$

In the following, we give the results on the existence and uniqueness of a solution of (6) with terminal condition (7). We give the sketch of the proof as it follows the same steps as in [Guéant \[2017\]](#).

Theorem 1. *There exists a unique solution $v : [0, T] \times \mathcal{Q} \rightarrow \mathbb{R}$, C^1 in time, solution of Equation (6) with terminal condition (7).*

Proof. It is clear that $H^{j,k}$ are Lipschitz continuous functions for all $j \in \{1, \dots, d\}$ and $k \in \{a, b\}$. Equation (6) with terminal condition (7) can be viewed as a backward Cauchy problem. According to the Cauchy-Lipschitz theorem, there exists some $\tau \in [0, T]$ such that the Equation (6) with the terminal condition (7) has a unique solution, C^1 in time, on the interval (τ, T) .

$\forall q \in \mathcal{Q}$, it is clear that the function $t \in (\tau, T) \mapsto v(t, q) + (-\mu \sum_{j=1}^d q^j \delta^j + \sum_{j=1}^d \frac{\kappa_j}{2} \sigma^2 (q^j \delta^j)^2 + \frac{\kappa}{2} \sigma^2 (\sum_{j=1}^d q^j \delta^j)^2)(T-t)$ is a decreasing function. Therefore, the only reason why there would not be a global solution on $[0, T]$ is because $\sup_{q \in \mathcal{Q}} v(t, q)$ blows up at $\tau > 0$. By the fact that $H^{j,k}$ is decreasing, a classical comparison principle can be obtained. It is clear that $\bar{v}(t, q) = \sum_{j=1}^d (H^{j,b}(0) + H^{j,a}(0) + |\mu \delta^j| Q^j)(T-t)$ defines a supersolution of (6) with terminal condition (7). Then by the comparison principle, we have

$$\sup_{q \in \mathcal{Q}} v(t, q) \leq \sum_{j=1}^d (H^{j,b}(0) + H^{j,a}(0) + |\mu \delta^j| Q^j)(T-t).$$

This shows that v is bounded for $t \in [0, T]$, which leads to its existence and uniqueness on $[0, T] \times \mathcal{Q}$. \square

By a classical verification argument, we can get the following result:

Theorem 2. *Considering the solution v of Equation (6) with terminal condition (7), the optimal market making decisions in the problem (5) are given by*

$$l_t^{j,k*}(q) = \mathbb{1}_{\{q_{t-}^j + \phi^k m^j \in \mathcal{Q}^j, \frac{v(t, q_{t-}) - v(t, q_{t-} + \phi^k m^j e^j)}{m^j} \leq \frac{D^j}{2}\}}, \quad \forall j \in \{1, \dots, d\}, k \in \{a, b\}.$$

Remark 1. *When the market maker controls only the portfolio's net risk, i.e. $\forall j \in \{1, \dots, d\}, \kappa_j = 0$ and $Q^j \rightarrow +\infty$, the d -dimensional variable q can be summarized with one variable. More precisely, by introducing the maximum risk bound $R > 0$ and $r_t := \sum_{j=1}^d q_t^j \delta^j$, now the value function $\theta : (t, r) \in [0, T] \times [-R, R] \mapsto \theta(t, R)$ associated with the problem can be written as*

$$\theta(t, r) = \sup_{\left(l_s^{1,a}, \dots, l_s^{d,a} \right)_{s \in [t, T]}} \mathbb{E}_{t,r} \left[\int_t^T \sum_{j=1}^d \sum_{k=a,b} \frac{D^j}{2} m^j l_s^{j,k} \Lambda^{j,k} \mathbb{1}_{\{r_s + \phi^k m^j \delta^j \in [-R, R]\}} + \mu r_s ds - \frac{\kappa}{2} \sigma^2 \int_t^T (r_s)^2 ds \right]. \quad (9)$$

And the associated HJB equation is given by

$$\begin{aligned} 0 = & -\frac{\partial \theta}{\partial t} - \mu r + \frac{\kappa}{2} \sigma^2 r^2 \\ & - \sum_{j=1}^d \sum_{k=a,b} \mathbb{1}_{\{r + \phi^k m^j \delta^j \in \mathcal{R}\}} m^j H^{j,k} \left(\frac{\theta(t, r) - \theta(t, r + \phi^k m^j \delta^j)}{m^j} \right), \end{aligned}$$

with terminal condition $\theta(T, r) = 0$.

3.2 Quadratic approximation

In our case with a portfolio made of numerous derivatives of SPX, Equation (6) becomes intricate to solve with classical numerical methods. Here we follow the idea of [Bergault et al. \[2021\]](#) that is approximating the Hamiltonian functions $H^{j,k}(\cdot)$ with quadratic ones. Then closed-form solution can be obtained for the new HJB equation. Moreover, the asymptotic of the solution can be derived, which are helpful for practical use. Note that most of the development in this part follows the same methodologies as in [Bergault et al. \[2021\]](#). We give some key results in the following to avoid any ambiguity.

The market making problem considered in [Bergault et al. \[2021\]](#) is more adapted to the over-the-counter market, where the optimal control is the quoting price and is taken to be continuous. The resulting Hamiltonian functions are C^2 and a natural choice is to use Taylor expansions as the approximations. Clearly, this does not apply for (8). In this work, we propose the following choice as the approximated Hamiltonian functions:

$$\hat{H}^{j,k}(p) = \Lambda^{j,k} \left(\frac{\alpha^j}{2} p^2 - p + \frac{D^j}{2} \right), \quad (10)$$

with $\alpha^j > 0$. Actually we have $\hat{H}^{j,k}(0) = H^{j,k}(0)$ and $(\hat{H}^{j,k})'(0) = (H^{j,k})'(0)$ for any $(j, k) \in \{1, \dots, d\} \times \{a, b\}$. We can choose $\alpha^j \sim \frac{1}{D^j}$ to make $\hat{H}^{j,\cdot}$ close to $H^{j,\cdot}$ around 0. In the case without maximal inventory constraints for individual assets, i.e. $Q^j = +\infty, j \in \{1, \dots, d\}$, following the same computations as in [Bergault et al. \[2021\]](#), the approximation of v associated with $(\hat{H}^{j,k})_{j \in \{1, \dots, d\}, k \in \{a, b\}}$, denoted by \hat{v} , can be written as a quadratic function of q , i.e.

$$\hat{v}(t, q) = -q^T A(t) q - q^T B(t) - C(t),$$

III. Multi-asset market making under quadratic rough Heston

where $A: [0, T] \mapsto S_d^+$, $B: [0, T] \mapsto \mathbb{R}^d$ and $C: [0, T] \mapsto \mathbb{R}$ are given by the unique solution of a system of ODEs¹². Since we are more interested in the behavior of \hat{v} as $T \rightarrow +\infty$, we focus mostly on the asymptotic formulas of A and B ¹³. Let $\delta := (\delta^1, \dots, \delta^d)^T$ and $\Sigma := \text{diag}(\frac{\kappa_1(\delta^1)^2}{\kappa}, \dots, \frac{\kappa_d(\delta^d)^2}{\kappa}) + \delta\delta^T$. We have

$$\begin{aligned} A(0) &\xrightarrow{T \rightarrow +\infty} \frac{\sigma}{2} \sqrt{\kappa} \Gamma, \\ B(0) &\xrightarrow{T \rightarrow +\infty} -D_+^{-\frac{1}{2}} \hat{A}^+ D_+^{\frac{1}{2}} \delta \mu - D_+^{-\frac{1}{2}} \hat{A} \hat{A}^+ D_+^{-\frac{1}{2}} (V_- + \frac{\sigma}{2} \sqrt{\kappa} D_- \mathcal{D}(\Gamma)), \end{aligned}$$

where

$$\begin{aligned} D_+ &= \text{diag}((\Lambda^{1,b} + \Lambda^{1,a})\alpha^1 z^1, \dots, (\Lambda^{d,b} + \Lambda^{d,a})\alpha^d z^d), \\ D_- &= \text{diag}((\Lambda^{1,b} - \Lambda^{1,a})\alpha^1 z^1, \dots, (\Lambda^{d,b} - \Lambda^{d,a})\alpha^d z^d), \\ V_- &= ((-\Lambda^{1,b} + \Lambda^{1,a})z^1, \dots, (-\Lambda^{d,b} + \Lambda^{d,a})z^d), \\ \Gamma &= D_+^{-\frac{1}{2}} (D_+^{\frac{1}{2}} \Sigma D_+^{\frac{1}{2}})^{\frac{1}{2}} D_+^{-\frac{1}{2}}, \quad \hat{A} = \sqrt{\kappa} V (D_+^{\frac{1}{2}} \Sigma D_+^{\frac{1}{2}})^{\frac{1}{2}}, \end{aligned}$$

\hat{A}^+ is the Moore-Penrose inverse of \hat{A} and \mathcal{D} is the linear operator mapping a matrix onto the vector of its diagonal coefficients. As we focus on the asymptotic of the approximated solution in the following, we use A, B and $\hat{v}(q)$ to represent the respective asymptotic formulas for ease of notation. As suggested in Bergault et al. [2021], for any $(j, k) \in \{1, \dots, d\} \times \{a, b\}$, one important application of \hat{v} is to make greedy decisions following

$$\begin{aligned} \tilde{l}^{j,k*}(q) &= \mathbb{1}_{\{q^j + \phi^k m^j \in \mathcal{Q}^j, \frac{\hat{v}(q) - \hat{v}(q + \phi^k m^j e^j)}{m^j} \leq \frac{D^j}{2}\}} \\ &= \mathbb{1}_{\{q^j + \phi^k m^j \in \mathcal{Q}^j, 2\phi^k (e^j)^T A q + m^j (e^j)^T A e^j + \phi^k (e^j)^T B \leq \frac{D^j}{2}\}}. \end{aligned}$$

It is clear that the set $\{q | \tilde{l}^{j,k*} = 1, q \in \mathcal{Q}\}$ and $\{q | \tilde{l}^{j,k*} = 0, q \in \mathcal{Q}\}$ are separated by an affine hyperplane defined on \mathcal{Q} . In the particular case with only portfolio-level net risk controlled, as suggested in Remark 1, and symmetric order execution intensities on the bid and ask sides, i.e. $\Lambda^{j,a} = \Lambda^{j,b} =: \Lambda^j$ for all $j \in \{1, \dots, d\}$, applying the above quadratic approximation on Equation (1), the optimal decisions can be simplified as

$$\tilde{l}^{j,k*}(r) = \mathbb{1}_{\{r + \phi^k m^j \delta^j \in [-R, R], 2Dr\delta^j + (\delta^j)^2 m^j \leq \frac{D^j}{2}\}},$$

where

$$D = \frac{\sigma}{2} \sqrt{\frac{\kappa}{2 \sum_j^d (\delta^j)^2 \Lambda^j m^j \alpha^j}}.$$

Remark 2. In the market making problem introduced in Section 2.2, the agent is restricted to sending limit orders only on the first bid and ask limits. The problem can be easily extended to cover other limits. The resulting Hamiltonian functions can also be approximated and then a closed-form asymptotic solution can be obtained in a similar manner as the above.

¹²We denote by S_d^+ the set of positive semi-definite symmetric matrices of dimension $d \times d$.

¹³We ignore C since it is irrelevant for the dependence of the value function on the states.

4 Numerical results

4.1 Daily hedging with SPX

In this part, we perform discrete hedging to evaluate the relevance of the QRH model and the idea of using constant δ_t during a short period, as suggested in Approximation III.1. For asset j , the cumulative profit of discrete hedging with a time step Δt is given by

$$\mathcal{J}_t^{j,\delta} = \sum_{k=0}^{\lfloor t/\Delta t \rfloor - 1} \hat{\delta}_{t_k}^j \Delta S_{t_{k+1}} + \hat{\delta}_{t_{\lfloor t/\Delta t \rfloor}}^j (S_t - S_{\lfloor t/\Delta t \rfloor \Delta t}),$$

where $t_k = k\Delta t$, $\Delta S_{t_k} := S_{t_k} - S_{t_{k-1}}$, and $\hat{\delta}_{t_k}^j$ is the hedging quantity at t_k computed by neural networks. We compare $\mathcal{J}_t^{j,\delta}$ with the price evolution of asset j observed in the market, defined by

$$\mathcal{J}_t^{j,P} = P_t^j - P_0^j.$$

Figure 1 and 2 show the results of two experiments of daily hedging on SPX call, VIX call, and VIX future with the following characteristics:

1. SPX call: maturity 2017-07-21, strike 2460; VIX call: maturity 2017-07-19, strike 11.5; VIX future: maturity 2017-07-19.
2. SPX call: maturity 2017-10-20, strike 2550; VIX call: maturity 2017-10-18, strike 13; VIX future: maturity 2017-10-18.

On day zero of each experiment, we calibrate the QRH model with market data and compute $(\hat{\delta}_{t_0}^j)_{j=1,\dots,d}$. For each following day, $(Z^i)_{i=1,\dots,n}$ are updated with respect to the evolution of SPX, and then $(\hat{\delta}^j)_{j=1,\dots,d}$ are recomputed. Remarkably in both tests, the hedging portfolios can follow well the market price's evolution with daily rebalancing for all assets. It implies that in practice high-frequency hedging is not mandatory to obtain satisfactory results, and verifies the rationality of Approximation III.1 over the time horizon of the considered market making problem, which is usually shorter than one day.




Figure 1: Daily hedging example 1.




Figure 2: Daily hedging example 2.

4.2 Market making approximation

4.2.1 Example 1

In this part, we assess the proposed quadratic approximation. We begin with a two-asset example with SPX and VIX future. The parameters of the QRH model are taken as follows¹⁴:

$$\lambda = 1.7, \quad \eta = 1.5, \quad a = 0.265, \quad b = 0.246, \quad c = 0.0001,$$

$$Z_0 = (-0.009, 0.015, 0.011, 0.036, 0.002, -0.011, -0.018, 0.074, 0.142, -0.171).$$

Then we have $V_0 = 0.18 \text{ year}^{-1}$. We set $S_0 = 3000\$$, and $\mu_0 = 0 \text{ year}^{-1}$. The maturity of VIX future is set to be one month. Along with $\delta^1 = 1$, using the neural network approximating the pricing function of VIX future, we have $\delta^2 = -0.028$.

As for the market making related parameters, we consider the following values:

- Tick size: $D^1 = 0.25\$$, $D^2 = 0.05\$$.
- Order size: $m^1 = 1$, $m^2 = 20$.
- Maximum inventory: $Q^1 = 15$, $Q^2 = 300$.
- Execution intensity: $\Lambda^{1,b} = \Lambda^{1,a} = 1 \text{ seconds}^{-1}$, $\Lambda^{2,b} = \Lambda^{2,a} = 0.1 \text{ seconds}$.
- Risk penalization: $\kappa^1 = 0.005\$^{-1}$, $\kappa^2 = 0.005\$^{-1}$, $\kappa = 0.01\$^{-1}$.
- Time horizon: $T = 300 \text{ seconds}$.

Here some characteristics of SPX refer to the front month E-mini S&P 500 future. In real life, the contract multiplier of E-mini S&P 500 future is 50, and that of VIX future is 1000. Here we normalize them against 50 to have order sizes. The execution intensities are selected such that on average the volume of SPX is 10 times that of VIX future.

First, we use a classical Euler scheme to approximate numerically the true value function $v(t, q)$. Figure 3 shows the results for several $t \in [0, T]$. Interestingly, the shapes of $v(t, q)$ at different t are very similar, implying that time does not play a crucial role for the optimal controls according to Theorem 2. This is confirmed by Figure 4, showing that for SPX or VIX future, the set $\{q | l^{i,b*} = 0\}$ does not change a lot with t when t is relatively small.

¹⁴The same multi-factor approximation method as that in Rosenbaum & Zhang [2022a] is used, leading to the same $(c_i)_{i=1,\dots,10}$ and $(\gamma_i)_{i=1,\dots,10}$




Figure 3: Snapshots of the value function $v(t, q)$ for $t=0, 50, 100, 150, 200, 250$.

Now we check the relevance of our quadratic approximation. As for the parameter α^j defined in Equation (10), we choose $\alpha^j = \frac{1}{4D^j}$, $j \in \{1, \dots, d\}$. The asymptotic approximated value function $\hat{v}(q)$ is shown in Figure 5. Its shape is very close to that of $v(0, q)$, and the resulting greedy decision $\tilde{l}^{j,k^*}(q)$, given in Figure 6, is very similar to $l_0^{j,k^*}(q)$ shown in Figure 4.

To see the impact of the quadratic approximation in cases with other values of $(\kappa^i)_{i=1,\dots,d}$ and κ , and to quantify the deviation of the market making decisions following the approximated solution from the ones based on the true solution, we conduct backtests on simulated data with different risk aversion preferences¹⁵. More precisely, we vary the inventory penalization parameter κ from 10^{-4} to 1 and fix $\kappa^1 = \kappa^2 = \frac{\kappa}{2}$ for simplicity. We test two strategies β^l and $\beta^{\bar{l}}$, based respectively on l^* and \bar{l}^* . For each pair (κ, β^l) and $(\kappa, \beta^{\bar{l}})$, we evaluate the mean and standard deviation of the profit of market making, defined as $\Pi_{T_b} - \Pi_0$ where T_b stands for the backtesting horizon, on N simulated paths. We use $T_b = \frac{T}{2} = 150$ seconds as the backtesting horizon since we are more interested in the results given by the asymptotic of v . With $N = 5000$, Figure 7 shows the resulting mean-risk profiles for β^l and $\beta^{\bar{l}}$. Even though for a given κ , slight deviations from β^l are observed for $\beta^{\bar{l}}$, the global mean-risk profile of its performance is not significantly degraded from that of β^l . This suggests the effectiveness of the proposed approximation method.

¹⁵Our simulations use neural networks to compute efficiently the price of SPX and VIX options, and VIX future as introduced in Rosenbaum & Zhang [2022a].

III. Multi-asset market making under quadratic rough Heston




Figure 4: Optimal decisions on whether sending bid order for SPX and VIX futures at different t and q . The black areas stand for the set $\{q | l_t^{j,b*} = 0\}$.




Figure 5: The asymptotic value function $\hat{v}(0, q)$ (right) as $T \rightarrow +\infty$, compared to $v(0, q)$ (left). Note that for ease of comparison $C(0)$ is chosen to have $\max_q \hat{v}(0, q) = \max_q v(0, q)$.




Figure 6: The greedy decisions $\tilde{l}^{j,b*}$ obtained with the asymptotic value function $\hat{v}(0, q)$. The black areas stand for the set $\{q | \hat{l}^{j,b*} = 0\}$.




Figure 7: Mean performance and risk of the market making strategies β^l and β^i with varying risk penalization parameter κ on two assets.

4.2.2 Example 2

Here we repeat the same experiment as the above for the case where only the portfolio's net inventory risk is controlled, as discussed in Remark (1). We consider the following assets:

- SPX option with maturity 10 days and strike 2950\$.
- SPX option with maturity 25 days and strike 3050\$.
- VIX option with maturity 10 days and strike 22\$.
- VIX option with maturity 25 days and strike 18\$.

And we have the following parameters:

- Tick size: $D^1 = 0.1$, $D^2 = 0.1$, $D^3 = 0.05$, $D^4 = 0.05$.
- Inventory risk: $\delta^1 = 0.533$, $\delta^2 = 0.134$, $\delta^3 = -0.014$, $\delta^4 = -0.013$.
- Order size: $m^j = 2$ for $j \in \{1, 2, 3, 4\}$.
- Inventory risk bound: $R = 10$.
- Execution intensity: $\Lambda^{j,b} = \Lambda^{j,a} = 0.05 \text{ second}^{-1}$ for $j \in \{1, 2, 3, 4\}$.
- Risk penalization: $\kappa = 1$.
- Time horizon: $T = 600$ seconds.

With four assets under consideration, the value function θ defined in Equation (9) can be solved with a classical Euler scheme based on grids. As shown in the left subfigure of Figure 8, the shape of the slices $\theta(\cdot, r)$ is very stable through time. We then apply the above quadratic approximation method and obtain its asymptotic form $\hat{\theta}(r)$ as $T \rightarrow +\infty$, which is shown to be very close to $\theta(0, r)$ in Figure 8.




Figure 8: The value function $\theta(t, r)$ (left), and the asymptotic solution $\hat{\theta}(0, r)$ as $T \rightarrow T$ (right). Note that the constant $C(0)$ in the expression of $\hat{\theta}(0, r)$ is chosen to have $\max_r \hat{\theta}(0, r) = \max_r \theta(0, r)$.

Similarly to the above example, we use backtests to perceive more directly the effectiveness of the approximated solution. We choose $T_b = \frac{T}{2} = 300$ seconds and for each pair (κ, β^l) and $(\kappa, \beta^{\bar{l}})$, the mean and standard deviation of the performance on $N = 5000$ paths are computed. Figure 9 shows the results with κ varied from 10^{-4} to 1. We remark again that the mean-risk profile of β^l is very close to that of $\beta^{\bar{l}}$.




Figure 9: Mean performance and risk of the market making strategies β^l and $\beta^{\bar{l}}$ with varying risk penalization parameter κ on four assets.

4.3 Example 3

In this part we consider a market making problem on the six assets mentioned above to exemplify the effect of risk mutualization among them. We test with $T_b = 2000$ seconds, κ varying from 10^{-4} to 10^{-1} and $\kappa^j = \frac{\kappa}{2}, j = 1, \dots, 6$. As indicated in the introduction, having a closed-form asymptotic solution allows us to recalibrate instantly the solution with the most recently observed parameters. To illustrate the idea, during the test, we update $(\delta^j)_{j=1, \dots, d}$ and δ every 100 seconds according to the running state of the QRH model, *i.e.* $X := (S, Z^1, \dots, Z^{10})$. As for the benchmarking case, we set $\kappa = 0$ so that the inventory risk of each asset is managed individually. A classical Euler scheme is used for solving the value function of each asset, and then the corresponding optimal decisions can be obtained. Figure 10 gives the results on 5000 simulated paths. Clearly, the performance of uni-asset market making is suboptimal compared to the other setting, which also considers the net risk at the portfolio level.




Figure 10: Mean performance and risk of the uni-asset and multi-asset market making strategies with varying risk penalization parameter κ . For the uni-asset case, optimal market making is conducted on each individual asset and the final profits are summed.

On the universality of the volatility formation process: when machine learning and rough volatility agree

Abstract

We train an LSTM network based on a pooled dataset made of hundreds of liquid stocks aiming to forecast the next daily realized volatility for all stocks. Showing the consistent outperformance of this universal LSTM relative to other asset-specific parametric models, we uncover nonparametric evidences of a universal volatility formation mechanism across assets relating past market realizations, including daily returns and volatilities, to current volatilities. A parsimonious parametric forecasting device combining the rough fractional stochastic volatility and quadratic rough Heston models with fixed parameters results in the same level of performance as the universal LSTM, which confirms the universality of the volatility formation process from a parametric perspective.

Keywords— Volatility formation, universality, forecast, LSTM, HAR, rough volatility, quadratic rough Heston, Zumbach

1 Introduction

The rough volatility paradigm, where the volatility dynamic is modelled with a fractional Brownian motion (fBm) with Hurst parameter of order 0.1, enables us to generate the stylized facts of realized volatilities [Bennedsen et al., 2021; Gatheral et al., 2018]. It is also very effective for derivatives pricing and hedging, see for example Bayer et al. [2016]; El Euch et al. [2019]; El Euch & Rosenbaum [2019] and El Euch & Rosenbaum [2018]. Under the rough fractional stochastic volatility (RFSV) model, one can predict the next realized volatilities following a simple formula that requires essentially only one parameter, namely the Hurst parameter [Gatheral et al., 2018]¹⁶. As it is observed that the Hurst parameter is around 0.1 for a wide range of financial assets [Bennedsen et al., 2021; Bolko et al., 2022; Gatheral et al., 2018; Wu et al., 2022], this actually results in a quasi-universal forecasting formula, suggesting some universality of the volatility formation process across assets. The recently introduced

¹⁶More precisely, only the Hurst parameter is required for forecasting log-volatility. When volatility forecasting is concerned, an additional parameter is needed, whose impact is marginal as we will show in Section 5.

IV. On the universality of the volatility formation process: when machine learning and rough volatility agree

quadratic rough Heston (QRH) model can reproduce at the same time the market observations of SPX and VIX options [Gatheral et al., 2020; Rosenbaum & Zhang, 2022a], which is known to be challenging with continuous stochastic volatility models. A key element to achieving this is the encoding of the feedback effect of price trends on volatility, or *strong* Zumbach effect as initially suggested in Zumbach [2010]. This motivates us to presume that the potentially universal volatility formation mechanism involves also past returns.

In this work, we aim at confirming this universality of volatility formation mechanism relating past volatilities and returns to current volatilities across hundreds of liquid stocks, *i.e.* the values of the involved parameters do not show significant differences among stocks. We are not suggesting that the volatility processes of different assets follow exactly the same genesis. The assets with different exposures to exogenous events such as earnings announcements will certainly show different volatility distributions. Here we focus on the mechanism of self-reflexivity of volatility processes. As the degree of endogeneity of financial markets is high [Filimonov & Sornette, 2012, 2015; Hardiman et al., 2013], we expect that a large part of current volatility levels can be explained by past volatilities and returns. One can imagine that this endogenous link between current volatilities and past market realizations emerges from some primary activities at the microstructural scale as suggested in Blanc et al. [2017]; Dandapani et al. [2021]; El Euch et al. [2018] and Jusselin & Rosenbaum [2020], leading to a universal volatility formation process across assets.

The universality of the volatility dynamics of stocks is considered in Blanc et al. [2014] and Chicheportiche & Bouchaud [2014] in the context of general quadratic autoregressive models, which are calibrated with return data, relying on a combination of the generalized method of moments and the maximum-likelihood estimation. In this work the access to intraday prices allows us to compute daily realized volatilities. Then the impact of past volatilities and returns on current volatilities can be inspected by fitting a volatility forecasting device dependent on histories of market realizations, by minimizing the error of forecast. The heterogeneous autoregressive (HAR) model introduced in Corsi [2009] is probably the most referenced one in this category. It offers a highly parsimonious way to reproduce the slow decay behavior of volatility autocorrelation. It is inspired by the heterogeneity and seasonality of market participants' activities, with a volatility prediction typically depending on past realized volatilities over one day, one week, and one month. Several extensions have been developed to reproduce additional features such as the leverage effect, see for example Corsi & Renò [2012] and Patton & Sheppard [2015]. Both the HAR and RFSV models presume some specific formulations for the effect of past market realizations on current volatilities. However, we prefer to be agnostic about the dynamic of the volatility formation process in the first place. Providing this dynamic is universal across assets, we can fit a forecasting device on the dataset mixing different assets to neutralize any asset-dependent effect and reduce the risk of overfitting, which should help the device uncover the "true" universal volatility formation process.

Because of our will to be as model-free as possible, in the first part of our work, we choose to apply a nearly nonparametric approach with machine learning. More particularly, we assess the volatility formation mechanism by trying to forecast the next daily realized volatility with a long short-term memory (LSTM) network, see Hochreiter & Schmidhuber [1997], based on past realized volatilities and returns. Technically speaking, LSTM is a parametric method. Nonetheless, being different from classical parametric methods where the mapping function between the inputs and the outputs is characterized by a handful of parameters, deep neural network involves much more parameters and can be seen as a universal nonparametric approximator Hornik et al. [1990]. It is a natural choice to use LSTM instead of multilayer perceptron or convolutional neural network for time series learning tasks given its recurrent structure. Our network is trained with a pooled dataset made of over 800 liquid stocks. We will see that this universal LSTM can consistently outperform the asset-specific parametric models. We

conduct further tests to confirm the targeted universality and its stationarity through time.

Our universal LSTM network can be viewed as a benchmark for volatility forecasting, capturing the universal volatility formation mechanism with a quasi-nonparametric approach. The question now is whether we can uncover this complex “black box” with a parsimonious parametric forecasting device, while similar performance can be held. Moreover, given the universality in question, we expect that one fixed set of parameter values should apply to all stocks. The RFSV and QRH models give parametric views about the impacts of past volatilities and returns on current volatility respectively, as introduced before. For each of them, we conduct comparison tests to confirm that a universal model with fixed parameters performs very similarly to those calibrated stock by stock. Then we obtain a simple combination of the two aiming at capturing the effect of both past volatilities and returns. Surprisingly, we show that we can apply this parsimonious forecasting device for all stocks using the same parameters, with the same level of performance that the LSTM obtained. This recalls the universality of the volatility formation process found by the universal LSTM network, while this time from a parametric perspective.

Machine learning is already widely used in works of volatility forecasting, from the ones applying classical statistical learning algorithms, such as LASSO [Audrino et al., 2020; Guyon & Lekeufack, 2022], to the others mostly based on deep-learning models, see for example Assaf et al. [2022]; Chen & Robert [2022]; Christensen et al. [2022]; Kim & Won [2018]; R. Liu et al. [2022]; Y. Liu [2019]; Petrozziello et al. [2022]; Rahimikia & Poon [2020]. Among the latter strand of research, LSTM has become a popular choice. Past returns and volatilities, measured through daily return series, are taken as predictors in Y. Liu [2019] and the superiority of LSTM to GARCH-type models is verified on S&P500 and AAPL series. Kim & Won [2018] present a hybrid design feeding the parameters of multiple GARCH-type models into LSTM and test it on data of the KOSPI 200 index. A rich set of explanatory variables collected from news and limit order book data is used to forecast realized volatility of 23 NASDAQ stocks in Rahimikia & Poon [2020]. Uni-asset volatility forecasting based on multi-assets panels, *i.e.* features of the same timestamp coming from multiple assets, are adopted in Assaf et al. [2022]; R. Liu et al. [2022]; Petrozziello et al. [2022], trying to capture some cross-sectional dependences between different assets. Note that in these works one need to train an LSTM network for each concerned asset. In our current work, through numerical tests on hundreds of stocks actively traded on the US and European markets, we present a highly performant universal LSTM-based volatility forecasting device applied to all stocks. We highlight the implied universal nature of the volatility formation process, which to our best knowledge is not mentioned in the previous works. Moreover, we then disclose this universality with a parsimonious parametric formulation. The forecasting in our work only involves return and volatility data since we focus on the endogenous volatility formation mechanism. In cases with additional predictors such as those extracted from news data, more complex forecasting methods can be investigated by taking the uncovered universality into account.

This chapter is organized as follows. In Section 2, we recall the forecasting formulas of HAR and RFSV models. We also consider the simple autoregressive (AR) model for reasons of comparison. The architecture of our LSTM network is then introduced and some key training hyperparameters are listed. We describe in Section 3 the dataset at hand and the evaluation metrics used for model assessment. The out-of-sample performances of all the models are then compared in Section 4. We conduct further tests to confirm the universality of the volatility formation process across assets. In Section 5, we evaluate the universal versions of the forecasting devices given by the RFSV and QRH models respectively, and combine them trying to approximate the performance of the above LSTM network. Finally, we conclude with our main findings in Section 6.

2 Description of the forecasting devices

2.1 Parametric methods

As the “true” volatility is unobservable, in our tests we use daily realized volatilities, which we will define in the next section, to inspect the volatility formation mechanism. We recall first the forecasting formulas of several parametric realized volatility models. In the AR and HAR model, the predicted volatility $\hat{\sigma}_t$ is given by a linear combination of past volatilities $(\sigma_{t-p})_{p=1,2,\dots}$. More specifically, we have:

- AR(p):

$$\hat{\sigma}_t = \alpha_0 + \sum_{j=1}^p \beta_j \sigma_{t-j},$$

- HAR:

$$\hat{\sigma}_t = \alpha_0 + \beta_1 \sigma_{t-1} + \beta_2 \frac{1}{5} \sum_{j=1}^5 \sigma_{t-j} + \beta_3 \frac{1}{22} \sum_{j=1}^{22} \sigma_{t-j},$$

where $\alpha_0, \beta_i \in \mathbb{R}$ can be estimated via ordinary least squares. We evaluate AR(22) in our tests to make it dependent on the same volatility histories as the HAR model.

In the RFSV model, the logarithm of volatility is modelled by a fBm W_t^H associated with the Hurst parameter H :

$$d \log \sigma_t = \nu dW_t^H, \quad (1)$$

with $\nu > 0$. For a fBm, we have the following property:

$$\mathbb{E}[|W_{t+\Delta}^H - W_t^H|^q] = \mathbb{E}[|Z|^q] \Delta^{qH}, \quad Z \sim N(0, 1), \quad (2)$$

for any $t \in \mathbb{R}$, $\Delta \geq 0$ and $q > 0$. For simplicity, we use the method suggested in [Gatheral et al. \[2018\]](#) to estimate H in this work. Of course other methods can be tested trying to get more accurate estimations, see for example [Bolko et al. \[2022\]](#), [Shi et al. \[2023\]](#), and [Wu et al. \[2022\]](#). Here we compute

$$m_q(l) = \frac{1}{T-l} \sum_{t=1}^{T-l} |\log \sigma_{t+l} - \log \sigma_t|^q, \quad (3)$$

over intervals of $l \in \{1, 2, \dots\}$, where T is the number of data points available. By (2), we have

$$m_q(l) \simeq \mathbb{E}[|Z|^q] \nu^q l^{qH}.$$

Then we can fit $\log(m_q(l))$ against $\log(l)$ using ordinary least squares (OLS) regression. Writing the results of OLS as $\widehat{\log(m_q(l))} = \beta_0 + \beta_1 \log(l)$, we use $\hat{H} = \frac{\beta_1}{q}$ and $\hat{\nu} = \exp(\frac{\beta_0 - \log(\mathbb{E}[|Z|^q])}{q})$ as estimators. As proposed in [Gatheral et al. \[2018\]](#), we then have the following predictors for $\hat{H} < \frac{1}{2}$:

$$\widehat{\log \sigma}_t = \frac{\cos(\hat{H}\pi)}{\pi} \int_{-\infty}^{t-1} \frac{\log \sigma_s}{(t-s+1)(t-s)^{\hat{H}+1/2}} ds, \quad (4)$$

$$\hat{\sigma}_t = \hat{\nu} \exp(\widehat{\log \sigma}_t),$$

with

$$\hat{\nu} = \exp\left(\frac{\Gamma(3/2 - \hat{H})}{2\Gamma(\hat{H} + 1/2)\Gamma(2 - 2\hat{H})} \hat{\nu}^2\right). \quad (5)$$

In practice, the integral in (4) is truncated to some length of history and then approximated through a Riemann sum. We refer to Wang et al. [2022] for a recent discussion on alternative forecasting methods of the fBm. In our tests we take $q = 2$, the approach actually coincides with the aggregated variance method, which is usually applied in estimating time series long-range dependence, see for example Taqqu et al. [1995]. The estimation for the volatility of log-volatility parameter ν appears only in the “scaling” term $\hat{\nu}$. We will see in Section 5 that with real data $\hat{\nu}$ is very close to 1 and shows a very small deviation across assets. The impact of ν on the forecasting is thus marginal.

2.2 Nonparametric forecasting with LSTM




Figure 1: Structure of an LSTM cell (left) and simplified computational graph of the network based on LSTM (right).

A recurrent neural network (RNN) is a type of neural networks suitable for sequential learning tasks, of which LSTM is a popular member. The left subfigure of Figure 1 illustrates the structure of the layer l of a multilayer LSTM cell. For each timestep t , the following computations are performed:

$$\begin{aligned} i_t^l &= s(W_{ii}^l x_t^l + b_{ii}^l + W_{hi}^l h_{t-1}^l + b_{hi}^l), \\ f_t^l &= s(W_{if}^l x_t^l + b_{if}^l + W_{hf}^l h_{t-1}^l + b_{hf}^l), \\ g_t^l &= \tanh(W_{ig}^l x_t^l + b_{ig}^l + W_{hg}^l h_{t-1}^l + b_{hg}^l), \\ o_t^l &= s(W_{io}^l x_t^l + b_{io}^l + W_{ho}^l h_{t-1}^l + b_{ho}^l), \\ c_t^l &= f_t^l * c_{t-1}^l + i_t^l * g_t^l, \\ h_t^l &= o_t^l * \tanh(c_t^l), \end{aligned}$$

where the superscript indexes the layer, c_t^l, h_t^l are the cell and hidden states, x_t^l equals the t -th element of the input sequence when $l = 1$ or h_{t-1}^{l-1} for $l > 1$, $s(\cdot)$ is the sigmoid function, $*$ is the Hadamard product, W^l and b^l are parameters to learn from data. The main idea of LSTM is to introduce the gates i, f and o to capture long-range dependence in the sequential data and to avoid the vanishing gradient problem of classical RNNs [Goodfellow et al., 2016; Hochreiter & Schmidhuber, 1997]. The right subfigure of Figure 1 gives the computational graph for predicting the next volatility with the most recent L histories. In the following, we test first with $x_t = (\sigma_t^2)$ and then with $x_t = (\sigma_t^2, r_t)$, where r_t is the daily return at time t . Table 1 summarizes the network architecture and the hyperparameters

IV. On the universality of the volatility formation process: when machine learning and rough volatility agree

of training. Since in practice neural networks are trained with stochastic gradient descent, to avoid artifacts generated by a particular random seed, we fix globally 10 random seeds with which we train 10 independent LSTM networks during each experiment. Their average is taken as the final forecast.

| | |
|--------------|--|
| Architecture | number of layers: 2, dimension of hidden state: 2, activation function: SiLU, output layer: linear |
| Training | optimizer: Adam, loss: mean squared error (MSE), batch size: 512, learning rate: 1e-3, number of epochs: 5 |

Table 1: Summary of the network architecture and training parameters.

3 Data and evaluation metrics

Our dataset contains 5-minutes intraday prices of Russell 1000 and STOXX Europe 600 constituents, for years between 2010 and 2020. After removing the ones with considerable missing or abnormal data, we get 862 names from the US market and 503 names from the European market. Figure 2 gives their distribution by associated sectors following the Global Industry Classification Standard (GICS)¹⁷, where we can see a diversity of involved sectors. The daily realized volatility is estimated by

$$\sigma_t = \sqrt{\sum_i r_{t,i}^2},$$

where $r_{t,i}$ is defined by the i -th 5-minutes intraday return after removing the first and last 30 minutes of the daily trading period. The return of day t is defined as

$$r_t = \frac{P_t - P_{t-1}}{P_{t-1}},$$

where P_t is the closing price of day t . To make the data of different stocks have similar scale, we perform the following scaling for each stock:

$$\tilde{\sigma}_t = \frac{\sigma_t}{\sqrt{\langle \sigma_t^2 \rangle}}, \quad \tilde{r}_t = \frac{r_t - \langle r_t \rangle}{\sqrt{\langle (r_t - \langle r_t \rangle)^2 \rangle}},$$

where $\langle \cdot \rangle$ refers to the average over t . For ease of notation, we use σ_t and r_t referring to the preprocessed quantities in the following.

We focus mostly on the US market, and the data of the European market is used for an out-of-sample double-check. We use the pooled dataset of 862 stocks over years 2010 - 2015 to train the LSTM network. The period 2016 - 2020¹⁸ is used for out-of-sample evaluation. For the parametric models introduced above, sliding window fit of size 1000 is applied, which means, the model is recalibrated with the last 1000 data points for every new forecast. We employ MSE as the model evaluation metric, which is defined as follows:

$$\text{MSE}(\sigma, \hat{\sigma}) = \frac{1}{T} \sum_{t=1}^T (\hat{\sigma}_t - \sigma_t)^2,$$

¹⁷<https://www.msci.com/our-solutions/indexes/gics>.

¹⁸As in this work we focus on the endogenous volatility feedback effect, the period 2020-02-01 - 2020-06-01 is excluded from the test set, which is conceived to be largely perturbed by exogenous information.




Figure 2: Count of stocks by GICS Sector.

where T is the number of trading days of the out-of-sample period. MSE is widely used in forecasting-like tasks. According to Patton [2011], it is a robust and homogeneous metric for model comparison of volatility forecasting. In this work, we focus more on each model's relative performance compared to that of the HAR model so that we compute instead $(\text{MSE}_m / \text{MSE}_{\text{HAR}})$, for $m \in \{\text{AR}(22), \text{RFSV}, \text{LSTM}\}$.

4 Capturing universality with LSTM

We estimate the Hurst parameter of each time series of realized volatility in the dataset following the approach described in Section 2.1. The resulting distributions by sector shown in Figure 3 via violin plot confirm the observations in the literature, with H of order 0.1. Moreover, we do not remark any significant differences across sectors and markets. As presented in the Introduction, this implies a quasi-universal volatility forecasting formula, suggesting a universal volatility formation mechanism relating past market realizations to current volatilities.

We use LSTM to uncover this universality without an *a priori* format. We train first an LSTM network, denoted by LSTM_{var}^{us} with $x_t = (\sigma_t^2)$ as input. Then we add past returns data, *i.e.* $x_t = (\sigma_t^2, r_t)$, to train another LSTM network LSTM_{ret}^{us} . We recall that the training is performed with the data of over 800 stocks of the US market. In the following, we compare this nonparametric approach with the other parametric methods fitted asset by asset in terms of out-of-sample MSE. The universal volatility formation mechanism across assets is then investigated through a series of experiments.

4.1 Parametric vs nonparametric

We recall that the parametric methods introduced before are all fitted on each stock trying to get more accurate forecasting, while the LSTM models are trained on the pooled dataset of over 800 stocks. Figure 4 summarizes the out-of-sample performances of the stocks considered. We perform also the

IV. On the universality of the volatility formation process: when machine learning and rough volatility agree




Figure 3: Empirical distribution of the estimated Hurst parameters inside each GICS Sector. The dashed gray line represents the overall mean.

model confidence set (MCS) procedure introduced in [Hansen et al., 2011] to select the best forecasting methods stock by stock with a significance level = 10%. We use the semi-quadratic statistic obtained from 2000 bootstrap iterations with a block length of 12. Among the 862 names, Table 2 reports the total times that each model is selected as one of the best candidates. We can make the following remarks:

- the simple AR(22) model, which imposes no structural constraint on the autoregressive coefficients, underperforms the HAR model. This is understandable as it is more likely to overfit,
- the RFSV outperforms the HAR model in terms of prediction error. It is remarkable as it involves only two free parameters, *i.e.* v and H as shown in (1). It benefits from the special parametrization of the weights of past volatilities through H , as shown in (4),
- the nonparametric models, $LSTM_{var}^{us}$ and $LSTM_{ret}^{us}$ outperform the other parametric models, especially when we incorporate past returns data. This indicates that the potential universal volatility formation mechanism across assets, relating past volatilities and returns to current volatilities, allows us to calibrate a universal model based on all assets, where the risk of overfitting is reduced due to enriched realized scenarios.

| AR(22) | HAR | RFSV | $LSTM_{var}^{us}$ | $LSTM_{ret}^{us}$ |
|--------|-----|------|-------------------|-------------------|
| 91 | 340 | 622 | 664 | 850 |

Table 2: Number of times that a model falls in the MCS with 10% as significance level.

The length of sequences for training $LSTM_{var}^{us}$ and $LSTM_{ret}^{us}$ is set here to be 22. The scheme in Figure 1 can adapt to sequences of variable lengths. We choose 22 to make it comparable with the HAR model. Figure 5 shows the forecasting errors of $LSTM_{ret}^{us}$ with different values of L relative to that with $L = 5$. We can remark that most of valuable information for the forecast comes from the most recent




Figure 4: Boxplot showing each model’s out-of-sample MSE relative to the HAR model for the stocks of the US market. LSTM^{US}_{var} is trained based on past realized variance, and LSTM^{US}_{ret} uses also past returns.

market realizations over a horizon of fewer than 15 days. The improvement with longer input sequences is marginal. We also tested with LSTM networks with more complicated architecture such as higher dimension of hidden state and more hidden layers, though no significant improvement is observed.




Figure 5: Out-of-sample MSE of LSTM^{US}_{ret} with different L relative to the one with $L = 5$.

4.2 Network inspection

We see from the above results that a data-driven universal model can outperform asset-specific models. However, the trained LSTM remains a “black box”. To shed some light on the universal mechanism captured by the network, we compute the average sensitivities of its output with respect to the inputs. Then we confirm empirically the existence of some universal volatility formation mechanism across assets following similar methods as those in [Sirignano & Cont \[2019\]](#). Note that when there is no clear indication, the LSTMs in the following are trained with volatility and return data.

4.2.1 Local sensitivity visualisation

We define the following local sensitivities:

$$\alpha_t(\tau) := \frac{\partial \hat{\sigma}_t}{\partial \sigma_{t-\tau}^2}, \quad \beta_t(\tau) := \frac{\partial \hat{\sigma}_t}{\partial r_{t-\tau}} \quad \text{for } \tau = 1, 2, \dots, 22, \quad (6)$$

where $\hat{\sigma}_t$ refers to the output of LSTM_{ret}^{us} . Note that these quantities can be computed efficiently with automatic differentiation. For each stock under study, we evaluate their average over t , *i.e.* $\alpha(\tau) := \langle \alpha_t(\tau) \rangle$ and $\beta(\tau) := \langle \beta_t(\tau) \rangle$. Figure 6 shows the average of $\alpha(\tau)$ and $\beta(\tau)$ across stocks, with their respective standard deviations. The shape of $\alpha(\cdot)$ reflects in particular the well-known slowly decaying autocorrelation function of volatility time series. The values of $\beta(\cdot)$ are in general one order of magnitude smaller than those of $\alpha(\cdot)$, showing particularly the negative correlation between last returns and current volatilities. Interestingly the deviations across stocks of both $\alpha(\cdot)$ and $\beta(\cdot)$ seem to be very small, implying some universality of the volatility formation mechanism. Of course other partial derivatives of higher orders can also be investigated such as $(\frac{\partial^2 \hat{\sigma}_t}{\partial r_{t-\tau_1} r_{t-\tau_2}})_{\tau_1, \tau_2=1, \dots, 22}$. However, their magnitude is relatively smaller and the noise involved in the computation is not negligible. In the following, we confirm this universality in another way, by showing mainly that the trained LSTM network can be applied to stocks not included in the training set without significant degradation of performance.




Figure 6: Average $\alpha(\tau)$ (left) and $\beta(\tau)$ (right) across stocks. The shadow regions represent the standard deviation across stocks.

4.2.2 Checking universality across assets

- Sector-level universality?

A natural question is whether this potential universality makes sense at the level of the whole market or rather at the level of the sectors. One way to verify it is to train sector-specific LSTMs and compare them with the universal LSTM trained with all the sectors. For a certain sector-specific LSTM, we are interested in its relative performance to the universal counterpart on its covered sector in comparison with the other sectors. If some important sector-dependent volatility formation mechanisms exist, we expect that this LSTM performs consistently better on its “local” sector than the others, where the performance is evaluated relative to that of the benchmark - the universal LSTM.




Figure 7: Out-of-sample MSE of the asset-specific LSTM trained with the stocks of the sector *Industrials*, relative to the universal LSTM covering all sectors.

Figure 7 and 8 give two examples, respectively, the sector *Industrials* and *Financials*. Interestingly, in both cases, the sector-specific LSTMs do not show consistently smaller forecasting errors on their covered sectors than the others. The deterioration of performance on all sectors compared to the universal LSTM is mainly caused by the shrinkage of training data. Figure 9 gives an additional example confirming our analysis, where we train an LSTM with multi sectors. Thanks to the enriched training set, the forecasting errors of all sectors are universally reduced. The LSTM with only a subset of sectors is slightly deteriorated compared to the universal LSTM. These observations conquer the idea of no significant difference on the volatility formation mechanism between sectors and support the universality across assets.

- Asset-dependent mechanisms?

The LSTM trained with all the stocks seems to capture the universal link between past market realizations and current volatilities. However, we may wonder whether there could be any asset-dependent features that the universal network failed to encode. For that, we fine-tune the universal LSTM on each

IV. On the universality of the volatility formation process: when machine learning and rough volatility agree




Figure 8: Out-of-sample MSE of the asset-specific LSTM trained with the stocks of the sector *Financials*, relative to the universal LSTM covering all sectors.




Figure 9: Out-of-sample MSE of an LSTM trained with the stocks of the sectors *Health care*, *Financials*, *Information technology*, *Communication services*, *Utilities*, and *Real estate*, relative to the universal LSTM.

individual stock to try to add some stock-specific components. More specifically, we train an LSTM, whose parameters are initialized with those of $LSTM_{ret}^{us}$, with only the data of one specific stock. In the deep learning community, it is also a common practice to fine-tune only the output layer of a given pre-trained network. Here we test both ideas allowing all the parameters or only those of the output layer to be updated. Figure 10 shows the out-of-sample performances of the models after 5-epochs fine-tuning relative to the original one. We do not remark any significant reduction of forecasting error after the parameters of the network being “adapted” locally to each stock. Longer retraining does not lead to any improvement of global performance. This seems to reject again any important asset-dependent mechanism relating past realizations to current volatilities, at least for those that can be captured nonparametrically by an LSTM network.




Figure 10: Empirical distribution of out-of-sample MSE of the fine-tuned asset-specific LSTMs relative to the universal LSTM. The left corresponds to the case where all the parameters of the network can be updated. The right is where we only fine-tune the parameters of the output layer.

- Universality across markets?

The above tests are all done on US market data. Does this universality hold even across markets? In other words, can the LSTM based on the US market be directly applied to stocks of the European market without deterioration of performance? To see that, with European market data, we fit the considered parametric forecasting devices and train two LSTM networks, $LSTM_{var}^{eu}$ and $LSTM_{ret}^{eu}$, similarly to the test as shown in Figure 4. Besides, we test also the $LSTM_{var}^{us}$ and $LSTM_{ret}^{us}$, which are trained with US market data, on the stocks of the European market. From Figure 11, the universal models depending on past volatilities and returns, *i.e.* $LSTM_{ret}^{eu}$ and $LSTM_{ret}^{us}$, obtain the best out-of-sample performance. Moreover, $LSTM_{var}^{us}$ and $LSTM_{ret}^{us}$ get slightly better results compared to their respective counterparts, $LSTM_{var}^{eu}$ and $LSTM_{ret}^{eu}$. This supports again the existence of a universal volatility formation mechanism across assets, which do not differ from one market to another. $LSTM_{var}^{us}$ and $LSTM_{ret}^{us}$ benefit from the abundance of data of the US market, and thus capture better the endogenous processes relating past volatilities and returns to current volatilities.

- Stability across time?

IV. On the universality of the volatility formation process: when machine learning and rough volatility agree




Figure 11: Boxplot showing models' out-of-sample performances relative to the HAR model for the stocks of the European market. Note that LSTM^{EU} and LSTM^{US} are trained on stocks of European and US markets respectively.

With financial time series, it is a common practice to recalibrate models based on a sliding window of the most recent data as it is perceived that there are changes in market regimes. To see whether the universal volatility formation mechanism is time-varying, we apply the idea of “dynamic” prediction. At the end of each year, we retrain an LSTM based on the data of last six years and use it for the forecasting of the following year. We monitor the out-of-sample performances of these “dynamic” models relative to the initial “static” one trained on data of 2010 - 2015. As shown in Figure 12, the “dynamic” models perform very closely to the “static” one, not showing any superiority in forecasting for the periods just following their training sets. This gives evidence that the universality is also stationary over time. One does not need to retrain the universal model frequently trying to adapt to varying market regimes.




Figure 12: Out-of-sample performance of the “dynamic” models relative to the “static” model.

5 Uncovering the universal volatility formation process

The above results confirm a universal volatility formation mechanism across assets, relating past volatilities and returns to current volatilities, from a nonparametric perspective. We wonder now whether we can disclose this universal nature with a parsimonious parametric formulation, while keeping a similar level of forecasting performance as that of the LSTM network. The RFSV model gives a parametric view of the universal link between past and current volatilities. The results in Figure 4 and 11 shows that the RFSV model can get similar out-of-sample forecasting performance as the LSTMs trained with past volatilities data. The recently introduced QRH model encodes *strong* Zumbach effect describing the feedback effect of price trends on volatility, which is missing in the RFSV model. In the following, we first verify the universality of both models. For each model, we will see that similar forecasting performance can be obtained with properly fixed parameters as in the case where parameters are calibrated on each stock. Then we evaluate a simple combination of these two universal parametric forecasting devices, comparing its performance with the universal LSTM.

- RFSV

As shown in (4), forecasting volatility in the RFSV model involves \hat{H} and \hat{v}^2 , where the impact of \hat{v}^2 is imposed only through the \hat{v} defined in (5). Figure 13 shows the fitting results of H and v with data of 2012 - 2015 of US market, by following the method introduced in Section 2.1. The deviations across stocks do not seem to be large for both \hat{H} and \hat{v} . Other methods can be tested in the future trying to reduce the estimation errors. For illustrating the idea of universality, we estimate H and v based on the entire training set of the US market and then apply them for the forecasting of all the stocks considered, including those of the European market. More precisely, the same procedure of estimation as introduced in Section 2.1 is applied, by using now $\bar{m}_q(l) := \frac{1}{N} \sum_{i=1}^N m_{q,i}(l)$, where $m_{q,i}(l)$ is the value of $m_q(l)$ defined in (3) for the i -th stock and $N = 862$ in our case. We get $\hat{H} = 0.059$ and $\hat{v} = 1.03$. Figure 14 gives the corresponding out-of-sample MSE relative to those fitted on each stock based on a sliding window of 1000 days. Interestingly, the device with predetermined parameters performs slightly better than the asset-specific ones, which could be impacted more by estimation bias due to limited data. This suggests the universality of the RFSV method in the sense that the same parameters work for all stocks¹⁹.




Figure 13: Distribution of \hat{H} and \hat{v} calibrated with data of 2012 - 2015 of US market.

¹⁹Similar results can be obtained when H and v are set with the medians of the estimations of all stocks.




Figure 14: Out-of-sample MSE of the RFSV forecast with $\hat{H} = 0.059$, $\hat{\nu} = 1.03$, relative to that of calibration on each stock.

- QRH

Following the idea on Zumbach effect of the QRH model, we propose the following forecasting device:

$$\begin{aligned}\hat{\sigma}_t^2 &= a(Z_{t-1} - b)^2 + c, \\ Z_t &= \sum_{i=1}^n c_i^d Z_{i,t},\end{aligned}\tag{7}$$

$$Z_{i,t} = e^{-\gamma_i^d} Z_{i,t-1} + r_t, \quad Z_{i,0} = z_{i,0}, \quad i = 1, \dots, n,\tag{8}$$

where $a, b, c > 0$ and $(c_i^d, \gamma_i^d)_{i=1, \dots, n}$ are given by the same multi-factor approximation of the rough kernel function $K(t) := \frac{t^{H-1/2}}{\Gamma(H+1/2)}$ as that used in Rosenbaum & Zhang [2022a]²⁰. Given some H and the number of factors n for the approximation, we recall that $(c_i^d, \gamma_i^d)_{i=1, \dots, n}$ are not free parameters to calibrate. Equations (7, 8) can be seen as a multi-factor discretization of the process $Z_t = \int_{-\infty}^t \frac{(t-s)^{H-\frac{1}{2}}}{\Gamma(H+\frac{1}{2})} \sigma_s dW_s$ ²¹, which is essentially a moving average of past returns. The above model is not very sensitive to $(z_{i,0})_{i=1, \dots, n}$ because of the exponential decay. In fact, by (7, 8), we have

$$Z_T = \sum_{i=1}^n \beta_T^i z_{i,0} + \sum_{t=1}^T r_t \sum_{i=1}^n \beta_{T-t}^i, \quad T = 1, 2, \dots,$$

where

$$\beta_t^i = c_i^d e^{-t\gamma_i^d}, \quad t = 1, 2, \dots.$$

One can easily verify that for any $i \in \{1, \dots, n\}$, the relative weight of $z_{i,0}$, given by $\tilde{\beta}_T^i := \frac{\beta_T^i}{\sum_{i=1}^n \sum_{t=0}^T \beta_t^i}$, is decreasing with T . In practice, given a time series $(r_t, \sigma_t)_{t=1, \dots, L}$ of length L , we can take $z_{i,0} = 0$ and

²⁰With daily data we use $c_i^d = \frac{c_i}{\sum_i c_i}, \gamma_i^d = \frac{\gamma_i}{250}$, where c_i and γ_i following the same definition as in Rosenbaum & Zhang [2022a]

²¹Applying the multi-factor approximation avoids choosing a predetermined Δ to compute $Z_t \simeq \int_{t-\Delta}^t \frac{(t-s)^{H-\frac{1}{2}}}{\Gamma(H+\frac{1}{2})} \sigma_s dW_s$. Long-range information can be well captured by the one-step update of $(Z_{i,\cdot})_{i=1, \dots, n}$ following (8).

IV. On the universality of the volatility formation process: when machine learning and rough volatility agree

compute $(Z_{i,t})_{t=1,\dots,L}$ following (8). Then we can simply discard the m initial samples $(Z_{i,t}, \sigma_t)_{t=1,\dots,m}$ by choosing

$$m = \max\{k | k \in \{1, \dots, L\}, \exists i \in \{1, \dots, n\}, \tilde{\beta}_k^i < \epsilon\},$$

with $0 < \epsilon \ll 1$. Then based on the remaining data, we take

$$(\hat{a}, \hat{b}, \hat{c}) = \operatorname{argmin} l(a, b, c) := \sum_{t=m+1}^L (\hat{\sigma}_t - \sigma_t)^2. \quad (9)$$

In our experiments we use SciPy optimize package with L-BFGS-B as the default algorithm for solving (9).

Figure 15 shows the empirical distributions of \hat{a} , \hat{b} and \hat{c} with data of 2012-2015 of the US market. The deviations across stocks seem to be large, especially for \hat{b} . In fact as indicated in Figure 6, the impact of past returns on current volatilities is not as apparent as that of past volatilities, therefore it is not surprising to get noisy estimations. We are interested in how the forecasting performance evolves when we use fixed $(\hat{a}, \hat{b}, \hat{c})$ for all stocks. Similarly, we fit them on the data of all the stocks of the US market. With now $\frac{1}{N} \sum_{i=1}^N l^i$ as the objective to minimize, where l^i representing the MSE corresponding to the i -th stock as defined in (9), we get $\hat{a} = 0.065$, $\hat{b} = 0.74$ and $\hat{c} = 0.55$. Figure 16 shows the resulting out-of-sample performance, along with that of stock-specific calibrations where (a, b, c) are fitted on each stock based on a sliding window of 1000 days. We do not remark that the forecasting with fixed parameters performs significantly worse than the other. Note that the QRH forecast is outperformed by the HAR as the former does not use past volatilities.




Figure 15: Distribution of \hat{a} , \hat{b} and \hat{c} based on the data of 2012-2015 of the US market.

- RFSV + QRH

We have seen that for the RFSV and QRH forecasts, properly fixing the parameters can obtain similar performance as allowing the parameters to be fitted on each stock. Since they do not depend on the same data for the forecast, we use the following combination as the final forecast:

$$(1 - \lambda)\hat{\sigma}^{RFSV} + \lambda\hat{\sigma}^{QRH},$$

where $\lambda \in (0, 1)$, $\hat{\sigma}^{RFSV}$ and $\hat{\sigma}^{QRH}$ stand for the RFSV and QRH forecasts respectively. We evaluate this mixed forecast for different λ , with the other parameters fixed to the same values as above, *i.e.*

$$\hat{H} = 0.059, \hat{v} = 1.03, \hat{a} = 0.065, \hat{b} = 0.74, \hat{c} = 0.55.$$




Figure 16: Out-of-sample MSE of the QRH forecast with calibration on each stock and with fixed parameters, relative to the HAR forecast.

Figure 17 and 18 show their forecasting performance relative to that of $LSTM_{ret}^{us}$ in US and European markets respectively. We remark that with $\lambda \sim 0.1$, the blended forecast can benefit from the complementary characteristic given by the RFSV and QRH forecasts, leading to the same level of performance as the universal nonparametric $LSTM_{ret}^{us}$.




Figure 17: Out-of-sample performance of the forecast $(1 - \lambda)\hat{\sigma}^{RFSV} + \lambda\hat{\sigma}^{QRH}$ relative to $LSTM_{ret}^{us}$ in the US market.

Basically, the future volatility is made of two components depending respectively on past volatilities and past price trends. The first one describes the expected future log-volatility from a linear combination of past log-volatilities. The associated coefficients depend only on the roughness parameter H . A universal H works for all stocks. The second component corresponds to the feedback effect of past price trends

IV. On the universality of the volatility formation process: when machine learning and rough volatility agree




Figure 18: Out-of-sample performance of the forecast $(1-\lambda)\hat{\sigma}^{RFSV} + \lambda\hat{\sigma}^{QRH}$ relative to $LSTM_{ret}^{eu}$ in the European market.

on volatility. The past price trends are built with a rough kernel, and their impact is expressed in a quadratic form, with essentially constant parameters across assets.

6 Conclusion

In this work, we shed a light on the universality of the volatility formation process relating past market realizations to current volatilities at the daily scale. For that, we apply the LSTM network to forecast the next daily volatility, based on past daily volatilities and returns. The universal model, trained on a pooled dataset of hundreds of stocks, outperforms consistently the asset-specific parametric models based on past volatilities. Interestingly, similar superior performances hold on assets that are not part of the training set, even on those of a different market. Fine-tuning the universal model with the data of each stock does not help improve the performance. These observations suggest the existence of a universal volatility formation mechanism from a nonparametric perspective. Then we uncover the learned universal volatility formation process with a parsimonious parametric formulation involving the RFSV and QRH forecasting devices. It turns out that a simple combination of the RFSV and QRH forecasts with fixed parameters perform similarly to our LSTM network. From a parametric perspective, this confirms again the “targeted” universality, showing that the main features of this universal volatility formation process can be well described by the rough volatility paradigm boosted with *strong* Zumbach effect. Our findings can also be helpful for volatility forecasting with richer predictor variables. For example, one can imagine a mixed forecast, containing the endogenous part induced by past market realizations via the uncovered universal mechanism and other temporal effects contributed by some exogenous variables.

Towards systematic intraday news screening: a liquidity-focused approach

Abstract

News can convey bearish or bullish views on financial assets. Institutional investors need to evaluate automatically the implied news sentiment based on textual data. Given the huge amount of news articles published each day, most of which are neutral, we present a systematic news screening method to identify the “true” impactful ones, aiming for more effective development of news sentiment learning methods. Based on several liquidity-driven variables, including volatility, turnover, bid-ask spread, and book size, we associate each 5-min time bin to one of two specific liquidity modes. One represents the “calm” state at which the market stays for most of the time and the other, featured with relatively higher levels of volatility and trading volume, describes the regime driven by some exogenous events. Then we focus on the moments where the liquidity mode switches from the former to the latter and consider the news articles published nearby impactful. We apply naive Bayes on these filtered samples for news sentiment classification as an illustrative example. We show that the screened dataset leads to more effective feature capturing and thus superior performance on short-term asset return prediction compared to the original dataset.

Keywords— News screening, intraday liquidity, mode fitting, sentiment learning, jump model, exogenous events.

1 Introduction

The price of a financial asset is driven by endogenous activities, such as self-reflexive trades, and also exogenous information. A main component of the latter comes from news releases. Nowadays, the financial market is becoming increasingly efficient, yet the embodiment of new information transferred by news in asset price is rarely accomplished instantaneously. Quick and effective estimation of news sentiment, *i.e.* whether the view given by a news article is bullish or bearish, can give profitable opportunities to investors. As said in Pedersen [2019], “financial markets are efficiently inefficient”, in the idea that professional investors with superior performance are compensated for their costs and risks, and the competition among them makes markets almost efficient. Given the large number of news publications every day, manual analysis of each piece of news is infeasible. To assess efficiently the impact of a news release on the market price of its associated financial asset, institutional investors then need to develop automatic news sentiment evaluation methods.

Numerous works using news data to predict financial assets’ price movements exist in the literature. In Chan [2003], it is documented that the monthly returns of the stocks associated with public news releases are less likely to reverse than those without identifiable news publication, suggesting that news can publish some information concerning the “fair” value of stocks. Jiang et al. [2021] extend the study with intraday returns in the same spirit. A long/short trading strategy exploiting the news-driven price drifts is shown to generate abnormal profit. In these works, the views transferred by certain news releases on the stocks of interest, whether bullish or bearish, are identified by the associated post-news market reactions. Thus to make an investment decision, investors have to wait until the emergence of some significant price drifts after news publications. To react more quickly, sentiment evaluation methods based on textual data are required. Tetlock [2007] applies the Harvard-IV psychosocial dictionary on the articles from *The Wall Street Journal* to estimate the pessimistic pressure on the Dow Jones Industrial Average index. Loughran & McDonald [2011] construct a customized dictionary more adapted to financial text using the term frequency-inverse document frequency (tf-idf) method. Stock market sentiment lexicons based on microblogging data are developed in Oliveira et al. [2016] and Renault [2017] by computing several statistical measures. Ke et al. [2019] estimate the tone weights of sentiment-charged words via a topic modeling method, and compute the article-level sentiment score using a regularized version of maximum likelihood estimation.

Another strand of research is using deep learning methods on news data in the same vein as current practices in natural language processing. Ding et al. [2014] transform news publications into structural events using Open Information Extraction techniques. A novel neural tensor network is used in Ding et al. [2015] to represent news text as dense vectors. Then the representations are put into a convolutional neural network for predicting future stock price movements. The dependence among sequential news releases is considered in Hu et al. [2018] by proposing an attention-based recurrent neural network. Chen [2021] develops a fine-tuned bidirectional encoder representations from transformers, see Kenton & Toutanova [2019], to produce contextualized word embeddings, which are then fed into a recurrent neural network to output news classification results.

To develop a news sentiment learner, one needs firstly labeled samples, *i.e.* bullish or bearish news releases, for model training. However, the sentimental nature of a news article is not explicitly marked in real life. In the case of Bloomberg News²², the first step of model development consists of manual labeling by human experts, in the idea to isolate “true” sentiment contained in news from realized price actions. However, this method can be exposed to subjective biases. More importantly, one has to repeat

²²<https://www.bloomberg.com/professional/product/event-driven-feeds/>

the manual labeling when moving to a new training set, which can be costly. More efficient systematic news labeling methods are preferred in this case. A common practice in the research community is to take the sign of the share price movement following a news publication as the ground truth label [Chen, 2021; Ding et al., 2014, 2015; Hu et al., 2018]. This is understandable since news sentiments are expected to predict positively the price drifts of associated assets. However, not every news article conveys a directional view of the underlying asset price. Labeling news only based on post-publication returns would result in a non-negligible proportion of falsely labeled samples in the training set, which causes extra uncertainties for model learning. Chen [2021] keeps only a small portion of samples with extreme market realized returns for model training to reduce the impact of neutral news. Still, we believe that relying solely on the posterior return is incomplete, given the low signal-to-noise ratio of return data.




Figure 1: Dynamics of stock price and volatility of Apple(left) and Nvidia(right) around two news publications, indicated by the arrows.

In this chapter, we aim at identifying the “true” impactful news releases, *i.e.* those imply positive or negative views on the associated assets, with a particular focus on the liquidity states of the assets of interest. Figure 1 illustrates our main motivation by examining stocks’ price and volatility dynamics around news publications with two particular examples. Both news releases would be considered positive if we look at only the short-term realized return. Yet the first one is more likely to be a piece of neutral news for Apple, which is consistent with the mediocre volatility fluctuation. As for the case of Nvidia, we observe a more significant volatility change when the message conveys a positive view on the company’s potential. Note that here we apply the model with uncertainty zones introduced by Robert & Rosenbaum [2011, 2012] for effective computation of high-frequency volatility. Thus, we can conceive a two-step news screening approach. First, we divide the daily core trading session into multiple nonoverlapping time bins of equal size, and for each bin, we estimate the realized price volatility. These estimations can be sorted out using clustering methods into two clusters representing two distinct liquidity modes. We designate the one with a lower volatility level as *Mode 1* and the other as *Mode 2*. Second, we focus on the jumps from *Mode 1* to *Mode 2*, and consider the news articles published around these jumps impactful. That is, these jumps are thought to have caused these volatility changes.

The above method is consistent with the findings in Groß-Klußmann & Hautsch [2011]. Based on news data concerning 40 stocks actively traded at the London Stock Exchange, preprocessed by the Reuters NewsScope Sentiment Engine, Groß-Klußmann & Hautsch [2011] suggests that the news with high relevance impact significantly volatilities and trading volumes. To get a more robust liquidity mode

fitting, the actual indicator set chosen in this work includes four common liquidity-driven variables, *i.e.* volatility, turnover, bid-ask spread, and book size [Bińkowski & Lehalle, 2022]. Instead of using classical clustering methods like K-means to distinguish two liquidity modes, we apply the jump model introduced in Bemporad et al. [2018] to take also the ordering of observations into account. In the jump model, we can easily penalize frequent mode switches, and thus some mode persistence is favored. This is relevant for time series describing intraday liquidity conditions.

Joulin et al. [2008]; Marcaccioli et al. [2022] investigate the volatility fluctuations around the price jumps induced by news releases, suggesting that they exhibit different dynamical patterns to those arising from endogenous activities. One can thus monitor the volatility dynamics around each news publication, and then decide whether the release has impacted the market according to the observed features. However, we have to consider the following inconveniences when applying this method in practice. First, news screening one by one is very time-consuming given the large size of the news dataset. Second, when some other liquidity-driven variables in addition to volatility are considered, as in the case of our current work, the dynamical features of these variables differentiating between the exogenous and endogenous events need to be specified explicitly, and then news classification rules should be modified accordingly.

Our approach is very efficient by structuring the task into two steps: liquidity mode fitting on a time-bin basis and associating news releases with detected liquidity mode jumps. The fitting results can also be used for news data from other providers or any other exogenous events/signals in a similar manner. The jump model fits data in a nonparametric way, allowing us to be agnostic about the dynamical properties of each measured variable. Other variables can be easily tested in the same manner. Note that only the news releases that happened during the daily core trading sessions are concerned by our screening approach. Once the impactful releases are targeted, we mark them as positive or negative by the signs of the post-news returns of the associated assets, similarly to Chen [2021]; Ding et al. [2014, 2015]; Hu et al. [2018]. Then various supervised learning methods can be applied to the labeled samples to learn sentiment-charged features. In this work, we focus on the effect of our news screening method instead of developments of alternative sentiment learning methods. To illustrate the idea, we fit two Bernoulli naive Bayes classifiers (NBCs) respectively on the original intraday news data and the dataset filtered by liquidity mode changes. Based on out-of-sample tests, the classification results given by the later classifier are more consistent with the post-news asset price movements than the former, *i.e.* in average the news articles with high probability to be positive (negative) assigned by the latter classifier show more significant and persistent positive (negative) post-news price drift than the ones sorted out by the former classifier. This indicates that the screened dataset includes fewer falsely labeled samples, and thus can lead to more effective sentiment learning.

This chapter is organized as follows. In Section 2, we describe firstly the data involved in this study and related preprocessing procedures. The application of the jump model on liquidity mode fitting is detailed in Section 3. Numerical results with market data are presented. We then present how to use the fitted mode sequences to identify impactful news releases in Section 4. The relevance of our method for more effective news sentiment learning is then illustrated through numerical experiments. Finally, we conclude with our main findings in Section 5.

2 Data and preprocessing

In this work, we use the Daily Trade and Quote (TAQ) dataset²³ for computing the considered liquidity-driven variables. The TAQ dataset covers all stocks traded in the US market. To avoid any results biased by small-cap names, we focus on the components of S&P 500. The following variables are measured on consecutive bins of five minutes during the daily core trading session:

- *Average bid-ask spread in ticks* (ϕ). For each second, we record the last observed bid-ask spread. The average value for each 5-minute bin is computed and its ratio against the tick size is kept.
- *Traded value/turnover* (V), is the total value traded during each 5-minute bin.
- *Volatility* (σ). Instead of using the Garman-Klass volatility as in Bińkowski & Lehalle [2022], here we take the one based on the model with uncertainty zones, whose effectiveness on high-frequency data is shown in [Robert & Rosenbaum, 2011, 2012].
- *Average book size* (B). We record the average volume available at the best bid and ask prices of each second inside each 5-minute bin.

After removing the data concerning the first and last 15 minutes of the daily core trading period²⁴, for each (stock, day) pair, we get a time series of length $T = 72$. Intraday seasonalities of liquidity variables, induced by certain fundamental reasons, are well known [Bińkowski & Lehalle, 2022; Lehalle & Laruelle, 2018], and our objective is to identify the short-term impacts of news beyond these effects. For each observation $o_t^{S,d} \in \{\phi_t^{S,d}, V_t^{S,d}, \sigma_t^{S,d}, B_t^{S,d}\}$ of stock S on day d with $t = 1, \dots, T$, we use the following stationarization procedure:

$$o_t^{S,d} = \frac{\log o_t^{S,d} - \frac{1}{D} \sum_{d'=1}^D \log o_t^{S,d'}}{\text{IQR}(\log o_t^{S,1}, \dots, \log o_t^{S,D})},$$

where $\text{IQR}(\cdot)$ means the difference between the 75th and 25th percentiles of the data, and D represents the total available days on the dataset under study. We use IQR instead of the ordinary standard deviation to reduce the impact of outliers. In this chapter, data covering 2017-01-01 ~ 2019-12-31 is selected for examining the association between intraday liquidities and news data, *i.e.* $D \simeq 750$. After this preprocessing, the intraday seasonalities can be mostly removed and all the variables of all stocks have similar scales, which is essential for the fitting of intraday liquidity states, as will be detailed in the following.

Table 1 gives several samples of the Bloomberg News data that we used in this work. Note that we evaluate the news sentiment only based on the headlines. The *Score* and *Confidence* fields give sentiment estimation of the news articles, which are based on Bloomberg's proprietary classification algorithm. The *Score* says whether a piece of news is bullish(1), bearish(-1), or neutral(0), with the *Confidence* indicating the reliability of this classification. We can thus compute an *composed score* := $\text{score} \times \text{confidence}$. Naturally, we expect on average the news releases with the highest *composed score* are followed with significantly positive price drift and the opposite for the ones with the lowest *composed score*. The predictive power of this score will be compared with the sentiment score given by the NBCs, as will be shown later.

Figure 2 shows the number of news releases related to our work. We screen the news headlines published during the years 2017~ 2019 based on the liquidity-driven variables as introduced above. Thus only the

²³<https://www.nyse.com/market-data/historical/daily-taq>

²⁴We recall that the regular trading hours for the US stock market are 9:30 a.m. to 4:00 p.m.

V. Towards systematic intraday news screening: a liquidity-focused approach

| Headline | TimeStamp | Ticker | Score | Confidence |
|---|-------------------------|---------|-------|------------|
| 1st Source Corp: 06/20/2015 - 1st Source announces the promotion of Kim Richardson in St. Joseph | 2015-06-20T05:02:04.063 | SRCE | -1 | 39 |
| Siasat Daily: Microsoft continues rebranding of Nokia Priority stores in India opens one in Chennai | 2015-06-20T05:14:01.096 | MSFT | 1 | 98 |
| Rosneft, Eurochem to cooperate on monetization at east urengoy | 2015-06-20T08:01:53.625 | ROSN RM | 0 | 98 |

Table 1: Several samples of the Bloomberg News data.

ones released during the daily trading hours is concerned for this step. We then fit a news classifier to output sentiment scores capturing the sentiment-charged features based on the samples picked out. The classifier is evaluated on all the intraday news headlines published during 2020-01-01 ~ 2021-12-31, with a particular focus on the predictive power of the sentiment scores on the short-term post-publication price drifts.




Figure 2: Count of the news releases concerned in current work.

3 Liquidity mode fitting

3.1 Jump model

Given a sequence of data pairs $(x_t, y_t)_{t=1, \dots, T}$ with $x_t \in \mathcal{X}, y_t \in \mathcal{Y}$ and the number of modes K , the jump model introduced in [Bemporad et al. \[2018\]](#) outputs a mode sequence $(m_t)_{t=0, \dots, T}$ with $m_t \in \{1, \dots, K\} := \mathcal{K}$, and the parameter $\theta_m \in \mathbb{R}^a$ associated with each $m \in \mathcal{K}$, with a some positive constant. The obtained $\Theta := (\theta_1, \dots, \theta_K)$ and $M := (m_0, \dots, m_T)$ minimize the following objective function

$$J(X, Y, \Theta, M) = \sum_{t=1}^T l(x_t, y_t, \theta_{m_t}) + \sum_{k=1}^K r(\theta_k) + \mathcal{L}(M),$$

where $X := (x_1, \dots, x_T), Y := (y_1, \dots, y_T), J(\cdot), l(\cdot), r(\cdot)$ and $\mathcal{L}(\cdot)$ are all functions taking value on \mathbb{R} . The first two parts to the right of the equal sign represent respectively the total fitting loss of the given data and a regularization term on the model parameters. For example, when $K = 1$, $l(x, y, \theta) = \|y - \theta'x\|_2^2$ and $r(\theta) = \lambda \|\theta\|_2^2$ with $\lambda > 0$, we get the standard setting for Ridge regression. The particularity of the jump model relies on the introduction of the loss $\mathcal{L}(M)$ taking the ordering of the mode sequence into account. It is defined as

$$\mathcal{L}(M) = \mathcal{L}^{init}(m_0) + \sum_{t=1}^T \mathcal{L}^{mode}(m_t) + \sum_{t=1}^T \mathcal{L}^{trans}(m_t, m_{t-1}),$$

where the loss is decomposed into three parts, the penalization on the initialization \mathcal{L}^{init} , the one linked to the fitting result of each timestamp \mathcal{L}^{mode} , and the cost concerning mode transitions \mathcal{L}^{trans} . As suggested in [Bemporad et al. \[2018\]](#), this definition generalizes popular models such as hidden Markov models. Various cost functions can be chosen to meet the needs of different applications. We refer to [Bemporad et al. \[2018\]](#) for more details.

In our case, we apply the jump model in an unsupervised learning manner. For each stock-day pair (S, d) with $S = 1, \dots, N$ and $d = 1, \dots, D$, we have a preprocessed sequence $X^{S,d} := (x_t^{S,d})_{t=1, \dots, T}$ where $x_t^{S,d} := (\phi_t^{S,d}, V_t^{S,d}, \sigma_t^{S,d}, B_t^{S,d}) \in \mathbb{R}^4$. The model consists in finding a mode sequence $M^{S,d} := (m_t^{S,d})_{t=1, \dots, T} \in \mathcal{K}^T$, under which the observations associated with the same mode are more similar to each other than to those linked with other modes. We denote the K mode represents by $\Theta^{S,d} := (\theta_k^{S,d})_{k=1, \dots, K} \in (\mathbb{R}^4)^K$. We have no prior knowledge about the initial mode $m_0^{S,d}$, and impose no mode-specific cost, that is, $\mathcal{L}^{init} = \mathcal{L}^{mode} = 0$. We penalize frequent mode switches, expecting some degree of persistence for the fitted mode sequence. Particularly, it leads to the same type of loss function as in [Nystrup et al. \[2021\]](#), which reads

$$J(X^{S,d}, \Theta^{S,d}, M^{S,d}) = \sum_{t=1}^T l(x_t^{S,d}, \theta_{m_t^{S,d}}^{S,d}) + \lambda \sum_{t=1}^{T-1} \mathbb{1}_{m_t^{S,d} \neq m_{t+1}^{S,d}}, \quad (1)$$

where $l(x, \theta) = \|x - \theta\|_2^2$, with $\|\cdot\|_2$ representing the L^2 norm, and λ is a hyperparameter trading off between clustering the given data and mode persistence. In practice, it can be chosen via cross-validation. Note that when $\lambda = 0$, we will get the classical K-means solution. As for the number of modes, larger K can bring better fit, while it becomes less evident to interpret and involves a higher risk of overfitting. Since we focus on the market liquidity regime switch caused by exogenous information, we simply take $K = 2$ in the following tests. We will see in the following that the resulting two modes are separated in terms of volatility level. We let the mode associated with lower volatility be *Mode 1*, and the other be *Mode 2*.

Therefore, the results of minimization of (1) are a sequence of liquidity modes associated with each time bin and two vectors of dimension four. Note that the latter are obtained based on $T = 72$ observations, which still implies some potential risk of overfitting. To reduce this risk, we can expand the fitting on sequences of multiple days. However, as the liquidity-driven variables are not homogeneous across time, *e.g.* certain periods are more volatile than the others, the fitted modes will not distribute uniformly over time. For example, we expect that *Mode 2* will concentrate on periods associated with relatively larger volatility. Thus in this case, the observed mode switches are mostly the results of some market-wide events such as monetary policy announcements, instead of stock-specific news releases. In this work, we expand the fitting space in the dimension of the asset. More precisely, on day d the model is fitted on the pooled set of sequences $X^d := \{(x_t^{S,d})_{t=1,\dots,T}\}_{S=1,\dots,N}$. As market-wide activities impact all the stocks to a similar extent, and are unlikely to change abruptly at the intraday scale, the observations $(x_t^{S,d})_{S=1,\dots,N,t=1,\dots,T}$ are influenced uniformly by these market-wide events across S and t . In this way, the fitting results can reflect the short-term liquidity fluctuations from the base level related to the market environment. Significant liquidity changes are thought to be induced by some exogenous events, *i.e.* news releases in our case. We ignore the superscript d for ease of notation in the following. The actual fitting objective reads thus

$$\arg \min_{\Theta, M} J(X, \Theta, M) = \sum_{S=1}^N \left(\sum_{t=1}^T l(x_t^S, \theta_{m_t^S}) + \lambda \sum_{t=1}^{T-1} \mathbb{1}_{m_t^S \neq m_{t+1}^S} \right), \quad (2)$$

where $\Theta := (\theta_1, \dots, \theta_K)$ and $M := \{(m_t^S)_{t=1,\dots,T}\}_{S=1,\dots,N}$. Problem (2) can be solved with a simple coordinate-descent optimization algorithm that alternates minimization with respect to Θ and M , which is detailed in Appendix V.A.

3.2 Fitting results

In this part, we give several statistics on the intraday liquidity mode fitting results for the period 2017-01-01 ~ 2019-12-31. We recall that in this study $K = 2$ and we test several different λ to see its effect. Given the set of all fitted mode sequences $\{M^{S,d}\}_{S=1,\dots,N,d=1,\dots,D}$, we count respectively the occurrences of *Mode 1* and *Mode 2*. We also estimate a 2×2 implicit transition matrix \mathcal{F} , whose entries are computed as follows:

$$\mathcal{F}_{ij} = \frac{\sum_{d=1}^D \sum_{S=1}^N \sum_{t=1}^{T-1} \mathbb{1}_{m_t^{S,d}=i, m_{t+1}^{S,d}=j}}{\sum_{d=1}^D \sum_{S=1}^N \sum_{t=1}^{T-1} \mathbb{1}_{m_t^{S,d}=i}}, \quad i, j \in \{1, 2\}.$$

Figure 3 gives the results with respect to different λ . Most of the time the market stays at *Mode 1*. With larger λ , we get a slightly smaller assignment ratio for *Mode 2* and more pronounced mode persistence. Note that even when we do not penalize the mode switch, *i.e.* $\lambda = 0$, the diagonal elements of \mathcal{F} are significantly larger than the off-diagonal ones. This is consistent with the observations in Bińkowski & Lehalle [2022] that all the selected liquidity-driven variables are positively autocorrelated, and thus the points of adjacent time bins are likely to be classified into the same mode. Accordingly, Figure 4 gives the average daily count of liquidity mode switches from *Mode 1* to *Mode 2* per stock.

Figure 5 plots the dynamics of fitted mode parameters during our testing period when $\lambda = 0.5$. As expected in Introduction, the volatility levels of the two modes are well distinct. Interestingly, it is also the case for the traded volume V . The differentiation of ϕ and B between the two modes are relatively less noticeable. It is unlikely that the detected *Mode 2* corresponds mostly to the time slots with endogenous volatility spikes, which are usually accompanied by increased bid-ask spread as suggested,




Figure 3: Mode distribution of all fitted 5-minute bins (top) and implicit mode transition matrix estimated from the fitting results.




Figure 4: Average daily count of jumps from *Mode 1* to *Mode 2* per stock.

for example, in [Wyart et al. \[2008\]](#). Of course considering the multivariate nature in (2), the resulting pattern depends on the set of variables that we chose at the beginning. Developing other features and selecting the most effective ones in the context of impactful news screening is out of the scope of the current work.




Figure 5: Historical evolution of fitted mode parameters when $\lambda = 0.5$.

4 News screening and learning

4.1 Methodology

It is often suggested in the literature that exogenously driven orders can generate a larger volatility footprint than endogenously mechanical ones, see for example [Huang et al. \[2015\]](#); [Rambaldi et al. \[2019\]](#). As *Mode 2* is characterized by higher volatility and increased trading volume, it is reasonable to assume that some exogenous events, news releases in this chapter, evoked the mode switches. Considering a piece of news arriving inside the t -th 5-minute bin as shown in Figure 6, with $t = 1, \dots, 72$, we consider it impactful if one of the following scenarios happens:

$$\begin{cases} \text{Jump from } \textit{Mode 1} \text{ to } \textit{Mode 2} \text{ at the moment } 5t - 5, & \text{for } t \in \{2, \dots, 72\}, \\ \text{Jump from } \textit{Mode 1} \text{ to } \textit{Mode 2} \text{ at the moment } 5t, & \text{for } t \in \{1, \dots, 71\}. \end{cases} \quad (3)$$

Let the set of original intraday news releases and the impactful ones selected with the above criterion




Figure 6: Possible scenarios where a news release is considered impactful.

be \mathcal{D} and \mathcal{D}_λ respectively, where λ is the mode switch penalization parameter defined in (2). Given a piece of news $a \in \mathcal{D}$, released at t^* and concerning the stock S , we denote the h -minute return of the stock S^* after the publication of a by r_a^h , i.e. $r_a^h := \frac{P_{t^*+h}^{S^*}}{P_{t^*}^{S^*}} - 1$, where P_t^S represents the price of stock S at time t . Since we are interested in the firm-specific price movements disentangled from market-wide activities, we replace r_a^h with its market-detrended version defined by

$$r_a^h := r_a^h - \frac{1}{N} \sum_{S=1}^N \left(\frac{P_{t^*+h}^S}{P_{t^*}^S} - 1 \right),$$

where for sake of simplicity, we follow the classical capital asset pricing model with the *betas* fixed to be one. With $\mathcal{R}^h := \{r_a^h | a \in \mathcal{D}\}$, let $r^{h,k}$ and $r^{h,100-k}$ be the k -th and $(100-k)$ -th percentile of \mathcal{R}^h respectively. We define the following sets

$$\mathcal{Z}_{h,k}^- := \{a | r_a^h \leq r^{h,k}, a \in \mathcal{D}\} \quad \text{and} \quad \mathcal{Z}_{h,k}^+ := \{a | r_a^h \geq r^{h,100-k}, a \in \mathcal{D}\}.$$

Thus, Z^- and Z^+ are respectively sets of bearish and bullish news releases according to the amplitude of post-publication stock returns, which is the criterion commonly used in works such as [Chen \[2021\]](#); [Ding et al. \[2014, 2015\]](#); [Hu et al. \[2018\]](#). The cases with $k \ll 50$ are in the spirit that “true” impactful news releases are more likely to be followed with significant price movements. Given λ, h and k , the sets of “true” bearish and bullish news in our approach are defined respectively by

$$\mathcal{N}_{\lambda,h,k}^- := \mathcal{D}_\lambda \cap \mathcal{Z}_{h,k}^- \quad \text{and} \quad \mathcal{N}_{\lambda,h,k}^+ := \mathcal{D}_\lambda \cap \mathcal{Z}_{h,k}^+.$$

Therefore, in addition to extreme post-publication return, the news publications selected by our method is also followed by noticeable changes in market liquidity conditions.

Given a set of labeled samples, such as $\mathcal{E} = \mathcal{Z}^+ \cup \mathcal{Z}^-$ or $\mathcal{E} = \mathcal{N}^+ \cup \mathcal{N}^-$, we firstly inquire the dependence of news sentiment, positive or negative, on the presence or absence of each individual word through measuring the mutual information between these two random variables. Then we fit a multi-variate Bernoulli NBC as a news sentiment predictor. Some key computational rules are recalled in Appendix V.B. We refer to for instance [Cover et al. \[1991\]](#) and [McCallum et al. \[1998\]](#) for more details of these methods. For a piece of news a , the classifier tells us the probabilities of a being positive and negative. Since there is no ground truth for the news sentiment, we evaluate the classification results based on

stocks’ post-news price movements, under the hypothesis that news sentiment is positively correlated with the direction of stocks’ future return. More precisely, for an NBC fitted on the training set \mathcal{E} , we define the sentiment score $F \in [-1, 1]$ of news a by

$$F(a)|\mathcal{E} := P_{\mathcal{E}}^{+}(a) - P_{\mathcal{E}}^{-}(a),$$

where $P_{\mathcal{E}}^{+}$ and $P_{\mathcal{E}}^{-}$ are the probabilities of a being bullish and bearish respectively given by the model. We expect naturally that a significantly high (low) F is more likely to be followed by a positive (negative) price drift for the associated stock.

4.2 Numerical results

In this part, we perform news sentiment learning on the screened news dataset. Figure 7 gives the resulting ratio of news releases selected with the criterion (3). Note that with $\lambda = 0.5$, only around 10% of total intraday releases are thought to have impacted the market in our approach.




Figure 7: Ratio of news releases with identified nearby jumps from *Mode 1* to *Mode 2*, i.e. $\frac{\#\mathcal{D}^{\lambda}}{\#\mathcal{D}}$.

4.2.1 Most informative words

Figure 8 shows the words reporting the most mutual information with news sentiment, measured respectively on $\mathcal{Z}_{15,10}^{+} \cup \mathcal{Z}_{15,10}^{-}$ and $\mathcal{N}_{0.5,15,10}^{+} \cup \mathcal{N}_{0.5,15,10}^{-}$. Visually, our news selection method boosted with liquidity-driven variables can reduce the weights of certain sentiment-neutral words, e.g. “boeing”, “737”, “stubhub”, “http”, etc. It also highlights some sentiment-charged words, e.g. “senseless”, “abandon”, “recall”, “poised”, “evil”, etc. Considering the limitation of uni-gram evaluation, it is not surprising that several neutral words still seem to be overvalued with our approach. For example, “headquarters” solely is not sentiment-charged, while “build new headquarters” is likely to drive some stock price movements.



Figure 8: Words showing the most mutual information with news class variable measured for the samples of $\mathcal{Z}_{15,10}^+ \cup \mathcal{Z}_{15,10}^-$ (left), and $\mathcal{N}_{0.5,15,10}^+ \cup \mathcal{N}_{0.5,15,10}^-$ (right). The font size of a word is proportional to the amount of its mutual information.




Figure 9: Average post-publication stock return for the news releases with significantly positive and negative sentiment scores assigned by the NBCs fitted on different training sets. The left represents the case with news screening solely based on return information. The right corresponds to our approach. Only the samples whose scores fall in the top/bottom 10% range are shown for ease of comparison. The shadow parts represent the empirical standard deviation of the average return.

4.2.2 Short-term return prediction

Now we evaluate the predictive power of the news sentiment scores, given by the fitted NBCs, for future price movement. Note that the following results are based on the news headlines published during the out-of-sample period, *i.e.* 2020-01-01 ~ 2021-12-31. We are more interested in the short-term returns following news publications since they reflect better the immediate impacts of news sentiment compared to price changes over longer horizons. We fit two NBCs on the training sets $\mathcal{Z}_{15,10}^+ \cup \mathcal{Z}_{15,10}^-$ and $\mathcal{N}_{0.5,15,10}^+ \cup \mathcal{N}_{0.5,15,10}^-$ respectively. They are then applied to the out-of-sample news data to output sentiment estimations. We plot the average post-publication price drifts of the news releases associated with distinct sentiment scores up to 150 minutes in Figure 9. Clearly, after screening the news dataset with our approach, NBC can better predict the short-term stock return. The pieces of news with significantly positive sentiment scores are followed by climbing prices, while the ones with negative scores drive the inverse phenomenon.




Figure 10: Average post-publication stock return for the news releases with significantly positive and negative sentiment scores assigned by the NBC fitted on the news samples associated with extreme realized return falling in the top/bottom 1% quantile (left), and the case with the sentiment scores given by Bloomberg (right).

With $\lambda = 0.5$, the size of $\mathcal{N}_{0.5,15,10}^+ \cup \mathcal{N}_{0.5,15,10}^-$ is only about one tenth of $\mathcal{Z}_{15,10}^+ \cup \mathcal{Z}_{15,10}^-$. To verify that our selection method is not equivalent to filtering news releases by the associated post-publication stock returns, we repeat the same test on $\mathcal{Z}_{15,1}^+ \cup \mathcal{Z}_{15,1}^-$. As shown in the left subfigure of Figure 10, with similar number of training samples to the case with $\mathcal{N}_{0.5,15,10}^+ \cup \mathcal{N}_{0.5,15,10}^-$, the prediction performance now is largely degraded. Keeping only the news releases with extreme post-publication stock returns can reduce the ratio of neutral samples in the training set to some extent, while it can also result in model underfitting because of data shortage. Our method presents a more efficient way to filter neutral news samples and identify the impactful ones, which helps the model learn more effectively. In Figure 10, we show also the results of Bloomberg *composed score* as defined in Section 2. Interestingly, despite its simplicity, the NBC fitted on the screened dataset performs even slightly better than the scores given by Bloomberg in terms of short-term return prediction. Moreover, we show in Appendix V.C that the

performance of our approach is not very sensitive to the value of λ , h and k .

5 Conclusion

In this work, we introduce a systematic method for identifying “true” impactful news releases that are conceived to contain unexpected information for the financial market. The identification method consists in associating significant changes in liquidity conditions during market open hours with nearby news publications. Four variables are used to monitor the intraday dynamics of liquidity mode, including volatility, turnover, bid-ask spread, and book size. Through numerical tests on S&P500 components, the two predetermined liquidity states are distinct from each other in terms of volatility and turnover level. We pick out the news releases with identifiable mode switch from that with lower volatility and less trading volume to the other one, and label them by the sign of post-news realized returns of the associated stocks. Experiments on news sentiment learning with NBC show that the proposed news screening method leads to more effective feature capturing and thus better model predictive performance.

The study can be extended in several directions. First, our news screening approach can be taken as a preprocessing procedure, and can be applied together with various news sentiment learning methods. Second, we focus on the Bloomberg News data in this study, while the same tests can be easily conducted on any other news dataset, or more generally on other types of exogenous events/signals. Last, it would be interesting to build some agent-based models to understand further the link between exogenous inputs and resulting dynamics of volatility/turnover.

Appendix

V.A Jump model fitting

In our work, all the daily calibrations are conducted following Algorithm 1. At each iteration the loss J is non-increasing, so Algorithm 1 can always terminate in a finite number of steps. However, similar to other classical clustering methods such as K-means or Gaussian mixture methods, the final solution depends on the initialization and may not be the global optimal one. In practice, one can run the above algorithm multiple times with different initial model sequences and keep the one with the smallest loss.

V.B Mutual information and naive Bayes classifier

Mutual information

Let $C \in \mathcal{C}$ and $W \in \{0, 1\}$ be two random variables denoting respectively the category of a news release, *i.e.* positive or negative in this chapter, and the presence or absence of word w in the news text. $W = 0$ represents the absence of the word, and $W = 1$ represents the presence of w . The mutual information between C and W is given by

$$I(C; W) = \sum_{c \in \mathcal{C}} \sum_{f_w \in \{0, 1\}} P_{(C, W)}(c, f_w) \log \left(\frac{P_{(C, W)}(c, f_w)}{P_C(c)P_W(f_w)} \right),$$

where P_C and P_W are the marginal probability mass functions of C and W respectively, and $P_{(C, W)}$ is their joint probability mass function. In this work, for a given labeled set of news headlines, $\mathcal{Z}^- \cup \mathcal{Z}^+$ or $\mathcal{N}^- \cup \mathcal{N}^+$, we estimate the above probability quantities by their classical empirical observations over all samples.

Multi-variate Bernoulli naive Bayes classifier

Let c be the news class variable, f_{w_1}, \dots, f_{w_n} indicate the presence or absence of word w_1, \dots, w_n in the news headlines, we are interested in the classification result given by

$$\hat{c} = \operatorname{argmax}_c P(c | f_{w_1}, \dots, f_{w_n}).$$

Following Bayes' theorem, we have

$$P(c | f_{w_1}, \dots, f_{w_n}) = \frac{P(c)P(f_{w_1}, \dots, f_{w_n} | c)}{P(f_{w_1}, \dots, f_{w_n})}.$$

The "naive" conditional independence assumption, we have

$$P(f_{w_1}, \dots, f_{w_n} | c) = \prod_{i=1}^n P(f_{w_i} | c),$$

Algorithm 1: Fitting algorithm for the problem (2)

Input: N observations sequences $(x_t^S)_{t=1,\dots,T}$, with $S = 1, \dots, N$, jump penalty λ and convergence tolerance ϵ .

1. Generate N initial mode sequences $\{(m_t^S)_{t=1,\dots,T}\}_{S=1,\dots,N}$ through applying K-means on the pooled set of observation sequences. Initial loss $J^0 = +\infty$.
2. Iterate for $l = 1, \dots$ until $|J^l - J^{l-1}| \leq \epsilon$:

- a) Model parameter fit: for $k = 1, \dots, K$, the optimal θ_k is given by the mean of all the samples assigned with mode k , *i.e.*

$$\theta_k = \frac{\sum_{S=1}^N \sum_{t=1}^T x_t^S \mathbb{1}_{m_t^S=k}}{\sum_{S=1}^N \sum_{t=1}^T \mathbb{1}_{m_t^S=k}}, \quad i = 1, \dots, 4$$

- b) For each $S = 1, \dots, N$, solve the optimal mode sequence with respect to θ_1, θ_2 :
 - i. Compute a matrix $F^S \in \mathbb{R}^{T \times K}$, which is defined by:

$$F^S(T, k) = \|x_T^S - \theta_k\|_2^2,$$

$$F^S(t, k) = \|x_t^S - \theta_k\|_2^2 + \min_j \{F^S(t+1, j) + \lambda \mathbb{1}_{k \neq j}\},$$

where $k = 1, \dots, K$.

- ii. Reconstruct the optimal mode sequence with

$$m_1^S = \arg \min_k F^S(1, k),$$

$$m_t^S = \arg \min_k \{F^S(t, k) + \lambda \mathbb{1}_{m_{t-1}^S \neq k}\}, \quad t = 2, \dots, T.$$

- c) Update the loss by $J^l = \sum_{S=1}^N F^S(1, m_1^S)$.

Output: Model parameters $\theta_1, \dots, \theta_K$, N mode sequences $\{(m_t^S)_{t=1,\dots,T}\}_{S=1,\dots,N}$.

where $P(f_{w_i}|c)$ can be easily estimated given the Bernoulli assumption.

V.C Robustness tests

Our approach involves mainly three parameters, *i.e.* λ , h and k . In the following we show their effects by varying their values.

- Effect of λ

As shown in Figure V.C.1, the results are relatively robust with intermediate λ . When $\lambda = 1$, the amount of news samples taken as impactful becomes too limited to accurate model learning.

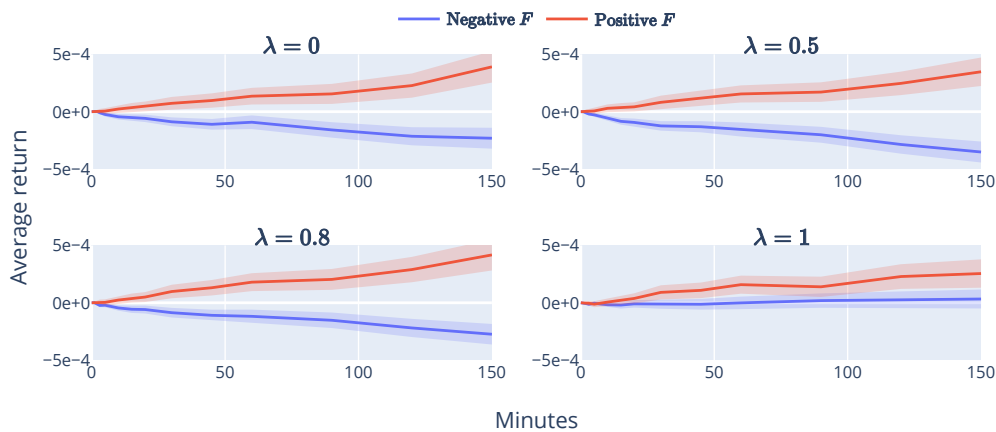


Figure V.C.1: Average post-news stock return corresponding to the cases with different jump penalization parameter λ . We set $h = 15$, $k = 10$.

- Effect of h

Figure V.C.2 gives the results when choosing the post-news realized return over different horizons for news sign labeling. We do not remark significant variation of performance for $h \in [5, 10, 15]$. Slight degradation is observed for $h = 30$, which is understandable given the increased volatility of realized return.

- Effect of k

When using only realized return information for news labeling, we are inclined to relatively small k to pick out only the news releases associated with significant post-publication price drifts. With our method, as shown in Figure V.C.3, the impact of k on the final results is very limited. Even with $k = 50$, which means that the magnitude of realized return is not considered, there is no significant




Figure V.C.2: Average post-news stock return corresponding to the cases with different h . We set $\lambda = 0.5$, $k = 10$.

deterioration for the predictive performance of the resulting model.




Figure V.C.3: Average post-news stock return corresponding to the cases with different k . We set $\lambda = 0.5$, $h = 15$.

Bibliography

- [“el2018perfect”, n.d.] (n.d.).
- [Abi Jaber, 2019] Abi Jaber, E. (2019). Lifting the Heston model. *Quantitative Finance*, 19(12), 1995–2013.
- [Abi Jaber & El Euch, 2019] Abi Jaber, E., & El Euch, O. (2019). Multifactor approximation of rough volatility models. *SIAM Journal on Financial Mathematics*, 10(2), 309–349.
- [Assaf et al., 2022] Assaf, O., Di Fatta, G., & Nicosia, G. (2022). Multivariate LSTM for stock market volatility prediction. In *Machine learning, optimization, and data science: 7th international conference, lod 2021, grasmere, uk, october 4–8, 2021, revised selected papers, part ii* (pp. 531–544).
- [Audrino et al., 2020] Audrino, F., Sigrist, F., & Ballinari, D. (2020). The impact of sentiment and attention measures on stock market volatility. *International Journal of Forecasting*, 36(2), 334–357.
- [Avellaneda & Stoikov, 2008] Avellaneda, M., & Stoikov, S. (2008). High-frequency trading in a limit order book. *Quantitative Finance*, 8(3), 217–224.
- [Baldacci, Bergault, & Guéant, 2021] Baldacci, B., Bergault, P., & Guéant, O. (2021). Algorithmic market making for options. *Quantitative Finance*, 21(1), 85–97.
- [Baldacci, Derchu, & Manziuk, 2021] Baldacci, B., Derchu, J., & Manziuk, I. (2021). An approximate solution for options market-making in high dimension. *Risk*, June.
- [Bayer et al., 2016] Bayer, C., Friz, P., & Gatheral, J. (2016). Pricing under rough volatility. *Quantitative Finance*, 16(6), 887–904.
- [Bayer et al., 2019] Bayer, C., Horvath, B., Muguruza, A., Stemper, B., & Tomas, M. (2019). On deep calibration of (rough) stochastic volatility models. *arXiv preprint arXiv:1908.08806*.
- [Bemporad et al., 2018] Bemporad, A., Breschi, V., Piga, D., & Boyd, S. P. (2018). Fitting jump models. *Automatica*, 96, 11–21.
- [Bennedsen et al., 2021] Bennedsen, M., Lunde, A., & Pakkanen, M. (2021). Decoupling the short- and long-term behavior of stochastic volatility. *Journal of Financial Econometrics*. doi: 10.1093/jjfinec/nbaa049

- [Bergault et al., 2021] Bergault, P., Evangelista, D., Guéant, O., & Vieira, D. (2021). Closed-form approximations in multi-asset market making. *Applied Mathematical Finance*, 28(2), 101–142.
- [Bergault & Guéant, 2021] Bergault, P., & Guéant, O. (2021). Size matters for OTC market makers: general results and dimensionality reduction techniques. *Mathematical Finance*, 31(1), 279–322.
- [Bińkowski & Lehalle, 2022] Bińkowski, M., & Lehalle, C.-A. (2022). Endogenous dynamics of intraday liquidity. *The Journal of Portfolio Management*, 48(6), 145–169.
- [Blanc et al., 2014] Blanc, P., Chicheportiche, R., & Bouchaud, J.-P. (2014). The fine structure of volatility feedback II: Overnight and intra-day effects. *Physica A: Statistical Mechanics and its Applications*, 402, 58–75.
- [Blanc et al., 2017] Blanc, P., Donier, J., & Bouchaud, J.-P. (2017). Quadratic Hawkes processes for financial prices. *Quantitative Finance*, 17(2), 171–188.
- [Bolko et al., 2022] Bolko, A., Christensen, K., Pakkanen, M., & Veliyev, B. (2022). A GMM approach to estimate the roughness of stochastic volatility. *Journal of Econometrics*.
- [Cartea et al., 2017] Cartea, Á., Donnelly, R., & Jaimungal, S. (2017). Algorithmic trading with model uncertainty. *SIAM Journal on Financial Mathematics*, 8(1), 635–671.
- [Cartea et al., 2015] Cartea, Á., Jaimungal, S., & Penalva, J. (2015). *Algorithmic and high-frequency trading*. Cambridge University Press.
- [Cartea et al., 2014] Cartea, Á., Jaimungal, S., & Ricci, J. (2014). Buy low, sell high: A high frequency trading perspective. *SIAM Journal on Financial Mathematics*, 5(1), 415–444.
- [Cartea et al., 2018] Cartea, Á., Jaimungal, S., & Ricci, J. (2018). Algorithmic trading, stochastic control, and mutually exciting processes. *SIAM review*, 60(3), 673–703.
- [Chan, 2003] Chan, W. S. (2003). Stock price reaction to news and no-news: drift and reversal after headlines. *Journal of Financial Economics*, 70(2), 223–260.
- [Chen, 2021] Chen, Q. (2021). Stock movement prediction with financial news using contextualized embedding from bert. *arXiv preprint arXiv:2107.08721*.
- [Chen & Robert, 2022] Chen, Q., & Robert, C.-Y. (2022). Multivariate realized volatility forecasting with graph neural network. In *Proceedings of the third acm international conference on ai in finance* (pp. 156–164).
- [Chicheportiche & Bouchaud, 2014] Chicheportiche, R., & Bouchaud, J.-P. (2014). The fine-structure of volatility feedback I: Multi-scale self-reflexivity. *Physica A: Statistical Mechanics and its Applications*, 410, 174–195.
- [Christensen et al., 2022] Christensen, K., Siggaard, M., & Veliyev, B. (2022). A machine learning approach to volatility forecasting. *Journal of Financial Econometrics*.

- [Corsi, 2009] Corsi, F. (2009). A simple approximate long-memory model of realized volatility. *Journal of Financial Econometrics*, 7(2), 174–196.
- [Corsi & Renò, 2012] Corsi, F., & Renò, R. (2012). Discrete-time volatility forecasting with persistent leverage effect and the link with continuous-time volatility modeling. *Journal of Business & Economic Statistics*, 30(3), 368–380.
- [Cover et al., 1991] Cover, T. M., Thomas, J. A., et al. (1991). Entropy, relative entropy and mutual information. *Elements of information theory*, 2(1), 12–13.
- [Dandapani et al., 2021] Dandapani, A., Jusselin, P., & Rosenbaum, M. (2021). From quadratic Hawkes processes to super-Heston rough volatility models with zumbach effect. *Quantitative Finance*, 21(8), 1235–1247.
- [Ding et al., 2014] Ding, X., Zhang, Y., Liu, T., & Duan, J. (2014). Using structured events to predict stock price movement: An empirical investigation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (emnlp)* (pp. 1415–1425).
- [Ding et al., 2015] Ding, X., Zhang, Y., Liu, T., & Duan, J. (2015). Deep learning for event-driven stock prediction. In *Twenty-fourth international joint conference on artificial intelligence*.
- [El Euch et al., 2018] El Euch, O., Fukasawa, M., & Rosenbaum, M. (2018). The microstructural foundations of leverage effect and rough volatility. *Finance and Stochastics*, 22(2), 241–280.
- [El Euch et al., 2019] El Euch, O., Gatheral, J., & Rosenbaum, M. (2019). Roughening Heston. *Risk*, 84–89.
- [El Euch & Rosenbaum, 2018] El Euch, O., & Rosenbaum, M. (2018). Perfect hedging in rough Heston models. *The Annals of Applied Probability*, 28(6), 3813–3856.
- [El Euch & Rosenbaum, 2019] El Euch, O., & Rosenbaum, M. (2019). The characteristic function of rough Heston models. *Mathematical Finance*, 29(1), 3–38.
- [Filimonov & Sornette, 2012] Filimonov, V., & Sornette, D. (2012). Quantifying reflexivity in financial markets: Toward a prediction of flash crashes. *Physical Review E*, 85(5), 056108.
- [Filimonov & Sornette, 2015] Filimonov, V., & Sornette, D. (2015). Apparent criticality and calibration issues in the Hawkes self-excited point process model: application to high-frequency financial data. *Quantitative Finance*, 15(8), 1293–1314.
- [Fodra & Pham, 2015] Fodra, P., & Pham, H. (2015). High frequency trading and asymptotics for small risk aversion in a markov renewal model. *SIAM Journal on Financial Mathematics*, 6(1), 656–684.
- [Fouque et al., 2011] Fouque, J.-P., Papanicolaou, G., Sircar, R., & Sølna, K. (2011). *Multiscale stochastic volatility for equity, interest rate, and credit derivatives*. Cambridge University Press.

- [Fukasawa et al., 2022] Fukasawa, M., Takabatake, T., & Westphal, R. (2022). Consistent estimation for fractional stochastic volatility model under high-frequency asymptotics. *Mathematical Finance*, 32(4), 1086–1132.
- [Gatheral & Jacquier, 2014] Gatheral, J., & Jacquier, A. (2014). Arbitrage-free SVI volatility surfaces. *Quantitative Finance*, 14(1), 59–71.
- [Gatheral et al., 2018] Gatheral, J., Jaisson, T., & Rosenbaum, M. (2018). Volatility is rough. *Quantitative Finance*, 18(6), 933–949.
- [Gatheral et al., 2020] Gatheral, J., Jusselin, P., & Rosenbaum, M. (2020). The quadratic rough Heston model and the joint S&P 500/VIX smile calibration problem. *Risk*, May.
- [Giles & Glasserman, 2006] Giles, M., & Glasserman, P. (2006). Smoking Adjoints: fast Monte Carlo Greeks. *Risk*, 19(1), 88–92.
- [Goodfellow et al., 2016] Goodfellow, I., Bengio, Y., & Courville, A. (2016). *Deep learning*. MIT press.
- [Groß-Klußmann & Hautsch, 2011] Groß-Klußmann, A., & Hautsch, N. (2011). When machines read the news: Using automated text analytics to quantify high frequency news-implied market reactions. *Journal of Empirical Finance*, 18(2), 321–340.
- [Guéant, 2016] Guéant, O. (2016). *The financial mathematics of market liquidity: From optimal execution to market making* (Vol. 33). CRC Press.
- [Guéant, 2017] Guéant, O. (2017). Optimal market making. *Applied Mathematical Finance*, 24(2), 112–154.
- [Guéant et al., 2013] Guéant, O., Lehalle, C.-A., & Fernandez-Tapia, J. (2013). Dealing with the inventory risk: a solution to the market making problem. *Mathematics and financial economics*, 7(4), 477–507.
- [Guéant & Manziuk, 2019] Guéant, O., & Manziuk, I. (2019). Deep reinforcement learning for market making in corporate bonds: beating the curse of dimensionality. *Applied Mathematical Finance*, 26(5), 387–452.
- [Guilbaud & Pham, 2013] Guilbaud, F., & Pham, H. (2013). Optimal high-frequency trading with limit and market orders. *Quantitative Finance*, 13(1), 79–94.
- [Guilbaud & Pham, 2015] Guilbaud, F., & Pham, H. (2015). Optimal high-frequency trading in a pro rata microstructure with predictive information. *Mathematical Finance*, 25(3), 545–575.
- [Guyon, 2020] Guyon, J. (2020). The joint S&P 500/VIX smile calibration puzzle solved. *Risk*, April.
- [Guyon & Lekeufack, 2022] Guyon, J., & Lekeufack, J. (2022). Volatility is (mostly) path-dependent. *Available at SSRN 4174589*.

- [Hansen et al., 2011] Hansen, P. R., Lunde, A., & Nason, J. M. (2011). The model confidence set. *Econometrica*, 79(2), 453–497.
- [Hardiman et al., 2013] Hardiman, S. J., Bercot, N., & Bouchaud, J.-P. (2013). Critical reflexivity in financial markets: a Hawkes process analysis. *The European Physical Journal B*, 86(10), 1–9.
- [Hendrycks & Gimpel, 2016] Hendrycks, D., & Gimpel, K. (2016). Gaussian error linear units (gelus). *arXiv preprint arXiv:1606.08415*.
- [Hernandez, 2016] Hernandez, A. (2016). Model calibration with neural networks. *Available at SSRN 2812140*.
- [Ho & Stoll, 1981] Ho, T., & Stoll, H. R. (1981). Optimal dealer pricing under transactions and return uncertainty. *Journal of Financial economics*, 9(1), 47–73.
- [Hochreiter & Schmidhuber, 1997] Hochreiter, S., & Schmidhuber, J. (1997). Long short-term memory. *Neural computation*, 9(8), 1735–1780.
- [Hornik et al., 1990] Hornik, K., Stinchcombe, M., & White, H. (1990). Universal approximation of an unknown mapping and its derivatives using multilayer feedforward networks. *Neural networks*, 3(5), 551–560.
- [Horvath et al., 2019] Horvath, B., Muguruza, A., & Tomas, M. (2019). Deep learning volatility. *Available at SSRN 3322085*.
- [Horvath et al., 2021] Horvath, B., Muguruza, A., & Tomas, M. (2021). Deep learning volatility: a deep neural network perspective on pricing and calibration in (rough) volatility models. *Quantitative Finance*, 21(1), 11–27.
- [Hu et al., 2018] Hu, Z., Liu, W., Bian, J., Liu, X., & Liu, T.-Y. (2018). Listening to chaotic whispers: A deep learning framework for news-oriented stock trend prediction. In *Proceedings of the eleventh acm international conference on web search and data mining* (pp. 261–269).
- [Huang et al., 2015] Huang, W., Lehalle, C.-A., & Rosenbaum, M. (2015). Simulating and analyzing order book data: The queue-reactive model. *Journal of the American Statistical Association*, 110(509), 107–122.
- [Huge & Savine, 2020] Huge, B. N., & Savine, A. (2020). Differential Machine Learning. *Risk*, October.
- [Jaisson & Rosenbaum, 2016] Jaisson, T., & Rosenbaum, M. (2016). Rough fractional diffusions as scaling limits of nearly unstable heavy tailed Hawkes processes. *Annals of Applied Probability*, 26(5), 2860–2882.
- [Jiang et al., 2021] Jiang, H., Li, S. Z., & Wang, H. (2021). Pervasive underreaction: Evidence from high-frequency data. *Journal of Financial Economics*, 141(2), 573–599.
- [Joulin et al., 2008] Joulin, A., Lefevre, A., Grunberg, D., & Bouchaud, J.-P. (2008). Stock price jumps: news and volume play a minor role. *Wilmott Magazine*(46).

- [Jusselin & Rosenbaum, 2020] Jusselin, P., & Rosenbaum, M. (2020). No-arbitrage implies power-law market impact and rough volatility. *Mathematical Finance*, 30(4), 1309–1336.
- [Ke et al., 2019] Ke, Z. T., Kelly, B. T., & Xiu, D. (2019). *Predicting returns with text data* (Tech. Rep.). National Bureau of Economic Research.
- [Kenton & Toutanova, 2019] Kenton, J. D. M.-W. C., & Toutanova, L. K. (2019). Bert: Pre-training of deep bidirectional transformers for language understanding. In *Proceedings of naacl-hlt* (pp. 4171–4186).
- [Kim & Won, 2018] Kim, H. Y., & Won, C. H. (2018). Forecasting the volatility of stock price index: A hybrid model integrating LSTM with multiple GARCH-type models. *Expert Systems with Applications*, 103, 25–37.
- [Lehalle & Laruelle, 2018] Lehalle, C.-A., & Laruelle, S. (2018). *Market microstructure in practice*. World Scientific.
- [R. Liu et al., 2022] Liu, R., Jiang, Y., & Lin, J. (2022). Forecasting the volatility of specific risk for stocks with LSTM. *Procedia Computer Science*, 202, 111–114.
- [Y. Liu, 2019] Liu, Y. (2019). Novel volatility forecasting using deep learning–long short term memory recurrent neural networks. *Expert Systems with Applications*, 132, 99–109.
- [Livieri et al., 2018] Livieri, G., Mouti, S., Pallavicini, A., & Rosenbaum, M. (2018). Rough volatility: evidence from option prices. *IJSE transactions*, 50(9), 767–776.
- [Loughran & McDonald, 2011] Loughran, T., & McDonald, B. (2011). When is a liability not a liability? textual analysis, dictionaries, and 10-ks. *The Journal of finance*, 66(1), 35–65.
- [Lu & Abergel, 2018] Lu, X., & Abergel, F. (2018). Order-book modeling and market making strategies. *Market Microstructure and Liquidity*, 4(01n02), 1950003.
- [Marcaccioli et al., 2022] Marcaccioli, R., Bouchaud, J.-P., & Benzaquen, M. (2022). Exogenous and endogenous price jumps belong to different dynamical classes. *Journal of Statistical Mechanics: Theory and Experiment*, 2022(2), 023403.
- [McCallum et al., 1998] McCallum, A., Nigam, K., et al. (1998). A comparison of event models for naive bayes text classification. In *Aaai-98 workshop on learning for text categorization* (Vol. 752, pp. 41–48).
- [Nystrup et al., 2021] Nystrup, P., Kolm, P. N., & Lindström, E. (2021). Feature selection in jump models. *Expert Systems with Applications*, 184, 115558.
- [Øksendal & Sulem, 2005] Øksendal, B., & Sulem, A. (2005). *Stochastic control of jump diffusions*. Springer.
- [Oliveira et al., 2016] Oliveira, N., Cortez, P., & Areal, N. (2016). Stock market sentiment lexicon acquisition using microblogging data and statistical measures. *Decision Support Systems*, 85, 62–73.

- [Patton, 2011] Patton, A. J. (2011). Volatility forecast comparison using imperfect volatility proxies. *Journal of Econometrics*, 160(1), 246–256.
- [Patton & Sheppard, 2015] Patton, A. J., & Sheppard, K. (2015). Good volatility, bad volatility: Signed jumps and the persistence of volatility. *Review of Economics and Statistics*, 97(3), 683–697.
- [Pedersen, 2019] Pedersen, L. H. (2019). *Efficiently inefficient: how smart money invests and market prices are determined*. Princeton University Press.
- [Petrozziello et al., 2022] Petrozziello, A., Troiano, L., Serra, A., Jordanov, I., Storti, G., Tagliaferri, R., & La Rocca, M. (2022). Deep learning for volatility forecasting in asset management. *Soft Computing*, 26(17), 8553–8574.
- [Rahimikia & Poon, 2020] Rahimikia, E., & Poon, S.-H. (2020). Machine learning for realised volatility forecasting. *Available at SSRN*, 3707796.
- [Rambaldi et al., 2019] Rambaldi, M., Bacry, E., & Muzy, J.-F. (2019). Disentangling and quantifying market participant volatility contributions. *Quantitative Finance*, 19(10), 1613–1625.
- [Renault, 2017] Renault, T. (2017). Intraday online investor sentiment and return patterns in the us stock market. *Journal of Banking & Finance*, 84, 25–40.
- [Robert & Rosenbaum, 2011] Robert, C. Y., & Rosenbaum, M. (2011). A new approach for the dynamics of ultra-high-frequency data: The model with uncertainty zones. *Journal of Financial Econometrics*, 9(2), 344–366.
- [Robert & Rosenbaum, 2012] Robert, C. Y., & Rosenbaum, M. (2012). Volatility and covariation estimation when microstructure noise and trading times are endogenous. *Mathematical Finance: An International Journal of Mathematics, Statistics and Financial Economics*, 22(1), 133–164.
- [Rosenbaum & Zhang, 2022a] Rosenbaum, M., & Zhang, J. (2022a). Deep calibration of the quadratic rough Heston model. *Risk*, September.
- [Rosenbaum & Zhang, 2022b] Rosenbaum, M., & Zhang, J. (2022b). On the universality of the volatility formation process: when machine learning and rough volatility agree. *arXiv preprint arXiv:2206.14114*.
- [Shi et al., 2023] Shi, S., Yu, J., & Zhang, C. (2023). Fractional gaussian noise: Spectral density and estimation methods.
- [Sirignano & Cont, 2019] Sirignano, J., & Cont, R. (2019). Universal features of price formation in financial markets: perspectives from deep learning. *Quantitative Finance*, 19(9), 1449–1459.
- [Sirignano & Spiliopoulos, 2018] Sirignano, J., & Spiliopoulos, K. (2018). DGM: A deep learning algorithm for solving partial differential equations. *Journal of computational physics*, 375, 1339–1364.

- [[Taqqu et al., 1995](#)] Taqqu, M. S., Teverovsky, V., & Willinger, W. (1995). Estimators for long-range dependence: an empirical study. *Fractals*, *3*(04), 785–798.
- [[Tetlock, 2007](#)] Tetlock, P. C. (2007). Giving content to investor sentiment: The role of media in the stock market. *The Journal of finance*, *62*(3), 1139–1168.
- [[Wang et al., 2023](#)] Wang, X., Xiao, W., & Yu, J. (2023). Modeling and forecasting realized volatility with the fractional ornstein-uhlenbeck process. *Journal of Econometrics*, *232*, 389-415.
- [[Wang et al., 2022](#)] Wang, X., Zhang, C., & Jun, Y. (2022). On the optimal forecast with the fractional brownian motion.
- [[Wu et al., 2022](#)] Wu, P., Muzy, J.-F., & Bacry, E. (2022). From rough to multifractal volatility: The log s-fbm model. *Physica A: Statistical Mechanics and its Applications*, *604*, 127919.
- [[Wyart et al., 2008](#)] Wyart, M., Bouchaud, J.-P., Kockelkoren, J., Potters, M., & Vettorazzo, M. (2008). Relation between bid-ask spread, impact and volatility in order-driven markets. *Quantitative finance*, *8*(1), 41–57.
- [[Zumbach, 2010](#)] Zumbach, G. (2010). Volatility conditional on price trends. *Quantitative Finance*, *10*(4), 431–442.

Titre : Quelques applications de l'apprentissage statistique en finance quantitative : calibration de modèles, mécanisme de formation de la volatilité et filtrage des nouvelles

Mots clés : quadratic rough Heston, volatilité rugueuse, market making, universalité, filtrage de nouvelles

Résumé : Nous commençons cette thèse au Chapitre I en présentant nos méthodologies basées sur l'apprentissage profond pour une calibration efficace du modèle *quadratic rough Heston*, qui est non-markovien et non-semimartingale. Une approximation multifactorielle du modèle est d'abord proposée. Deux réseaux neuronaux profonds sont ensuite entraînés sur des données simulées afin d'apprendre les fonctions de tarification des options SPX et VIX respectivement. Les résultats de la calibration jointe peuvent être obtenus instantanément. Des tests empiriques montrent que le modèle calibré peut très bien reproduire les surfaces de volatilité implicite du SPX et du VIX. Grâce à l'application de la différenciation adjointe automatique, ces deux réseaux neuronaux peuvent être utilisés pour calculer efficacement les quantités de couverture.

Dans le cadre de quadratic rough Heston, nous formulons ensuite un problème de *market making* optimale sur un panier composé de plusieurs dérivés de SPX, tels que des contrats à terme SPX et VIX, des options SPX et VIX. Le *market maker* maximise son profit en capturant les écarts entre l'achat et la vente et pénalise le risque d'inventaire du portefeuille, qui peut être expliqué principalement par la variation de SPX. Nous abordons la haute dimensionnalité du problème avec plusieurs approximations pertinentes. Des solutions asymptotiques à forme fermée peuvent ensuite être déduites. Motivés par le fait largement établi que la volatilité est *rough* dans de nombreuses classes d'actifs et par les résultats pro-

metteurs grâce à l'incorporation de l'effet *Zumbach*, nous nous intéressons à l'universalité des mécanismes de formation de la volatilité endogène au Chapitre III. Nous commençons par une approche quasi-nonparamétrique. Un réseau LSTM entraîné sur un ensemble de données regroupées couvrant des centaines d'actions liquides montre une performance de prévision de la volatilité supérieure aux autres dispositifs spécifiques à un actif ou à un secteur, ce qui suggère l'existence de l'universalité en question. Nous proposons ensuite une méthode de prévision paramétrique parcimonieuse combinant les principales caractéristiques des modèles *rough fractional stochastic volatility* et quadratic rough Heston. Avec des paramètres fixes, cette approche présente le même niveau de performance que le LSTM universel, confirmant à nouveau l'universalité d'un point de vue paramétrique.

Enfin, nous nous concentrons sur le lien entre les publications des nouvelles et les changements des conditions de liquidité intrajournalière dans le marché financier. En particulier, en détectant les sauts significatifs de la volatilité et des volumes, nous proposons une approche systématique pour distinguer les nouvelles impactantes des neutres présentés massivement. Un prédicteur du sentiment des nouvelles apprendit sur les échantillons sélectionnés s'avère beaucoup plus efficace que celui calibré sur l'ensemble des données brutes en termes de prédiction des mouvements de prix à court terme pour les actifs associés.

Title : Some applications of machine learning in quantitative finance: model calibration, volatility formation mechanism and news screening

Keywords : quadratic rough Heston, rough volatility, market making, universality, news screening

Abstract : We begin this thesis in Chapter I by introducing our deep learning-based methodologies for efficient calibration of the quadratic rough Heston model, which is non-Markovian and non-semimartingale. A multi-factor approximation of the model is first proposed. Two deep neural networks are then trained on simulated data to learn the resulting pricing functions of SPX and VIX options respectively. Given SPX/VIX smiles, joint calibration results can be instantaneously obtained. Empirical tests show that the calibrated model can reproduce both SPX and VIX implied volatility surfaces very well. Through the application of automatic adjoint differentiation, the neural networks can be used to compute efficiently related hedging quantities.

Under quadratic rough Heston, we then formulate an optimal market making problem on a basket composed of multiple derivatives of SPX, such as SPX and VIX futures, SPX and VIX options. The market maker maximizes its profit from spread capturing and penalizes the portfolio's inventory risk, which can be mostly explained by the variation of SPX. We tackle the high dimensionality of the problem with several relevant approximations. Closed-form asymptotic solutions can then be obtained.

Motivated by the widely established fact that volatility is rough across many asset classes and the promising results of the qua-

dratic rough Heston model thanks to the incorporation of the Zumbach effect, we are interested in the universality of the endogenous volatility formation mechanism in Chapter III. We take first a quasi-nonparametric approach. An LSTM network trained on a pooled dataset covering hundreds of liquid stocks shows superior volatility forecasting performance than other asset-specific or sector-specific devices, suggesting the existence of the universality in question. We propose then a parsimonious parametric forecasting method combining the main features of the rough fractional stochastic volatility and quadratic rough Heston models. With fixed parameters, this approach presents the same level of performance as the universal LSTM, confirming again the universality from a parametric perspective.

Finally, we focus on the link between informative news releases and switches of intraday liquidity conditions. Particularly, by detecting significant jumps of volatility and trading volumes inside daily trading sessions, we propose a systematic approach to distinguish impactful news releases from the massively presented neutral ones. A news sentiment predictor fitted on the identified impactful data is shown to be much more effective than the one calibrated on the raw dataset in terms of short-term price movement prediction for the associated assets.