

Arbitrage-Free Neural-SDE Market Models of Traded Options



InFoMM
Industrially Focused
Mathematical Modelling

Sheng (Victor) Wang
University College
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy
Trinity 2022

To my parents, wife and this fabulous planet.

Acknowledgements

The greatest difference that the DPhil journey has made to me is that I will always enjoy solving challenging problems, and will never get bored with thinking differently and independently. The two key figures who have enlightened me are undoubtedly my academic supervisors Samuel N. Cohen and Christoph Reisinger. I am grateful for their illuminating research ideas and supports in all aspects.

Many professors in the Mathematical Institute have kindly offered advice for my research in various occasions, such as our internal seminars, random corridor chats and my transfer/confirmation exams. In particular, I would like to thank Prof. Rama Cont, Álvaro Cartea, Jan Obłój and Dr. Katia Babbar for carefully assessing my research progress. They have provided invaluable feedback, of which some is critical for shaping my research directions.

My research has received generous supports from my previous employer, CME Group, in terms of finance, data, computational resources and unique business insights. Particularly, Florian Huchedé and Dr. David Nicolay had spared no effort to make the collaboration between Oxford and CME not only possible but also successful. In addition, I appreciate the encouraging conversations with Udesch Jha and Sunil Cutinho (CIO), who acknowledge the potentials of my research outcomes for much more efficient option clearing solutions.

In the journey of approaching the doctoral degree at Oxford, I have never felt alone, thanks to the company of my excellent cohorts at EPSRC InFoMM CDT, and especially my officemates Lingyi Yang, Ambrose Yim and Zhen Shao. Importantly, we all benefit from the rich and highly structured curriculum of InFoMM, which is impossible without the passion and efforts of our directors Prof. Chris Beward and Colin Please.

Finally, I would like to thank my two housemates at Oxford, Jamie Saw (at Univ and Rowland Hill) and James H. Morrill (at Osney Mew), who are not necessarily helpful for completing this thesis by offering me dinners, beers, cocktails and film nights, but definitely help me to get through a few hard moments.

Abstract

Modelling joint dynamics of liquid vanilla options is crucial for arbitrage-free pricing of illiquid derivatives and managing risks of option trade books. This thesis develops a nonparametric model for the European options book respecting underlying financial constraints while being practically implementable. We derive a state space for prices which are free from static (or model-independent) arbitrage and study the inference problem where a model is learnt from discrete time series data of stock and option prices. We use neural networks as function approximators for the drift and diffusion of the modelled SDE (stochastic differential equation) system, and impose constraints on the neural nets such that no-arbitrage conditions are preserved. In particular, we give methods to calibrate *neural-SDE* models which are guaranteed to satisfy a set of linear inequalities.

We validate our approach with numerical experiments using both (i) synthetic data generated from a Heston stochastic local volatility model and (ii) real world data. In both cases, we assess how well the model captures option price dynamics, and therefore its capacity as a realistic option market simulator. However, the presence of arbitrage in the real world option price data causes difficulties in the calibration of our models, making pre-processing of the data to eliminate arbitrage necessary. We formulate the arbitrage repair as a linear programming (LP) problem, where the no-arbitrage relations are constraints, and the objective is to minimise prices' changes within their bid and ask price bounds. Through empirical studies, we show that the proposed arbitrage repair method gives sparse perturbations on data, and is fast when applied to real world large-scale problems due to the LP formulation. In addition, we show that removing arbitrage from price data by our repair method can improve model calibration with enhanced robustness and reduced calibration error.

Finally, we explore the capacity of the model in two important applications for managing risks of option books, namely the Value-at-Risk (VaR) calculation and hedging. In terms of calculating VaR, we show that our models are more computationally efficient and accurate for evaluating risks of option portfolios, with better coverage performance and less procyclicality than standard filtered historical simulation approaches. In the hedging application, we derive sensitivity and minimum variance based hedging ratios under neural-SDE market models, and compare the performance of hedging various option portfolios with standard Greeks-based hedging under the Black–Scholes model, demonstrating favourable performance of our approach.

Contents

1	Introduction	1
1.1	Option arbitrage theory	2
1.2	Martingale and market model approaches	3
1.3	Factor-based arbitrage-free market models	5
1.4	Model inference and constrained neural-SDE	7
1.5	Key contributions and structure of this thesis	8
2	Static arbitrage constraints	10
2.1	Assumptions	11
2.2	Static arbitrage	12
2.3	Shape constraints of the call price surface	13
2.4	Constraints reduction	15
	Appendices	15
2.A	Localisation of static arbitrage constraints	15
2.B	Proof of Proposition 2.2	17
3	Application: detecting and repairing arbitrage	21
3.1	Introduction	21
3.2	Arbitrage repair	24
3.3	Design of the objective without liquidity consideration	25
3.4	Inclusion of bid and ask prices	26
3.5	Empirical studies	28
4	Arbitrage-free neural-SDE market model	42
4.1	Introduction	42
4.2	Assumptions	44

4.3	Liquid options	44
4.4	Factor representation of call option prices	46
4.5	Absence of arbitrage	49
4.6	Constrained state space for static arbitrage	50
4.7	Model inference	51
	Appendices	54
4.A	HJM drift restriction for dynamic arbitrage	54
5	Estimating neural-SDE constrained by a polytope	58
5.1	Nonattainability of a diffusion process	59
5.2	The likelihood function approximation	60
5.3	Deep learning as a function approximation algorithm	61
	Appendices	66
5.A	Adaptive diffusion shrinking scales	66
5.B	The ρ^* -interior points for drift correction	67
6	Calibrating to synthetic data: Heston SLV model	68
6.1	Input synthetic data	68
6.2	Implementation details of the model estimation strategy	69
6.3	Factor construction and no-arbitrage boundaries	71
6.4	Neural network training results	74
6.5	In-sample test	75
6.6	Out-of-sample simulation test	78
6.7	Concluding remarks on the numerical example	82
	Appendices	82
6.A	Implied volatility simulation	83
6.B	Necessity of drift and diffusion operators	83
6.C	Neural network sensitivity analysis	85
6.D	Higher-dimensional models	87
7	Calibrating to real data: EURO STOXX 50 and DAX	91
7.1	Option price data description and pre-processing	91
7.2	Revisiting the neural-SDE market models	93
7.3	Factor decoding	94
7.4	Factor dynamics estimation	97
7.5	Underlying index dynamics estimation	99
7.6	Analysis of historical residuals	100
7.7	Out-of-sample simulation of factors and implied volatility surface	101

7.8	VIX-like volatility index simulation	102
	Appendices	103
7.A	Data pre-processing	104
7.B	Computing z_t for estimating dynamic arbitrage	106
7.C	Alternative model configurations	106
8	Application: estimating risks of option books	110
8.1	Introduction	110
8.2	Measuring risk	113
8.3	Value-at-Risk backtesting analysis	115
8.4	Conclusions and further extensions	121
	Appendices	121
8.A	Heston model calibration results and historical scenarios	121
8.B	Tested option portfolios for risk simulation	122
8.C	More VaR backtesting results	124
8.D	In-sample VaR backtesting	128
9	Application: hedging risks of option books	130
9.1	Introduction	130
9.2	Option hedging problem	133
9.3	Hedging with neural-SDE market models	135
9.4	Empirical results	141
9.5	Conclusion	155
10	Conclusions and extensions	157
	Bibliography	159

Introduction

“予尝求古仁人之心，或异二者之为，何哉？不以物喜，不以己悲，居庙堂之高则忧其民，处江湖之远则忧其君。”

— 范仲淹《岳阳楼记》

“I have tried to study the minds of people of lofty ideals in ancient times. Perhaps they were different from the people I mentioned above. Why is this? The reason is that they were not thrown into ecstasies over their success, nor felt depressed over their failures. When they were in high positions at court, they were concerned about people. When they were in remote places, they were concerned about their emperor.”

— Zongyan Fan, *The Yueyang Tower* (translated by Jingguo Luo)

Consider a financial market where the following assets are liquidly traded: a stock¹ S and a collection of European call options $C(T, K)$ on S with various expiries $T \in \mathcal{T}^e$ and strikes $K \in \mathcal{K}$. Pricing and risk management generally require a statistical model for these assets' prices, which are known to be related to each other in complex ways. For example, the volatility of the stock price through time is typically related to the value of call options, in a monotone but nonlinear fashion. These stylised relationships suggest that there is significant statistical information captured in the interrelated prices of the stock and options, however modelling these jointly is challenging. Furthermore, there are various constraints on prices which must hold in the absence of arbitrage, and these should be reflected in any statistical model.

We aim to construct a class of models for the stock and options that should

- (i) permit no arbitrage;
- (ii) allow exact cross-sectional calibration; and

¹Although we refer to the asset as “stock” throughout the thesis, our methods are equally applicable to other asset classes such as currencies or commodities.

(iii) reflect stylised facts observed from market price dynamics.

Importantly, it should be practically convenient to estimate these models, given that observations are discrete time series of prices for a large but finite collection of options. We aim to exploit the recent successes in the use of neural networks as function approximators in order to give a flexible class of models. However, the no-arbitrage constraints on option prices imply that we need to fit neural networks where the resulting behaviour (when used in an *stochastic differential equation*, or SDE in short) will satisfy a family of linear (in)equalities. In this thesis we will develop methods to solve this challenge, and demonstrate their effectiveness when building a financial model in real world applications such as calculating Value-at-Risk and deriving hedging ratios.

The remainder of this chapter reviews the relevant literature of this field of research. This inspires the broad research direction and refines many modelling choices of the research. Meanwhile, it also provides definitions of the concepts and mathematical structures under consideration.

1.1 Option arbitrage theory

Arbitrage refers to a costless trading strategy that has a positive probability of earning risk-free profit. A *static* arbitrage is an arbitrage exploitable by fixed positions in options and the underlying stock at initial time, while the position of the underlying stock can be modified at a finite number of trading times in the future. Any other arbitrage is called *dynamic* arbitrage. As an example of a static arbitrage, for call options, it must hold that $C(T, K_1) \geq C(T, K_2)$ for $K_1 < K_2$, otherwise by going long one (T, K_1) option and short one (T, K_2) option, we make an immediate profit of $C(T, K_2) - C(T, K_1)$ with non-negative terminal payoff². An example of dynamic arbitrage is a continuously delta-hedged short position on an over-priced option in the perfect Black–Scholes world [22].

Now we associate no-arbitrage with mathematical models. A *model* \mathbb{M} is a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \in \mathcal{T}}, \mathbb{P})$, that carries an adapted price process $\{(S_t, \mathbf{C}_t)\}_{t \in \mathcal{T}}$, where \mathbf{C}_t gives the prices of all options at time t . Here \mathcal{T} denotes the set of times at which the asset can be traded so that $0 \in \mathcal{T}$ and the set of option expiries $\mathcal{T}^e \subset \mathcal{T}$, and $\mathcal{F}_0 = \{\Omega, \emptyset\}$ augmented with all null sets of $\mathcal{F}_{\sup \mathcal{T}^e}$. The First Fundamental Theorem of Asset Pricing (FFTAP) establishes an equivalence relation between no-arbitrage (static and dynamic) and the existence of an equivalent martingale measure (EMM). After the landmark work of Harrison and Kreps [81], various versions of the FFTAP and extensions

²A call option, denoted its time- t value as $C_t(T, K)$, gives its buyer the right but not the obligation to buy the underlying asset at K on the expiration date T . Therefore, its payoff at expiry is $C_T(T, K) = (S_T - K)^+ := \max(S_T - K, 0)$.

of the no-arbitrage concept have been developed (e.g. no free lunch by Kreps [107], no free lunch with vanishing risk by Delbaen and Schachermayer [57]). In this thesis, we work with a simplified version of FFTAP as follows by implicitly using a strong definition of arbitrage. Given a model \mathbb{M} , we say that there is no arbitrage if and only if $\exists \mathbb{Q} \sim \mathbb{P}$, such that, for all $t < s \leq T$ where $t, s \in \mathcal{T}$,

$$\forall \text{ tradable } (T, K), D_0(t)C_t(T, K) = D_0(s)\mathbb{E}^{\mathbb{Q}}[C_s(T, K)|\mathcal{F}_t], \quad (1.1)$$

where $D_t(T)$ is the time t price of the zero-coupon bond expiring at time T .

No static arbitrage corresponds to a much smaller set of conditions than (1.1), since the path dynamics governed by \mathbb{Q} no longer matter. As discussed by Carr, Géman, Madan and Yor [38, 39] and Davis and Hobson [56], static arbitrage is present if and only if no \mathbb{Q} exists such that $C_0(T, K) = D_0(T)\mathbb{E}^{\mathbb{Q}}[C_T(T, K)|\mathcal{F}_0]$. Therefore, static arbitrage constraints are consequences of relations between terminal payoffs, projected to the present time. In contrast, dynamic arbitrage relies on dynamics and path properties of the tradable assets. Nevertheless, static arbitrage constraints establish the prerequisites that the price data have to satisfy at time zero for admitting a dynamically arbitrage-free model.

1.2 Martingale and market model approaches

Many models in the literature are derived using the *martingale approach*. In this approach, one specifies the dynamics of the underlying S (usually in the form of an SDE) under some pricing measure \mathbb{Q} , and uses no-arbitrage arguments (1.1) and Itô calculus to derive option prices, written as discounted conditional expectation of options' payoffs under \mathbb{Q} . For example, $C_t(T, K) = D_t(T) \cdot \mathbb{E}^{\mathbb{Q}}[(S_T - K)^+|\mathcal{F}_t]$. Prestigious examples of martingale models are the Black–Scholes (BS) [22] model, the local volatility (LV) model of Dupire [62] and Derman and Kani [58], stochastic volatility (SV) models such as Heston [85] and SABR [79, 80], and stochastic local volatility (SLV) models [98]. A data-driven approach to learn an overparametrised martingale model with SDE coefficients expressed as neural nets is taken by Gierjatowicz, Sabate-Vidales, Šiška, Szpruch and Žurič [74].

The martingale approach guarantees no-arbitrage by the FFTAP [81]. Therefore, the martingale approach eliminates arbitrage by construction, and the LV model and SV models can calibrate exactly³ for *static* price surfaces. However, some challenges remain:

- (i) martingale models do not give explicit expressions for option prices, often requiring heavy, model-specific numerical methodology to calibrate these models to market data;

³In general, exact calibration of stochastic volatility models is not possible as there is a limited number of model parameters to parametrise many more data points on the price surface.

- (ii) the values of calibrated model parameters are observed to change over time, even though they are assumed to be constant by the posited model;
- (iii) martingale models are naturally posed under the (risk-neutral) measure \mathbb{Q} , and additional steps are often needed to use them to model dynamics under the historical measure \mathbb{P} , as is needed for risk management.

An alternative family of models admitting exact calibration are *market models*, where one specifies the joint dynamics of all liquid tradables simultaneously, and imposes additional conditions to exclude arbitrage, as option prices are no longer automatically discounted conditional expectations, as specified by (1.1). This idea originates from the framework of Heath–Jarrow–Morton (HJM) [82] for interest rate modelling, which establishes no-arbitrage through drift restrictions. The case for options is more complicated, as the contract specifications at maturity enforce convoluted relationships between option prices and the underlying price; see, for example, Carr and Madan [39], Davis and Hobson [56]. In other words, the state space of these processes is heavily constrained.

Most work on arbitrage-free market models for option prices has sought a convenient parametrisation, or *codebook* in the language of Carmona [33], such that the codebook processes have a simple state space and yet capture all the static arbitrage constraints. Specifying the dynamics of the codebook then leads to a tractable arbitrage-free dynamic model for options.

A proper choice of codebook processes improves the descriptive capability of the model in terms of the observed strikes and expiries (\mathcal{K} and \mathcal{T}^e). Lyons [113], Schönbucher [96] and Babbar [8] propose a market model in terms of the BS *implied volatility* σ_{imp} for a single option, i.e. the case $\mathcal{K} = \{K\}$ and $\mathcal{T}^e = \{T\}$. The model is free from static arbitrage if $S > 0$ and $\sigma_{\text{imp}} \geq 0$. However, as discussed by Schweizer and Wissel [135], market models of implied volatilities cannot be easily extended to the general case with more than one option, because the absence of static arbitrage between different options enforces awkward constraints between the corresponding implied volatilities (see, for example, Lemma 2.2 in [72]). For the term structure case, $\mathcal{K} = \{K\}$ and $\mathcal{T}^e = [0, +\infty)$, Schönbucher [96] and Schweizer and Wissel [136] use the *forward implied volatilities*, defined as $\sigma_{\text{fw}}^2(T) := \partial((T-t)\sigma_{\text{imp}}^2(T))/\partial T$, as the codebook. Their model is statically arbitrage-free if $S > 0$ and $\sigma_{\text{fw}}^2(T) \geq 0$ for all $T \in \mathcal{T}^e$. Jacod and Protter [97] directly use *call option prices* as the codebook and work in a more general setting of jump processes. Static arbitrage is ruled out by imposing a non-decreasing term structure on the call option prices. For the multi-strike case $\mathcal{K} = [0, +\infty)$ and $\mathcal{T}^e = \{T\}$, Schweizer and Wissel [135] introduce a new parametrisation of option prices called *local implied volatilities* $X(K)$

and *price level* Y , see Definition 4.1 and 4.4 of [135]. Their model does not admit static arbitrage if $X(K) > 0$ for all $K \in \mathcal{K}$.

The surface case $\mathcal{K} = [0, +\infty)$ and $\mathcal{T}^e = [0, +\infty)$ has been considered by Derman and Kani [59], Carmona and Nadtochiy [34, 35] and Kallsen and Krühner [100]. Carmona and Nadtochiy use Dupire’s local volatilities $\sigma_{\text{loc}}(T, K)$ as the codebook processes and rigorously analyse the dynamic arbitrage conditions derived by Derman and Kani. Later, these authors independently built market models relying on time-inhomogeneous Lévy processes, which allow jumps and may hence be particularly suitable for short-term options. Both parametrisations make the static arbitrage constraints naturally hold.

The non-trivial construction of codebook processes, i.e. a PDE for local volatility [34] and Fourier transforms for Lévy process [35, 100], add complexities for model calibration. By considering a continuous spectrum of strikes and expiries, one is usually forced to consider an infinite-dimensional problem, where it is difficult to prove the existence of models and computationally infeasible to implement them exactly. The only work, to our knowledge, that considers market models for a finite family of strikes and expiries is Wissel [147]. Wissel combines ideas from the LV model of Dupire with the market model of Schweizer and Wissel [135] and uses *local implied volatilities* and *price level* as a codebook for parametrising statically arbitrage-free option prices on a *rectangular* lattice⁴. The rectangular lattice setup is necessary for the codebook construction but is a rather restrictive assumption in practice. In addition, there is not a straightforward model calibration method for general cases.

1.3 Factor-based arbitrage-free market models

To allow practical calibration, we will focus on constructing *market models* for *finitely* many options that are liquidly traded; after all, price data are only readily available for liquid options.

Typically, the range of strikes actively quoted on the market is much smaller for short-expiry options than for long-expiry options, as seen in Figure 1.1⁵. This causes problems for methods based on a rectangular lattice. One might augment market data using arbitrage-free interpolation techniques (see e.g. Kahalé [99] and Fengler [67]), but these are prone to extrapolate unrealistic prices for deep ITM and OTM options. Our models allow an arbitrarily shaped lattice of options, which can be specified according to the stylised liquidity profile of the options market.

⁴The lattice of strikes and expiries is called rectangular if prices are specified for a finite collection of discrete strikes and, for each strike, for a common finite set of expiries.

⁵Unsurprisingly, the ranges of quoted Black–Scholes deltas appear quite similar across expiries, which is consistent with the OTC FX convention of quoting option price by delta (see Reiswich and Wystup




Figure 1.1: The quoted strikes/log-moneynesses/Black-Scholes deltas (horizontal axis) for all expiries (vertical axis) of CME weekly- and monthly-listed EURUSD European call options as of 31st May, 2018. Each dot represents a market quote.

Unlike the neat static arbitrage constraints for options on a rectangular lattice, e.g. see Carr and Madan [39], the flexibility of modelling options on an arbitrarily-shaped lattice induces challenges for deriving static arbitrage constraints. Based on the work of Davis and Hobson [56] and Cousot [52], we will produce an efficient construction of static arbitrage constraints for large-scale practical problems.

As seen in Carr and Madan [39] and Davis and Hobson [56], the static arbitrage constraints are *linear* with respect to option prices, and enforce simple geometric shape constraints on the surface $(T, K) \mapsto C(T, K)$, including positivity, monotonicity and convexity. This effectively restricts the evolution of the option price surface to a low-dimensional submanifold; see for example Cont, Fonseca and Durrleman [49, 50] for the use of Karhunen–Loève decomposition on implied volatility (IV) surfaces to extract low-dimensional factor models.

We then assume that the evolution of the collection of option prices is driven by a small number of *latent factors*, for which we specify diffusion-like dynamics. In other words, rather than parametrising call option price with an equidimensional codebook as in the traditional market models [59, 34, 35, 100, 147, 135], we reduce the call price object to a low-dimensional factor object. Thereafter, from the no-arbitrage conditions, we derive a class of low-dimensional diffusion models for the factor with HJM-type drift restrictions, and a statically arbitrage-free state space. We shall see that this low-dimensional representation has various advantages, in particular it is possible to calibrate non-parametric models for the factor dynamics, the estimated models are noticeably more stable, and the number of constraints that need to be verified is significantly reduced.

[127]).

1.4 Model inference and constrained neural-SDE

For our models, inference consists of two independent steps: decoding factors from call prices and calibrating SDE models from the decoded factors.

The factor decoding step is to extract a smaller number of market factors from prices of finitely many options. These factors are built to reflect the joint goals of eliminating static and dynamic arbitrage in reconstructed prices and guaranteeing statistical accuracy. Expressing option prices as affine functions of market factors, we use a modified form of principal component analysis (PCA) to decompose option prices into a small number of driving factors.

In the SDE model calibration step, we represent the drift and diffusion functions by neural networks, referred to as *neural-SDE* by Cuchiero, Khosrawi and Teichmann [53] and Gierjatowicz et al. [74]. By leveraging deep learning algorithms, we train the neural networks by maximising the likelihood of observing the factor paths, subject to the derived arbitrage constraints. As argued by Gierjatowicz et al. in their framework of neural martingale models, this allows for calibration and model selection to be done simultaneously.

1.4.1 Neural networks as functional approximators

Since the universal approximation theorems by Hornik, Stinchcomber and White [86, 87] establish the theoretical foundation of using neural networks as approximators for continuous functions on compact domains, there is a vast literature on approximation capabilities in terms of error bounds (e.g. [14, 149]) and network width and depth (e.g. [64, 118, 132]). As for finance applications, neural nets have gained fast-growing popularity (and success) since Hutchinson et al. [92] in 1990s; see the surveys by Ruf and Wang [129] and Cohen et al. [48].

Broadly speaking, neural nets are mainly used to approximate the following types of functions in finance applications, classified by use cases. The first type is parametric pricing and hedging operators, which map parameters from some handcrafted models (for examples, Black–Scholes, Heston or rough volatility models) to the resultant prices and hedging strategies ([6, 18, 88, 92]). This moves the computationally expensive option pricing offline, thus significantly accelerating, for example, calibration tasks which require repetitively evaluating the pricing operator. The second type is solutions of high-dimensional and/or non-linear PDEs ([15, 19, 131]), and the third type is value functions in temporal optimal control problems ([29, 105]). These use cases for neural networks boil down to their ability to learn complex, high dimensional, non-linear relationships.

Finally, when using neural nets to represent drift and diffusion coefficients of an SDE, we get a neural-SDE ([53, 74]).

Neural-SDEs are hybrid models where some universal neural networks are constrained by adding domain knowledge, i.e. the SDE structure, to the architecture for more efficiency and better interpretability. Compared with the neural-SDE *martingale* models [53, 74], our models are neural-SDE *market* models, which are subject to different loss functions for training and additional arbitrage constraints.

1.4.2 Constrained neural networks

No-arbitrage conditions are embedded as part of model inference. Specifically, static arbitrage constraints are characterised by a compact state space for the latent factors; we identify sufficient conditions on the drift and diffusion to restrict the factors to their arbitrage-free domain, using the classic results of Friedman and Pinsky [71]. Consequently, the neural network that is used to parametrise the drift and diffusion functions needs to be constrained.

Chatatigner, Crépey and Dixon [42] compare the *hard constraint* and *soft constraint* approaches for imposing no-arbitrage constraints on put option prices. To enforce hard constraints, one modifies the network structure to embed the constraints (see Dugas, Benigo, Bélisle, Dadeau and Garcia [61]), while the soft constraint approach introduces penalty terms favouring the satisfaction of the constraints (see Itkin [94]). Though all these works are concerned with enforcing no-arbitrage conditions on neural networks that generate option prices, they are all solving cross-sectional calibration problems, rather than learning dynamic models.

For ensuring no-arbitrage in our neural-SDE market models, we constrain the drift and diffusion neural networks by developing appropriate transformations for the output of the neural networks.

1.5 Key contributions and structure of this thesis

The principal objective of this thesis is to construct nonparametric factor-based and arbitrage-free market models for European options, and establish an efficient inference method for estimating the models by leveraging deep learning algorithms. The key results of the thesis have been previously appeared as the three papers [45], [46] and [47].

Chapter 2 and 3 are essentially the article [45]. There are two major contributions. First, we establish an efficient construction of empirically verifiable, model-independent static arbitrage constraints for finitely many European call options, mainly based on Carr, Géman, Madan and Yor [38, 39], Davis and Hobson [56], and Cousot [52]. These

arbitrage constraints are crucial for constraining the market models that we will build later in Chapter 4 to ensure the absence of static arbitrage. Second, we study the option price data *repair* problem where only necessary and minimal changes are made to the given data in the presence of arbitrage. Specifically, we formulate data repair as a constrained linear programming problem, and the design of the objective function is carefully discussed. We shall see, in Chapter 4 and 7, that the developed arbitrage repair algorithm has wide applications in pre-processing of data for model calibration and simulation from the calibrated model.

Chapter 4, 5 and 6 are mainly the article [46], where the key contributions are summarised as follows.

- (i) We construct a family of factor-based market models, where the factor representation in principle allows exact static calibration, and the joint dynamics of options are straightforwardly available once the dynamics of the factors are specified. The models are given by a finite system of SDEs for the factors and the stock price.
- (ii) We derive an HJM-type drift condition on the factor SDEs which guarantees freedom from dynamic arbitrage, and the state space of the market factor processes is constrained where the models are free from static arbitrage.
- (iii) To calibrate our models, we need to tackle the problem of calibrating a neural-SDE with a convex polytope state space. We propose a novel hard constraint approach that modifies the network to respect sufficient conditions on the drift and diffusion to restrain the process within the polytope. Readers who are interested in this calibration problem, rather than market models, can directly jump to Chapter 5.

Chapter 7 and 8 present the work in [47]. While in Chapter 6 we have validated the market models using synthetic data generated from a known model, here we apply the modelling approach using real world data, and explore the capacity of the model as a realistic option market simulator. In addition, we show that our models are more computationally efficient and accurate for evaluating risks of option portfolios, with better coverage performance and less procyclicality than standard filtered historical simulation approaches [78].

Chapter 9 considers deriving sensitivity and minimum variance based hedging ratios under market models. Using the same dataset and calibrated market model as in Chapter 7 and 8, we compare the hedging performance for portfolios of European vanilla options with standard Greeks-based hedging under the Black–Scholes model, demonstrating favourable performance of our approach.

Chapter 10 concludes the thesis and outlines future extensions. Some chapters have been supplied appendices at the end of the corresponding chapter.

Static arbitrage constraints

“故曰：‘天下熙熙，皆为利来；天下攘攘，皆为利往。’夫千乘之王，万家之侯，百室之君，尚犹患贫，而况匹夫编户之民乎。”

— 司马迁 《史记·货殖列传》

“So it is said,

Jostling and joyous,
The whole world comes after profit;
Racing and rioting,
After profit the whole world goes!

If even the king of a land of a thousand chariots, the lord of ten thousands households, or the master of a hundred dwellings must worry about poverty, how much more so the common peasant whose name is enrolled in the tax collector’s list? ”

— Sima Qian, *The Biographies of the Money-Makers, Records of the Grand Historian* (translated by Burton Watson)

Ensuring no static arbitrage is a crucial element in developing our market models. It is therefore important to derive static arbitrage constraints in terms of option prices. Our derivation is mainly based on Carr, Géman, Madan and Yor [38] [39], Davis and Hobson [56], and Cousot [52]. In particular, we consider a finite collection of traded European call options¹ written on the same asset. These options can have arbitrary expiry and strike parameters rather than a rectangular grid of parameters, a restrictive prerequisite for many arbitrage detection [39] and spline-type smoothing methods [67] [99] to work. In practice, it is uncommon to have price data on a rectangular grid, see, e.g. Figure 2.1.

Consider N European call options that have expiries $0 < T_1 < T_2 < \dots < T_m$. For a

¹We focus on European style vanilla options in this study. Specifically, we only consider call options, since the static arbitrage constraint between call and put options is the put-call parity, which can be easily incorporated in our approach.

given expiry T_i , available strikes are $0 < K_1^i < K_2^i < \dots < K_{n_i}^i$. The (i, j) -th option has reference price C_j^i at present time 0, and terminal payoff $(S_{T_i} - K_j^i)^+$, where S_t denotes the price of the underlying asset at time t . Hence $N = \sum_{i=1}^m n_i$. Denote $\mathcal{T}^e = \{T_i\}_{1 \leq i \leq m}$ and $\mathcal{P}^{T,K} = \{(T_i, K_j^i)\}_{1 \leq j \leq n_i, 1 \leq i \leq m}$. In Figure 2.1, we show how (T, K) are distributing for traded call options on a typical trading day. A detailed description of the data used can be seen in Section 3.5.




Figure 2.1: Distributions of (K, T) for traded EURUSD call options in the OTC market (Bloomberg data) and at the CME market, observed as of 31st May, 2018.

2.1 Assumptions

The arbitrage constraints that price data should satisfy are derived under a frictionless market assumption. As a consequence, when the price data break these constraints, it may not be possible in practice to exploit the apparent arbitrage, given practical market barriers and transaction costs. However, the assumption that prices should be arbitrage free is justified by the fact that the single-price data² are not executable prices in the market, but are designed to be reference or benchmark prices for tradable assets, which are useful inputs to a variety of models.

We allow for non-zero but deterministic interest and dividends³. At present time 0, we use $D(T)$ to denote the market discount factor for time T , and $\Gamma(T)$ to denote the number of shares which will be owned by time T if dividend income is invested in shares. Then there is a model-independent, arbitrage-free forward price, $F(T) = S_0/(\Gamma(T)D(T))$, for delivery of the asset at T .

We assume that zero-coupon bonds and forward contracts on the risky asset, with the same expiries as the options, are traded in the market. In addition, they are sufficiently liquid that we can neglect their bid-ask spreads (e.g. usually one or two ticks). Therefore,

²Nevertheless, we will use bid and ask prices as soft bounds for guiding our arbitrage repair mechanism in Chapter 3.

³When applying our method to other asset classes, dividends of stock shares are comparable to foreign currency interest rates for FX rates, or convenience yields for commodities.

we observe market discount factors $D_i := D(T_i)$ and forward prices $F_i := F(T_i)$ for $1 \leq i \leq m$. However, when the underlying (spot or forward) trades at a sufficiently large bid-ask spread, then any arbitrage strategy can become impossible (see the discussion by Gerhold and Gülüm [73]).

2.2 Static arbitrage

Arbitrage refers to a costless trading strategy that has a positive probability of earning risk-free profit. A *static arbitrage* is an arbitrage exploitable by fixed positions in options and the underlying stock at initial time, while the position of the underlying stock can be modified at a finite number of trading times in the future. Any other arbitrage is called *dynamic arbitrage*. As an example of a static arbitrage, it must hold the condition that $C_1^1 \geq C_2^1$ for $K_1^1 < K_2^1$, otherwise by going long one (T_1, K_1^1) option and short one (T_1, K_2^1) option, we make immediate profit of $C_2^1 - C_1^1$ with non-negative terminal payoff. An example of dynamic arbitrage is a continuously delta-hedged short position on an over-priced option in the perfect Black–Scholes world.

A *model* \mathbb{M} is a filtered probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t \in \mathcal{T}}, \mathbb{P})$, that carries an adapted price process $\{(S_t, \mathbf{C}_t)\}_{t \in \mathcal{T}}$, where \mathbf{C}_t gives the prices of the N options at time t , and we observe \mathbf{C}_0 . Here \mathcal{T} denotes the set of times at which the asset can be traded so that $0 \in \mathcal{T}$ and $\mathcal{T}^e \subset \mathcal{T}$, and $\mathcal{F}_0 = \{\Omega, \emptyset\}$ augmented with all null sets of \mathcal{F}_{T_m} .

The First Fundamental Theorem of Asset Pricing (FFTAP) establishes an equivalence relation between no-arbitrage (static and dynamic) and the existence of an equivalent martingale measure (EMM). We work with a simplified version of FFTAP as follows. Given a model \mathbb{M} , there is no arbitrage if and only if $\exists \mathbb{Q} \sim \mathbb{P}$, such that

$$\forall (T, K) \in \mathcal{P}^{T, K} \cup (\mathcal{T}^e \times \{0\}), \quad D(t)C_t(T, K) = D(s)\mathbb{E}^{\mathbb{Q}}[C_s(T, K)|\mathcal{F}_t] \quad (2.1)$$

for all $t < s \leq T$ where $t, s \in \mathcal{T}$. No static arbitrage corresponds to a much smaller set of conditions, since the path dynamics governed by \mathbb{Q} no longer matter. As discussed by Carr, Géman, Madan and Yor [38], [39] and Davis [56], static arbitrage is present if no \mathbb{Q} exists such that $C_0(T, K) = D(T)\mathbb{E}^{\mathbb{Q}}[C_T(T, K)|\mathcal{F}_0]$. Therefore, static arbitrage constraints are consequences of relations between terminal payoffs, projected to the present time.

2.3 Shape constraints of the call price surface

Let us define $M_{T_i} = S_{T_i}/F_i$, $k_j^i = K_j^i/F_i$, $c_j^i = C_j^i/(D_i F_i)$, for all i, j . To have no static arbitrage, there must exist \mathbb{Q} such that

$$c_j^i = \mathbb{E}^{\mathbb{Q}} \left[\left(\frac{S_{T_i}}{F_i} - \frac{K_j^i}{F_i} \right)^+ \middle| \mathcal{F}_0 \right] = \mathbb{E}^{\mathbb{Q}} \left[(M_{T_i} - k_j^i)^+ \middle| \mathcal{F}_0 \right], \quad \forall i, j.$$

We will work on “normalised” quantities M , k , c in the rest of this section. We define the *normalised call function* $c(T, k)$ as

$$c(T, k) := \mathbb{E}^{\mathbb{Q}} \left[(M_T - k)^+ \middle| \mathcal{F}_0 \right], \quad \text{where } T \in \mathbb{R}_{>0}, k \in \mathbb{R}_{\geq 0}. \quad (2.2)$$

Given the specific structure in (2.2), a probability measure \mathbb{Q} exists only when the call function satisfies some shape constraints. For arbitrary but fixed T , using Breeden and Litzenberger’s analysis [26], the marginal measure $\mathbb{Q}_T := \mathbb{Q}(\cdot | \mathcal{F}_T)$ exists if

$$\forall k_3 > k_2 > k_1 \geq 0, \quad -1 \leq \frac{c(T, k_2) - c(T, k_1)}{k_2 - k_1} \leq \frac{c(T, k_3) - c(T, k_2)}{k_3 - k_2} \leq 0.$$

If a family of marginal measures $\{\mathbb{Q}_T\}_{T \in \mathcal{T}^e}$ on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ exists with time-independent mean, and $\mathbb{Q}_{T_1} \geq_{\text{cx}} \mathbb{Q}_{T_2}$ whenever $T_1 > T_2$, then, by Kellerer’s theorem [102], there exists a Markov martingale measure with these marginals. Here we write $\mathbb{Q}_{T_1} \geq_{\text{cx}} \mathbb{Q}_{T_2}$ if $\int_{\mathbb{R}} \phi d\mathbb{Q}_{T_1} \geq \int_{\mathbb{R}} \phi d\mathbb{Q}_{T_2}$ for each convex function $\phi : \mathbb{R} \rightarrow \mathbb{R}$, and we say $\{\mathbb{Q}_T\}_{T \in \mathcal{T}^e}$ is Non-Decreasing in Convex Order (NDCO). The convex order can be equivalently characterised in terms of the call function [137]:

$$\mathbb{Q}_{T_1} \geq_{\text{cx}} \mathbb{Q}_{T_2} \iff \begin{cases} \mathbb{Q}_{T_1} \text{ and } \mathbb{Q}_{T_2} \text{ have equal means;} \\ \int_{\mathbb{R}} (x - k)^+ d\mathbb{Q}_{T_1} \geq \int_{\mathbb{R}} (x - k)^+ d\mathbb{Q}_{T_2} \quad \forall x \in \mathbb{R}. \end{cases}$$

Given that $\mathbb{E}^{\mathbb{Q}_T}[M_U] = \mathbb{E}^{\mathbb{Q}_T}[S_U/F_T(U)] = 1$ is time-independent for any $T < U$ where $T, U \in \mathcal{T}^e$, it is then sufficient to conclude that $\{\mathbb{Q}_T\}_{T \in \mathcal{T}^e}$ is NDCO if $c(T, \cdot) \leq c(U, \cdot)$. Also note that $\lim_{k \downarrow 0} c(T, k) = \mathbb{E}^{\mathbb{Q}}[M_T | \mathcal{F}_0] = M_0 = 1$ for any T , and by monotonicity we have $0 \leq c(T, k) \leq 1$. Hence, if we define a set of functions $s(x, y) : X \times Y \rightarrow \mathbb{R}$, where $X, Y \subseteq \mathbb{R}_{\geq 0}$ are compact sets, by

$$\mathcal{S}(X \times Y) = \left\{ (x, y) \mapsto s(x, y) : \begin{aligned} &\forall x_1 < x_2 \in X, y_1 < y_2 < y_3 \in Y, \\ &0 \leq s \leq 1, \quad s(x_1, \cdot) \leq s(x_2, \cdot), \\ &-1 \leq \frac{s(\cdot, y_2) - s(\cdot, y_1)}{y_2 - y_1} \leq \frac{s(\cdot, y_3) - s(\cdot, y_2)}{y_3 - y_2} \leq 0 \end{aligned} \right\}, \quad (2.3)$$

then no arbitrage can be constructed on the static surface $(T, k) \mapsto c(T, k)$ if $c \in \mathcal{S}(\mathbb{R}_{>0} \times \mathbb{R}_{\geq 0})$. Consequently, no static arbitrage can be constructed from the finite collection of prices if

$$\exists c \in \mathcal{S} \left(\mathcal{T}^e \times [0, \max_{i,j} k_j^i] \right), \text{ s.t. } \forall (T_i, k_j^i) \in \mathcal{P}^{T,k}, c(T_i, k_j^i) = c_j^i, \quad (2.4)$$

where $\mathcal{P}^{T,k} = \{(T_i, k_j^i)\}_{1 \leq j \leq n_i, 1 \leq i \leq m}$.

Condition (2.4) can be characterised by practically verifiable constraints of prices \mathbf{c} . We slightly revise Cousot's construction (Definition 2.1 – 2.2 in [52]). We augment the given price data with the price that corresponds to a call struck at 0 for each expiry. This means $\forall i \in \{1, \dots, m\}$ we add $K_0^i = 0$ and $C_0^i = F_i$, or equivalently $k_0^i = 0$ and $c_0^i = 1$. This augmentation is necessary to check arbitrage relationships between call options and forwards. Define, for any $k_{j_1}^{i_1} > k_{j_2}^{i_2}$, where $1 \leq i_1, i_2 \leq m$, $0 \leq j_1 \leq n_{i_1}$, and $0 \leq j_2 \leq n_{i_2}$,

$$\beta(i_1, j_1; i_2, j_2) := \frac{c_{j_1}^{i_1} - c_{j_2}^{i_2}}{k_{j_1}^{i_1} - k_{j_2}^{i_2}}, \quad (2.5)$$

which can be viewed as the slope of the straight line passing through the two points $(k_{j_1}^{i_1}, c_{j_1}^{i_1})$ and $(k_{j_2}^{i_2}, c_{j_2}^{i_2})$, if we plot all prices on the (k, c) plane. We will employ $\beta(\cdot)$ to define the price of some *test strategies*.

Definition 2.1. A test *spread* strategy is defined $\forall 1 \leq i_1 \leq i_2 \leq m$, and $\forall 0 \leq j_1 \leq n_{i_1}$, $0 \leq j_2 \leq n_{i_2}$ such that $k_{j_1}^{i_1} \geq k_{j_2}^{i_2}$, by

$$S_{j_1, j_2}^{i_1, i_2} = \begin{cases} -\beta(i_1, j_1; i_2, j_2) & \text{if } k_{j_1}^{i_1} > k_{j_2}^{i_2}, \\ c_{j_2}^{i_2} - c_{j_1}^{i_1} & \text{if } k_{j_1}^{i_1} = k_{j_2}^{i_2}. \end{cases}$$

In particular, there are three types of test spread strategies:

- (i) *Vertical spread*: $VS_{j_1, j_2}^i = S_{j_1, j_2}^{i, i}$ with $k_{j_1}^i > k_{j_2}^i$.
- (ii) *Calendar spread*: $CS_j^{i_1, i_2} = S_{j_1, j_2}^{i_1, i_2}$ with $k_{j_1}^{i_1} = k_{j_2}^{i_2}$ and $i_1 < i_2$.
- (iii) *Calendar vertical spread*: $CVS_{j_1, j_2}^{i_1, i_2} = S_{j_1, j_2}^{i_1, i_2}$ with $k_{j_1}^{i_1} > k_{j_2}^{i_2}$ and $i_1 < i_2$.

Definition 2.2. A test *butterfly* strategy is defined $\forall i, i_1, i_2 \in [1, m]$ s.t. $i \leq i_1$ and $i \leq i_2$, $\forall j \in [0, n_i]$, $j_1 \in [0, n_{i_1}]$, $j_2 \in [0, n_{i_2}]$ such that $k_{j_1}^{i_1} < k_j^i < k_{j_2}^{i_2}$, by

$$B_{j, j_1, j_2}^{i, i_1, i_2} = -\beta(i, j; i_1, j_1) + \beta(i_2, j_2; i, j).$$

In particular, there are two types of test butterfly strategies:

- (i) *Vertical butterfly*: $VB_{j, j_1, j_2}^i = B_{j, j_1, j_2}^{i, i, i}$.
- (ii) *Calendar butterfly*: $CB_{j, j_1, j_2}^{i, i_1, i_2} = B_{j, j_1, j_2}^{i, i_1, i_2}$ where i, i_1 and i_2 are not all equal.

Based on these definitions of test strategies, we restate Cousot's constraints for no-arbitrage in the following proposition.

Proposition 2.1 (Cousot [52]). *All test strategies are non-negative, and all test vertical spreads are not greater than 1, if and only if there exist m risk-neutral measures $\{\mathbb{Q}_{T_i}\}_{1 \leq i \leq m}$ corresponding to all option expiries, that are NDCO. In addition, all their means are equal to $M_0 = 1$.*

Together with Kellerer's theorem [102], Proposition 2.1 gives sufficient conditions for the existence of a \mathbb{Q} -martingale (thus no static arbitrage), in terms of constraints on prices of the test strategies. Those constraints are also necessary for no static arbitrage if semi-static strategies are allowed to exploit arbitrage opportunities, as proved by Cousot in Appendix A of [52].

2.4 Constraints reduction

Cousot's constraints contain redundancies. For instance, if two vertical spreads VS_{j_2, j_1}^i and VS_{j_3, j_2}^i (where $k_{j_1}^i < k_{j_2}^i < k_{j_3}^i$) are non-negative, then $VS_{j_3, j_1}^i \geq 0$ holds automatically. We will reduce the number of constraints from $\mathcal{O}(N^3)$ to $\mathcal{O}(m^2N)$ by localisation on the surface. Localisation can successfully reduce the amount of constraints because the shape constraints specified in (2.3) include only boundedness, positivity, monotonicity and convexity, which are all local properties. The reduced set of constraints is listed in Table 2.1, where the (order of the) number of constraints in each category is also indicated.

We give details of the localisation method in Appendix 2.A. We claim that the reduced set of constraints listed in Table 2.1 are sufficient to imply Cousot's constraints, thus are sufficient and necessary to guarantee no-arbitrage, as stated in Proposition 2.2.

Proposition 2.2. *If the constraints C1 – C6 are satisfied, then all test strategies are non-negative, and all test vertical spreads are not greater than 1.*

Proof. See Appendix 2.B. □

Appendix

2.A Localisation of static arbitrage constraints

To localise calendar butterfly constraints, we use a sequential build-up of local constraints from the shortest expiry to the longest expiry. Define $\mathcal{D}_i := \{(k_j^i, c_j^i) : 1 \leq j \leq n_i\}$ as price data for options of expiry $T_i \in \mathcal{T}^e$. Given arbitrage-free \mathcal{D}_{i^*} , we construct constraints such that adding price data of any longer-expiry option should not introduce arbitrage. This

Category	Constraints	Number
C1 Outright	$\forall i \in [1, m], c_{n_i}^i \geq 0$	m
C2 Vertical spread	$\forall i \in [1, m], j \in [1, n_i],$ $VS_{j,j-1}^i \geq 0$ and $VS_{1,0}^i \leq 1$	$N + m$
C3 Vertical butterfly	$\forall i \in [1, m], j \in [1, n_i - 1], VB_{j,j-1,j+1}^i \geq 0$	$N - m$
C4 Calendar spread	$\forall 1 \leq i_1 < i_2 \leq m, j_1 \in [0, n_{i_1}], j_2 \in [0, n_{i_2}],$ $CS_{j_1,j_2}^{i_1,i_2} \geq 0$	$\mathcal{O}(mN)$
C5 Calendar vertical spread	$\forall i^* \in [1, m], j^* \in [1, n_{i^*}],$ define $\mathcal{I} := \{i, j : T_i > T_{i^*}, k_{j^*-1}^{i^*} < k_j^i < k_{j^*}^{i^*}\},$ then $\forall i, j \in \mathcal{I}, CVS_{j^*,j}^{i^*,i} \geq 0$	$\mathcal{O}(mN)$
C6.1 Calendar butterfly I (Absolute location convexity)	$\forall i^* \in [1, m], j^* \in [1, n_{i^*} - 1],$ define $\mathcal{I} := \{i, j : T_i > T_{i^*}, k_{j^*-1}^{i^*} < k_j^i < k_{j^*}^{i^*}\},$ then $\forall i, j \in \mathcal{I}, CB_{j^*,j^*+1}^{i^*,i^*} \geq 0;$ <hr/> $\forall i^* \in [1, m], j^* \in [2, n_{i^*}],$ define $\mathcal{I} := \{i, j : T_i > T_{i^*}, k_{j^*-1}^{i^*} < k_j^i < k_{j^*}^{i^*}\},$ then $\forall i, j \in \mathcal{I}, CB_{j^*-1,j^*-2,j}^{i^*,i^*} \geq 0;$ <hr/> $\forall i^* \in [1, m],$ define $\mathcal{I} := \{i, j : T_i > T_{i^*}, k_j^i > k_{n_{i^*}}^{i^*}\},$ then $\forall i, j \in \mathcal{I}, CB_{n_{i^*},n_{i^*}-1,j}^{i^*,i^*} \geq 0$	$\mathcal{O}(m^2N)$
C6.2 Calendar butterfly II (Relative location convexity)	$\forall i^* \in [1, m], j^* \in [1, n_{i^*} - 1],$ define $\mathcal{I}_1 := \{i, j : T_i > T_{i^*}, k_{j^*-1}^{i^*} < k_j^i < k_{j^*}^{i^*}\},$ $\mathcal{I}_2 := \{i, j : T_i > T_{i^*}, k_{j^*}^{i^*} < k_j^i < k_{j^*+1}^{i^*}\},$ $\forall i_1, j_1 \in \mathcal{I}_1, \forall i_2, j_2 \in \mathcal{I}_2, CB_{j^*,j_1,j_2}^{i^*,i_1,i_2} \geq 0;$ <hr/> $\forall i^* \in [1, m],$ define $\mathcal{I}_1 := \{i, j : T_i > T_{i^*}, k_{n_{i^*}-1}^{i^*} < k_j^i < k_{n_{i^*}}^{i^*}\},$ $\mathcal{I}_2 := \{i, j : T_i > T_{i^*}, k_j^i > k_{n_{i^*}}^{i^*}\},$ $\forall i_1, j_1 \in \mathcal{I}_1, \forall i_2, j_2 \in \mathcal{I}_2, CB_{n_{i^*},j_1,j_2}^{i^*,i_1,i_2} \geq 0$	$\mathcal{O}(m^2N)$

Table 2.1: The reduced set of static arbitrage constraints

is done locally in two steps, where we scan a neighbourhood of each $k_j^{i^*}$. The first step, we call “absolute location convexity” C6.1, finds constraints ensuring that adding *any single* data point (k_j^i, c_j^i) where $i > i^*$ will not introduce arbitrage. In Figure 2.A.1 we indicate the regions where adding a single data point will not introduce arbitrage for four types of strike neighbourhood. In the second step “relative location convexity” C6.2, we find constraints making sure that adding *all* data points (k_j^i, c_j^i) where $i > i^*$ will not introduce arbitrage for two types of strike neighbourhood. As shown in Figure 2.A.2, if we draw line segments by linking each added point and the reference point $o = (k_{j^*}^{i^*}, c_{j^*}^{i^*})$, we require the slope of any line on the left $\{l_i\}$ to be not greater than the slope of any line on the right $\{r_j\}$.




Figure 2.A.1: Absolute location convexity constraint, discussed in four cases. Points falling in the green region satisfy the absolute location convexity constraint.




Figure 2.A.2: Relative location convexity constraint, discussed in two cases. Points within the green region satisfy the absolute location convexity constraint.

2.B Proof of Proposition 2.2

We prove Proposition 2.2 by establishing Lemma 2.1, 2.2 and 2.3.

Lemma 2.1. *If C1, C2 and C3 are satisfied, then all outright, vertical spreads and vertical butterflies are non-negative. In addition, all test vertical spreads are not greater than 1.*

Proof. We consider the prices of call options with the same expiry T_i where $i \in [1, m]$.

First, we prove that any vertical spread is non-negative, i.e. $\forall 0 \leq j_1 < j_2 \leq n_i$, $c_{j_1}^i \geq c_{j_2}^i$. This is true by the vertical spread constraint C2, as $c_{j_1}^i \geq c_{j_1+1}^i \geq \dots \geq c_{j_2}^i$.

Second, we show that all outright are non-negative, i.e. $\forall j \in [0, n_i]$, $c_j^i \geq 0$. Given the outright constraint C1 and that any vertical spread is non-negative, we have $c_j^i \geq c_{n_i}^i \geq 0$.

Next, we show that any vertical butterfly is non-negative, i.e. $\forall 0 \leq j_1 < j < j_2 \leq n_i$, $\beta(i, j_2; i, j) - \beta(i, j; i, j_1) \geq 0$. To do that, we claim

$$\beta(i, j_1 + 1; i, j_1) \leq \beta(i, j; i, j_1 + 1), \quad \text{if } j_1 < j - 1, \quad (2.6a)$$

$$\beta(i, j_2; i, j_2 - 1) \geq \beta(i, j_2 - 1; i, j), \quad \text{if } j_2 > j + 1. \quad (2.6b)$$

These two claims can be proved by induction. Here we only show the proof for (2.6a). It is true that $\beta(i, j - 1; i, j - 2) \leq \beta(i, j; i, j - 1)$ (the $j_1 = j - 2$ case for (2.6a)) by the vertical butterfly constraint C3. Assume that (2.6a) holds for $j_1 = l < j - 2$, i.e.

$$\begin{aligned} c_j^i &\geq c_{l+1}^i + (k_j^i - k_{l+1}^i) \frac{c_{l+1}^i - c_l^i}{k_{l+1}^i - k_l^i} = c_{l+1}^i + [(k_j^i - k_l^i) - (k_{l+1}^i - k_l^i)] \frac{c_{l+1}^i - c_l^i}{k_{l+1}^i - k_l^i} \\ &= c_l^i + (k_j^i - k_l^i) \beta(i, l + 1; i, l). \end{aligned}$$

This leads to $\beta(i, j; i, l) \geq \beta(i, l + 1; i, l)$. Again by C3 we have $\beta(i, l + 1; i, l) \geq \beta(i, l; i, l - 1)$. Hence, $\beta(i, j; i, l) \geq \beta(i, l; i, l - 1)$, which is the stated inequality by setting $j_1 = l - 1$. Therefore, (2.6a) holds by induction in reverse order from $j_1 = j - 2$ to $j_1 = 0$. Thereafter, (2.6a) implies

$$\begin{aligned} -c_{j_1}^i &\leq -c_{j_1+1}^i + (k_{j_1+1}^i - k_{j_1}^i) \frac{c_j^i - c_{j_1+1}^i}{k_j^i - k_{j_1+1}^i} \\ &= -c_{j_1+1}^i + [(k_{j_1+1}^i - k_j^i) - (k_{j_1}^i - k_j^i)] \frac{c_j^i - c_{j_1+1}^i}{k_j^i - k_{j_1+1}^i} \\ &= -c_j^i + (k_{j_1+1}^i - k_j^i) \beta(i, j; i, j_1 + 1), \end{aligned}$$

which leads to $\beta(i, j; i, j_1) \leq \beta(i, j; i, j_1 + 1)$. Similarly we have $\beta(i, j; i, j_1) \leq \beta(i, j; i, j_1 + 1) \leq \dots \leq \beta(i, j; i, j - 1)$. In similar fashion we can prove (2.6b) by induction from $j_1 = j + 2$ to $j_1 = n_i$, and deduce $\beta(i, j_2; i, j) \geq \beta(i, j_2 - 1; i, j) \geq \dots \geq \beta(i, j + 1; i, j)$. Therefore, with C3, we can conclude $\beta(i, j_2; i, j) \geq \beta(i, j; i, j_1)$.

Finally, we show that any vertical spread is bounded by 1, i.e. $0 \leq j_1 < j_2 \leq n_i$, $-\beta(i, j_2; i, j_1) \leq 1$. For the case $j_1 > 1$, given that any butterfly spread is non-negative, we have $-\beta(i, j_2; i, j_1) \leq -\beta(i, j_1; i, 1) \leq -\beta(i, 1; i, 0) \leq 1$, where the last inequality holds due to the vertical spread constraint C2. If $j_1 = 1$, then $-\beta(i, j_2; i, j_1 = 1) \leq -\beta(i, 1; i, 0) \leq 1$. Otherwise $j_1 = 0$, applying (2.6a) by assigning $j_1 = 0$, $j = j_2$ yields $-\beta(i, j_2; i, j_1 = 0) \leq -\beta(i, 1; i, 0) \leq 1$. \square

Lemma 2.2. *If C2, C4 and C5 are satisfied, then any calendar spread or calendar vertical spread is non-negative.*

Proof. We would like to prove that $\forall 0 \leq i_1 < i_2 \leq m$ and $\forall j_1 \in [0, n_{i_1}]$, $j_2 \in [0, n_{i_2}]$ where $k_{j_1}^{i_1} \geq k_{j_2}^{i_2}$, we have $c_{j_1}^{i_1} \leq c_{j_2}^{i_2}$.


First consider the calendar spread case when $k_{j_1}^{i_1} = k_{j_2}^{i_2}$. The calendar spread constraint C4 immediately leads to $c_{j_1}^{i_1} \leq c_{j_2}^{i_2}$.

Otherwise $k_{j_1}^{i_1} > k_{j_2}^{i_2}$, which implies that j_1 must be greater than 0. Given $i_1 \in [1, m]$, $j_1 \in [1, n_{i_1}]$, there must be $k_{j_2}^{i_2} \in [k_{j_1-p-1}^{i_1}, k_{j_1-p}^{i_1})$ for some $p \in [0, j_1 - 1]$. By the calendar vertical spread constraint C5, we have $c_{j_2}^{i_2} \geq c_{j_1-p}^{i_1}$. In addition, $c_{j_1-p}^{i_1} \geq c_{j_1}^{i_1}$ due to the vertical spread constraint C2. Hence, $c_{j_2}^{i_2} \geq c_{j_1}^{i_1}$. \square

Lemma 2.3. *If C3 and C6 are satisfied, then any calendar butterfly is non-negative.*

Proof. We would like to prove that $\forall i, i_1, i_2 \in [1, m]$ where $i \leq i_1, i \leq i_2$ and $\forall j \in [1, n_i]$, $j_1 \in [0, n_{i_1}]$, $j_2 \in [0, n_{i_2}]$ where $k_{j_1}^{i_1} < k_j^i < k_{j_2}^{i_2}$, we have $\beta(i, j; i_1, j_1) \leq \beta(i_2, j_2; i, j)$.

Given $i \in [1, m]$, $j \in [1, n_i]$, it must be that $k_{j_1}^{i_1} \in [k_{j-p-1}^i, k_{j-p}^i]$ for some $p \in [0, j - 1]$ and either $k_{j_2}^{i_2} \in [k_{j+q}^i, k_{j+q+1}^i]$ for some $q \in [0, n_i - j - 1]$ or $k_{j_2}^{i_2} \in (k_{n_i}^i, \infty)$. See Figure 2.B.1.



(a) $k_{j_1}^{i_1} \in [k_{j-p-1}^i, k_{j-p}^i]$ for some $p \in [0, j - 1]$

(b) $k_{j_2}^{i_2} \in [k_{j+q}^i, k_{j+q+1}^i]$ for some $q \in [0, n_i - j - 1]$

Figure 2.B.1: Locations of $k_{j_1}^{i_1}$ and $k_{j_2}^{i_2}$ relative to k_j^i .

Let us consider the case when $k_{j_2}^{i_2} \leq k_{n_i}^i$ (which implies that $j < n_i$). If $p = q = 0$, then by the calendar butterfly relative location constraints C6.2 we conclude $\beta(i, j; i_1, j_1) \leq \beta(i_2, j_2; i, j)$. Otherwise, we claim that if $p > 0$

$$\beta(i, j; i_1, j_1) \leq \beta(i, j; i, j - p), \quad (2.7a)$$

$$\beta(i, j; i, j - p) \leq \beta(i, j; i, j - p + 1) \leq \dots \leq \beta(i, j; i, j - 1); \quad (2.7b)$$

and similarly if $q > 0$

$$\beta(i_2, j_2; i, j) \geq \beta(i, j + q; i, j), \quad (2.8a)$$

$$\beta(i, j + q; i, j) \geq \beta(i, j + q - 1; i, j) \geq \cdots \geq \beta(i, j + 1; i, j). \quad (2.8b)$$

We will show the proof for the four claims later. If $p > 0$ and $q > 0$, the four claims and the vertical butterfly constraint C3 lead to the stated result. If $p > 0$ but $q = 0$, then (2.7) and the calendar butterfly absolute location convexity constraint C6.1 lead to the stated result. If $p = 0$ but $q > 0$, then (2.8) and C6.1 lead to the stated result.

Next we would like to prove the claims (2.7) and (2.8). First of all, (2.7b) and (2.8b) hold because of the convexity of the set of points $\{(k_l^i, c_l^i)\}_{l \in [j-p, j+q]}$ resulted from the vertical butterfly constraint C3. The calendar butterfly absolute location convexity constraint C6.1 results in $\beta(i, j - p; i_1, j_1) \leq \beta(i, j - p + 1; i, j - p)$. In addition, the vertical butterfly constraint C3 results in $\beta(i, j; i, j - p + 1) \geq \beta(i, j - p + 1; i, j - p)$, then

$$\begin{aligned} c_j^i &\geq c_{j-p+1}^i + (k_j^i - k_{j-p+1}^i) \frac{c_{j-p+1}^i - c_{j-p}^i}{k_{j-p+1}^i - k_{j-p}^i} \\ &= c_{j-p+1}^i + [(k_j^i - k_{j-p}^i) - (k_{j-p+1}^i - k_{j-p}^i)] \frac{c_{j-p+1}^i - c_{j-p}^i}{k_{j-p+1}^i - k_{j-p}^i} \\ &= c_{j-p}^i + (k_j^i - k_{j-p}^i) \beta(i, j - p + 1; i, j - p). \end{aligned}$$

Hence $\beta(i, j, i, j - p) \geq \beta(i, j - p + 1, i, j - p) \geq \beta(i, j - p, i_1, j_1)$. Then

$$\begin{aligned} c_{j_1}^{i_1} &\leq c_{j-p}^i + (k_{j_1}^{i_1} - k_{j-p}^i) \frac{c_j^i - c_{j-p}^i}{k_j^i - k_{j-p}^i} = c_{j-p}^i + [(k_{j_1}^{i_1} - k_j^i) - (k_{j-p}^i - k_j^i)] \frac{c_j^i - c_{j-p}^i}{k_j^i - k_{j-p}^i} \\ &= c_j^i + (k_{j_1}^{i_1} - k_j^i) \beta(i, j; i, j - p), \end{aligned}$$

which indicates that $\beta(i, j; i_1, j_1) \leq \beta(i, j; i, j - p)$, i.e. (2.7a).

The calendar butterfly spread absolute location convexity constraint C6.1 results in $\beta(i_2, j_2; i, j + q) \geq \beta(i, j + q; i, j + q - 1)$. In addition, the vertical butterfly constraint C3 results in $\beta(i, j + q - 1; i, j) \leq \beta(i, j + q; i, j + q - 1)$, then

$$\begin{aligned} -c_j^i &\leq -c_{j+q-1}^i + (k_{j+q-1}^i - k_j^i) \frac{c_{j+q}^i - c_{j+q-1}^i}{k_{j+q}^i - k_{j+q-1}^i} \\ &= -c_{j+q-1}^i + [(k_{j+q-1}^i - k_{j+q}^i) - (k_j^i - k_{j+q}^i)] \frac{c_{j+q}^i - c_{j+q-1}^i}{k_{j+q}^i - k_{j+q-1}^i} \\ &= -c_{j+q}^i + (k_{j+q}^i - k_j^i) \beta(i, j + q - 1; i, j + q). \end{aligned}$$

Hence $\beta(i, j + q; i, j) \leq \beta(i, j + q; i, j + q - 1) \leq \beta(i_2, j_2; i, j + q)$. Then

$$\begin{aligned} c_{j_2}^{i_2} &\geq c_{j+q}^i + (k_{j_2}^{i_2} - k_{j+q}^i) \frac{c_j^i - c_{j+q}^i}{k_j^i - k_{j+q}^i} = c_{j+q}^i + [(k_{j_2}^{i_2} - k_j^i) - (k_{j+q}^i - k_j^i)] \frac{c_j^i - c_{j+q}^i}{k_j^i - k_{j+q}^i} \\ &= c_j^i + (k_{j_2}^{i_2} - k_j^i) \beta(i, j; i, j + q), \end{aligned}$$

which indicates that $\beta(i_2, j_2; i, j) \geq \beta(i, j + q; i, j)$, i.e. (2.8a).

Now we consider the case when $k_{j_2}^{i_2} > k_{n_i}^i$. The same proof as above applies if we let $q = n_i - j$ and allow j to take the value n_i .

□

Application: detecting and repairing arbitrage

“爱人不亲，反其仁；治人不治，反其智；礼人不答，反其敬。行有不得者皆反求诸己，其身正而天下归之。”

— 孟子《孟子·离娄》

“If a man love others, and no responsive attachment is shown to him, let him turn inwards and examine his own benevolence. If he is trying to rule others, and his government is unsuccessful, let him turn inwards and examine his wisdom. If he treats others politely, and they do not return his politeness, let him turn inwards and examine his own feeling of respect. When we do not, by what we do, realise what we desire, we must turn inwards, and examine ourselves in every point. When a man’s person is correct, the whole kingdom will turn to him with recognition and submission.”

— Mencius, *Li Lou*, *Mencius* (translated by James Legge)

The neat and efficient construction of static arbitrage constraints in Chapter 2 not only is necessary for developing market models, but also enables the application of detecting and repairing arbitrage in market-observed option prices. As we shall see later, this application is useful for pre-processing input option price data that are used for calibrating the established market models. In this chapter, we show how we develop arbitrage detection and repairing algorithms on top of the established static arbitrage constraints, and apply the algorithms on cleansing real world data.

3.1 Introduction

Price data of vanilla options are widely used in various financial applications such as calibrating models for pricing and hedging, and computing risk-neutral densities (RND) of the underlying. The presence of arbitrage in option price data can lead to poor or

even failed model calibration, as well as erroneous RND estimation. Derivative pricing models in nearly all applications are constructed to be arbitrage-free, as it is economically meaningless to have a model that has the potential to make risk-free profits. Exact calibration is impossible for any arbitrage-free model or RND function with any size of input data if the price data contain arbitrage. For example, calibration of the local volatility model of Dupire [62] and Derman and Kani [58] would fail given arbitrageable data. It is also not possible to have any arbitrage-free interpolation such as Kahale [99] and Wang et al. [144], because the data to be interpolated are not arbitrage-free from the beginning. Though inexact calibration methods are available for many models, it seems natural to expect that removing arbitrage from input data can improve calibration of arbitrage-free models, such as enhancing robustness or reducing calibration error.

Therefore, it is important to remove arbitrage (if present) from option price data. Most attention in the relevant literature has been devoted to the *smoothing* and *filtering* of data. Notable works on smoothing include Aït-Sahalia and Duarte [3], Fengler et al. [67] [68] [69], Gatheral and Jacquier [72], and Lim [111]. In fact, the calibration of many pricing models, such as stochastic volatility models, is essentially arbitrage-free smoothing. Arbitrage-free data is only a byproduct of smoothing, since the main goal of smoothing is to produce a $C^{1,2}$ call price function $(T, K) \mapsto C(T, K)$ (or equivalently implied volatility function $(T, K) \mapsto \sigma_{\text{imp}}(T, K)$). For smoothing, usually an ℓ^2 -norm optimisation is used when searching over polynomial, spline or kernel parameters that produce values as close to the given price data as possible. This method leads to changes for nearly *all* data. Though liquidity considerations can be included by adding weights in the optimisation, it remains unclear what should be an effective way to set weight values for different options. The filtering of data refers to simply removing suspiciously low-quality price data according to criteria in terms of moneyness, expiry, trading volume, intra-day activity, etc. A good survey of popular empirical filtering criteria can be found in Ivanovas [95] and Meier [116]. Filtering can be quite subjective, can cause information loss, and might not even be feasible as many criteria are based on order-book level data, which are not always available (for example, in OTC markets).

In contrast to smoothing, which typically changes nearly all data, or filtering, which truncates data, we propose to *repair* data in the sense that only necessary and minimal changes are made to the given data in the presence of arbitrage. If arbitrages in data are mainly consequences of infrequent price updates of illiquid options rather than noncompetitive market, it is better to only perturb as few data points as possible. In addition, when making changes, we use bid and ask prices as soft bounds such that liquidity profiles of different options are considered in an objective way. Bid-ask spread is a measure of liquidity, i.e. the narrower the spread, the easier a market order can be matched and

executed. Since the “fair” price could lie anywhere within the bid-ask price bounds, the width of the bid-ask spread represents the degree of certainty in the market prices. Empirically, deep out-of-the-money (OTM) and in-the-money (ITM) options are thinly traded with wide bid-ask spreads, leading to less trustworthy price data compared with more liquid options. We therefore formulate the data repair as a constrained optimisation problem, where the no-arbitrage relations are written as constraints, and the objective is to minimise price changes within soft bounds. By carefully choosing the objective function, we can rewrite the formulation as a linear programming (LP) problem, so that we can take advantage of efficient solution techniques and software for large-scale LPs. A related approach using quadratic programming was proposed in [110], focusing on options with a single maturity.

Our method is to repair single-price data. At any moment during the trading day, each tradable asset has multiple prices, i.e. bid price and ask price. However, most applications require single-price inputs. There is a need to construct some “fair” reference price from the market-quoted multiple prices. Examples of a reference price are the mid-price, the quantity-weighted price, the last trade price or the micro-price by Stoikov [138]. In this article, we do not discuss the construction of reference price, and use the mid-price by default, however other reference prices could easily be considered.

We envisage further applications of our methodology in repairing data generated by models which do not themselves rule out arbitrage. Included in this class are prices predicted by deep learning methods, which have gained substantial popularity recently, as documented by the survey paper by Ruf and Wang [129]. Typically, there is no guarantee for arbitrage-free predicted option prices even if the training set is arbitrage free; see also a more detailed discussion of this point in the introduction of Dixon, Crépey and Chataigner [42], which goes on to use the local volatility code book for arbitrage free vanilla prices as a means of guaranteeing arbitrage-free interpolation of prices. The arbitrage repair method from our paper can provide a simple post-processing step of potentially arbitrageable learned prices. By repairing a discrete set of input prices directly without extra assumptions, using linear constraints only, the method distinguishes itself by versatility, transparency, and speed, making it particularly well-suited to online computations.

The rest of the chapter is structured as follows. Given the set of empirically verifiable model-independent, static arbitrage constraints for European call options in Chapter 2, we formulate data repair¹ as a constrained LP problem in Section 3.2, and the design of the objective function is carefully discussed². Finally in Section 3.5, we apply our arbitrage

¹The framework of our arbitrage repair method is applicable to a mixture of a wider range of options, as long as their arbitrage constraints can be defined by feasible linear inequalities of prices.

²Our implementation of this algorithm in Python is available in the repository <https://github.com/vicaws/arbitragerepair>.

repair method to FX option data to justify why arbitrage repair is needed for real data, and demonstrate how our method performs empirically on various metrics, especially on the improvement of model calibration. We also show an example of how we can use our repair method for identifying the formation and disappearance of executable arbitrage in intra-day price data.

3.2 Arbitrage repair

It is worth addressing that we focus on repairing *static* rather than *dynamic* arbitrage in option price data. Dynamic arbitrage relies on dynamics and path properties of the tradable assets. From the data repair perspective, we should minimise model dependence, because the repaired data are to be used in more generic applications. Hence, data should only be adjusted by model-independent constraints, so we restrict ourselves to static arbitrage in which no dynamics need to be modelled. Static arbitrage constraints establish the prerequisites that the price data have to satisfy at time zero for admitting a dynamically arbitrage-free model.

The static arbitrage constraints in Table 2.1 are linear inequalities of at most three call prices. Therefore, we can write these constraints in the form $A\mathbf{c} \geq \mathbf{b}$, where $\mathbf{c} = [c_1^1 \ \dots \ c_{n_1}^1 \ \dots \ c_{n_m}^m]^\top \in \mathbb{R}^N$, and $A = (a_{ij}) \in \mathbb{R}^{R \times N}$ and $\mathbf{b} = (b_j) \in \mathbb{R}^R$ are a constant matrix and a vector corresponding to coefficients and bounds of the inequalities, respectively, that are completely determined by the expiries and strikes of observed options. Here, R is the number of no-arbitrage constraints, where $R \sim \mathcal{O}(m^2N)$. These constraints are *feasible* by construction, i.e. $\{\mathbf{x} \in \mathbb{R}^N : A\mathbf{x} \geq \mathbf{b}\} \neq \emptyset$, because $\mathcal{S}(\mathcal{T}^e \times [0, \max_{i,j} k_j^i]) \neq \emptyset$, (for example, the prices under a Black–Scholes model satisfy the requirements).

When some row of the system of inequalities $A\mathbf{c} \geq \mathbf{b}$ is not satisfied, there is arbitrage. We define $\boldsymbol{\varepsilon}$ to be the vector of perturbations added to the vector of call prices \mathbf{c} such that the perturbed prices are arbitrage-free, i.e. $A(\mathbf{c} + \boldsymbol{\varepsilon}) \geq \mathbf{b}$. Hence, to remove arbitrage from the call price data, we seek the “minimal” repair subject to no-arbitrage constraints:

$$\min_{\boldsymbol{\varepsilon} \in \mathbb{R}^N} f(\boldsymbol{\varepsilon}), \quad \text{subject to } A\boldsymbol{\varepsilon} \geq \mathbf{b} - A\mathbf{c}, \quad (3.1)$$

where the objective $f : \mathbb{R}^N \rightarrow \mathbb{R}$ measures how much the perturbation deviates from zero. The formulation (3.1) is *feasible* because its constraints are feasible.

3.3 Design of the objective without liquidity consideration

We start from the simple case where there is no liquidity difference among options. It seems natural to use the ℓ^2 -norm for measuring the size of perturbations due to its convexity and computational efficiency when optimising by gradient-based methods. The ℓ^2 -norm has been widely used in data smoothing algorithms, such as [3], [67], [68] and [69].

However, the ℓ^2 -norm usually leads to small perturbations for *all* prices, while in our application *sparse* perturbation is desirable. An alternative is the ℓ^0 -norm³, which is a natural way of comparing difference, and produces sparse solutions. Nevertheless, the ℓ^0 -norm is nonconvex and in general leads to an NP-hard [121] optimisation problem. Hence, it is natural to consider the ℓ^1 -norm, which is well known as a convex relaxation of the ℓ^0 -norm. In fact, optimal solutions of the ℓ^0 and ℓ^1 norms objectives are equivalent under certain conditions, see [32], [77] and [60].

Choosing the ℓ^1 -norm has other benefits. When minimising a convex continuous objective function like the ℓ^1 -norm, every local minimum is a global minimum, see Chapter 4 of [25]. In addition, our repair problem is a Linear Programming (LP) problem with the ℓ^1 -norm objective, which can be solved fairly quickly even for large-scale problems. Finally, compared with the ℓ^2 -norm, the ℓ^1 -norm is more robust to outliers because the ℓ^2 -norm squares values, which increases the cost of outliers quadratically, see Huber [90].

Consequently, the ℓ^1 -norm is a natural candidate for the objective function. Blacque-Florentin and Missaoui [23] also choose the ℓ^1 -norm as objective when fitting tensor polynomials to sparse data, as inspired by the compressed sensing framework. The differences between our work and theirs are that they are concerned with smoothing data rather than repairing data, and assume a rectangular grid of strikes and expiries. The ℓ^1 -norm optimisation with linear constraints can be expressed as an LP problem. We write the objective function as $f(\boldsymbol{\varepsilon}) := \|\boldsymbol{\varepsilon}\|_{\ell^1} = \sum_{i=j}^N |\varepsilon_j| = \sum_{i=j}^N (\varepsilon_j^+ + \varepsilon_j^-)$, where $\varepsilon_j^+ = \max(\varepsilon_j, 0)$, $\varepsilon_j^- = -\min(\varepsilon_j, 0)$ for each j . We denote $\boldsymbol{\varepsilon}^+ = [\varepsilon_1^+ \cdots \varepsilon_N^+]$ and $\boldsymbol{\varepsilon}^- = [\varepsilon_1^- \cdots \varepsilon_N^-]$ so that $\boldsymbol{\varepsilon} = \boldsymbol{\varepsilon}^+ - \boldsymbol{\varepsilon}^-$. We define $B = [-A \ A]$ and $\boldsymbol{\theta} = [\boldsymbol{\varepsilon}^+ \boldsymbol{\varepsilon}^-]^\top$. Hence, the repair problem with the ℓ^1 -norm minimisation is equivalent to the following LP in canonical form:

$$\min_{\boldsymbol{\theta}} \mathbf{1}^\top \boldsymbol{\theta}, \quad \text{subject to } B\boldsymbol{\theta} \leq A\mathbf{c} - \mathbf{b}, \quad \boldsymbol{\theta} \geq \mathbf{0}. \quad (3.2)$$

After solving for an optimal $\hat{\boldsymbol{\theta}} = [\hat{\boldsymbol{\varepsilon}}^+ \ \hat{\boldsymbol{\varepsilon}}^-]^\top$, the optimal perturbation vector is recovered by $\hat{\boldsymbol{\varepsilon}} = (\hat{\boldsymbol{\varepsilon}}^+ - \hat{\boldsymbol{\varepsilon}}^-)^\top$.

³Note that the ℓ^0 -norm is not actually a “norm” as it violates the homogeneity and triangle inequality properties that a vector norm must satisfy.

3.4 Inclusion of bid and ask prices

The reference prices will typically lie within their corresponding bid-ask price bounds. In the presence of arbitrage, we not only want minimal repair, but also wish to have as many perturbed prices falling within the bid-ask price bounds as possible. Specifically, a reference price with wider bid-ask spread shall be given more freedom to be perturbed. The sparsity of the solution of the ℓ^1 -norm optimisation is less desirable if perturbing a larger number of prices can keep more perturbed prices within the bid-ask price bounds.

3.4.1 Design of the objective with bid and ask prices

We consider using the best bid/ask prices for data repair. To incorporate bid-ask price constraints into the repair problem, we revise the objective function f rather than adding extra constraints. In other words, we treat bid-ask price bounds as *soft constraints* rather than *hard constraints* like the arbitrage constraints. There may not be arbitrage-free prices within the bid-ask price bounds, and adding bid-ask price bounds as hard constraints may cause the repair problem to be infeasible.

We choose an objective function of the form $f(\varepsilon) = \sum_{j=1}^N f_j(\varepsilon_j)$ with $f_j(x) \geq 0$ for $x \in \mathbb{R}$. Then $f_j(x)$ can be naturally interpreted as the cost of perturbing the j -th option price, and $df_j(x)/d|x| > 0$ (if defined) gives the marginal cost. The ℓ^1 -norm objective sets $f_j(x) = |x| = \max(-x, x)$ and any perturbation x , where $|x| > 0$, has marginal cost 1 for all j . Let $\delta_j^a, \delta_j^b > 0$ be ask-reference spread and bid-reference spread for the j -th price, respectively. To incorporate these spreads into the objective, we require that $f_j(x)$ should have the following properties, for all $j \in [1, N]$:

- (i) $f_j(0) = \inf_x f_j(x) = 0$. The minimum is attained when there is no perturbation, which is costless to the objective;
- (ii) $f_j(x)$ is monotonically increasing (decreasing) for $x > 0$ ($x < 0$);
- (iii) $f_j(-\delta_j^b) = f_j(\delta_j^a) = \delta_0$, where $\delta_0 \geq 0$ is a constant. The cost of perturbing a price to its bid or ask price is the same for all options;
- (iv) $df_j(x)/d|x| = 1$ for $x \in (-\infty, -\delta_j^b) \cup (\delta_j^a, +\infty)$. The marginal cost of perturbing a price out of the bid-ask price bounds is the same for all options.

We therefore propose the following objective that meets all the properties, and, with particular merit, retains the ability to be expressed as an LP:

$$f_j(x) = \max \left(-x - \delta_j^b + \delta_0, -\frac{\delta_0}{\delta_j^b}x, \frac{\delta_0}{\delta_j^a}x, x - \delta_j^a + \delta_0 \right),$$

with $\delta_0 \leq \min(\delta_j^a, \delta_j^b)$ for all $j \in [1, N]$, as such the marginal cost of perturbing a price within the bid-ask price band is not greater than the marginal cost of perturbing mid prices outside the bid-ask price bounds.




Figure 3.4.1: Plot of the objective function component $f_j^{\ell^1}(x)$ and $f_j(x)$

Denote $f_j^{\ell^1}$ as the j -th component of the ℓ^1 -norm objective. We visualise the difference between $f_j^{\ell^1}$ and f_j in Figure 3.4.1. Note that $f_j^{\ell^1}$ is a special case of f_j when $\delta_j^a = \delta_j^b = \delta_0 > 0$ for all j . Choosing smaller δ_0 makes it relatively more costly to move prices outside of their bid-ask price bounds. Nevertheless, letting $\delta_0 = 0$ causes the optimisation problem to be ill-posed as it admits infinitely many solutions. For example, if $\varepsilon_j^* = 0$ is optimal, then so is $\varepsilon_j^* = \omega \min(\delta_j^a, \delta_j^b)$ for all $\omega \in [0, 1]$. In practice, we will choose

$$\delta_0 = \frac{1}{N} \wedge \min_{j=1, \dots, N} (\delta_j^a \wedge \delta_j^b). \quad (3.3)$$

This means we prefer to move all options (by ε) within the bid-ask, rather than moving one option outside its bid-ask bounds.

Hence, the objective function taking into account bid-ask spread is

$$f(\varepsilon) = \sum_{j=1}^N \max \left(-\mathbf{e}_j^\top \varepsilon - \delta_j^b + \delta_0, -\frac{\delta_0}{\delta_j^b} \mathbf{e}_j^\top \varepsilon, \frac{\delta_0}{\delta_j^a} \mathbf{e}_j^\top \varepsilon, \mathbf{e}_j^\top \varepsilon - \delta_j^a + \delta_0 \right), \quad (3.4)$$

where \mathbf{e}_j is the standard basis vector for \mathbb{R}^N with its j -th element being 1 and others being 0. With objective (3.4), we can rewrite the repair problem (3.1) as the following LP by introducing auxiliary variables $\mathbf{t} = [t_1 \ \cdots \ t_N]^\top$:

$$\begin{aligned} & \underset{\varepsilon, \mathbf{t}}{\text{minimise}} && \sum_{j=1}^N t_j \\ & \text{subject to} && -\varepsilon_j - \delta_j^b + \delta_0 \leq t_j, \quad \varepsilon_j - \delta_j^a + \delta_0 \leq t_j, \quad \forall j \in [1, N], \\ & && -\frac{\delta_0}{\delta_j^b} \varepsilon_j \leq t_j, \quad \frac{\delta_0}{\delta_j^a} \varepsilon_j \leq t_j, \quad \forall j \in [1, N], \\ & && -A\varepsilon \leq -\mathbf{b} + A\mathbf{c}. \end{aligned} \quad (3.5)$$

After solving for the optimal perturbation vector $\hat{\varepsilon}$, we get the arbitrage-free normalised call price $\hat{\mathbf{c}} = \mathbf{c} + \hat{\varepsilon}$. For each i, j , the arbitrage-free call price is $\hat{C}_j^i = \hat{c}_j^i D_i F_i$.

3.4.2 Executable arbitrage opportunities

We refer to the objective function taking into account bid-ask spread with δ_0 as in (3.3) as the ℓ^1 -BA objective. We define the *effectively perturbed prices* as those that are perturbed outside of the bid-ask price bounds. We denote the number of perturbed (resp. effectively perturbed) prices by N^ε (resp. $N^{\varepsilon, \delta}$), thus

$$N^\varepsilon = \sum_{j=1}^N \mathbb{1}_{\{|\varepsilon_j| > 0\}}, \quad N^{\varepsilon, \delta} = \sum_{j=1}^N \mathbb{1}_{\{\varepsilon_j > \delta_j^a\} \cup \{\varepsilon_j < -\delta_j^b\}}. \quad (3.6)$$

We say an arbitrage is *executable* if we can realise it by buying and selling its components at their ask and bid quotes, respectively. The arbitrage detected in options' reference prices is not necessarily executable. However, if the ℓ^1 -BA repair results in effective perturbations, i.e. $N^{\varepsilon, \delta} > 0$, then there must exist executable arbitrages. To see this, let $E_j = [c_j - \delta_j^b, c_j + \delta_j^a]$, and we can characterise $N^{\varepsilon, \delta} > 0$ as

$$\text{if } \forall i \in [1, R], \sum_{j=1}^N a_{ij} \hat{c}_j \geq b_i, \text{ then } \exists j \in [1, N] \text{ s.t. } \hat{c}_j \notin E_j.$$

Equivalently, its contrapositive statement is

$$\text{if } \forall j \in [1, N], \hat{c}_j \in E_j, \text{ then } \exists i^* \in [1, R] \text{ s.t. } \sum_{j=1}^N a_{i^*j} \hat{c}_j < b_{i^*}.$$

Therefore, it holds that

$$\sum_{j=1}^N a_{i^*j} [(c_j + \delta_j^a) \mathbb{1}_{\{a_{i^*j} \geq 0\}} + (c_j - \delta_j^b) \mathbb{1}_{\{a_{i^*j} < 0\}}] < b_{i^*}.$$

By going long on the left-hand side and going short on the right side of the inequality, we construct a portfolio that makes immediate positive profit, while the portfolio has non-negative future payoffs. The left-hand side of the inequality consists of positions in options, for which we buy at ask price $(c_j + \delta_j^a)$ and sell at bid price $(c_j - \delta_j^b)$.

3.5 Empirical studies

We carry out a series of empirical studies. We show that arbitrage is frequently present in historical price data, so repairing data is important. We also demonstrate empirical performances of the repair method in terms of sparsity, speed and improvement to model

calibration. Last, we use the ℓ^1 -BA repair for identifying the formation and disappearance of arbitrage in the intra-day S&P 500 options market on a day when the market underwent a regime switch.

3.5.1 Presence of arbitrage in historical option price data

We collect daily close (bid, ask and mid) prices from 1st November, 2007 to 31st May, 2018 for OTC FX options from Bloomberg. Bloomberg provides price quoted as implied volatility given in terms of delta. We choose 13 benchmark tenors (expiries) from overnight (one-day) to two-year. For each tenor, a list of standard *instruments* are available: at-the-money (ATM), risk-reversal (RR) and butterfly (BF). We choose the liquid 10-delta, 15-delta, 25-delta and 35-delta instruments, and construct a vanilla volatility smile of 9 moneynesses for each tenor. Following the OTC FX market conventions [148], we compute strike and time-to-expiry for each IV mid quote, and generate vanilla IV spreads from the bid/ask quotes for the instruments⁴. Thereafter, we calculate mid call prices and vanilla call price spreads using the mid vanilla IVs and generated vanilla IV spreads, together with Bloomberg FX mid forward curves. There are $117 = 13$ (tenors) \times 9 (deltas and ATM) data points on the call price surface for each day. In Figure 3.5.1, we show an example of OTC call option mid-prices and IV mid-quotes on one day.

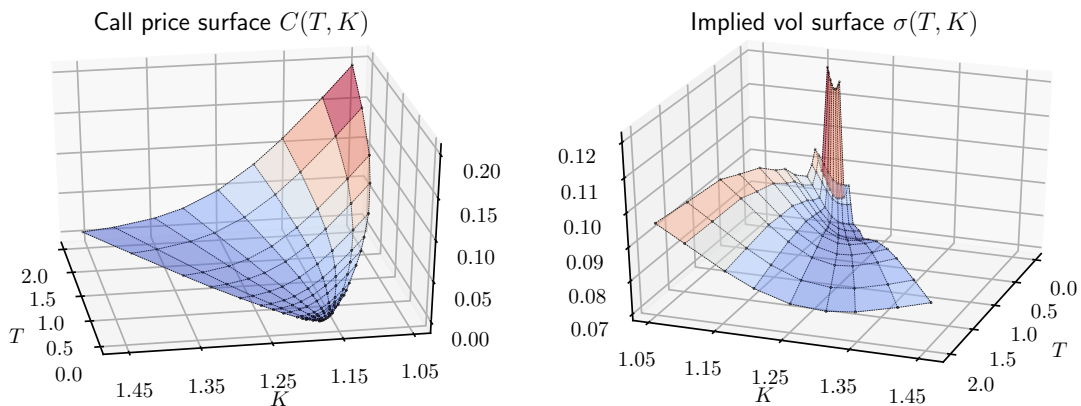


Figure 3.5.1: An example of observed OTC-traded call option prices. These are end of day prices settled by Bloomberg for EURUSD European call options as of 31st May, 2018.

We count violations of arbitrage constraints in raw daily close mid-prices over time for some major currencies and emerging market (EM) currencies. In Figure 3.5.2, we see that

⁴Given the *instrument bid-ask spreads* for ATM, RR and BF, one cannot uniquely determine the corresponding *vanilla spreads* without specifying some rule. For example, in practice, trading desks may estimate vanilla spreads only using ATM spreads, which makes the spread of each option at the same expiry equal, see Section 4.2.1 of [148]. Since vanilla IVs are linear transformations of instrument IVs, we conservatively assume that vanilla spreads are weighted sums of instrument spreads. This does not take into account that delta-symmetric vanilla spreads are dependent on each other, and generates the widest possible bid-ask spreads for vanilla IVs.

there are more arbitrages in the EM currency markets. We also see persistent clustering of (mild) arbitrages from early 2007 to mid 2012 in major currency markets. Further investigation suggests that these are caused by over-priced 1-day options, which result in calendar arbitrages with longer-dated options. We conjecture that the systematic appearance of the same type of arbitrage is due to Bloomberg's legacy data cleansing method.




Figure 3.5.2: Time series of number of daily violated arbitrage constraints in OTC FX option market, during the period from 1st November, 2007 to 31st May, 2018.

Calendar arbitrage (especially CVS C5 and CBS C6) is more difficult and costly to exploit than non-calendar arbitrages, as it requires rebalancing the hedging portfolio over time. Most arbitrage-free smoothing algorithms in the literature only remove calendar arbitrage of C4 type, because they assume a rectangular grid of expiries and strikes. However, calendar arbitrage can be a major source of arbitrage. In Figure 3.5.3, we consider what fraction of the arbitrages are of calendar type for different currency pairs. Comparing medians (and overall distributions), as shown in the plot, the proportion of calendar arbitrages for major currencies (AUD, EUR, GBP, CAD, CHF, and JPY) is larger than that for EM currencies (BRL, KRW, and MXN), though the cross pair EURGBP is an exception. In fact, the medians are very close to 100% for almost all major currencies except sterling. In other words, nearly all arbitrages in major currencies' option markets are calendar ones.

We examine the day when the EURUSD option price data have the most occurrences of calendar arbitrages over our observation period, and plot the call price curves for the first three expiries in Figure 3.5.4. There is no non-calendar arbitrage on that day since each




Figure 3.5.3: Fraction of calendar arbitrages on a given day, for different currency pairs during the period from 1st November, 2007 to 31st May, 2018. The light blue shadow is a violin plot which indicates the kernel density of the percentages, and the red notched box is a box plot. The horizontal short bar shows the median of each sample.

curve is non-increasing and convex. After the repair, the T_1 -curve is pushed downwards until it does not lie beyond the other two curves, which ensures NDCO marginal risk-neutral measures.




Figure 3.5.4: An example of arbitrage repair for EURUSD call options on 2nd April, 2015. *Left* – raw call price curves for the first three expiries. *Middle* – repaired arbitrage-free call price curves. *Right* – perturbations added to each data point.

However, when calendar and non-calendar arbitrages are mixed, the perturbations added to ensure no arbitrage tend to be more varied in signs. For instance, in Figure 3.5.5 we plot the call price curves for the first four expiries on the day when USDBRL options had the most occurrences of calendar arbitrage, however, there are also many non-calendar arbitrages. Unlike the above EURUSD example, the repair does not simply translate any curve. Therefore, the perturbations are not systematically negative.

3.5.2 Properties of the repair method

Sparse solution of the ℓ^1 -norm objective

The ℓ^1 -norm objective leads to sparse perturbations. We show the fraction of perturbed prices in Figure 3.5.6. The medians are very close to zero for all currency pairs, indicating




Figure 3.5.5: An example of arbitrage repair for USDBRL call options on 28th October, 2008.

that very few data points need to be perturbed on average to remove arbitrage. This is especially true for major currencies, as their distributions collapse almost entirely to zero.




Figure 3.5.6: Number of perturbed prices as a percentage of all prices, for different currency pairs during the period from 1st November, 2007 to 31st May, 2018.

Computational time

Our data repair method is designed to be fast due to the LP formulation. In addition, the reduction of arbitrage constraints shrinks the scale of the LP and speeds up the repair. We investigate the computational time of our repair method when applied to a few practical cases. All of the following studies were carried out on a quadcore Intel Core i7-8650U CPU with 32GB RAM. All LPs are solved using the GLPK (GNU Linear Programming Kit) solver wrapped by the CVXOPT [114] Python package.

In Figure 3.5.7, we plot histograms of (1) the number of constraints $R \sim \mathcal{O}(m^2N)$, (2) the fraction of violated constraints, and (3) the elapsed times for constructing the constraints (Table 2.1) and solving the LP (3.2). We take EURUSD and USDBRL as representatives of major currencies and EM currencies, respectively. The number of constraints rarely exceeds 4000, while it takes less than 0.4 seconds to transform them to the matrix form in most cases. Solving the LP with $N = 117$ variables and $R < 4000$ constraints takes less than 0.05 seconds for EURUSD options and less than 0.1 seconds for USDBRL options. More violated constraints lead to higher computational time.

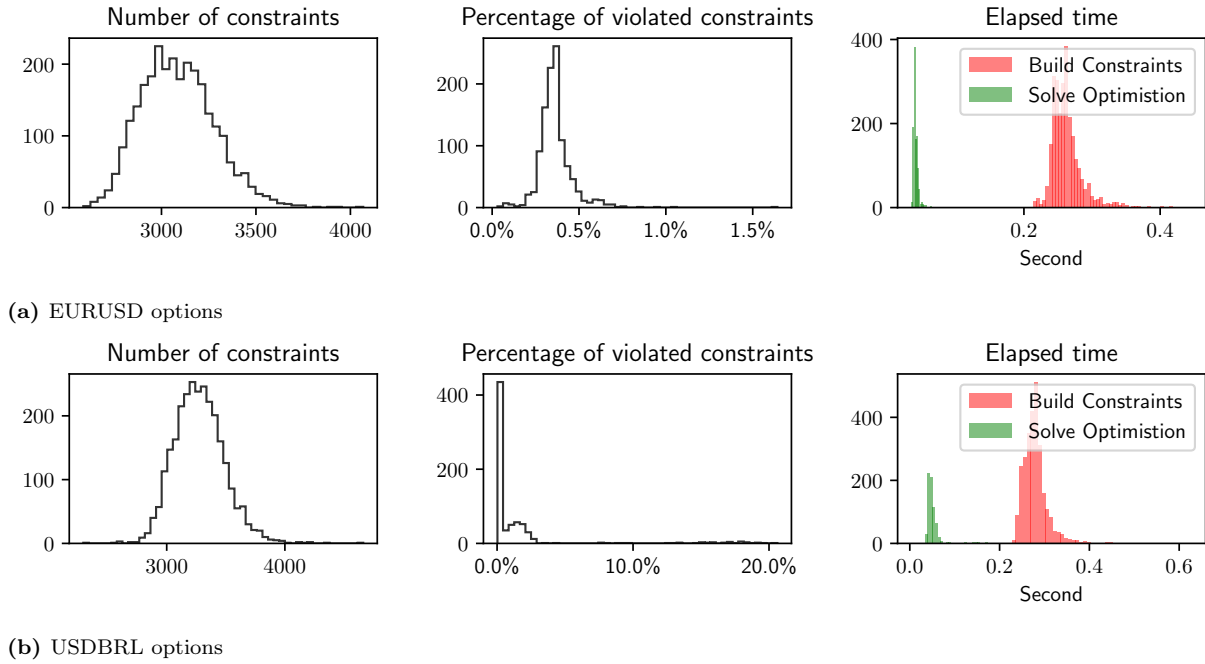


Figure 3.5.7: Histograms of various statistics for repairing data of EURUSD options and USDBRL options, during the period from 1st November, 2007 to 31st May, 2018.

To see how the repair method performs for larger-scale problems, we collect daily settled prices for all *traded* EURUSD call options listed by CME from 1st January, 2013 to 31st December, 2018. The number of traded options varies from one day to another, see Figure 3.5.8a. We show the distribution of traded expiry and strikes on a typical day in Figure 3.5.8b. In Figure 3.5.8c, we plot similar repair statistics to those for the OTC data. On average, there are 500 call prices per day, which result in on average 25000 arbitrage constraints to verify. Though the number of constraints is observed as high as 90000, it takes less than 1 second to construct them in the matrix form. Solving the LP now can take up to 6 seconds, but on average it only takes 1.44 seconds.

Stress testing the ℓ^1 -norm objective repair

We test how our repair method works in hypothetical extreme scenarios when there is massive arbitrage. First, we collect arbitrage-free call prices for a day, denote these data by $\mathbf{c} \in \mathbb{R}^N$, and let $\mathcal{I} = \{1, \dots, N\}$ be the set of its indices. Next, we simulate noises and add them to a portion $\lambda \in (0, 1]$ of the price data, where we denote $\mathcal{I}_\xi \subset \mathcal{I}$ as the set of indices of those polluted prices. Here, \mathcal{I}_ξ is randomly sampled without replacement such that $|\mathcal{I}_\xi| = \lceil \lambda |\mathcal{I}| \rceil$. Constructing the noise $\boldsymbol{\xi} = (\xi_j) \in \mathbb{R}^N$ by taking $\xi_j = \zeta_j \mathbb{1}_{\{j \in \mathcal{I}_\xi\}}$ where $\boldsymbol{\zeta} = (\zeta_j)$ are i.i.d., we then define the noisy price $\tilde{\mathbf{c}} \in \mathbb{R}^N$ by

$$\tilde{c}_j = c_j e^{\xi_j}, \quad \forall 1 \leq j \leq N.$$

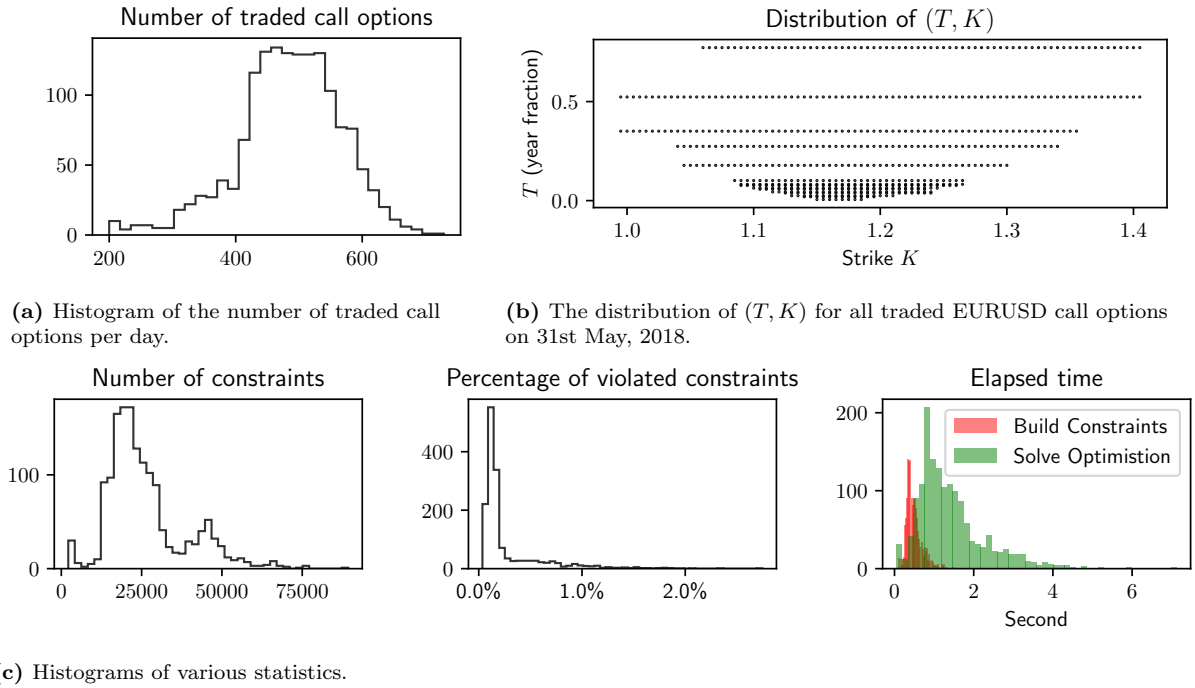


Figure 3.5.8: Statistics for repairing data of CME-listed EURUSD options during the period from 1st January, 2013 to 31st December, 2018.

The noisy price vector contains arbitrage in general, which we repair by seeking an optimal perturbation $\varepsilon \in \mathbb{R}^N$. The perturbed arbitrage-free price vector is $\hat{\mathbf{c}} = \tilde{\mathbf{c}} + \varepsilon$. An example of \mathbf{c} , $\tilde{\mathbf{c}}$, and $\hat{\mathbf{c}}$ is given in Figure 3.5.9.




Figure 3.5.9: An example used for stress testing the repair method. There are $N = 591$ prices in total. The data used are CME-traded EURUSD options' prices on 31st May, 2018.

We assess how well \mathbf{c} is recovered by examining the quantities $\ln(\hat{c}_j/c_j)$ and $\hat{\lambda} := \frac{1}{N} \sum_{j=1}^N \mathbb{1}_{\{\hat{c}_j \neq c_j\}}$. Note that $\ln(\hat{c}_j/c_j) \approx (\hat{c}_j - c_j)/c_j$ (when $\hat{c}_j \approx c_j$). In addition, $\hat{\lambda}$ counts the portion of different prices per call surface. It is unrealistic to expect any repair method to fully recover a price as it is unlikely to know the exact marginal that generates the price. However, given the ground truth that a portion λ of the surface prices has been polluted by noise, a desirable data repair method should leave as many unpolluted prices unchanged as possible, i.e. $\hat{\lambda} - \lambda$ should be small.

Assuming Gaussian noises $\zeta \sim \mathcal{N}(\mathbf{0}, \sigma_\xi I)$, we simulate noises M times, and compute the average value of $\hat{\lambda}$ and plot the histograms of $\ln(\hat{c}_j/c_j)$, conditional on non-zero values, as shown in Figure 3.5.10. For a fixed noise magnitude σ_ξ , the gap between $\hat{\lambda}$ and λ widens as λ increases, i.e. the repair method adjusts a larger number of prices to remove arbitrage if there are more noisy prices. The same observation holds for different values of σ_ξ , though larger noise magnitude σ_ξ results in more arbitrages. Note that taking $\sigma_\xi = 1$ and $\lambda = 25\%$ already results in, on average, $\hat{\lambda} = 30.80\%$ of the price data being perturbed, an extremely large fraction that has rarely been seen in our data, see Figure 3.5.6. Hence, in practice our repair method seems to only perturb a few additional (i.e. $\hat{\lambda} - \lambda \approx 5\%$ for $\lambda = 25\%$) prices to ensure no arbitrage.

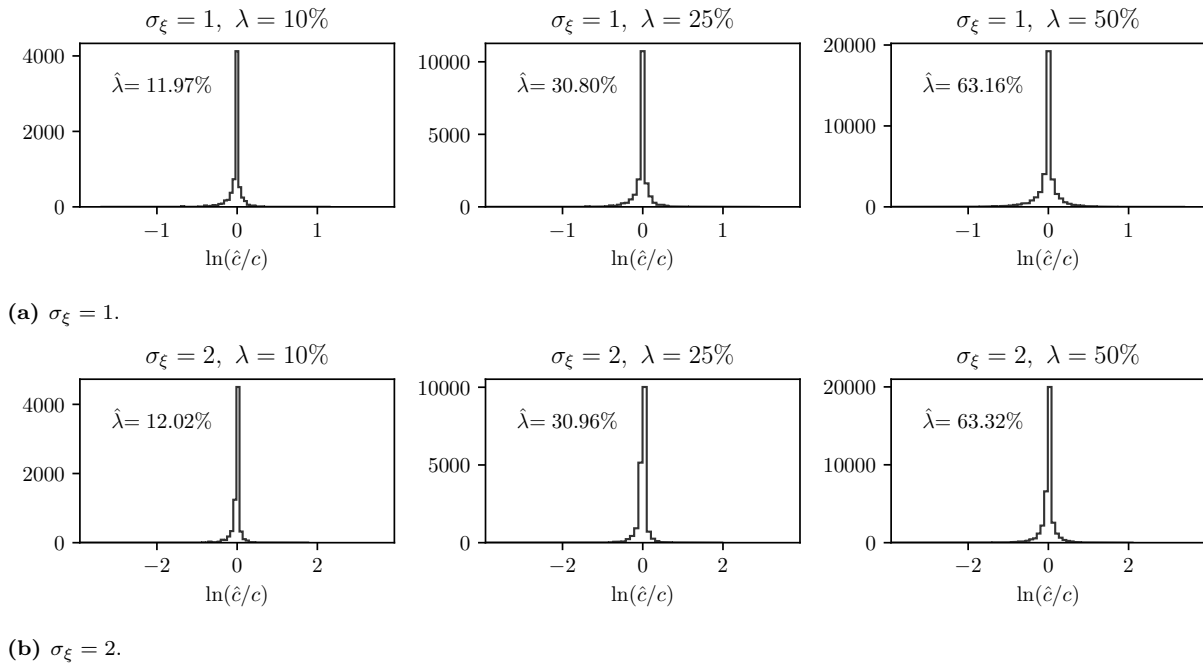


Figure 3.5.10: Histograms of $\ln(\hat{c}_j/c_j) \approx (\hat{c}_j - c_j)/c_j$, conditional on non-zero values, computed under differently valued noise simulation parameters (λ, σ_ξ). We simulate $M = 100$ times.

Comparing the objectives: ℓ^1 -norm and ℓ^1 -BA

The ℓ^1 -BA repair is designed to perturb more prices (larger N^ε) than the ℓ^1 -norm repair does, but fewer of them are effective (smaller $N^{\varepsilon, \delta}$), if possible. To verify this, we apply the ℓ^1 -BA repair method to the same OTC FX option price data. In Figure 3.5.11, we show the histograms of the difference in these two statistics N^ε and $N^{\varepsilon, \delta}$ that are produced by the two repair methods.

A detailed example showing how the two repair methods work in reality is given in Figure 3.5.12. From left to right, the displayed data are ordered by increasing strikes, grouped by expiry. The light blue areas are confined by bid-ask spread as a percentage of

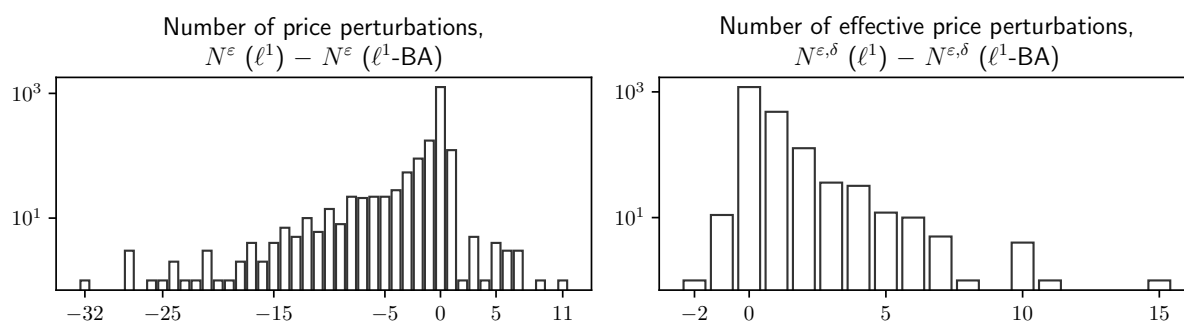


Figure 3.5.11: Histograms of the difference in N^ε and $N^{\varepsilon,\delta}$ that are produced by the ℓ^1 -norm repair method and the ℓ^1 -BA repair method. These two methods are separately applied to the same set of OTC FX data as in Figure 3.5.2, and the histograms are plotted by stacking data of all ten currency pairs and historical dates.

option prices (green lines). We see that ITM and OTM options have wider bid-ask spreads than ATM options do. The ℓ^1 -BA repair method results in fewer effective perturbations. First, there is one less effective perturbation of 1M option prices, at the cost of perturbing a few 2W, 3W and 1M option prices to their bid or ask prices. Second, all four effective perturbations of 4M option prices by the ℓ^1 -norm repair are replaced by six ineffective perturbations of 6M option prices by the ℓ^1 -BA repair.

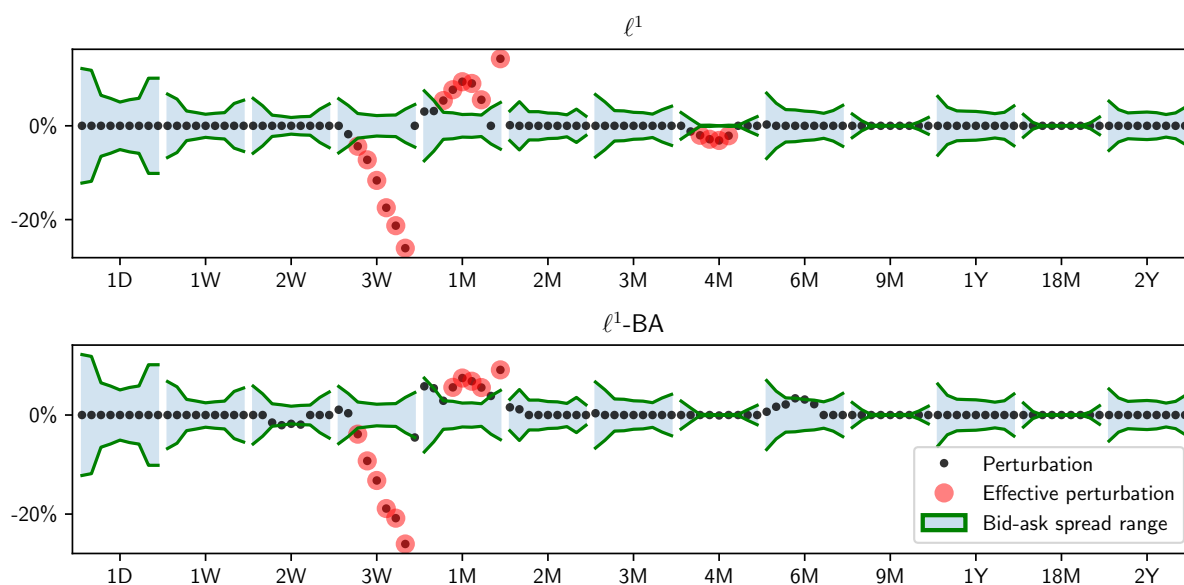


Figure 3.5.12: Perturbations (as percentages of the raw price data) resulted from the ℓ^1 -norm and the ℓ^1 -BA objectives. Data used are bid, ask and mid prices for OTC-traded USDBRL options on 18th September, 2008.

For a given set of prices, if none of the perturbations is effective, then the bid and ask quotes given by the market admit some arbitrage-free prices that fall within the bid-ask price bounds. In contrast, effective perturbations imply the existence of *executable*

arbitrages that are exploitable through matching existing bid or ask orders in the market, see Section 3.4. In Table 3.5.1, we count the number of days when there is arbitrage in mid-prices ($N^\varepsilon > 0$) and the number of days when there is executable arbitrage ($N^{\varepsilon, \delta} > 0$) in historical data for the four currency pairs that have been seen to have the most occurrences of arbitrages.

Currency pair	EURGBP	USDBRL	USDKRW	USDMXN
#days when $N^\varepsilon > 0$	470	708	623	577
#days when $N^{\varepsilon, \delta} > 0$	285	89	163	144

Table 3.5.1: Number of days when there is arbitrage in mid-prices ($N^\varepsilon > 0$) and when there is executable arbitrage ($N^{\varepsilon, \delta} > 0$).

3.5.3 Improvement to model calibration

We verify that our repair method improves model calibration with more robust parameter estimates and smaller calibration error.

Test framework

Let Θ be model parameters. We specify $\Theta = \bar{\Theta}$ and generate model prices \mathbf{c} for call options on a set of expiries and strikes. Then we carry out the following steps M times. For the m -th time:

- (i) Simulate noises to create synthetic arbitrageable price data $\tilde{\mathbf{c}}^{(m)}$, following the method in Section 3.5.2. Recall that $\lambda \in (0, 1]$ portion of prices are polluted by Gaussian noises of variance σ_ξ^2 .
- (ii) Repair arbitrage in $\tilde{\mathbf{c}}^{(m)}$ to get arbitrage-free data $\hat{\mathbf{c}}^{(m)}$.
- (iii) Calibrate model parameters Θ to $\tilde{\mathbf{c}}^{(m)}$ and $\hat{\mathbf{c}}^{(m)}$ separately⁵, and get calibrated parameters $\tilde{\Theta}^{(m)}$ and $\hat{\Theta}^{(m)}$, respectively. Defining the calibration objective as $G(\Theta; \mathbf{c}) = \sum_{j=1}^N (c_j^\Theta - c_j)^2$ where c_j^Θ is the model price for the j -th option, we have

$$\tilde{\Theta}^{(m)} = \arg \min_{\Theta} G(\Theta; \tilde{\mathbf{c}}^{(m)}), \quad \hat{\Theta}^{(m)} = \arg \min_{\Theta} G(\Theta; \hat{\mathbf{c}}^{(m)}).$$

We measure model calibration performance by two metrics, which are (a) the *robustness* defined by variations in the parameter estimates, and (b) the calibration error defined as the square root of the minimal objective value. Since we have parameter

⁵Note that we must apply exactly the same numerical procedure for these two separate calibrations, i.e. the same optimisation algorithm, terminal criteria, lower and upper bounds, and initial values.

estimates $\{\tilde{\Theta}^{(m)}\}_{1 \leq m \leq M}$ and $\{\hat{\Theta}^{(m)}\}_{1 \leq m \leq M}$, we can compare the variations in them for assessing robustness. For each m , we define the (relative) reduction of calibration error as

$$\Delta G^{(m)} = 1 - \sqrt{\frac{G(\hat{\Theta}^{(m)}; \hat{\mathbf{c}}^{(m)})}{G(\tilde{\Theta}^{(m)}; \tilde{\mathbf{c}}^{(m)})}}.$$

Heston model calibration

We carry out a test on calibration of the Heston model [85]. Recall that the Heston model is described by the SDEs with model parameters $\Theta = (\nu_0, \theta, k, \sigma, \rho)$:

$$\begin{aligned} dS_t &= r_t S_t dt + \sqrt{\nu_t} S_t dW_t^S, \\ d\nu_t &= k(\theta - \nu_t) dt + \sigma \sqrt{\nu_t} dW_t^\nu, \quad d\langle W_t^S, W_t^\nu \rangle = \rho dt, \end{aligned}$$

where the Feller condition $2k\theta > \sigma^2$ is sufficient to ensure strict positivity of the instantaneous variance process ν_t .

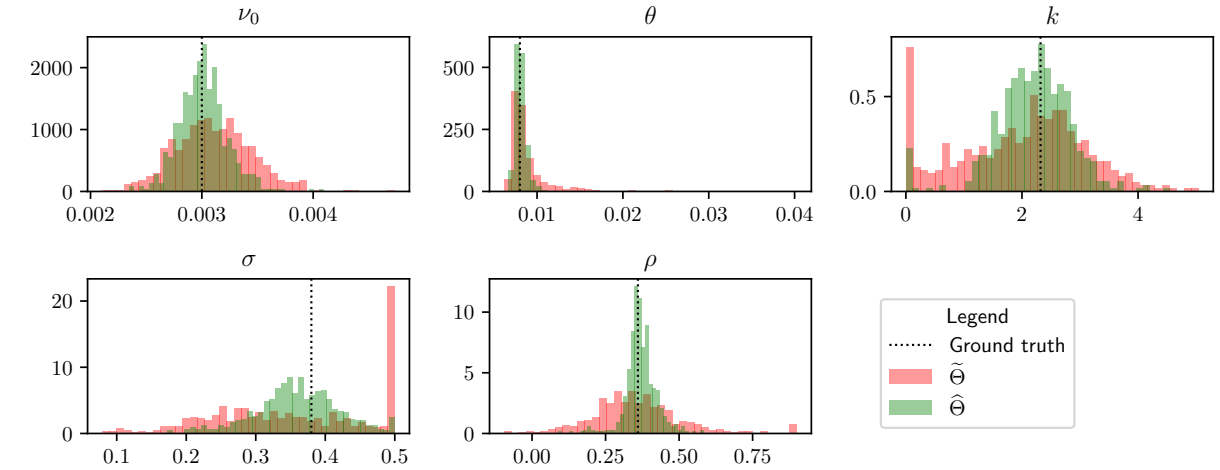
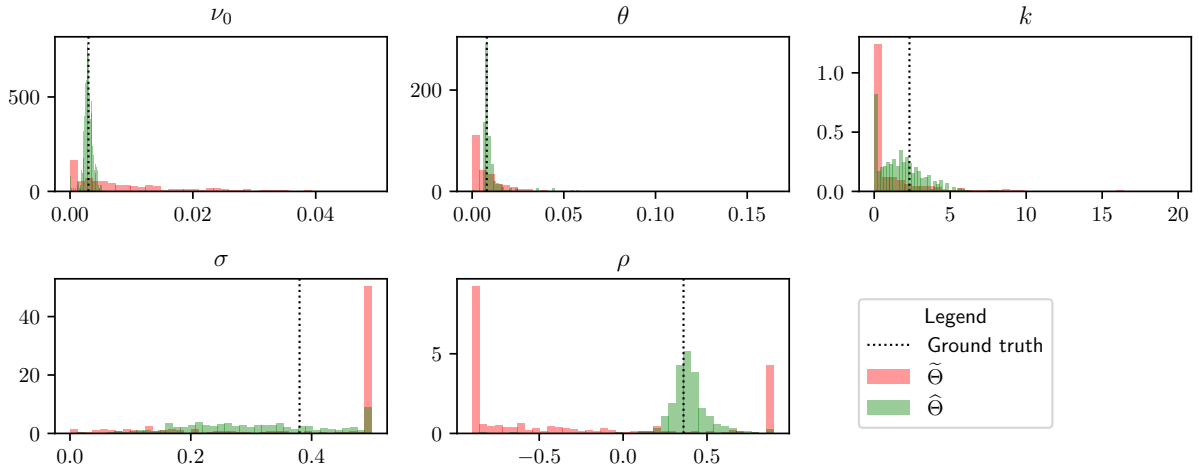
We specify a typical set of expiries and strikes that is observed on a day in the OTC market, such as the one shown in Figure 3.5.1. Other simulation parameters and ground truth model parameters⁶ are listed in Table 3.5.2.

Parameter	Heston model					Simulation			
	ν_0	θ	k	σ	ρ	N	M	λ	σ_ξ
Value	0.003	0.008	2.32	0.38	0.36	117	500	0.25	0.1 or 1

Table 3.5.2: Parameter values

Next, we follow the test framework and evaluate $\tilde{\Theta}^{(m)}$, $\hat{\Theta}^{(m)}$, $\Delta G^{(m)}$ and $\Delta t^{(m)}$ for $m = 1, \dots, M$. In Figure 3.5.13, we plot and compare the normed histograms of calibrated Heston parameters $\tilde{\Theta}$ (using noisy data) and $\hat{\Theta}$ (using repaired data) given different choices of σ_ξ . The ground truth parameter values are also indicated by vertical dotted lines. Repairing data does make the model calibration more robust, as supported by two types of evidence. First, there are apparently more variations in $\tilde{\Theta}$ than in $\hat{\Theta}$. Second, $\tilde{\Theta}$ tends to hit the bounds set in the numerical optimisation procedure (e.g. 0 for k , 0.5 for σ , 1 for ρ) much more often than $\hat{\Theta}$ does. Moreover, when the price data are more noisy (larger σ_ξ) so that more prices with arbitrage are present, the robustness improvement of model calibration by the repair method becomes more significant.

⁶Heston model parameters are chosen as those that reproduce a typical call price surface for USDBRL options. Noise simulation parameters λ and σ_ξ are chosen to mimic severe but not extreme arbitrage scenarios (measured by the fraction of perturbed prices by the repair method) observed in real world data.

(a) $\sigma_\xi = 0.1$.(b) $\sigma_\xi = 1$.**Figure 3.5.13:** Sample (normed) histograms of $\tilde{\Theta}$ and $\hat{\Theta}$, where $\Theta = (\nu_0, \theta, k, \sigma, \rho)$.

In Figure 3.5.14, we plot the histograms of ΔG and indicate their means by vertical dotted lines. Repairing arbitrage in data reduces the calibration errors in all M simulations with no exception. Moreover, the more noisy the raw data are, the arbitrage repair method reduces relatively more calibration errors. On average, repairing data can reduce the calibration error by more than 70% for $\sigma_\xi = 0.1$, and more than 95% for $\sigma_\xi = 1$.

Hence, for model calibration task, there is more benefit of repairing data by removing arbitrage when the data contain larger noise.

3.5.4 Identifying intra-day executable arbitrage

We can use the ℓ^1 -BA repair method on order book data for identifying executable arbitrage. An example is given in Figure 3.5.15. We collect the order book data for all E-mini

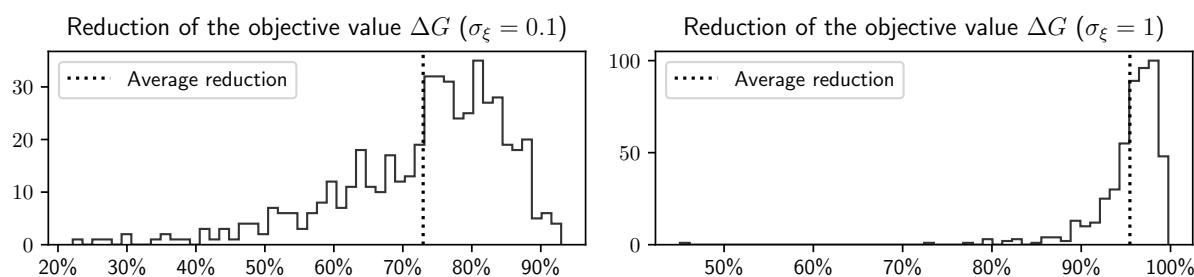


Figure 3.5.14: Sample histograms of (relative) reductions in calibration error.

S&P 500 monthly European call options from 12:00 ET to 16:10 ET on 12th June, 2020.




Figure 3.5.15: *Top* – The formation and disappearance of intra-day executable arbitrage opportunities in the E-mini S&P 500 monthly European call option market on 12th June, 2020. *Bottom* – front-month futures’ prices and 15-minute return.

We extract the active best ask and best bid prices for all quoted call options from the order book at the end of every minute. Then we compute mid prices, apply the ℓ^1 -BA repair method to the mid prices, and count the number of effective perturbations $N^{\varepsilon, \delta}$. Recall from Section 3.4.2 that, given δ_0 small, there exists executable arbitrage if the ℓ^1 -BA repair method results in effective perturbations. In the top plot of Figure 3.5.15, the black line gives $N^{\varepsilon, \delta}$ over time, while we also indicate the ATM implied volatility of the front-month option (which would expire on 19th June) by the grey line. The bottom plot gives the prices and the 15-minute returns of the front-month futures contract. The downward trend of the futures market was inverted around 14:00 ET, after when the implied volatility also falls gradually from its peak.

There is a large spike of $N^{\varepsilon, \delta}$ at around 15:52 ET, a few minutes before the close of the S&P 500 index market at 16:00 ET. This spike coincided with rallies in the futures market, while the IV maintained its relatively low level. There are some clusters of smaller spikes

of $N^{\varepsilon, \delta}$ outside of the US trading hours. Apart from this arbitrage outbreak preceding the close of the underlying market, which lasted for around 15 minutes, there seems to be trivial executable arbitrage during the rest of the afternoon trading hours, even when the market underwent regime switch (from downward trend to upward trend) at around 14:00 ET.

Arbitrage-free neural-SDE market model

“尝一脔肉，知一镬之味；悬羽与炭，而知燥湿之气；以小明大。见一叶落，而知岁之将暮；睹瓶中之冰，而知天下之寒；以近论远。”

— 刘安《淮南子·说山训》

“By tasting one piece of meat, you can know the flavor of a potful; By suspending feathers and charcoal [in a balance beam], you can know the humidity of the air. One uses the small to illuminate the large. By seeing one leaf fall, you can know that the year will soon end; By noticing ice in a jug, you can know the temperature throughout the world. One uses the near to assess the far.”

— An Liu, *A Mountain of Persuasions, Huainanzi* (translated by John S. Major et al.)

4.1 Introduction

As already outlined in Chapter 1, we consider a financial market where the following assets are liquidly traded: a stock S and a collection of European call options $C(T, K)$ on S with various expiries $T \in \mathcal{T}^e$ and strikes $K \in \mathcal{K}$. We aim to construct a class of models for the stock and options that should (i) permit no arbitrage, (ii) allow exact cross-sectional calibration, and (iii) reflect stylised facts observed from market price dynamics. Importantly, it should be practically convenient to estimate these models, given that observations are discrete time series of prices for a large but finite collection of options.

After comparing the pros and cons of the martingale approach and the market model approach (see Section 1.2), in terms of achieving the abovementioned properties, and acknowledging the low-dimensional nature of statically arbitrage-free call price surface (see Section 1.3), we will focus on constructing factor-based market models for finitely many options that are liquidly traded. In addition, we aim to exploit the recent successes

in the use of neural networks as function approximators in order to give a flexible class of models. Nevertheless, we distinguish our use case from existing literature on using neural networks for finance applications as neural-SDE market models (see Section 1.4). In particular, the no-arbitrage constraints on option prices imply that we need to fit neural nets where the resulting behaviour (when used in an SDE) will satisfy a family of linear (in)equalities. In this chapter we will develop methods to solve this challenge, and demonstrate their effectiveness when building a financial model.

Imposing no-arbitrage conditions is a key component in our models. To achieve this, our models first ensure that for each fixed time, prices take values in a constrained state space such that options with different strikes (moneynesses) and expiries are *statically* arbitrage-free. Second, we aim both to guarantee the statistical accuracy of our model (by training with historical data) and to minimise the opportunities for dynamic arbitrage – the discounted price process of each traded asset, i.e. stock and all options, should be a martingale under *some* joint risk-neutral measure (which is not generally the measure from which our training data are sampled). This involves compatibility restrictions on the model coefficients, analogous to the HJM [82] drift restrictions.

There is a useful distinction to be drawn, in this context, between “model-free” and “model-based” arbitrage. If a market model produces prices which display model-free arbitrages, then these are easily exploited, as they do not depend on the dynamic model used. A counterparty simply needs to observe prices inconsistent with a set of restrictions (in practice given by a finite set of linear inequalities), and can quickly identify arbitrage opportunities. On the other hand, if a market model only has the risk of producing model-based arbitrage, then this may be less of a concern; to exploit this opportunity, a counterparty would typically need to know the model being used (including all conventions around interpolation of prices). Indeed, if no model-free arbitrages are present, then it is known that there exists a dynamic model which is arbitrage free and replicates these prices (see, for example, Carr and Madan [39], Davis and Hobson [56]).

Conveniently, the model-free arbitrage restrictions correspond to *static* arbitrage constraints, that is, to restrictions on the state-space of the price process. In Chapter 2 (drawing on Cousot [52]), we have reduced these constraints, for a collection of arbitrary strikes and maturities, to an efficient set of linear inequalities. In this thesis we will focus our attention, therefore, on estimating (under the historical measure \mathbb{P}) a market model where the model-free arbitrage constraints will be satisfied. We will see that this restriction of the state-space of our model yields significant benefits to the task of statistical calibration.

To summarise, our models are friendly for implementation for the following reasons:

- (i) They are flexible enough to model options on an arbitrarily-shaped lattice of strikes and expiries. One can specify the lattice according to the stylised liquidity profile of the target option market.
- (ii) The model inference procedure includes both cross-sectional calibration (i.e. the factor decoding step) and time-series estimation (i.e. the neural-SDE calibration step). Neural-SDEs enable a data-driven model selection approach, and mitigate model risk arising from restrictive parametric forms.
- (iii) Further economic considerations (for example bounds on statistical arbitrage) can be incorporated into the training of these models through penalization.

4.2 Assumptions

We assume a frictionless market in continuous-time, with a distant fixed finite horizon T^* . We allow non-zero interest and dividends¹, and assume a deterministic interest rate r_t and dividend yield q_t at time t . We use $D_t(T) = \exp(-\int_t^T r_s ds)$ to denote the market discount factor for time T , and $\Gamma_t(T) = \exp(\int_t^T q_s ds)$ to account for reinvestment of dividends. There is a model-independent, arbitrage-free forward price, $F_t(T) = S_t/(\Gamma_t(T)D_t(T))$, for delivery of the asset at T , where S_t is the price of the underlying asset. The primary traded securities in the market are zero coupon bonds, the stock, and a collection of forwards and European call options written on the stock. We assume no derivatives expire after T^* .

4.3 Liquid options

While it is theoretically reasonable to search for a model to produce arbitrage-free option prices for every possible T and K at any time t , doing so can be both practically unnecessary and restrictive. It is of limited use to ensure absence of arbitrage with regards to deep in or out-of-the-money options, given these are not liquidly traded, and modelling these would impose significantly more constraints when calibrating to data. Our models focus on the joint dynamics of *liquid* options, rather than options of *any* strikes and expiries. Liquid options are usually those that are neither too far from the money nor too close to expiry (typically a few days). In addition, the range of liquid strikes typically broadens for longer expiries, as seen in Figure 1.1.

¹When applying our methods to other asset classes, dividends of stocks are comparable to foreign currency interest for FX, or convenience yield for commodities.

Over time, the range of liquid options in (T, K) -coordinates tends to change *stochastically*. For example, a liquid at-the-money option as of today will become deep out-of-the-money if the underlying price plunges quickly. It is empirically advantageous, therefore, to cast the liquid range into relative coordinates (τ, m) , where $\tau = T - t$ is time-to-expiry and m denotes moneyness. There is usually a stable range of time-to-expiries and moneynesses for which the options are actively quoted and therefore price data are most readily available. A convenient parameterisation is given by the *forward log-moneyness* m_t , which is

$$m_t = M(K; F_t(T)) := \ln \left(\frac{K}{F_t(T)} \right). \quad (4.1)$$

The function $M(\cdot)$ is referred to as the moneyness function. In Figure 4.3.1, we highlight the trajectories of (τ_t, m_t) for three differently struck options (with the same expiry). The option that is very far from the money, i.e. $K = 1.3$ (green line), moved out of the liquid range (which could be approximated by the gray dotted area) and was never traded after its time-to-expiry became shorter than 0.4 years.




Figure 4.3.1: An example of the trajectory of (τ_t, m_t) for fixed (T, K) 's. Grey dots indicate the traded strikes for CME EURUSD option contracts with expiry date 6th March, 2020, from their first listing day until expiry. We highlight the trajectories for three differently struck contracts.

We assume the range of liquid options in the (τ, m) -coordinates is fixed over time, and develop arbitrage-free models on the the liquid range. We denote by $\underline{\tau}$ and $\bar{\tau}$ the minimal and maximal liquid time-to-expiries. For fixed τ , we denote by $\underline{m}(\tau)$ and $\bar{m}(\tau)$ the minimal and maximal liquid moneynesses. We then define the constant set $\mathcal{R}_{\text{liq}} = \{(\tau, m) \in \mathbb{R}^2 : \tau \in [\underline{\tau}, \bar{\tau}], m \in [\underline{m}(\tau), \bar{m}(\tau)]\}$ as the range of moneynesses and time-to-expiries where options are liquid. A typical market's range exhibits that $0 < \underline{\tau} < \bar{\tau} < T^*$; $\underline{m}(\tau) < 0 < \bar{m}(\tau)$ for all $\tau \in [\underline{\tau}, \bar{\tau}]$; $\underline{m}(\tau_1) \geq \underline{m}(\tau_2)$ and $\bar{m}(\tau_1) \leq \bar{m}(\tau_2)$ for any $\underline{\tau} \leq \tau_1 < \tau_2 \leq \bar{\tau}$. An example can be seen in Figure 4.3.2.

Our models construct option price surfaces on \mathcal{R}_{liq} through smooth interpolation of a finite collection of option prices, with parameters given by the set $\mathcal{L}_{\text{liq}} = \{(\tau_i, m_j^i) \in \mathcal{R}_{\text{liq}} :$




Figure 4.3.2: An example of liquid range and lattice.

$\underline{m}(\tau_i) = m_1^i < \dots < m_{n_i}^i = \overline{m}(\tau_i), \underline{\tau} = \tau_1 < \dots < \tau_p = \overline{\tau}, 1 \leq j \leq n_i, 1 \leq i \leq p\}$. We call \mathcal{L}_{liq} an *option lattice*. There are two reasons for us to have this model design. First, given the static arbitrage requirements concern monotonicity and convexity of the price surface, our models are free of static arbitrage on \mathcal{R}_{liq} if options on \mathcal{L}_{liq} are statically arbitrage-free and these prices are extended to \mathcal{R}_{liq} using a shape-preserving interpolation method (in particular, a method which preserves monotonicity in time to maturity and monotonicity and convexity in moneyness). Second, options with very close strike and expiry will have strongly dependent price processes, and so contain little new statistical information. We use options on \mathcal{L}_{liq} to represent local dynamics of different sub-areas of the surface on range \mathcal{R}_{liq} ; this is analogous to a common practice for quoting OTC data. For example, in the OTC FX market, traders quote options on the lattice [1-day, 1-week, 1-month, ...] \times [10-delta, 15-delta, ...] and smoothly interpolate them for pricing a specific contract.

4.4 Factor representation of call option prices

Even after reducing our model to the finite set \mathcal{L}_{liq} of liquid options, one usually finds that these prices contain significant redundancies. For this reason, it is efficient to express them through a low-dimensional factor representation, for which a reliable model can be built. These factors will be built from data, but are simple enough to be intuitively understood and plotted, as they have a linear relationship with option prices.

We will build a model based on the randomness provided by $d + 1$ standard independent Brownian motions W_0, W_1, \dots, W_d under the objective measure \mathbb{P} . Writing $W = [W_1 \dots W_d]^\top$ and $\overline{W} = [W_0 \ W_1 \ \dots \ W_d]^\top$, we describe the information available at time t by the σ -algebra \mathcal{F}_t , which is the completion of $\cap_{s>t} \mathcal{F}_s^0$ with \mathbb{P} -null sets, where

$\mathcal{F}_t^0 := \Sigma(\{\overline{W}_s : 0 \leq s \leq t\}) \vee \mathcal{F}_0$ (here we use $\Sigma(\cdot)$ to denote the sigma-algebra generator). Taking Ω to be the space of paths of \overline{W} , we have a filtered probability space $(\Omega, \{\mathcal{F}_t\}_{0 \leq t \leq T^*}, \mathbb{P})$.

Suppose that the market is described by the underlying price $S \in \mathbb{R}$ and some latent *market factor* $\xi \in \mathbb{R}^d$. We model the joint dynamics of the underlying price and the market factor by a $(d+1)$ -dimensional diffusion process solving the following time-homogeneous SDE:

$$\begin{cases} dS_t = (\alpha(\tilde{\xi}_t) - q_t) S_t dt + \gamma(\tilde{\xi}_t) S_t dW_{0,t}, & S_0 = s_0 \in \mathbb{R}; \\ d\xi_t = \mu(\tilde{\xi}_t) dt + \sigma(\tilde{\xi}_t) dW_t, & \xi_0 = \zeta_0 \in \mathbb{R}^d, \end{cases} \quad (4.2)$$

where $\tilde{\xi}_t = (S_t, \xi_t)$. We use α_t , γ_t , μ_t and σ_t to denote $\alpha(\tilde{\xi}_t)$, $\gamma(\tilde{\xi}_t)$, $\mu(\tilde{\xi}_t)$ and $\sigma(\tilde{\xi}_t)$, respectively. We denote by $L_{\text{loc}}^p(\mathbb{R}^d)$ the space of all \mathbb{R}^d -valued, progressively measurable, and locally p -integrable (in t , \mathbb{P} -a.s.) processes on $[0, T^*]$, and assume that $\alpha S \in L_{\text{loc}}^1(\mathbb{R})$, $\mu \in L_{\text{loc}}^1(\mathbb{R}^d)$, $\gamma S \in L_{\text{loc}}^2(\mathbb{R})$ and $\sigma \in L_{\text{loc}}^2(\mathbb{R}^d)$.

We now associate the latent market factor with call option prices. We denote by $C_t(T, K)$ the market price at time t of an European call option with expiry T and strike K . We call the map $(T, K) \mapsto C_t(T, K)$ a *call price surface*. After casting the surface into relative coordinates (τ, m) , the surface is parametrised over the range $\{(\tau, m) : \tau \in (0, T^* - t], m \in \mathbb{R}\}$. For some fixed T and K , we define the following transformations of call option prices for all $t \in [0, T^*]$:

$$\hat{c}_t(T, K) = \frac{C_t(T, K)}{D_t(T)F_t(T)}, \quad \tilde{c}_t(T - t, M(K; F_t(T))) = \hat{c}_t(T, K), \quad (4.3)$$

where \hat{c}_t and \tilde{c}_t are the *normalised* prices but with different arguments. The parametric surface we shall model is the normalised call price surface $(\tau, m) \mapsto \tilde{c}_t(\tau, m)$ rather than the original call price surface $(T, K) \mapsto C_t(T, K)$, though $C(T, K)$ can be implied from a family of $\tilde{c}(\tau, m)$. A key advantage of $\tilde{c}(\tau, m)$ over $C(T, K)$ is that it is more reasonable to suppose $t \mapsto \tilde{c}_t(\tau, m)$ to be a time-homogeneous process, whereas $C_t(T, K) \rightarrow (S_T - K)^+$ as $t \uparrow T$, making a statistical model more difficult.

For some fixed τ and m , we assume that the normalised option price $\tilde{c}_t(\tau, m)$ is determined by the latent factor ξ through a time-independent linear transformation

$$\tilde{c}_t(\tau, m) = g(t, \tau, m, \xi_t) := G_0(\tau, m) + \sum_{i=1}^d G_i(\tau, m) \xi_{it}, \quad (4.4)$$

where $G_i \in C^{1,2}(\mathcal{R}_{\text{liq}})$ is called a *price basis function*, for $i = 0, \dots, d$. The smoothness requirements on $G_i(\cdot)$ are necessary for deriving the dynamics of $C(T, K)$, for given T and K , from the dynamics of a family of $\tilde{c}(\tau, m)$. Though this representation does not show explicit dependence between the normalised call price $\tilde{c}(\tau, m)$ and the underlying price S , it does relate the two variables through the moneyness m .

Combining (4.2) (4.3) and (4.4), we have a model for the joint dynamics of S and $C(T, K)$ for all relevant strikes and expiries. We observe that the drift and diffusion coefficients of ξ are functions of S . In addition, given the potential nonlinearity in (4.2), we argue that a linear transformation in (4.4) is not a restrictive assumption, as a smooth nonlinear relation with respect to one factor can be approximated by a linear combination of polynomials of that factor.

Remark 4.1. The use of a Markovian model for ξ may seem a little restrictive. In theory, any observed data could be included in the model for ξ . In practice, we shall build factors ξ to minimize reconstruction errors in (4.4). As this involves all liquid option prices, it determines the law of S under \mathbb{Q} in the future (through the Breeden–Litzenberger formula [26]); in other words, the joint process given by S and all call option prices is guaranteed to be Markovian under the pricing measure. This is closely related to the endogenous completeness of market models with options, see Davis and Obłój [55], Schwarz [133] and Wissel [147].

4.4.1 Convergence to terminal payoff

To ensure prices of call options converge to correct payoffs at expiry, independently of the behavior of ξ , by examining (4.4) it is clearly sufficient (and generally necessary) that

$$G_0(0, m) = (1 - e^m)^+, \text{ and } G_i(0, m) = 0, \forall i = 1, \dots, d. \quad (4.5)$$

Under this assumption, $\tilde{c}_t(0, m) = (1 - e^m)^+$ and, for any T and K ,

$$\begin{aligned} \lim_{t \uparrow T} C_t(T, K) &= \lim_{t \uparrow T} D_t(T) F_t(T) \tilde{c}_t \left(T - t, \ln \left(\frac{K}{F_t(T)} \right) \right) \\ &= S_T \tilde{c}_T \left(0, \ln \left(\frac{K}{S_T} \right) \right) = (S_T - K)^+. \end{aligned}$$

4.4.2 Connections with the Black–Scholes model

We can also see that the Black–Scholes model is within the class of factor models we consider (in particular, it is a trivial model where no additional factors ξ are needed). The Black–Scholes model assumes deterministic drift and constant volatility for the underlying asset price, i.e. $\gamma_t \equiv \gamma$. In fact, with $d_{1,2}(\tau, m) = -\frac{m}{\gamma\sqrt{\tau}} \pm \frac{1}{2}\gamma\sqrt{\tau}$,

$$\tilde{c}_t(\tau, m) = G(\tau, m) := \Phi(d_1(\tau, m)) - e^m \Phi(d_2(\tau, m)). \quad (4.6)$$

More generally, we can see that $\tilde{c}(\tau, m)$ has a deterministic bijective relation with the Black–Scholes implied volatility, for arbitrary but fixed τ and m . With implied volatility

$\sigma_t^{\text{imp}}(\tau, m)$, the transformation (4.3) indicates that

$$\begin{aligned}\tilde{c}_t(\tau, m) &= \frac{C_t(t + \tau, e^m F_t(t + \tau))}{D_t(t + \tau) F_t(t + \tau)} \\ &= \Phi\left(-\frac{m}{\sigma_t^{\text{imp}} \sqrt{\tau}} + \frac{1}{2} \sigma_t^{\text{imp}} \sqrt{\tau}\right) - e^m \Phi\left(-\frac{m}{\sigma_t^{\text{imp}} \sqrt{\tau}} - \frac{1}{2} \sigma_t^{\text{imp}} \sqrt{\tau}\right),\end{aligned}\quad (4.7)$$

where $\Phi(\cdot)$ is the cumulative density of the standard normal distribution. In addition,

$$\frac{d\tilde{c}_t}{d\sigma_t^{\text{imp}}}(\sigma_t^{\text{imp}}; \tau, m) = \sqrt{\tau} \phi\left(-\frac{m}{\sigma_t^{\text{imp}} \sqrt{\tau}} + \frac{1}{2} \sigma_t^{\text{imp}} \sqrt{\tau}\right) > 0,$$

where $\phi(\cdot)$ is the density of the standard normal distribution. Hence, there is a time-independent and deterministic bijective relation between \tilde{c}_t and σ_t^{imp} for arbitrary but fixed τ and m . Consequently, by modelling the normalised call prices, we are essentially modelling the implied volatility surface as a function of τ and m . This aligns with the conventional “sticky delta” rule (see Daglish, Hull and Suo [54]), where the process for implied volatility depends on T , K , F and t only through its dependence on τ and m .

4.5 Absence of arbitrage

Arbitrage refers to a costless trading strategy that has zero risk and a positive probability (under \mathbb{P}) of profit. The First Fundamental Theorem of Asset Pricing (FFTAP) establishes an equivalence relation between no-arbitrage and the existence of an equivalent martingale measure (EMM). We work with a simplified version of FFTAP as stated in (2.1). Given the model, there is no arbitrage if and only if $\exists \mathbb{Q} \sim \mathbb{P}$, such that the discounted price processes for all tradable assets are martingales. These tradable assets include the stock and all forwards and options written on it.

If the model is arbitrage-free, then, for *any* fixed T and K , the discounted option price $D_0(t)C_t(T, K)$ should be a \mathbb{Q} -martingale. We derive an HJM-type drift restriction for this no-arbitrage condition to hold in Appendix 4.A, by analyzing the existence of a *market price of risk* process $\psi_t = [\psi_{1,t} \ \cdots \ \psi_{d,t}]^\top$. Specifically, suppose we model $N = |\mathcal{L}_{\text{liq}}|$ options, with prices given in (τ, m) -coordinates. Then we look for a solution ψ_t to

$$\begin{cases} \mathbf{G}^\top \sigma_t \psi_t = \mathbf{G}^\top \mu_t - z_t, & \text{where} \\ z_t = \left(-\frac{\partial}{\partial \tau} - \frac{1}{2} \gamma_t^2 \frac{\partial}{\partial m} + \frac{1}{2} \gamma_t^2 \frac{\partial^2}{\partial m^2}\right) \begin{bmatrix} \tilde{c}_t(\tau_1, m_1) \\ \vdots \\ \tilde{c}_t(\tau_N, m_N) \end{bmatrix}. \end{cases}\quad (4.8)$$

Here, we define \mathbf{G} as the $d \times N$ matrix with i -th row \mathbf{G}_i , where $\mathbf{G}_i = (G_{ij})_{j=1}^N \in \mathbb{R}^N$ for $i = 0, \dots, d$, and $G_{ij} = G_i(\tau_j, m_j)$. We call \mathbf{G}^\top a *price basis* and each of its columns a *price basis vector*.

If z_t lives in the column space of the price basis \mathbf{G}^\top , and given that σ_t is invertible, then (4.8) always admits a solution for ψ_t , regardless of what the drift and diffusion functions are. In other words, the HJM conditions are principally related to the choice of price basis \mathbf{G} , rather than to the specific choice of calibrated model (i.e. to μ, σ).

This encourages us to construct factors from principal components of z_t , as will be explored in detail in Algorithm 1. Unless $z_t \in \mathbb{R}^N$ lives in an \mathbb{R}^d submanifold, equality (4.8) always yields an over-determined linear system when the number of options N is greater than the number of latent market factors d , which is usually the case in practice. Given a price basis \mathbf{G} , we will then use (4.8) to encourage choices of μ and σ that do not suggest unreasonable values for the market price of risk ψ , evaluated on our liquid lattice \mathcal{L}_{liq} .

Remark 4.2. If we modelled *all* τ and m and enforced the drift restrictions at *all* times, then this would be enough to guarantee no-arbitrage. However, our models will be built on finitely many options on a fixed liquid lattice in the (τ, m) -coordinates, as specified by \mathcal{L}_{liq} . These options do not correspond to fixed contracts (i.e. fixed T and K) over time. Therefore, the HJM-type drift restrictions cannot be implemented in practice, and further consideration of no-arbitrage conditions is needed.

4.6 Constrained state space for static arbitrage

To ensure that our model does not generate static/model-free arbitrage, we introduce static arbitrage constraints on the underlying factors.

4.6.1 Static arbitrage-free relations

We study the static arbitrage-free relations at time $t = 0$ without loss of generality. For some T and K , we define $M_T = S_T/F_0(T)$, $k = K/F_0(T)$, and $\check{c}(\tau, k) = \check{c}_0(\tau, \ln k)$. We assume that $\check{c}(T, k) \in C^{1,2}(\mathbb{R}_{\geq 0}^2)$.

Similar to the analysis in Section 2.3, if we define the set of surface functions $s(x, y) : D \rightarrow \mathbb{R} \cap [0, 1]$, where $D \subseteq \mathbb{R}_{\geq 0}^2$ are compact sets, by

$$\mathcal{S}(D) = \left\{ s(x, y) \in C^{1,2}(D) : 0 \leq s \leq 1, \frac{\partial s}{\partial x} \geq 0, -1 \leq \frac{\partial s}{\partial y} \leq 0, \frac{\partial^2 s}{\partial y^2} \geq 0 \right\}, \quad (4.9)$$

then no arbitrage can be constructed on the surface $(T, k) \mapsto \check{c}(T, k)$ if $\check{c} \in \mathcal{S}(\mathbb{R}_{\geq 0}^2)$.

Our models aim to ensure absence of static arbitrage for the liquid range \mathcal{R}_{liq} defined in the (τ, m) -coordinates, or equivalently the liquid range $\check{\mathcal{R}}_{\text{liq}}$ defined on the (τ, e^m) -coordinate, where $\check{\mathcal{R}}_{\text{liq}} = \{(\tau, k) : (\tau, \ln k) \in \mathcal{R}_{\text{liq}}\}$. Accordingly, the corresponding discrete lattice \mathcal{L}_{liq} in (τ, e^m) -coordinates is denoted by $\check{\mathcal{L}}_{\text{liq}} = \{(\tau_i, k_j^i) : (\tau_i, \ln k_j^i) \in \mathcal{L}_{\text{liq}}\}$. We will construct models to ensure that $\check{c} \in \mathcal{S}(\check{\mathcal{R}}_{\text{liq}})$. We achieve this in two steps:

- (i) Ensure that no static arbitrage can be constructed from the finitely many liquid options on $\check{\mathcal{L}}_{\text{liq}}$, i.e.

$$\exists s \in \mathcal{S}(\check{\mathcal{R}}_{\text{liq}}), \text{ s.t. } \forall (\tau, k) \in \check{\mathcal{L}}_{\text{liq}}, s(\tau, k) = \check{c}(\tau, k). \quad (4.10)$$

- (ii) Given statically arbitrage-free prices on $\check{\mathcal{L}}_{\text{liq}}$, apply some shape-preserving smooth interpolation method which ensures monotonicity (on τ and k axes) and convexity (on k axis) of the surface. The interpolant function then belongs to $\mathcal{S}(\check{\mathcal{R}}_{\text{liq}})$. For example, Fengler and Hin [69] use bivariate tensor-product B-spines to estimate a statically arbitrage-free call price surface from observed prices of options on an arbitrarily shaped lattice.

4.6.2 Static arbitrage constraints on the liquid lattice

For notational simplicity, we reindex our liquid lattice $\mathcal{L}_{\text{liq}} = \{(\tau_i, m_j^i)\}_{i,j}$ to a single sequence $\{(\tau_j, m_j)\}_{j=1}^N$, and denote $c_j = \check{c}(\tau_j, k_j) = \tilde{c}(\tau_j, m_j = \ln k_j)$.

Cousot [52] finds practically verifiable conditions for (4.10), which we have simplified in Section 2.4. We can write these constraints in the form $\mathbf{A}\mathbf{c} \geq \widehat{\mathbf{b}}$, where $\mathbf{c} = [c_1 \cdots c_N]^\top \in \mathbb{R}^N$ is the vector of normalised call prices, $\mathbf{A} = (A_{ij}) \in \mathbb{R}^{R \times N}$ and $\widehat{\mathbf{b}} = (\hat{b}_j) \in \mathbb{R}^R$ are a known constant matrix and vector. Here R is the number of static arbitrage constraints. The construction of \mathbf{A} and $\widehat{\mathbf{b}}$ depends on the choice of \mathcal{L}_{liq} , and is given in Chapter 2. We assume for simplicity that \mathcal{L}_{liq} is fixed over time.

Using our factor representation (4.4) of option prices, we have $\mathbf{c}_t = \mathbf{G}_0 + \mathbf{G}^\top \xi_t$, with \mathbf{G} as in (4.8) and \mathbf{G}_0 the vector with entries $G_0(\tau_j, m_j)$. Consequently, the market model allows no static arbitrage among options on the liquid lattice \mathcal{L}_{liq} if ξ_t satisfies

$$\mathbf{A}\mathbf{G}^\top \xi_t \geq \mathbf{b} := \widehat{\mathbf{b}} - \mathbf{A}\mathbf{G}_0^\top, \text{ for all } t \in [0, T^* - \bar{\tau}]. \quad (4.11)$$

The dynamics of the diffusion process ξ are characterised by its drift and diffusion coefficients, which are assumed to be functions of $\tilde{\xi} = (S, \xi)$. We will address how to enforce this constraint on ξ_t when training a model in Section 5.1.

4.7 Model inference

Given an observable normalised option price surface, we estimate the model following two steps:

- (i) we decode the factor representation $\tilde{c}_t(\tau, m) = G_0(\tau, m) + \sum_{i=1}^d G_i(\tau, m)\xi_{it}$ by calibrating the price basis functions $\{G_i\}_{i=0}^d$;
- (ii) we infer the drift and diffusion functions of the decoded factor process ξ .

While these models are written in a continuous-time setting, inference is made from discretely observed data. Let us assume that we observe a discrete time series of normalised call option prices $\mathbf{c}_t = [\tilde{c}_t(\tau_1, m_1) \cdots \tilde{c}_t(\tau_N, m_N)]^\top$ at times $0 = t_0 < t_1 < \cdots < t_L = T$, where $\{(\tau_j, m_j)\}_{1 \leq j \leq N} = \mathcal{L}_{\text{liq}}$, assuming $d \ll N \ll L$.

4.7.1 Decoding the factor representation

Projecting the data from \tilde{c} to ξ reduces dimensions by a linear transformation. The matrix representation of this linear projection gives an estimate of the price basis vectors $\{\mathbf{G}_i\}_{i=0, \dots, d}$. Define matrices $\mathbf{C} = (C_{lj}) \in \mathbb{R}^{(L+1) \times N}$ and $\mathbf{\Xi} = (\Xi_{li}) \in \mathbb{R}^{(L+1) \times d}$, where $C_{lj} = \tilde{c}_{t_{l-1}}(\tau_j, m_j)$ and $\Xi_{li} = \xi_{i, t_{l-1}}$, respectively. We represent the observed data \mathbf{C} , with residuals $\mathbf{\Upsilon} \in \mathbb{R}^{(L+1) \times N}$, by

$$\mathbf{C} = \mathbf{1}_{L+1} \otimes \mathbf{G}_0 + \mathbf{\Xi} \mathbf{G} + \mathbf{\Upsilon}, \quad (4.12)$$

where $\mathbf{1}_{L+1}$ is an $(L+1)$ -vector of ones, and \otimes denotes the outer product of two vectors. We call $\mathbf{C} - \mathbf{\Upsilon}$ the *reconstructed prices*, and the Frobenius norm $\|\mathbf{\Upsilon}\|_F$ the *reconstruction error*. For some fixed $d < N$, an ideal factor representation should achieve the following three objectives:

- (i) **Statistical accuracy:** The reconstruction error $\|\mathbf{\Upsilon}\|_F$ should be as small as possible. Among linear dimension reduction techniques, PCA applied to the observed prices \mathbf{C} gives the optimal reconstruction error.
- (ii) **Minimal dynamic arbitrage:** There should be as few violations of the HJM-type drift restrictions (4.8) as possible. For this purpose, we estimate z_t from our observed prices \mathbf{C} , then construct factors from principal components of z_t , to minimize the discrepancy between z_t and the space spanned by the factors.
- (iii) **No static arbitrage:** The reconstructed prices should violate the static arbitrage constraints (4.11) as seldom as possible.

Suppose we construct d^{st} , d^{da} and d^{sa} factors for the above-mentioned objectives, respectively, where $d^{\text{st}} + d^{\text{da}} + d^{\text{sa}} = d$. Algorithm 1 gives the details of the factor decoding process.

Remark 4.3. We discuss here some practical aspects of Algorithm 1:

- (i) The order of fulfilling the three objectives could be altered. However, if the factor construction for no static arbitrage was not placed in the last step, it is possible that any factors defined afterwards could introduce more static arbitrage.

Algorithm 1: Decoding the factor representation

Input : Matrix of price data $\mathbf{C} \in \mathbb{R}^{(L+1) \times N}$; number of factors d ; number of statistical accuracy factors d^{st} ; number of dynamic arbitrage factors d^{da} ; number of static arbitrage factors d^{sa} ; data for z_t in (4.8), denoted as $\mathbf{Z} \in \mathbb{R}^{(L+1) \times N}$; constraint matrices \mathbf{A}, \mathbf{b} from (4.11).

Output: Factor data $\Xi \in \mathbb{R}^{(L+1) \times d}$; price basis vectors $\mathbf{G}_i \in \mathbb{R}^N$ for $i = 0, \dots, d$.

- 1 Let $G_{0j} = \frac{1}{L+1} \sum_{l=1}^{L+1} C_{lj}$ for $j = 1, \dots, N$;
/* Construct factors to minimize dynamic arbitrage. */
- 2 Compute the residual data $\mathbf{R}_0 = \mathbf{C} - \mathbf{1}_{L+1} \otimes \mathbf{G}_0$;
- 3 Compute the principal component decomposition of \mathbf{Z} and assign, $\forall i = 1, \dots, d^{\text{da}}$ and $\forall l = 1, \dots, L + 1$,
$$\mathbf{G}_i = i\text{-th principal component of } \mathbf{Z}, \Xi_{li} = \langle (\mathbf{R}_0)_l, \mathbf{G}_i \rangle,$$
where $(\mathbf{R}_0)_l$ is the l -th row of \mathbf{R}_0 ;
/* Construct factors to maximize statistical accuracy. */
- 4 Compute the residual data $\mathbf{R}_{d^{\text{da}}} = \mathbf{R}_0 - \sum_{k=1}^{d^{\text{da}}} \Xi_k \otimes \mathbf{G}_k$, where $\Xi_k = [\Xi_{1k} \dots \Xi_{L+1,k}]$;
- 5 Compute the principal component decomposition of $\mathbf{R}_{d^{\text{da}}}$ and assign, $\forall i = d^{\text{da}} + 1, \dots, d^{\text{da}} + d^{\text{st}}$ and $\forall l = 1, \dots, L + 1$,
$$\mathbf{G}_i = (i - d^{\text{da}})\text{-th principal component of } \mathbf{R}_{d^{\text{da}}}, \Xi_{li} = \langle (\mathbf{R}_{d^{\text{da}}})_l, \mathbf{G}_i \rangle,$$
where $(\mathbf{R}_{d^{\text{da}}})_l$ is the l -th row of $\mathbf{R}_{d^{\text{da}}}$;
/* Construct factors to minimize static arbitrage. */
- 6 **foreach** $i = d^{\text{da}} + d^{\text{st}}, \dots, d - 1$ **do**
- 7 Compute the residual data $\mathbf{R}_i = \mathbf{R}_0 - \sum_{k=1}^i \Xi_k \otimes \mathbf{G}_k$, where $\Xi_k = [\Xi_{1k} \dots \Xi_{L+1,k}]$;
- 8 Compute the covariance of the residual data $\Sigma_i = \mathbf{R}_i^\top \mathbf{R}_i$;
- 9 Eigen-decompose $\Sigma_i = \mathbf{Q}_i \Lambda_i \mathbf{Q}_i^{-1}$ (where \mathbf{Q}_i is a unitary matrix) ;
- 10 For some $\mathbf{w} \in \mathbb{R}^d$ such that $\|\mathbf{w}\|_2 = 1$, define
$$\mathbf{q}_i(\mathbf{w}) = \mathbf{Q}_i \mathbf{w}, \mathbf{s}_i(\mathbf{w}) = \mathbf{R}_i \mathbf{q}_i(\mathbf{w}).$$

Then solve the following minimization problem:

$$\mathbf{w}^* = \arg \min_{\|\mathbf{w}\|_2=1} \left(- \sum_{l=1}^{L+1} \mathbb{1}_{\{(\mathbf{A} \mathbf{q}_i(\mathbf{w}), \mathbf{s}_i(\mathbf{w})) \geq \mathbf{b} - \mathbf{A}(\sum_{k=1}^{i-1} \Xi_k \mathbf{G}_k)\}} + \lambda \|\mathbf{R}_i - \mathbf{s}_i(\mathbf{w}) \otimes \mathbf{q}_i(\mathbf{w})\|_2 \right), \quad (4.13)$$
where we require $\lambda < \min\left(1, \frac{1}{\|\mathbf{R}_i\|}\right)$ so that the penalty term is strictly less than 1;
- 11 Let $\mathbf{G}_{i+1} = \mathbf{q}_i(\mathbf{w}^*)$ and $\Xi_{i+1} = \mathbf{s}_i(\mathbf{w}^*)$;
- 12 **end**

(ii) The objective function (4.13) calculates the number of arbitrage inequalities which are violated in the reconstructed data, with an additional penalty term concerning the reconstruction errors. The penalty term plays a role in ensuring uniqueness of the preferred weights \mathbf{w} in the case that multiple weights produce the same number of arbitrage-free samples.

(iii) The decoding algorithm depends on the choice of the parameters d^{st} , d^{da} and d^{sa} , which should be determined through some preliminary analysis of the given data.

We show examples in the applications given in Section 6.3 and 7.3.

- (iv) For numerical convenience, it is useful to ensure that the factors constructed are orthonormal (i.e. \mathbf{G} is a unitary matrix). This can be done through a simple Gram–Schmidt process once an initial set of factors has been constructed.

4.7.2 Estimating the constrained diffusion process

Our goal is to estimate the functions $\alpha : \mathbb{R}^{d+1} \rightarrow \mathbb{R}$, $\mu : \mathbb{R}^{d+1} \mapsto \mathbb{R}^d$, $\gamma : \mathbb{R}^{d+1} \rightarrow \mathbb{R}$, and $\sigma : \mathbb{R}^{d+1} \mapsto \mathbb{R}^{d \times d}$ of the diffusion process $\tilde{\xi} = (S, \xi)$, given observations of the solution of the SDE (under the objective measure \mathbb{P})

$$dS_t = \left(\alpha(\tilde{\xi}_t) - q_t \right) S_t dt + \gamma(\tilde{\xi}_t) S_t dW_{0,t}, \quad (4.14a)$$

$$d\xi_t = \mu(\tilde{\xi}_t) dt + \sigma(\tilde{\xi}_t) dW_t. \quad (4.14b)$$

To ensure reconstructed prices do not generate static arbitrage, μ, σ in (4.14b) should be chosen to guarantee that ξ takes values in the (deterministic) polytope given by (4.11).

The SDE system (4.14) can be divided into two subsystems and estimated independently, as the constraints are only needed in (4.14b), and do not depend on S . In Chapter 5, we will discuss an neural network-based estimation method for the model of ξ (4.14b) that is constrained a polytope. The estimation method is generic enough for estimating the model of S (4.14a).

Appendix

4.A HJM drift restriction for dynamic arbitrage

In order to understand and avoid dynamic arbitrage, we will find a condition on the drifts α, μ in (4.2) such that the resulting discounted prices $\bar{c}_t(T, K) := C_t(T, K)/D_t(T) = D_0(t)C_t(T, K)/D_0(T)$ are martingales under some equivalent measure.

From the price transformation (4.3) and the latent factor representation (4.4), we can write $\bar{c}_t(T, K) = F_t(T)g(t, T - t, M(K; F_t(T)), \xi_t)$. For some fixed τ and m , we write $g(t, y) = g(t, \tau, m, y)$ for notational simplicity. By Itô's lemma, we have

$$\begin{aligned} d\bar{c}_t &= dg(t, \xi_t) = \left[\frac{\partial g}{\partial t}(t, \xi_t) + \sum_{i=1}^d \mu_{i,t} \frac{\partial g}{\partial y_i}(t, \xi_t) + \frac{1}{2} \sum_{i=1}^d \sum_{k=1}^d \sum_{j=1}^d \sigma_{ij,t} \sigma_{kj,t} \frac{\partial^2 g}{\partial y_i \partial y_k}(t, \xi_t) \right] dt \\ &\quad + \sum_{i=1}^d \sum_{j=1}^d \sigma_{ij,t} \frac{\partial g}{\partial y_i}(t, \xi_t) dW_{j,t} \\ &=: \tilde{\eta}_t dt + \tilde{\nu}_t dW_t. \end{aligned} \quad (4.15)$$

Let $F_t(T) = F(t, S_t; r_t, q_t)$. The dynamics of the T -maturity forward price are then

$$dF_t = dF(t, S_t) = (\alpha_t - r_t)F_t dt + \gamma_t F_t dW_{0,t}.$$

Since moneyness is a function of forward price, we use Itô's lemma again to derive the dynamics of moneyness. With $M(f) = M(K; f)$ for some fixed strike K , we have

$$dm_t = dM(F_t) = \left[(\alpha_t - r_t)F_t \frac{\partial M}{\partial f}(F_t) + \frac{1}{2} \gamma_t^2 F_t^2 \frac{\partial^2 M}{\partial f^2}(F_t) \right] dt + \gamma_t F_t \frac{\partial M}{\partial f}(F_t) dW_{0,t}. \quad (4.16)$$

Now, for some fixed expiry T and strike K , $\hat{c}_t(T, K) = \tilde{c}_t(\tau_t, m_t)$ where m_t has the dynamics given by (4.16) and $\tau_t = T - t$. Using a generalised Itô–Wentzell formula [143],

$$\begin{aligned} d\hat{c}_t(T, K) &= d\tilde{c}_t(T - t, m_t) \\ &= \left\{ \tilde{\eta}_t(T - t, m_t) - \frac{\partial g}{\partial \tau}(t, T - t, m_t) + \frac{1}{2} \left[\gamma_t F_t \frac{\partial M}{\partial f}(F_t) \right]^2 \frac{\partial^2 g}{\partial m^2}(t, T - t, m_t) \right. \\ &\quad \left. + \left[(\alpha_t - r_t)F_t \frac{\partial M}{\partial f}(F_t) + \frac{1}{2} \gamma_t^2 F_t^2 \frac{\partial^2 M}{\partial f^2}(F_t) \right] \frac{\partial g}{\partial m}(t, T - t, m_t) \right\} dt \\ &\quad + \gamma_t F_t \frac{\partial M}{\partial f}(F_t) \frac{\partial g}{\partial m}(t, T - t, m_t) dW_{0,t} + \tilde{\nu}_t(T - t, m_t) dW_t \\ &=: \eta_t(T, K) dt + \nu_t(T, K) d\bar{W}_t. \end{aligned} \quad (4.17)$$

Finally, let $\bar{c}_t = h(F_t, \hat{c}_t) = F_t \hat{c}_t$ and applying Itô's lemma to h , we have

$$d\bar{c}_t = dh(F_t, \hat{c}_t) = [\eta_t F_t + (\alpha_t - r_t)\bar{c}_t + \nu_{0,t} \gamma_t F_t] dt + (\nu_t F_t + \gamma_t \bar{c}_t \mathbf{e}_1) d\bar{W}_t, \quad (4.18)$$

where $\mathbf{e}_i \in \mathbb{R}^{(d+1) \times 1}$ is the unit vector with i -th entry one and other entries zero.

Drift restriction

We wish to connect the dynamics of \bar{c} with a no-arbitrage condition. Let $\bar{W} = [W_0 \cdots W_d]^\top$ and $\widetilde{W}_t = [\widetilde{W}_{0,t} \cdots \widetilde{W}_{d,t}]^\top$ have dynamics given by

$$\widetilde{W}_t = \bar{W}_t + \int_0^t \varphi_t dt,$$

where $\varphi_t = [\psi_{0,t} \cdots \psi_{p,t}]^\top \in L_{\text{loc}}^2(\mathbb{R}^{d+1})$ is a $(d+1)$ -dimensional progressively measurable process satisfying Novikov's² condition $\mathbb{E}^\mathbb{P}[\exp(\frac{1}{2} \int_0^{T^*} \varphi_t^\top \varphi_t dt)] < \infty$. Then Girsanov's theorem enables us to construct a measure $\mathbb{Q} \sim \mathbb{P}$ with the Radon–Nikodym derivative

$$\left. \frac{d\mathbb{Q}}{d\mathbb{P}} \right|_{\mathcal{F}_t} = \exp \left(- \int_0^t \varphi_s^\top d\bar{W}_s - \frac{1}{2} \int_0^t \varphi_s^\top \varphi_s ds \right). \quad (4.19)$$

²If Novikov's condition is not satisfied, but the process φ is locally square-integrable in time, then some no-arbitrage results are still possible; see Karatzas and Kardaras [101].

Under \mathbb{Q} the process \widetilde{W}_t is a $(d+1)$ -dimensional standard Brownian motion. The process ψ_i is interpreted as the market price of risk associated with the source of randomness W_i for each $i = 0, \dots, d$.

To ensure no arbitrage, we require the existence of a measure $\mathbb{Q} \sim \mathbb{P}$ under which the discounted price processes of all tradable assets are martingales. Since \mathbb{Q} can be found through (4.19) from the market price of risk process φ , we see that no-arbitrage requires the existence of an appropriate φ . Letting

$$\psi_{0,t} = \frac{\alpha_t - r_t}{\gamma_t}, \quad (4.20)$$

under the measure \mathbb{Q} , the discounted stock price $\widetilde{S}_t = S_t \exp(\int_0^t (q_s - r_s) ds)$ and the forward price are \mathbb{Q} -local martingales, i.e. $d\widetilde{S}_t = \gamma_t \widetilde{S}_t d\widetilde{W}_{0,t}$ and $dF_t = \gamma_t F_t d\widetilde{W}_{0,t}$. Since we have assumed that $\gamma S \in L^2_{\text{loc}}(\mathbb{R}^{d+1})$, it follows that \widetilde{S} and F are also \mathbb{Q} -martingales.

We choose the rest of the d components of φ such that \bar{c} is also a \mathbb{Q} -local martingale, i.e. it is driftless under \mathbb{Q} . Using (4.18) and (4.20), we have, under \mathbb{Q} ,

$$d\bar{c}_t = (\eta_t + \nu_{0,t}\gamma_t - \nu_t \cdot \varphi_t) F_t dt + (\nu_t F_t + \gamma_t \bar{c}_t \mathbf{e}_1) d\widetilde{W}_t. \quad (4.21)$$

Hence, the no-arbitrage condition, also known as the *drift restriction*, is that there must exist a process φ satisfying

$$\nu_t(T, K) \cdot \varphi_t(T, K) = \eta_t(T, K) + \nu_{0,t}(T, K)\gamma_t. \quad (4.22)$$

Given the parametrisation in terms of moneyness m_t and the market factor representation (4.4), we can simplify the drift restriction (4.22). For some fixed T and K , we let $m_t = M(K; F_t(T))$ and $\tau_t = T - t$. Then $\hat{c}_t(T, K)$ has drift and diffusion

$$\begin{aligned} \eta_t(T, K) &= \sum_{i=1}^d \left\{ \left[\mu_{i,t} + \xi_{i,t} \left(-\frac{\partial}{\partial \tau} + \left(r_t - \alpha_t + \frac{1}{2} \gamma_t^2 \right) \frac{\partial}{\partial m} + \frac{1}{2} \gamma_t^2 \frac{\partial^2}{\partial m^2} \right) \right] G_i(\tau_t, m_t) \right\} \\ &\quad + \left[-\frac{\partial}{\partial \tau} + \left(r_t - \alpha_t + \frac{1}{2} \gamma_t^2 \right) \frac{\partial}{\partial m} + \frac{1}{2} \gamma_t^2 \frac{\partial^2}{\partial m^2} \right] G_0(\tau_t, m_t), \end{aligned} \quad (4.23)$$

$$\nu_{j,t}(T, K) = \begin{cases} -\sum_{i=1}^d \xi_{i,t} \gamma_t \frac{\partial G_i(\tau_t, m_t)}{\partial m} - \gamma_t \frac{\partial G_0(\tau_t, m_t)}{\partial m}, & j = 0; \\ \sum_{i=1}^d \sigma_{ij,t} G_i(\tau_t, m_t), & 1 \leq j \leq d. \end{cases} \quad (4.24)$$

Substituting η_t and ν_t into the drift restriction (4.22) yields

$$\begin{aligned} \sum_{j=1}^d \left(\sum_{i=1}^d \sigma_{ij,t} G_i(\tau_t, m_t) \right) \psi_{j,t} &= \sum_{i=1}^d \left\{ \mu_{i,t} G_i(\tau_t, m_t) + \mathcal{Z} [G_i(\tau_t, m_t)] \xi_{i,t} \right\} \\ &= \sum_{i=1}^d \mu_{i,t} G_i(\tau_t, m_t) + \mathcal{Z} [\bar{c}_t] (\tau_t, m_t), \end{aligned} \quad (4.25)$$

where \mathcal{Z} is the infinitesimal operator defined by

$$\mathcal{Z} := -\frac{\partial}{\partial \tau} - \frac{1}{2}\gamma_t^2 \frac{\partial}{\partial m} + \frac{1}{2}\gamma_t^2 \frac{\partial^2}{\partial m^2}.$$

It follows that a necessary and sufficient condition for our models to be dynamic arbitrage-free is that (4.25) admits solutions with sufficient integrability.

Estimating neural-SDE constrained by a polytope

“孔子曰：‘君子有三戒：少之时，血气未定，戒之在色；及其壮也，血气方刚，戒之在斗；及其老也，血气既衰，戒之在得。’”

— 孔子弟子及再传弟子 《论语·季氏》

“Confucius said, ‘There are three things which the superior man guards against. In youth, when the physical powers are not yet settled, he guards against lust. When he is strong and the physical powers are full of vigour, he guards against quarrelsomeness. When he is old, and the animal powers are decayed, he guards against covetousness.’ ”

— Disciples of Confucius, *Ji Shi, Analects of Confucius* (translated by James Legge)

Consider a stochastic process $Y \in \mathbb{R}^d$ which is a solution of the SDE

$$dY_t = \mu(Y_t) dt + \sigma(Y_t) dW_t, \quad (5.1)$$

where W is a standard d -dimensional Brownian motion under the measure \mathbb{P} . In addition, μ, σ satisfy the usual Lipschitz regularity conditions such that a unique strong solution Y of the SDE exists. Given discrete observations of Y , we study the problem of estimating the drift $\mu : \mathbb{R}^d \rightarrow \mathbb{R}^d$ and the diffusion $\sigma : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$ in the model (5.1), such that paths of Y are guaranteed to lie within the deterministic convex polytope generated by a finite matrix-vector pair (\mathbf{V}, \mathbf{b}) , that is,

$$Y_t \in \mathcal{P} := \{x \in \mathbb{R}^d : \mathbf{V}x \geq \mathbf{b}\} \quad \text{for all } t \geq 0 \quad \mathbb{P}\text{-a.s.}$$

Without loss of generality (by rescaling), we assume the rows of \mathbf{V} are unit vectors.

5.1 Nonattainability of a diffusion process

Let \mathbf{v}_k denote the k -th row of \mathbf{V} . For each k , the diffusion process Y should not enter the domain $\mathcal{P}_k^c := \{x \in \mathbb{R}^d : \mathbf{v}_k^\top x < b_k\}$ at any time $t > 0$, \mathbb{P} -a.s. Writing $\mathcal{P}^c := \cup_k \mathcal{P}_k^c = \mathbb{R}^d \setminus \mathcal{P}$, we formulate the desired nonattainability of \mathcal{P}^c from its exterior by the solution Y of the SDE (5.1) as

$$1 - \mathbb{P}(Y_t \in \mathcal{P} \text{ for all } t > 0 | Y_0 \in \mathcal{P}) = \mathbb{P}(\exists t > 0 \text{ s.t. } Y_t \in \mathcal{P}^c | Y_0 \in \mathcal{P}) = 0. \quad (5.2)$$

Friedman and Pinsky [71] give a set of sufficient conditions for (5.2) in terms of μ and σ . It suffices that the normal components of the diffusion and the drift vanish on any boundaries $\partial\mathcal{P}_k := \{x \in \mathbb{R}^d : \mathbf{v}_k^\top x = b_k\}$, and that an additional ‘‘convexity’’ relation between the drift and diffusion coefficients on $\partial\mathcal{P}_k$ is imposed. Conveniently, this results in separate conditions for each of our k constraints, which are easily analysed.

For each k , let $\rho_k(y) := \text{dist}(y, \mathcal{P}_k^c)$ be the distance function defined for $y \notin \mathcal{P}_k^c$, and let $\nu : \mathbb{R}^d \rightarrow \mathbb{R}^d$ be the outward normal to $\partial\mathcal{P}_k$ (relative to \mathcal{P}_k^c , that is, a vector towards the interior of \mathcal{P}). Then, restricting to the closure of \mathcal{P} , we have $\nu_i = \partial\rho_k/\partial y_i$ on $\partial\mathcal{P}_k$. Defining $a := \sigma\sigma^\top$, to ensure that \mathcal{P}_k^c is nonattainable we verify

$$\sum_{i,j=1}^d a_{ij}\nu_i\nu_j = 0 \quad \text{on } \partial\mathcal{P}_k, \quad (5.3a)$$

$$\sum_{i=1}^d \mu_i\nu_i + \frac{1}{2} \sum_{i,j=1}^d a_{ij} \frac{\partial^2 \rho_k}{\partial y_i \partial y_j} \geq 0 \quad \text{on } \partial\mathcal{P}_k, \quad (5.3b)$$

assuming that ρ_k has second-order derivatives in a neighbourhood of $\partial\mathcal{P}_k$. This is formally written as Theorem 5.1.

Theorem 5.1 (Friedman and Pinsky (1973) [71]). *Suppose μ and σ satisfy Lipschitz and linear growth conditions, and let Y_t be any solution to (5.1) with initial data $Y_0 \in \mathcal{P}$. If (5.3) holds for all k , then \mathcal{P}^c is nonattainable by Y_t , i.e. $\mathbb{P}(\exists t > 0 \text{ s.t. } Y_t \in \mathcal{P}^c | Y_0 \in \mathcal{P}) = 0$.*

In our case of interest, the distance function $\rho_k(y)$ measures the (Euclidean) distance between a point $y \in \mathbb{R}^d$ and a hyperplane $\mathbf{v}_k^\top x = b_k$, that is, $\rho_k(y) = \mathbf{v}_k^\top y - b_k$. (Note that $\rho_k(y)$ is guaranteed to be non-negative for all y within the interior of \mathcal{P} .) Therefore, $\partial\rho_k(y)/\partial y = \mathbf{v}_k$ and $\partial^2\rho_k(y)/\partial y_i\partial y_j = 0$ on $\partial\mathcal{P}_k$. Therefore, we are able to re-write the conditions (5.3) as, for all k ,

$$\mathbf{v}_k^\top a(y)\mathbf{v}_k = 0 \quad \text{if } \mathbf{v}_k^\top y = b_k, \quad (\text{no diffusion over the boundary}); \quad (5.4a)$$

$$\mathbf{v}_k^\top \mu(y) \geq 0 \quad \text{if } \mathbf{v}_k^\top y = b_k, \quad (\text{inward-pointing drift at the boundary}). \quad (5.4b)$$

By enforcing these simple linear (in)equality constraints on $a = \sigma\sigma^\top$ and μ , we ensure that our process cannot leave the desired region \mathcal{P} .

5.2 The likelihood function approximation

Denote by $y_{0:T}$ the observed discrete sample path of Y_t at times $0 = t_0 < t_1 < \dots < t_L = T$. The problem of estimating the model (5.1) is formulated as the constrained optimization problem:

$$\max_{\mu, \sigma} p(y_{0:T}; \mu, \sigma), \text{ subject to (5.4a) and (5.4b), for all } k. \quad (5.5)$$

Here, $p(y_{0:T}; \mu, \sigma)$ is the likelihood (under measure \mathbb{P}) of observing the time series $y_{0:T}$ given μ and σ . The formulation (5.5) is an infinite-dimensional optimization problem, where the optimization variables are real-valued Lipschitz functions that are subject to boundary conditions.

We can compute the likelihood $p(y_{0:T}; \mu, \sigma)$ using the marginal density of the finite dimensional observations $y_{0:T}$. However, the transition density is generally not available in closed form. A simple approximation is to use the Euler–Maruyama scheme for the SDE of Y_t . Suppose we have a uniform time mesh and let $\Delta t = t_i - t_{i-1} = T/L$ for all $1 \leq i \leq L$. We have the single step approximation

$$Y_{t_{i+1}} - Y_{t_i} \approx \mu(Y_{t_i})\Delta t + \sigma(Y_{t_i})(W_{t_{i+1}} - W_{t_i}),$$

where $W_{t_{i+1}} - W_{t_i} \sim \mathcal{N}(\mathbf{0}, \mathbf{I}_d)$ and $0 \leq i < L$. For notional simplicity, let $\mu(i) = \mu(y_{t_i})$, $\sigma(i) = \sigma(y_{t_i})$ and $a(i) = \sigma(i)\sigma(i)^\top$. The likelihood of observing $y_{t_{i+1}}$ given y_{t_i} is then approximately

$$p(y_{t_{i+1}}|y_{t_i}; \mu, \sigma) \propto \frac{1}{\sqrt{|a(i)|}} \exp \left\{ -\frac{\|y_{t_{i+1}} - y_{t_i} - \mu(i)\Delta t\|_{a(i)}}{2\Delta t} \right\},$$

where $|a(i)| = \det(a(i))$, and for $u, \nu \in \mathbb{R}^d$ and a positive definite matrix $a \in \mathbb{R}^{d \times d}$,

$$(u, \nu)_a := u^\top a^{-1} \nu, \quad \|u\|_a^2 := (u, u)_a.$$

Due to the Markov property, we have

$$\begin{aligned} p(y_{0:T}|y_0; \mu, \sigma) &= \prod_{i=0}^{L-1} p(y_{t_{i+1}}|y_{t_i}; \mu, \sigma) \\ &\propto \left(\prod_{i=0}^{L-1} |a(i)|^{-\frac{1}{2}} \right) \times \exp \left\{ -\frac{1}{2\Delta t} \sum_{i=0}^{L-1} \|y_{t_{i+1}} - y_{t_i}\|_{a(i)}^2 \right\} \\ &\quad \times \exp \left\{ -\frac{1}{2} \sum_{i=0}^{L-1} \left[\|\mu(i)\|_{a(i)}^2 \Delta t - 2(\mu(i), y_{t_{i+1}} - y_{t_i})_{a(i)} \right] \right\}. \end{aligned}$$

Taking the log of the approximate likelihood p , we have

$$\begin{aligned} &\ln p(y_{0:T}|y_{t_0}0; \mu, \sigma) \\ &\propto -\frac{1}{2} \sum_{i=0}^{L-1} \left[\ln |a(i)| + \frac{1}{\Delta t} \|y_{t_{i+1}} - y_{t_i}\|_{a(i)}^2 + \|\mu(i)\|_{a(i)}^2 \Delta t - 2(\mu(i), y_{t_{i+1}} - y_{t_i})_{a(i)} \right]. \end{aligned} \quad (5.6)$$

5.3 Deep learning as a function approximation algorithm

Given the approximated likelihood function (5.6), the optimization (5.5) can be viewed as a *supervised learning* problem, where the input variable is y_t and the output variable is the increment $y_{t+\Delta t} - y_t$.

We will use a neural network to represent functions $\mu(\cdot)$ and $\sigma(\cdot)$ in a non-parametric way, limited only by properties such as continuity and boundedness. We design a neural network such that $\mu(\cdot)$ and $\sigma(\cdot)$ always satisfy the constraints listed in (5.4). This is achieved by viewing the *target* functions $\mu(\cdot)$ and $\sigma(\cdot)$ as transformations of some *underlying* functions $\hat{\mu}(\cdot)$ and $\hat{\sigma}(\cdot)$ through some operators \mathcal{G}_μ and \mathcal{G}_σ , respectively:

$$\mu = \mathcal{G}_\mu[\hat{\mu}], \quad \sigma = \mathcal{G}_\sigma[\hat{\sigma}].$$

We construct the operators \mathcal{G}_μ and \mathcal{G}_σ such that (5.4) holds for any functions $\hat{\mu}(\cdot)$ and $\hat{\sigma}(\cdot)$, but impose essentially no further restrictions. We then proceed by modelling the functions $\hat{\mu}(\cdot)$ and $\hat{\sigma}(\cdot)$ using a standard neural network. In Figure 5.3.1, we show an illustration of the proposed constrained neural network model.

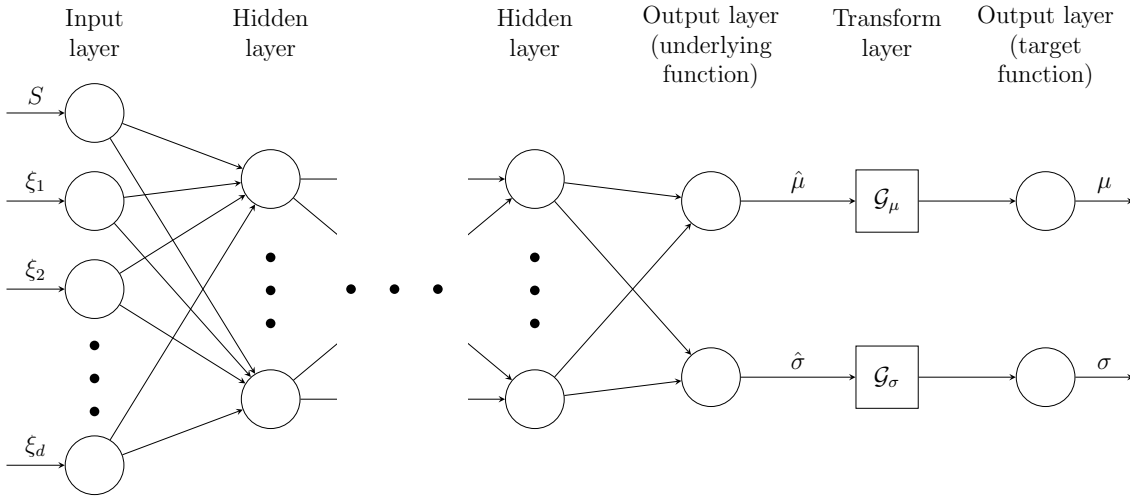


Figure 5.3.1: Constrained neural network.

5.3.1 Operator for diffusion constraints

The diffusion function σ is unconstrained except when y approaches a boundary of \mathcal{P} . More specifically, if y is on the k -th boundary, i.e. $\rho_k(y) = \mathbf{v}_k^\top y - b_k = 0$, the diffusion matrix $\sigma(y)$ will live in the kernel of \mathbf{v}_k^\top , as $(\mathbf{v}_k^\top \sigma)^2 = \mathbf{v}_k^\top a \mathbf{v}_k = 0$. We want to construct a Lipschitz operator \mathcal{G}_σ such that, for any non-degenerate locally (Lipschitz) continuous

$$\hat{\sigma} : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d},$$

$$\sigma = \mathcal{G}_\sigma[\hat{\sigma}] \in \ker(\mathbf{v}_k^\top), \text{ whenever } \rho_k(y) = 0.$$

We propose to achieve this by *shrinking* $\mathbf{v}_k^\top \sigma$ to zero as $\rho_k \downarrow 0$. We accordingly construct an operator which is able to act when near multiple boundaries, and should be unaffected by boundaries distant from y . Specifically, we define

$$\sigma(y) = \mathcal{G}_\sigma[\hat{\sigma}](y) := (\mathbf{P}(y))^\top \hat{\sigma}(y), \quad (5.7)$$

where $\mathbf{P}(y)$ is a linear shrinking matrix constructed using Algorithm 2.

Algorithm 2: Constructing diffusion shrinking transformation matrix

Input : A non-empty d -polytope $\mathcal{P} = \{x \in \mathbb{R}^d : \mathbf{V}x \geq \mathbf{b}, \mathbf{V} \in \mathbb{R}^{R \times d}, \mathbf{b} \in \mathbb{R}^d\}$; a Lipschitz continuous and monotonically increasing function $h^\sigma : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R} \cap [0, 1]$ satisfying $h^\sigma(0) = 0$ and $h^\sigma(x) \rightarrow 1$ for $x \rightarrow \infty$; some fixed $y \in \mathbb{R}^d$.

Output: The shrinking transformation $\mathbf{P}(y) \in \mathbb{R}^{d \times d}$.

/ In our experiments, we use the simple example $h^\sigma(x) = 1 - 1/(1+x)$. */*

1 **foreach** $k = 1, \dots, R$ **do**

2 Construct the normalised distances $\varepsilon_k^\sigma = h^\sigma \circ \rho_k(y)$, where ρ_k is the Euclidean distance to the k -th boundary of \mathcal{P} , as defined in Section 5.1.

3 **end**

4 Sort $\{\varepsilon_k^\sigma\}_k$ in ascending order, and denote the corresponding permutation of the indices of ε^σ as π . That is, $\varepsilon_k^\sigma = \varepsilon_{(\pi(k))}^\sigma$;

5 Keep the first d minimal ε^σ s, denoted by $\{\varepsilon_{(k)}^\sigma\}_{1 \leq k \leq d}$, and the associated hyperplane boundary coefficients $\{\mathbf{v}_{(k)}\}_{1 \leq k \leq d}$. Let $\mathbf{U} = [\mathbf{v}_{(1)} \ \dots \ \mathbf{v}_{(d)}]$ and assume it has full rank. (To ensure \mathbf{U} is of full rank, we can exclude from the ordering any k such that $\mathbf{v}_{(k)} \in \text{span}\{\mathbf{v}_{(1)}, \dots, \mathbf{v}_{(k-1)}\}$);

/ The set $\{\mathbf{v}_{(k)}\}_k$ represents the d directions on which the projected diffusion components most need to vanish. Therefore, we consider shrinking the diffusion along the orthogonal basis that is formed by these directions according to the order and scale of the corresponding distances $\{\varepsilon_{(k)}^\sigma\}_k$. */*

6 Let $\mathbf{q}_1 = \mathbf{v}_{(1)}$;

7 **foreach** $k = 2, \dots, d$ **do**

8 Following the Gram–Schmidt process, let

$$\mathbf{q}_k = \frac{1}{\sqrt{1 - \sum_{j=1}^{k-1} \langle \mathbf{q}_j, \mathbf{v}_{(k)} \rangle^2}} \left(\mathbf{v}_{(k)} - \sum_{j=1}^{k-1} \langle \mathbf{q}_j, \mathbf{v}_{(k)} \rangle \mathbf{q}_j \right), \quad (5.8)$$

then \mathbf{q}_k is the unit component of $\mathbf{v}_{(k)}$ that is orthogonal to $\text{span}\{\mathbf{v}_{(1)}, \dots, \mathbf{v}_{(k-1)}\}$;

9 **end**

10 Let $\mathbf{Q} = [\mathbf{q}_1 \ \dots \ \mathbf{q}_d]$, be the resulting change of basis matrix

/ A geometric illustration of the construction of \mathbf{Q} is shown in Figure 5.3.2 */*

11 Define

$$\mathbf{P}(y) = \text{diag}(\sqrt{\varepsilon_{(1)}^\sigma}, \dots, \sqrt{\varepsilon_{(d)}^\sigma}) \times \mathbf{Q}; \quad (5.9)$$

/ Unless the directions $\{\mathbf{v}_{(k)}\}_k$ are orthogonal, shrinking the diffusion along \mathbf{q}_j will shrink it along all directions in $\{\mathbf{v}_{(k)}\}_{k \geq j}$. This may cause the diffusion along $\mathbf{v}_{(k)}$, for any $k \geq 2$, to be shrunk by a much smaller factor than $\sqrt{\varepsilon_{(k)}^\sigma}$. To mitigate this compounded shrinking issue, we propose an alternative method to compute adaptive shrinking scales in Appendix 5.A. */*




Figure 5.3.2: Orthogonal directions for shrinking the diffusion function. As y is closer to boundary 4 (bottom right) than the others, direction \mathbf{q}_1 points to this boundary, and will be shrunk the most by the linear shrinking matrix $\mathbf{P}(y)$.

Proposition 5.1. *Given the operator \mathcal{G}_σ defined by (5.7) and (5.9), we have $\sigma = \mathcal{G}_\sigma[\hat{\sigma}] \in \ker(\mathbf{v}_k^\top)$ whenever $\rho_k(y) = 0$, for each k and any function $\hat{\sigma} : \mathbb{R}^d \rightarrow \mathbb{R}^{d \times d}$.*

Proof. Suppose that y is on the k -th boundary. Recall that we define π as the permutation of the indices of ε^σ when sorting $\{\varepsilon_j^\sigma\}_j$ in ascending order. If $\pi(k) \geq d$, then y is on a vertex of \mathcal{P} with $\varepsilon_{(j)} = 0$ for all $j = 1, \dots, d$. Hence,

$$\mathbf{v}_k^\top a \mathbf{v}_k = \sum_{i,j=1}^d a_{ij} (\mathbf{v}_k^\top \mathbf{q}_i) (\mathbf{v}_k^\top \mathbf{q}_j) \sqrt{\varepsilon_{(i)} \varepsilon_{(j)}} = 0.$$

Otherwise, when $1 \leq \pi(k) < d$, the Gram–Schmidt process ensures that $\mathbf{v}_k = \mathbf{v}_{(\pi(k))} \in \text{span}\{\mathbf{q}_1, \dots, \mathbf{q}_{\pi(k)}\} \perp \mathbf{q}_j$ for all $j > \pi(k)$. Hence,

$$\mathbf{v}_k^\top a \mathbf{v}_k = \sum_{i,j=1}^{\pi(k)} a_{ij} (\mathbf{v}_k^\top \mathbf{q}_i) (\mathbf{v}_k^\top \mathbf{q}_j) \sqrt{\varepsilon_{(i)} \varepsilon_{(j)}} = 0,$$

as $\varepsilon_{(j)} = 0$ for all $j = 1, \dots, \pi(k)$. Therefore, $\sigma \in \ker(\mathbf{v}_k^\top)$ as $(\mathbf{v}_k^\top \sigma)^2 = \mathbf{v}_k^\top a \mathbf{v}_k = 0$. \square

5.3.2 Operator for drift constraints

As with the diffusion function, the drift function μ should be unconstrained except when y approaches a boundary of \mathcal{P} . The geometric interpretation of the constraint $\mathbf{v}_k^\top \mu \geq 0$ is that the drift is pointing inwards whenever y is on the boundary \mathcal{P}_k .

As \mathcal{P} is convex, we know that the vector from any point on the boundary y towards any interior point ζ_k will satisfy $\mathbf{v}_k^\top (\zeta_k - y) > 0$. Therefore, we construct μ as a unconstrained locally (Lipschitz) continuous function $\hat{\mu}$, corrected by some \mathcal{P} -inward-pointed d -vectors, i.e.

$$\mu(y) = \mathcal{G}_\mu[\hat{\mu}](y) := \hat{\mu}(y) + \sum_k \lambda_k(y) (\zeta_k - y), \quad (5.10)$$

where $\{\zeta_k\}_k$ is a collection of interior points in \mathcal{P} , and λ_k is a scalar weight chosen such that $\zeta_k - y$ will dominate $\hat{\mu}$ when y moves close to the k -th boundary.

In order to ensure our correction acts smoothly as we approach a boundary, for some small $\rho^* > 0$, we select $\{\zeta_k\}_k$ from the ρ^* -interior points, defined as $\mathcal{P}^{\rho^*} := \{x \in \mathbb{R}^d : \mathbf{V}x \geq \mathbf{b} - \rho^* \mathbf{1}\}$, provided that $\mathcal{P}^{\rho^*} \neq \emptyset$ (which is guaranteed for ρ^* sufficiently small, as long as our inequality constraints are not degenerate). We pre-compute one ρ^* -interior point of \mathcal{P} for each boundary (see Algorithm 3 in Appendix 5.B for a method for constructing interior points). Ideally, an interior point should be as far from its corresponding boundary as possible, so that the correction vector $\zeta_k - y$ aligns more with the inward orthogonal vector to the corresponding boundary, where the correction is more efficient. An example is shown in Figure 5.3.3.




Figure 5.3.3: An example of ρ^* -interior points for calculating drift correction directions

Since the drift constraint is only required on the boundary, we require the correction vectors $\{\zeta_k - y\}_k$ to vanish quickly when y moves away from the boundary. To achieve this, it is sufficient to define

$$\lambda_k(y) = \left(\frac{\mathbf{v}_k^\top \hat{\mu} - \varepsilon_k^\mu(y)}{\mathbf{v}_k^\top (\zeta_k - y)} \right)^+, \quad (5.11)$$

where $\varepsilon_k^\mu = h^\mu \circ \rho_k$, for some continuous and monotonically increasing function $h^\mu : \mathbb{R}_{\geq 0} \rightarrow \mathbb{R}_{\geq 0}$ with $h^\mu(0) = 0$ and $h^\mu(x) \geq \varepsilon^{\mu*}$ for $x \geq \rho^*$ and some sufficiently large real number $\varepsilon^{\mu*}$. For example, we define $h^\mu(x) = \varepsilon^{\mu*}(e^x - 1)/(e^{\rho^*} - 1)$.

Proposition 5.2. *Given the operator \mathcal{G}_μ defined by (5.10) and (5.11), we have $\mathbf{v}_k^\top \mu \geq 0$ whenever $\rho_k(y) = 0$, for each k and any function $\hat{\mu} : \mathbb{R}^d \rightarrow \mathbb{R}^d$.*

Proof. For each k and any ρ^* -interior point ζ , whenever $\rho_k(y) < \rho^*$, we have

$$\mathbf{v}_k^\top (\zeta - y) = (\mathbf{v}_k^\top \zeta - b_k) - (\mathbf{v}_k^\top y - b_k) \geq \rho^* - \rho_k(y) > 0. \quad (5.12)$$

Consequently, by the definition of λ_k in (5.11),

$$\mu = \hat{\mu} + \sum_{\rho_k(y) < \rho^*} \lambda_k(\zeta_k - y) + \sum_{\rho_k(y) \geq \rho^*} 0 \cdot (\zeta_k - y).$$

If $y \in \mathcal{P}^{\rho^*}$, we have $\rho_i(y) \geq \rho^*$ for all k , which leads to $\mu = \hat{\mu}$. Therefore, no direction correction vectors are imposed on the drift when y is in the ρ^* -interior region of \mathcal{P} . For any k such that $\rho_k(y) < \rho^*$,

$$\mathbf{v}_k^\top \mu = \mathbf{v}_k^\top \hat{\mu} + \sum_{\rho_i(y) < \rho^*} \lambda_i \mathbf{v}_k^\top (\zeta_i - y) \geq \mathbf{v}_k^\top \hat{\mu} + \lambda_k \mathbf{v}_k^\top (\zeta_k - y) \geq -\varepsilon_k^\mu(y),$$

where we use the property of ρ^* -interior points as shown in (5.12) and the definition of λ_k in (5.11). It follows that, when $\rho_k(y) = 0$, we have $\mathbf{v}_k^\top \mu \geq -\varepsilon_k^\mu(y) = 0$. \square

5.3.3 Loss functional

With the approximated likelihood function and the transformed drift and diffusion functions, we are able to re-define the model estimation problem from the constrained optimization problem (5.5) to the following unconstrained optimization problem:

$$\begin{aligned} \min_{\hat{\mu}, \hat{\sigma}} J[\hat{\mu}, \hat{\sigma}] = & \sum_{i=0}^{L-1} \left[\ln |a(i)| + \frac{1}{\Delta t} \|y_{t_{i+1}} - y_{t_i}\|_{a(i)}^2 + \|\mu(i)\|_{a(i)}^2 \Delta t - 2(\mu(i), y_{t_{i+1}} - y_{t_i})_{a(i)} \right] \\ & + \lambda \mathcal{R}(\hat{\mu}, \hat{\sigma}), \end{aligned} \quad (5.13)$$

where $\mathcal{R}(\hat{\mu}, \hat{\sigma})$ is a regularization term added to the loss function to encourage useful properties such as sparsity; $\lambda > 0$ is a tuning parameter which controls the importance of the regularization term. Here, μ and a are constructed from $\hat{\mu}$ and $\hat{\sigma}$ by $\mu = \mathcal{G}_\mu[\hat{\mu}]$ and $a = \mathcal{G}_\sigma[\hat{\sigma}](\mathcal{G}_\sigma[\hat{\sigma}])^\top$.

Finally, to prevent the likelihood becoming degenerate, it is convenient to ensure $\hat{\sigma}$ (and hence a) is of full rank everywhere within \mathcal{P} . We do this by parametrising $\hat{\mu}$, $\hat{\sigma}$ using $\phi^\theta : \mathbb{R}^d \rightarrow \mathbb{R}^{\frac{1}{2}d(d+3)}$, a neural network with weight parameters θ , and let $\phi^\theta = (\phi_1^\theta, \dots, \phi_{\frac{1}{2}d(d+3)}^\theta)$. We construct

$$\hat{\sigma} = \begin{bmatrix} \exp(\phi_1^\theta) & 0 & \cdots & 0 \\ \phi_2^\theta & \exp(\phi_3^\theta) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \phi_{\frac{1}{2}d(d-1)+1}^\theta & \phi_{\frac{1}{2}d(d-1)+2}^\theta & \cdots & \exp(\phi_{\frac{1}{2}d(d+1)}^\theta) \end{bmatrix}, \quad \hat{\mu} = \begin{bmatrix} \phi_{\frac{1}{2}d(d+1)+1}^\theta \\ \phi_{\frac{1}{2}d(d+1)+2}^\theta \\ \vdots \\ \phi_{\frac{1}{2}d(d+3)}^\theta \end{bmatrix}. \quad (5.14)$$

Here we choose a lower triangular form for $\hat{\sigma}$ without loss of generality, and its diagonal terms have been exponentiated to ensure positivity. In addition, the lower triangular

structure of $\hat{\sigma}$ greatly simplifies the computation of the likelihood part of the loss functional J . Specifically,

$$\ln |a| = \ln |\mathbf{P}^\top \hat{\sigma} \hat{\sigma}^\top \mathbf{P}| = \ln \left(|\mathbf{P}|^2 \prod_{i=1}^d \hat{\sigma}_{ii}^2 \right) = 2 \ln |\mathbf{P}| + 2 \sum_{i=1}^d \phi_{\frac{1}{2}i(i+1)}^\theta,$$

and for any $u, v \in \mathbb{R}^d$,

$$(u, v)_a = u^\top (\mathbf{P}^\top \hat{\sigma} \hat{\sigma}^\top \mathbf{P})^{-1} v = \left[(\hat{\sigma}^\top)^{-1} (\mathbf{P}^{-1} u) \right]^\top \left[(\hat{\sigma}^\top)^{-1} (\mathbf{P}^{-1} v) \right]$$

can be computed efficiently, as $\hat{\sigma}$ is a triangular matrix.

Appendix

5.A Adaptive diffusion shrinking scales

In Algorithm 2, we would ideally like to shrink the diffusion component projected on $\mathbf{v}_{(k)}$ by $\sqrt{\varepsilon_{(k)}}$, for each k . However, since $\{\mathbf{v}_{(k)}\}_k$ is, in general, not orthogonal, shrinking along one direction will also shrink along all other non-orthogonal directions.

To mitigate this compounded shrinking issue, a modification of the definition of $\mathbf{P}(y)$ (5.9) in Algorithm 2 is as follows:

Similarly to 2, we define

$$\mathbf{P}(y) = \text{diag}(\sqrt{\varepsilon_1}, \dots, \sqrt{\varepsilon_d}) \times \mathbf{Q},$$

where $\{\sqrt{\varepsilon_k}\}_k$ is a set of adaptive shrinking scales. The novelty here is that we compute these scales recursively, together with the Gram–Schmidt process in (5.8).

Let $\varepsilon_1 = \varepsilon_{(1)}$. For each $k \geq 2$, suppose $\{\varepsilon_j\}_{j \leq k-1}$ is known. By the Gram–Schmidt process, we have

$$\mathbf{v}_{(k)} = \sum_{j=1}^{k-1} \langle \mathbf{q}_j, \mathbf{v}_{(k)} \rangle \mathbf{q}_j + \sqrt{1 - \sum_{j=1}^{k-1} \langle \mathbf{q}_j, \mathbf{v}_{(k)} \rangle^2} \mathbf{q}_k.$$

Since we shrink along \mathbf{q}_j by $\sqrt{\varepsilon_j}$ for all j , \mathbf{v}_k will shrink to

$$\hat{\mathbf{v}}_{(k)} = \sum_{j=1}^{k-1} \sqrt{\varepsilon_j} \langle \mathbf{q}_j, \mathbf{v}_{(k)} \rangle \mathbf{q}_j + \sqrt{\varepsilon_k} \sqrt{1 - \sum_{j=1}^{k-1} \langle \mathbf{q}_j, \mathbf{v}_{(k)} \rangle^2} \mathbf{q}_k.$$

In addition, we want the rescaled vector to have norm $\|\hat{\mathbf{v}}_{(k)}\|^2 = \varepsilon_{(k)}$, which yields

$$\varepsilon_k = \frac{1}{1 - \sum_{j=1}^{k-1} \langle \mathbf{q}_j, \mathbf{v}_{(k)} \rangle^2} \left(\varepsilon_{(k)} - \sum_{j=1}^{k-1} \langle \mathbf{q}_j, \mathbf{v}_{(k)} \rangle^2 \varepsilon_j \right).$$

If using this method, it is then natural to re-sort the boundary directions, so that at each step we shrink the direction corresponding to the largest remaining value of ϵ_k (rather than the largest value of the original shrinking scale ε_k). The matrix \mathbf{Q} is then built accordingly.

5.B The ρ^* -interior points for drift correction

We apply Algorithm 3 to define the interior points for drift correction. The correction is more efficient when the correction vectors formed by the interior points align with the inward orthogonal vector to the corresponding boundaries. Therefore, for each boundary, we first compute its midpoint (as a boundary of the polytope \mathcal{P}) and then, starting from the midpoint, we find the furthest possible ρ^* -interior point along the inward orthogonal direction. We can then pre-compute the set of drift correction vectors $\{\zeta_k - y_{t_i}\}_{k=1, \dots, R}$ for every observation y_{t_i} , where $i = 0, \dots, L$.

Algorithm 3: Computation of ρ^* -interior points

Input : A non-empty d -polytope $\mathcal{P} = \{x \in \mathbb{R}^d : \mathbf{V}x \geq \mathbf{b}, \mathbf{V} \in \mathbb{R}^{R \times d}, \mathbf{b} \in \mathbb{R}^d\}$, and a small positive constant ρ^* .

Output: One ρ^* -interior point corresponding to each of the R hyperplane boundaries of \mathcal{P} .

- 1 Apply the double description algorithm (Motzkin, Raiffa, Thompson and Thrall [119]) to identify the set of vertices of \mathcal{P} , denoted by \mathcal{V} ;
 - 2 **foreach** $k = 1, \dots, R$ **do**
 - 3 Find the set of vertices of \mathcal{P} that are passed through by the k -th boundary, denoted by $\mathcal{V}^k = \mathcal{V} \cap \{x \in \mathbb{R}^d : \mathbf{v}_k^\top x = b_k\}$;
 - 4 Compute the midpoint of the k -th boundary as

$$m_k = \frac{1}{|\mathcal{V}^k|} \sum_{v \in \mathcal{V}^k} v;$$
 - 5 Solve $c^* = \arg \max_{c \in \mathbb{R}_{>0}} \{c : \mathbf{V}(c\mathbf{v}_k + m_k) \geq \mathbf{b} - \rho^* \mathbf{1}\}$. The ρ^* -interior point is given by $\zeta_k = c^* \mathbf{v}_k + m_k$;
 - 6 **end**
-

Calibrating to synthetic data: Heston SLV model

We validate our modelling and inference approach with numerical experiments using data generated from a stochastic local volatility model (Jex, Henderson and Wang [98]), which is a state-of-the-art model frequently used on equity and FX derivatives desks in investment banks. We will here use simulated data, as our aim is to study the effectiveness of these methods in a situation with a well understood ground truth model, where we are not limited by any lack of data. Application to market data has further market-specific complexities, which we explore in Chapters 7, 8 and 9.

6.1 Input synthetic data

The Heston-type Stochastic Local Volatility (Heston-SLV) model we use is given by:

$$\begin{aligned} dS_u &= (r_u - q_u) S_u du + \mathcal{L}_t(u, S_u) \sqrt{\nu_t} S_u dW_u^S, \\ d\nu_u &= \kappa(\theta - \nu_u) du + \sigma \sqrt{\nu_u} dW_u^\nu, \\ d\langle W_u^S, W_u^\nu \rangle &= \rho du, \end{aligned} \quad u \in (t, T^*). \quad (6.1)$$

This model is popular in industry for its realistic dynamic properties through the stochastic volatility component $\sqrt{\nu_t}$, and its ability to exactly reprice vanilla options through fitting a local volatility component, the so-called *leverage function* \mathcal{L} . Here, the subscript t of \mathcal{L} indicates that calibration of \mathcal{L} is performed from derivative prices observed at t . The calibration typically involves an initial optimisation of the fit over the Heston parameters, followed by a calibration of the leverage function using Markovian projections (see Piterbarg [124] for a review of early approaches and the references in Bain, Mariapragassam and Reisinger [9] for more recent contributions).

We obtain the parameters and the leverage function used for generating synthetic input data in our tests by calibrating a Heston-SLV model to Bloomberg OTC USDBRL

option price data on 28th October, 2008, using the QuantLib library [1]. The Heston parameters are listed in Table 6.1.1, and we plot the leverage function \mathcal{L} in Figure 6.1.1.

We simulate discrete time series data for S , ν and call options with N fixed values of moneyness and maturity, over a time period of $(L + 1)\Delta t$ (see Table 6.1.1 for the values used). These values are chosen to broadly emulate the observation, at moderate frequency (of the order of every 10 minutes), of a single year's prices for the underlying asset and liquid European options book. Here, a single path of (S, ν) is sampled by an Euler–Maruyama approximation to (6.1) on $[0, T^*]^1$, and, for each t , option values at (S_t, ν_t) are produced using the calibrated Heston parameters and the initial portion, i.e. $\mathcal{L}_t(u, \cdot) = \mathcal{L}_0(u - t, \cdot)$, $u \in (t, T^*)$.²

This step is only required to generate the simulated input data, but not when real market data are used. The SDE (6.1) is formulated in the risk-neutral measure used for pricing options for each t , started at (S_t, ν_t) along the chosen trajectory. We also use (6.1) for generating the (S_t, ν_t) trajectory on $[0, T]$. Here, any drift could be used, and we choose 0 for simplicity, but note that this does not imply simulation under the risk-neutral measure.³

	Heston-SLV model								Simulation		
Parameter	r_t	q_t	S_0	ν_0	θ	κ	σ	ρ	L	Δt	N
Value	0	0	100	0.0083	0.0085	8.3	0.32	-0.42	10000	0.0001	46

Table 6.1.1: Parameter values

6.2 Implementation details of the model estimation strategy

Approximation of z_t in (4.8). We need $\mathbf{Z} = (z_t)_t$, as inputs for Algorithm 1 to decode our factors. However, neither γ_t (see (4.2) and thereafter) nor the partial derivatives of call option prices, as seen in (4.8), are readily available for computing z_t .

To estimate γ_t , we fit an *initial* model⁴ for S with state variable $\tilde{\xi} = (S, \xi^{(0)})$, where $\xi^{(0)}$ is the factor projected from the first $d^{(0)}$ principal components of the input call prices.

¹Since the leverage function is only defined on a bounded interval of S , we choose a path where the trajectory of S is always within the interval in order to avoid extrapolating the input leverage function. The trajectory of S can be found on the top right graph of Figure 6.6.1.

²This is consistent with the re-calibration approach in financial practice, which lends justification to our assumption of a stationary model for each fixed maturity.

³In fact, selection of a path within a certain range implicitly introduces a drift, and the use of the leverage function translated forward at different t , means that the simulation measure differs from the pricing measure.

⁴We use the neural SDE estimation method described in Chapter 5, where the neural network architecture is defined in (6.3).

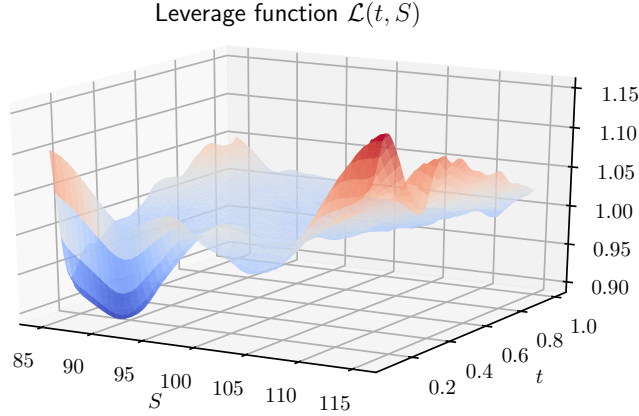


Figure 6.1.1: The leverage function $\mathcal{L}(t, S)$ of the Heston-SLV model used for simulation.

We use a large $d^{(0)}$, such that the reconstruction error is small. This initial calibration gives us a candidate volatility $\hat{\gamma}_t^{(0)}$.

Next, we need to evaluate the partial derivatives $\partial\tilde{c}/\partial\tau$, $\partial\tilde{c}/\partial m$ and $\partial^2\tilde{c}/\partial m^2$ on the option lattice $\mathcal{L}_{\text{liq}} = \{(\tau_j, m_j)\}_{j=1, \dots, N}$. We interpolate these values using $C^{1,2}$ basis functions⁵, from which we can approximate the required derivatives. In Figure 6.2.1, we show the approximated partial derivatives at $\{(\tau_j, m_j)\}_{j=1, \dots, N}$ for some fixed t . With $\hat{\gamma}_t^{(0)}$ and the approximated partial derivatives, we can compute z_t .




Figure 6.2.1: Partial derivatives of the normalised call price surface approximated by interpolation at $t = 0$.

Decorrelation and normalisation of factors. Though the price basis decoded from Algorithm 1 is orthogonal by construction, i.e. $\mathbf{G}\mathbf{G}^\top = \mathbf{I}_d$, the corresponding factors Ξ could still have high correlations. To enhance the explanatory capability of factors, we

⁵One could also use shape-preserving interpolation or convex regression to ensure monotonicity and convexity of the surface. This is likely to induce significant additional computational costs, as the interpolation method would depend on the state of the risk factors ξ .

decorrelate the factor data by considering the full principal component decomposition $\Xi^{(1)} = \Xi \mathbf{H}$, where columns of \mathbf{H} are eigenvectors of $\Xi^\top \Xi$.

In addition, data of different factors could distribute over very different scales. To improve numerical stability of model training, we further normalise the decorrelated factor data $\Xi^{(1)}$ as $\Xi^{(2)} = \Xi^{(1)} \mathbf{\Lambda}^{-1}$, where

$$\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_d), \text{ with } \lambda_j = \lambda_0 \left[\max_l \left(\Xi_{lj}^{(1)} \right) - \min_l \left(\Xi_{lj}^{(1)} \right) \right].$$

Here λ_0 is a positive constant; we choose $\lambda_0 = 10$ in our numerical example so that the min-max ranges of all normalised factor data are $1/\lambda_0 = 0.1$. Consequently, the factor representation of call price data as described in (4.12) can be re-written as

$$\mathbf{C} = \mathbf{1}_{L+1} \otimes \mathbf{G}_0 + \Xi^{(2)} \mathbf{G}^{(2)} + \Upsilon, \text{ where } \mathbf{G}^{(2)} = \mathbf{\Lambda} \mathbf{H}^\top \mathbf{G}.$$

We will build models for the decorrelated and normalised factors. In the following sections, for notational simplicity, we let $\Xi \leftarrow \Xi^{(2)}$ and $\mathbf{G} \leftarrow \mathbf{G}^{(2)}$.

Parametrisation of the drift μ_t . To improve the efficiency of the maximum likelihood estimator of the drift μ_t for the model of ξ , we provide a baseline estimator of the drift and use neural network only as a modification to the baseline model. Specifically, rather than constructing $\hat{\mu}$ with the neural network ϕ^θ directly as in (5.14), we define, for some baseline drift function $\bar{\mu}(\tilde{\xi}) : \mathbb{R}^{d+1} \rightarrow \mathbb{R}^d$,

$$\hat{\mu} = \text{diag} \left(\phi_{\frac{1}{2}d(d+1)+1}^\theta, \dots, \phi_{\frac{1}{2}d(d+3)}^\theta \right) \times \bar{\mu}.$$

A sensible baseline drift model can be built based on the assumption that the input option data do not permit dynamic arbitrage, i.e, the HJM drift restriction (4.8) holds, and that the market price of risk ψ_t is small. In particular, assuming zero market price of risk $\psi_t \equiv 0$, the drift restriction (4.8) implies⁶ that $\mu_t = \mathbf{\Lambda}^{-2} \mathbf{G} z_t$. In the following tests, we first perform a regression on the data $\{(\tilde{\xi}_t, \mathbf{\Lambda}^{-2} \mathbf{G} z_t)\}_t$ to obtain a baseline drift model $\tilde{\xi} \mapsto \bar{\mu}(\tilde{\xi})$.

6.3 Factor construction and no-arbitrage boundaries

We apply Algorithm 1 to decode factors from the normalised call prices that are generated from the Heston-SLV model. To see how well the decoded factors reconstruct the input data, we examine the following metrics:

⁶The normalisation of factors results in $\mathbf{G} \mathbf{G}^\top = \mathbf{\Lambda}^2$, rather than an identity matrix.

- Mean absolute percentage error (MAPE):

$$\text{MAPE} = \frac{1}{(L+1)N} \sum_{l=0}^L \sum_{j=1}^N \frac{\left| \tilde{c}_{t_l}(\tau_j, m_j) - G_{0j} - \sum_{i=1}^d G_{ij} \xi_{it_l} \right|}{\tilde{c}_{t_l}(\tau_j, m_j)}.$$

- Proportion of dynamic arbitrage (PDA):

$$\text{PDA} = 1 - \frac{\text{tr} \left(\mathbf{\Lambda}^{-1} \mathbf{G} \tilde{\mathbf{Z}}^{\top} \tilde{\mathbf{Z}} \mathbf{G}^{\top} \mathbf{\Lambda}^{-1} \right)}{\text{tr} \left(\tilde{\mathbf{Z}}^{\top} \tilde{\mathbf{Z}} \right)},$$

where $\tilde{\mathbf{Z}} = (\tilde{z}_{lj}) \in \mathbb{R}^{(L+1) \times N}$, with $\tilde{z}_{lj} = z_{t_{lj}} - \frac{1}{L+1} \sum_k z_{t_{kj}}$. This metric gives the fraction of variance of z_t that is unexplained by the constructed factors.

- Proportion of statically arbitrageable samples (PSAS):

$$\text{PSAS} = 1 - \frac{\sum_{l=0}^L \mathbf{1}_{\{\mathbf{A} \mathbf{G}^{\top} \xi_{t_l} \geq \mathbf{b}\}}}{L+1}.$$

We show the three metrics for a few combinations of factors in Table 6.3.1. Using only two factors, i.e. one dynamic arbitrage factor and one static arbitrage factor⁷ (last row in the table), can represent the whole collection of call prices with reasonable accuracy. Note that the use of a static arbitrage factor is significant in reducing the number of violations of the static arbitrage constraints, as evidenced by the corresponding reduction in PSAS.

Factors	MAPE	PDA	PSAS
Dynamic arb.	24.37%	3.51%	60.67%
Dynamic arb. + Statistical acc.	5.11%	3.21%	28.11%
Dynamic arb. + Static arb.	3.85%	3.04%	0.37%

Table 6.3.1: MAPE, PDA and PSAS metrics when including different combinations of factors.

As a proof-of-concept, we will restrict our attention to this simple three-factor model (i.e. two vectors ξ , in addition to the stock price), as it also allows us to demonstrate qualitative features of the model easily. We plot the price basis functions of these two factors, denoted as G_1 and G_2 , as well as G_0 , the constant term of \tilde{c} , in Figure 6.3.1. The points in the liquid lattice (in (τ, m) coordinates)⁸ are also shown on this plot. Here the real-valued functions G_1 and G_2 are obtained by interpolating the price basis vectors \mathbf{G}_1

⁷The minimization problem (4.13) in the factor decoding algorithm has a discrete objective function, where gradient-based optimisation methods do not apply. We use Py-BOBYQA (Cartis, Fiala, Marteau and Roberts [40]), a derivative-free optimization solver, to find its global optimum heuristically.

⁸These lattice points are chosen to mimic a typical liquid range of EURUSD options traded at CME, as seen in Figure 1.1.

and \mathbf{G}_2 , and G_0 is obtained by interpolating the normalised call prices averaged over time. Since the minimal option expiry in the input data is nonzero, there is no guarantee that the terminal payoff convergence condition (4.5) will be satisfied by simply extrapolating the input data along expiry. Hence, we introduce a few artificial data points at $\tau = 0$ for G_i , where $i = 0, 1, 2$, that satisfy (4.5) before interpolation.

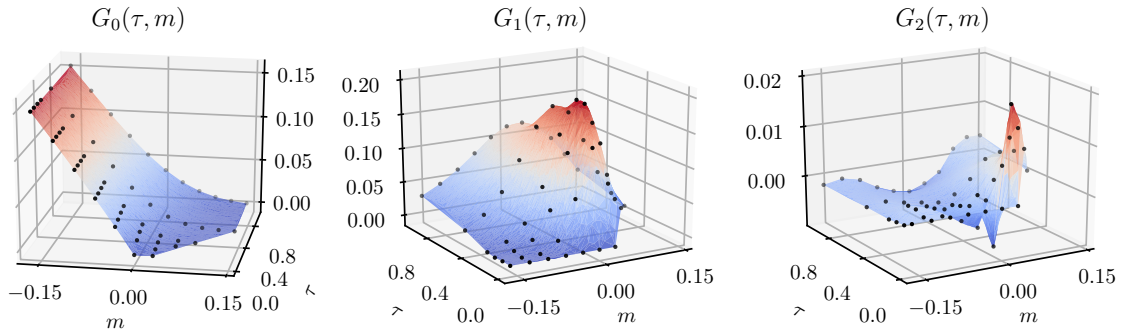


Figure 6.3.1: Price basis functions of the normalised call price surface.

Given these factors, we use the linear programming method (Caron, McDonald and Ponic [36]) to eliminate redundant constraints in the system $\mathbf{A}\mathbf{G}^\top \xi \geq \mathbf{b}$, which is the projection of the original no-arbitrage constraints, constructed in price space, to the \mathbb{R}^2 factor space. This results in only 7 constraints, which we indicate as red dashed lines in Figure 6.3.2. The convex polygonal domain surrounded by these constraints (light green area) is the statically arbitrage-free zone for the factors, that is, provided the factor process remains in this region, we are guaranteed to have no static arbitrage in the reconstructed call prices on our liquid lattice \mathcal{L}_{liq} .




Figure 6.3.2: Trajectory (black dots) of the \mathbb{R}^2 factors and the corresponding static arbitrage constraints (red dashed lines) projected to the \mathbb{R}^2 factor space.

The data we use for training the market model should be in the interior of the arbitrage-free factor space, otherwise we either have data exhibiting model-free arbitrage, or (if

points are on the boundary) the covariance matrix a will be singular, causing a failure in computing the likelihood part of the loss function J . Therefore, we truncate the factor data by removing those observations outside or on the boundary of the no-arbitrage region⁹.

6.4 Neural network training results

We apply the method in Chapter 5 to estimate the dynamics of the decoded factors. Specifically, we let $y = \xi$ and $\mathbf{V} = \mathbf{A}\mathbf{G}^\top$ and, by slight abuse of notation, μ and σ take $\tilde{\xi} = (S, \xi)$ as their argument rather than ξ . This does not alter the estimation method, except that the input layer of the neural network shall consist of $d + 1$, instead of d , neurons.

For the neural-SDE model of ξ , we use a simple architecture¹⁰ that is a composition of fully-connected layers and activation functions in the following orders:

$$\phi^\theta = \mathcal{F}_{d+1} \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{256} \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{256} \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{256}, \quad (6.2)$$

where \mathcal{F}_x is a fully connected layer, or affine transformation, with x units, and \mathcal{A}_{xxx} is an activation function. Each layer \mathcal{F} is parametric, but we omit the parameters for notational simplicity. We use a smaller network $\phi^{S,\theta} : \mathbb{R}^{d+1} \rightarrow \mathbb{R}^2$ for the neural-SDE model of S :

$$\phi^{S,\theta} = \mathcal{F}_{d+1} \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{128} \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{128} \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{128}. \quad (6.3)$$

We train the model for S separately from that for ξ . In addition, to mitigate over-fitting problems, we train both networks with 50% sparsity, meaning that the 50% smallest weights are pruned to zero. We implement and train our model using the standard tools within the Tensorflow [2] environment.

In Figure 6.4.1, we show the evolution of training losses and validation losses¹¹ over epochs during the training of the models for S and ξ , respectively. For the model of S , the loss value quickly drops during the first 10 epochs, and slowly declines and converges within 100 epochs. The loss function for the model of ξ is much more complex than that for the model of S (due to the higher dimensionality and transformations near boundaries), so we train it for significantly more epochs. Similar to the loss history for the model of S , the loss value for the model of ξ has a rapid decline for the first a few epochs, and then gradually converges.

⁹We are solving a supervised learning problem with the input-label data as $(x_l, y_l) = (\tilde{\xi}_{t_l}, \xi_{t_{l+1}} - \xi_{t_l})$ for $l = 1, \dots, L$. Hence, if ξ_{t_l} is an arbitrageable data point, we need to remove two samples $(\tilde{\xi}_{t_{l-1}}, \xi_{t_l} - \xi_{t_{l-1}})$ and $(\tilde{\xi}_{t_l}, \xi_{t_{l+1}} - \xi_{t_l})$ from the training data. For our training data, this leads to removing 49 samples given that we have observed 37 arbitrageable data points (after factor reconstruction) in our simulation.

¹⁰The robustness of the estimated neural network is assessed via sensitivity analysis in Appendix 6.C.

¹¹The first 90% of the dataset is used for training and the last 10% is reserved as validation data.

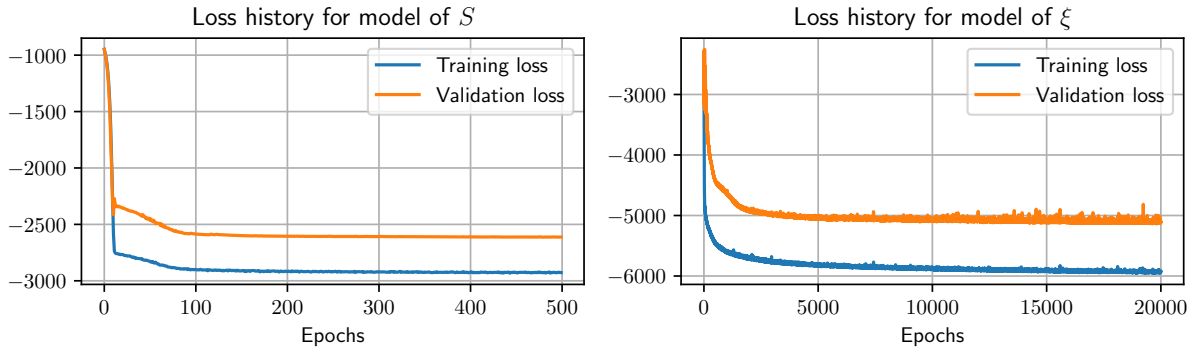


Figure 6.4.1: Evolution of training losses and validation losses.

In addition to observing the convergence of the loss functions, we need to demonstrate that the model has been trained sensibly compared with the ground truth Heston-SLV model, and that the learnt model is capable of generating data similar to the input data. We consider the following in-sample and out-of-sample tests.

6.5 In-sample test

We assess how well the learnt model recovers the ground truth Heston-SLV model (6.1). Specifically, we study $\sigma_t^S = \mathcal{L}(t, S_t)\sqrt{\nu_t}S_t$ (the diffusion coefficient of S), $\mu_t = \kappa(\theta - \nu_t)$ (the drift coefficient of ν) and $\sigma_t = \eta\sigma\sqrt{\nu_t}$ (the diffusion coefficient of ν), and compare values deduced from the input data against those generated by the learned model.


Model for S . The neural network for the model of S is written as $\phi^{S,\theta} = (\phi_\mu^{S,\theta}, \phi_\sigma^{S,\theta})$, where $\phi_\mu^{S,\theta}, \phi_\sigma^{S,\theta} : \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ give the approximation to the drift and diffusion coefficients of S , respectively. Using in-sample data $\{\tilde{\xi}_{t_l} = (S_{t_l}, \xi_{t_l})\}_{l=0,\dots,L}$, we can compute $\{\phi_\mu^{S,\theta}(\tilde{\xi}_{t_l})\}_l$ and $\{\phi_\sigma^{S,\theta}(\tilde{\xi}_{t_l})\}_l$.

It is well known that the estimation of drift terms in an SDE, from observation of a single time series over a short period, is prone to significant error. We hope that the drift estimates will be close to 0, the ground truth drift of S . However, as seen in Figure 6.5.1a, there are some deviations from 0, as well as some polynomial relations with the factors. Nevertheless, these deviations from 0 are small when compared to the size of the estimated diffusion, as will be seen later, and the polynomial relations with regards to the factors are possibly also artifacts introduced by selecting a path where S does not leave the domain of the leverage function.

The diffusion estimates $\{\phi_\sigma^{S,\theta}(\tilde{\xi}_{t_l})\}_l$ approximate $\{\sigma_{t_l}^S\}_l$, which is generated from the ground truth Heston-SLV model. As we see in Figure 6.5.1b, the volatility of S has been essentially correctly captured, with some noise and a slight upward bias for very low

volatilities. We observe a 4.96% mean absolute percentage error in this volatility estimate, computed as $\frac{1}{L+1} \sum_l |\phi_\sigma^{S,\theta}(\tilde{\xi}_{t_l}) - \sigma_{t_l}^S|/\sigma_{t_l}^S$.

We could also use these methods to perform an out-of-sample calibration check, provided that we can simulate $\tilde{\xi}_t$ realistically – this is indeed verified in Section 6.6; here we present in-sample performance to avoid any potential biases introduced by simulation methods.



(a) Scattergrams of $\phi_\mu^{S,\mu}(\tilde{\xi}_{t_l})$ against S , ξ_1 and ξ_2 . The estimated drift $\phi_\mu^{S,\mu}(\tilde{\xi}_{t_l})$ has values around 0, the ground-truth drift. (b) Scattergram of $\phi_\sigma^{S,\theta}(\tilde{\xi}_{t_l})$ against σ_t^S .

Figure 6.5.1: Estimated drift and diffusion coefficients for S .

Model for ξ . It is non-trivial to derive the ground-truth model for the factor ξ , given that the input data are generated from a Heston-SLV model. Nevertheless, the first calibrated factor $\{\xi_{1t_l}\}_{l=0,\dots,L}$ has a very strong linear relationship with the Heston-SLV simulated variance process $\{\nu_{t_l}\}_{l=0,\dots,L}$, as shown in Figure 6.5.2. Using this, an approximated ground-truth model for ξ_1 can be used to benchmark the learnt neural network model.




Figure 6.5.2: The linear relationship between ξ_1 and ν .

Assuming a simple linear model $\nu = \beta_0 + \beta_1 \xi_1$ with $\beta_1 \neq 0$, we apply Itô's lemma to

the SDE for ν in (6.1) to get

$$d\xi_{1t} = \mu^{\xi_1}(\xi_{1t}) dt + \sigma^{\xi_1}(\xi_{1t}) dW_t^\nu, \quad \mu^{\xi_1}(x) = \frac{\kappa(\theta - \beta_0 - \beta_1 x)}{\beta_1}, \quad \sigma^{\xi_1}(x) = \frac{\eta\sigma\sqrt{\beta_0 + \beta_1 x}}{\beta_1}. \quad (6.4)$$

Let $\phi_\mu^{\xi_1, \theta} : \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ and $\phi_\sigma^{\xi_1, \theta} : \mathbb{R}^{d+1} \rightarrow \mathbb{R}$ be the components of the neural network model that approximate the drift and diffusion functions of ξ_1 , respectively. Using in-sample data $\{\tilde{\xi}_{t_l} = (S_{t_l}, \xi_{t_l})\}_{l=0, \dots, L}$, we can compute $\{\phi_\mu^{\xi_1, \theta}(\tilde{\xi}_{t_l})\}_l$ and $\{\phi_\sigma^{\xi_1, \theta}(\tilde{\xi}_{t_l})\}_l$, which are supposed to approximate $\{\mu^{\xi_1}(\xi_{1t_l})\}_l$ and $\{\sigma^{\xi_1}(\xi_{1t_l})\}_l$, which were generated from the ground truth model (6.4). We plot and compare these data against $\{\xi_{1t_l}\}_l$ in Figure 6.5.3. We observe that the neural network model has captured the ground truth model well for $\xi_1 < 0.04$, where there are rich data for training. The square-root behavior in the diffusion of ν has been particularly well captured, as has the decreasing drift, despite the presence of noise.

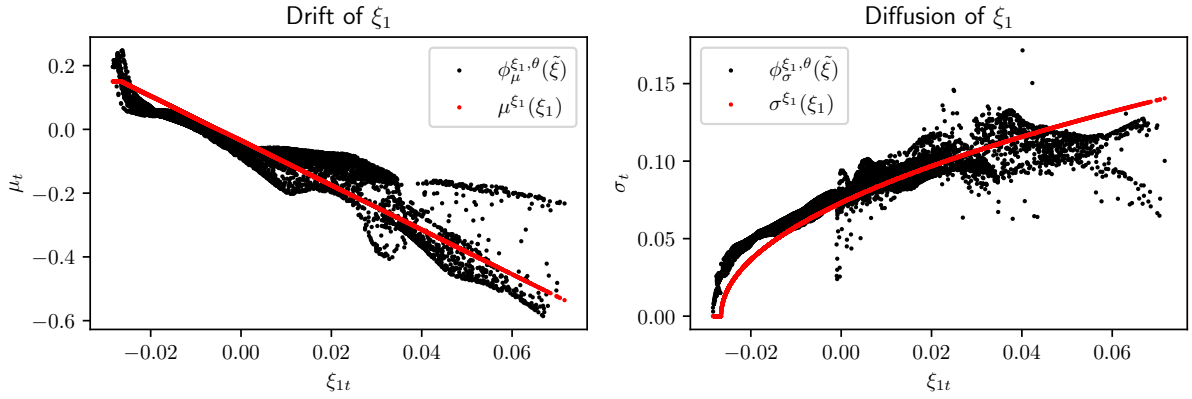


Figure 6.5.3: Comparison of the estimated drift and diffusion functions and the (approximated) ground-truth drift and diffusion functions for ξ_1 .

Market price of risk As discussed in Section 4.5, the no-arbitrage HJM-type drift restriction implies an over-determined linear system (4.8). Given the already estimated stock volatility γ_t and the approximated partial derivatives of \tilde{c} , we compute z_t and find an approximate solution to (4.8) using the ordinary least squares method,

$$\hat{\psi}_t = (\sigma_t^\top \mathbf{G} \mathbf{G}^\top \sigma_t)^{-1} \sigma_t^\top \mathbf{G} (\mathbf{G}^\top \mu_t - z_t). \quad (6.5)$$

The violation to the drift restriction can be measured by

$$\chi_t = \left\| \mathbf{G}^\top \mu_t - z_t - \mathbf{G}^\top \sigma_t \hat{\psi}_t \right\|_2. \quad (6.6)$$

We estimate the market price of risk ψ_t from the learnt drift and diffusion coefficients using (6.5) for all in-sample data points. In Figure 6.5.4, we show the heatmaps of the

violation of the HJM-type drift restriction χ , as defined in (6.6), and the size of the market price of risk $\|\hat{\psi}\|_2$. The calibrated market price of risk $\hat{\psi}$ initially appears large, given we have simulated under a risk-neutral measure; however, the use of a factor representation, forward translation of the leverage function, interpolation of prices, and the selection of a training path where S does not leave the region of the leverage function (as discussed in Section 6.1) may lead to a non-vanishing market price of risk in the underlying ground truth model. We also note, comparing with Figures 6.3.2, 6.6.1 and 6.6.3, that the most common values of ξ_1 are slightly negative, coinciding with the region where the market price of risk is low.




Figure 6.5.4: Statistics of the estimated market price of risk.

6.6 Out-of-sample simulation test

Many potential applications of our models involve the simulation of S and ξ , which can then be used for pricing or risk management. We take the trained model and simulate sample paths using a tamed¹² Euler scheme (see Hutzenthaler, Jentzen and Kloeden [93] or Szpruch and Zhang [140]), which is given by

$$\begin{aligned} S_{t+\Delta t} &= S_t + \frac{\mu^S(\tilde{\xi}_t)}{1 + |\mu^S(\tilde{\xi}_t)|\sqrt{\Delta t}}\Delta t + \frac{\sigma^S(\tilde{\xi}_t)}{1 + |\sigma^S(\tilde{\xi}_t)|\sqrt{\Delta t}}(W_{0,t+\Delta t} - W_{0,t}), \\ \xi_{t+\Delta t} &= \xi_t + \frac{\mu(\tilde{\xi}_t)}{1 + |\mu(\tilde{\xi}_t)|\sqrt{\Delta t}}\Delta t + \frac{\sigma(\tilde{\xi}_t)}{1 + \|\sigma(\tilde{\xi}_t)\|\sqrt{\Delta t}}(W_{t+\Delta t} - W_t). \end{aligned} \quad (6.7)$$

The values of μ^S , σ^S , μ and σ are approximated by the trained neural networks (along with the diffusion scaling \mathcal{G}_σ and drift corrections \mathcal{G}_μ).

¹²We include the taming method simply to ensure stability of simulations if our neural networks were to produce unusually large values for drifts or volatilities. In our examples, the tamed Euler scheme has very similar performance to the classical Euler–Maruyama scheme.

To demonstrate the learnt model's ability to simulate time series data that are alike the input data, we show an independent sample path of 10000 steps for $\tilde{\xi} = (S, \xi)$ in Figure 6.6.1, right.¹³ In the scatter plot on the left of Figure 6.6.1, we see that the dependence structure between ξ_1 and ξ_2 is well captured. In addition, the simulated factors remain within the no-arbitrage region, due to the hard constraints imposed on the drift and diffusion functions.

As expected, when the factor process is close to any no-arbitrage boundary, its diffusion component that is normal to the boundary will tend to vanish, while its drift will point inwards the no-arbitrage region. To illustrate this, we sample a few simulated factor data, and visualize their drift and diffusion coefficients in Figure 6.6.2.

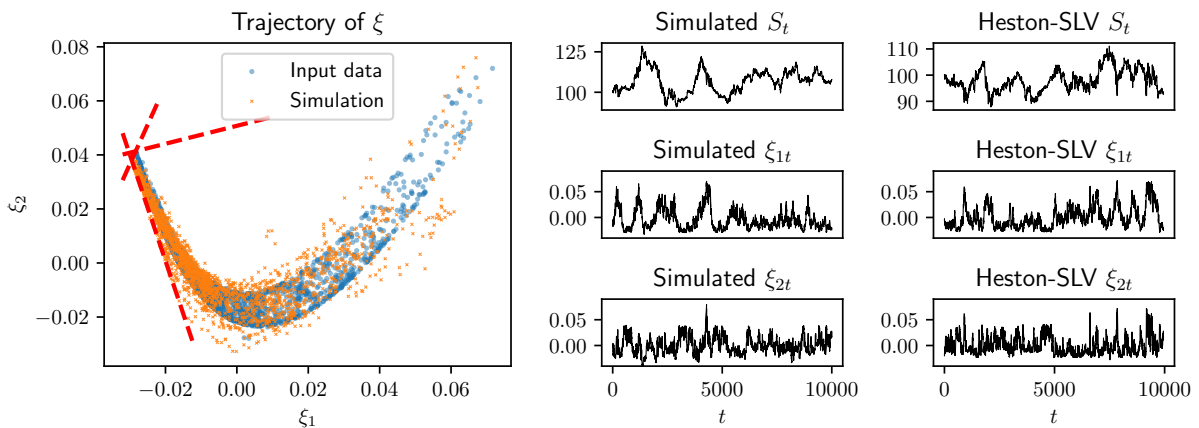


Figure 6.6.1: Simulation of S and ξ from the learnt neural network model, compared with the Heston-SLV model generated data.

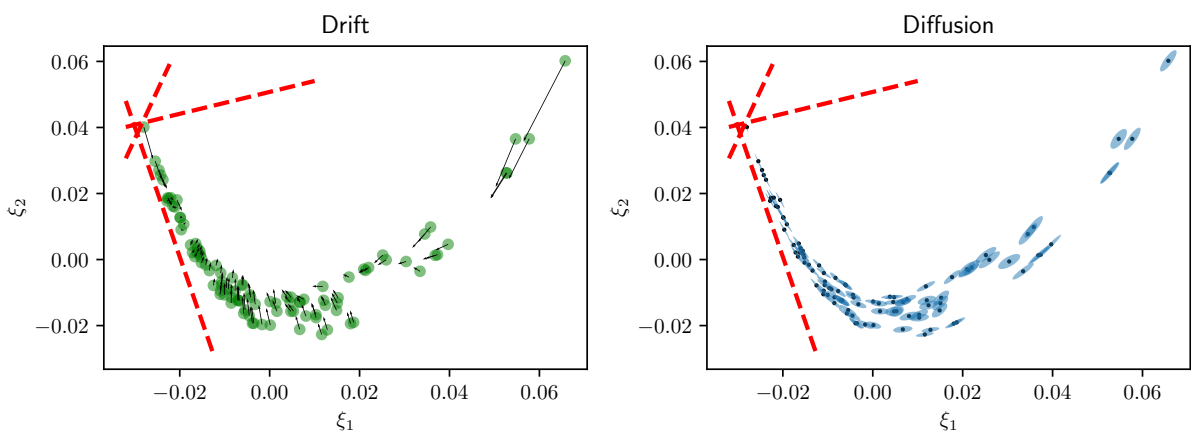


Figure 6.6.2: Drift vectors (arrows on the left plot) and diffusion matrices (ellipses representing the principal components of the diffusion on the right plot) for some randomly selected factor data points.

¹³There is no pathwise similarity here, as our learned factor model is based on a larger number of Brownian motions than the ground truth model.

In Figure 6.6.3, we compare the empirical distributions of the simulated log-return of S , ξ_1 and ξ_2 with those of the input data. We see that the learnt model is capable of generating realistic long time series data that is similar to the input data. More simulation results for implied volatilities can be found in Appendix 6.A, including the simulated paths of implied volatilities for a variety of option specifications.




Figure 6.6.3: Comparison of the marginal and joint distributions of the simulated ξ_1 and ξ_2 with the real distributions (generated from the Heston-SLV model).

VIX simulation In addition, we follow the CBOE VIX calculation methodology [41] to compute a *volatility index* from the simulated option prices. Specifically, the square of VIX is a linear combination of OTM call and put option prices, which can be further written as a linear combination of call prices only, provided that put-call parity holds under no-arbitrage.

Suppose the time t VIX index can be written as $\text{VIX}_t := \mathbf{h}^\top \mathbf{c}_t$ for some constant vector $\mathbf{h} \in \mathbb{R}^N$. Given simulated factors ξ_t^s , we first reconstruct prices \mathbf{c}_t^s and then compute VIX as

$$\text{VIX}_t^s = \mathbf{h}^\top \mathbf{c}_t^s = \mathbf{h}^\top \mathbf{G}_0 + \mathbf{h}^\top \mathbf{G}^\top \xi_t^s. \quad (6.8)$$

In Figure 6.6.4a, we plot the marginal and joint distributions of the log-return of S and the VIX index, for both the input data and the simulation data. While the distributions of the log-return of S are very close, our simulation gives a VIX of much lower kurtosis. This error is principally due to our factor representation, which does not align with the weights involved in the VIX calculation.

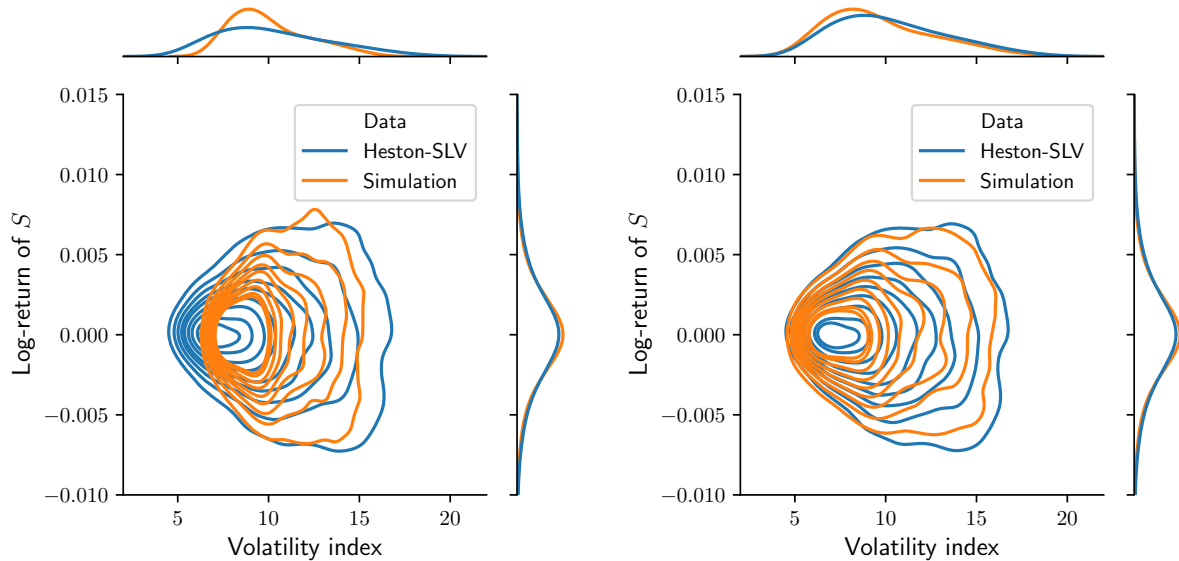
To improve the performance of our existing model for VIX simulation, we seek to minimise the impact of the factor reconstruction error. Consider fitting the linear regression model

$$\text{VIX}_t = \beta_0 + \boldsymbol{\beta}^\top \boldsymbol{\xi}_t + \varepsilon_t, \quad (6.9)$$

Using this relationship, we can compute VIX directly from the simulated factors $\boldsymbol{\xi}_t^s$ using

$$\text{VIX}_t^r = \hat{\beta}_0 + \hat{\boldsymbol{\beta}} \boldsymbol{\xi}_t^s. \quad (6.10)$$

We plot the distributions for this factor-regression based VIX in Figure 6.6.4b. Both its marginal distribution and the joint distribution with the log-return of S look reasonably similar to those of the Heston-SLV input data. This demonstrates that our model is capturing the dependence structure between the volatility index and the underlying S . In addition, the simulated time series of VIX and log-return of S are plotted in Figure 6.6.5. We see several occurrences of volatility clustering in the return series, which always coincide with high VIX values.



(a) Heston-SLV VIX v.s. VIX calculated from the simulated option prices.

(b) Heston-SLV VIX v.s. VIX^r calculated from the regression model on the simulated factors.

Figure 6.6.4: Joint distribution of the log-return of S and the VIX-like volatility index.

Remark 6.1. The nontrivial improvement in VIX simulation from (6.8) to (6.10) indicates that the decoded factors and their corresponding price basis do not optimise the representation of VIX. If one values a more accurate replication of VIX, an additional “VIX factor” could be included as a price basis vector in \mathbf{G} , with the objective of minimising the VIX reconstruction error and included in the model. (As VIX is here given by a linear combination of option prices, it can be directly included as a basis vector.)




Figure 6.6.5: Simulated time series of VIX^r , together with log-returns of the underlying stock.

6.7 Concluding remarks on the numerical example

We have chosen to use three factors (S, ξ_1, ξ_2) to represent a whole collection of call prices, while the ground-truth Heston-SLV model is a two-factor model with states (S, ν) . Our model adopts an extra factor to compensate for the restrictive linear representation (4.4) of normalised call prices in terms of latent factors. In fact, we see that ξ_1 has a strong linear relationship with ν (Figure 6.5.2), where most of the residuals from this linearity has been captured by ξ_2 . In addition, as seen in Figure 6.6.1, there seems to be a dominating polynomial relationship between ξ_2 and ξ_1 . For higher dimensional models, this polynomial relationship (which is forced by the use of linear factors) seems likely to cause the main difficulty in calibration, and appropriate methods to address this may be valuable. We develop an \mathbb{R}^{13} model and assess its simulation performance in Appendix 6.D.

While our focus here is on understanding the presence of true arbitrage in our model, in practice it may also be of concern if the model exhibits statistical arbitrage, that is, where a combination of options moves with very high probability in a given direction. This would correspond to the estimated market price of risk process $\hat{\psi}$ providing a perfect solution (so (6.6) is zero) but taking very large values. In practice, therefore, an additional regularization term to discourage large $\hat{\psi}$ may also be useful (for example the addition of $\|\hat{\psi}_t\|_2$ to (6.6)).

Appendix

6.A Implied volatility simulation

As a further description of the performance of our model, we consider the simulated time series of implied volatilities (IV) and IV smiles for multiple expiries and moneynesses. These are given in Figure 6.A.1 and Figure 6.A.2.

We see that the long-maturity and short-maturity-ITM implied volatilities are similar, but the short-maturity-OTM volatility processes have noticeably different behavior. This is also reflected in the varied shapes of the volatility smiles at different points in time.

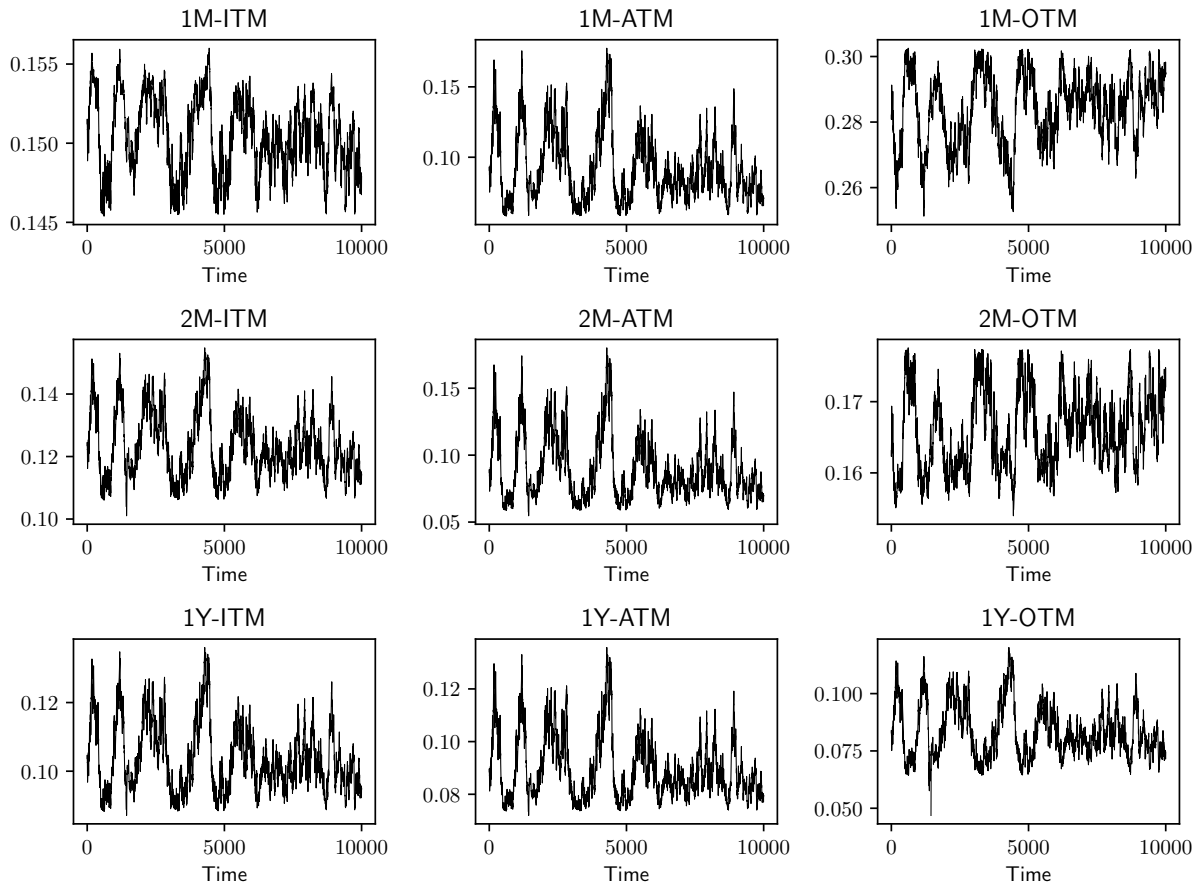


Figure 6.A.1: Simulated evolution of implied volatilities for multiple expiries and moneynesses. ATM: $m = 0$; ITM: $m = -0.0943$; OTM: $m = 0.1133$.

6.B Necessity of drift and diffusion operators

Given the underlying drift and diffusion functions are represented by neural networks, our approach is to shrink the diffusion and correct the drift near boundaries such the Friedman–Pinsky conditions (5.4) are satisfied. These conditions are sufficient for constraining the factor process within the state space where there is no static arbitrage.




Figure 6.A.2: Simulated implied volatility smiles for 1M, 2M, 6M and 1Y options on multiple days.

The question is whether it is necessary in practice to do the diffusion shrinking and drift correction. In particular, since the input data are already arbitrage-free, is the neural network capable of learning the correct boundary behaviours for the drift and diffusion functions without enforcing extra constraints?

To address this question, we estimate a model using exactly the same input data, neural network architecture and other technical details, except for neither shrinking the diffusion nor correcting the drift. The evolution of training and validation losses over epochs is shown in Figure 6.B.1. We find that, compared with Figure 6.4.1,

- (i) The losses converge faster. This is probably because that diffusion shrinking and drift correction constrain the gradients of the loss function when using backpropagation.
- (ii) The initial loss value is much larger. By shrinking diffusion and correcting drift at every epoch, we force the neural network training to start from and stay within the correct function subspace.




Figure 6.B.1: Evolution of training losses and validation losses (without imposing drift correction and diffusion shrinking).

We investigate how well the learnt model can simulate out-of-sample data. As shown in Figure 6.B.2, the simulated factors can easily break the static arbitrage constraints (the

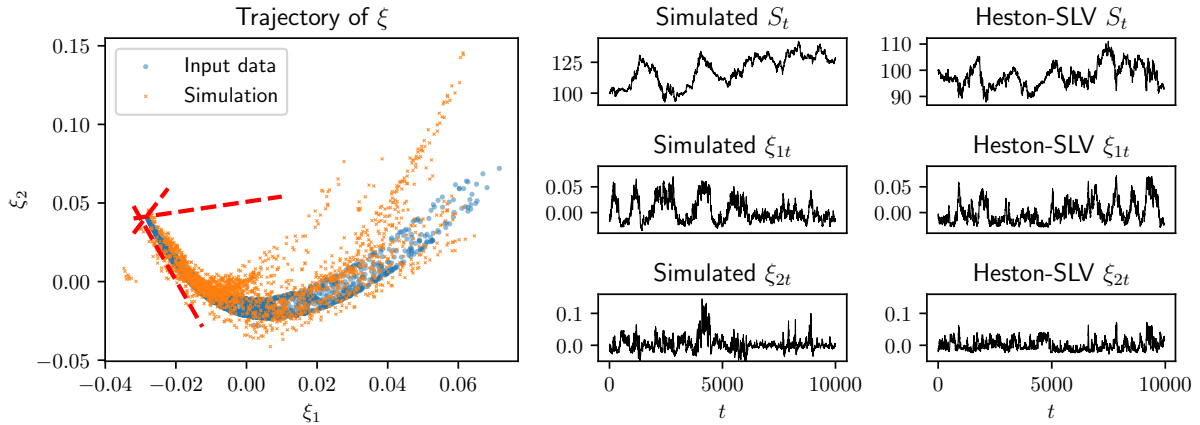


Figure 6.B.2: Simulation of S and ξ from the learnt neural network model (without imposing drift correction and diffusion shrinking).

dashed red lines), and generate very unrealistic values. Compared with the simulation results shown in Figure 6.6.1 (where the same random seed is used), drift correction and diffusion shrinking are effective at producing the correct boundary behaviours, which is crucial for an arbitrage-free model.

6.C Neural network sensitivity analysis

We demonstrate the sensitivity of the model estimation outcomes against a few neural network hyperparameters, including:

- Neural network depth: the number of hidden layers plus one (the output layer);
- Neural network width: the number of neurons per layer;
- Sparsity ratio: the fraction (between 0 and 1) of the network layers' weights that are pruned to zero;
- Activation function.

We call the network given in Section 6.4 the *benchmark* network. The values of the hyperparameters for the benchmark network for ξ are listed in Table 6.1.1. To carry out sensitivity analysis, we will train a collection of modified networks by varying one hyperparameter from the benchmark network.

Hyperparameters	NN depth	NN width	Sparsity ratio	Activation function
Values	3	256	0.5	ReLU

Table 6.C.1: Hyperparameters used for the benchmark model.

We compare the evolution of training losses and simulation performance between the benchmark network and the modified networks. The training losses are shown in Figure 6.C.1. We have shown in Figure 6.6.3 that the trained benchmark network simulates data with a similar distribution to the training data. To quantify this similarity between the distributions of simulated data and input data, we compute the Wasserstein distance of order one between the two distributions. We compare the computed distances for the benchmark network and the modified networks and list the results in Table 6.C.2.

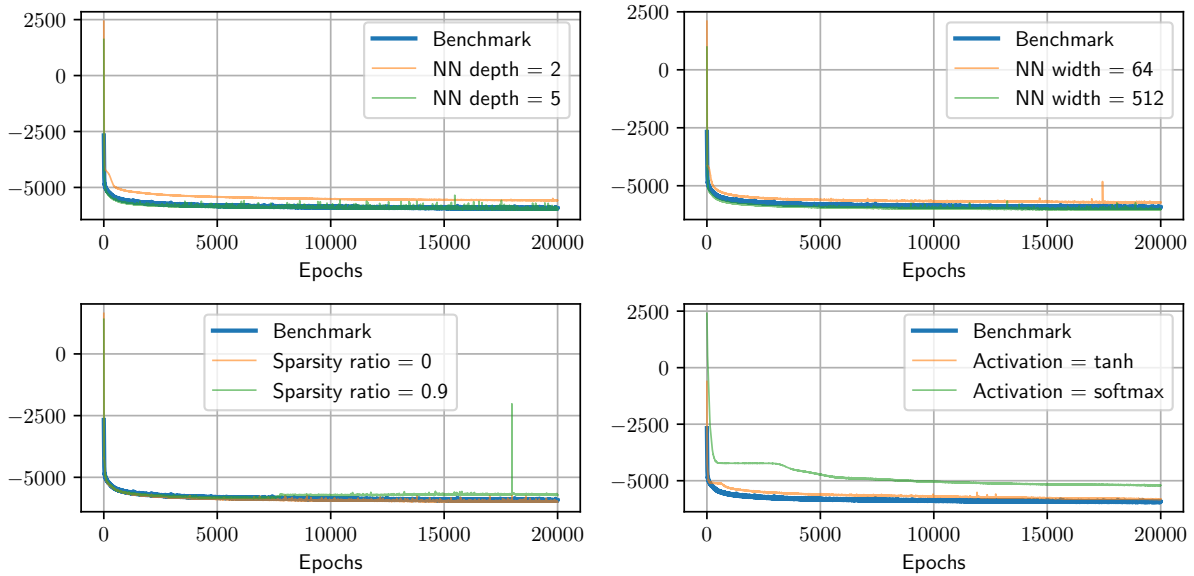


Figure 6.C.1: Evolution of training losses under different neural network architectures.

Neural network	Log-return of S	ξ_1	ξ_2
Benchmark	1.71	40.44	20.30
NN depth = 2	3.79	176.89	71.99
NN depth = 5	2.83	55.07	52.14
NN width = 64	2.32	52.46	29.17
NN width = 512	1.82	58.90	24.97
Sparsity ratio = 0	1.51	53.92	57.47
Sparsity ratio = 0.9	3.68	69.69	105.61
Activation = tanh	2.17	46.57	35.26
Activation = softmax	2.97	148.58	462.74

Table 6.C.2: Wasserstein distances ($\times 10^{-4}$) between the empirical distributions of simulated data and the input data.

The training losses of all networks drop fairly quickly during the first 1000 epochs, and then gradually converge. Together with the simulation performance, we find that:

- The 2-layer shallow neural network gives higher convergent losses and much worse simulation results. Increasing the depth of the neural network, or varying its width,

has a relatively small effect on performance, with mixed direction.

- The ReLU activation function generally produces smaller convergent losses and more similar simulations (i.e. smaller Wasserstein distances). The performance using a softmax activation function (as implemented by Tensorflow) was surprisingly poor.
- Pruning the network with moderate sparsity improves the extrapolation capability of the learnt neural network, thus resulting in better simulation results for ξ .

6.D Higher-dimensional models

In our numerical example, we observe that representing the call price data with more factors can further reduce reconstruction error, which improves the accuracy of simulations for option portfolios, such as VIX. However, a few challenges arise:

- The inclusion of more factors leads to higher-dimensional models; and the number of unknown functions to estimate increases quadratically with the number of factors, increasing the chance of overfitting and making the estimates less robust.
- There are dominating polynomial relationships between factors, specifically because the call price data are generated from a two-factor model.

These issues led to poor performance of a high dimensional version of our model, in particular, highly unrealistic simulations of prices.

To overcome these issues, we assume that some *secondary* factors could be expressed as polynomials of some *primary* factors plus noise. Specifically, suppose there are d' primary factors and $d - d'$ secondary factors. For each secondary factor ξ_i , where $i = d' + 1, \dots, d$, we have an independent Brownian motion W_i , and the model

$$\xi_i = f_i(\xi_1, \dots, \xi_{d'}) + \varepsilon_i, \text{ and } d\varepsilon_{it} = \kappa_i(\theta_i - \varepsilon_{it}) dt + \varsigma_i dW_{it}. \quad (6.11)$$

In other words, noises are assumed to be mean-reverting Ornstein-Uhlenbeck (OU) processes with constant parameters κ , θ and ς . Here f is calibrated to the observed dominating polynomial relations between factors. Let ξ^p be the collection of all primary factors, with drift $\mu^p : \mathbb{R}^{d'+1} \rightarrow \mathbb{R}^{d'}$ and diffusion $\sigma^p : \mathbb{R}^{d'+1} \rightarrow \mathbb{R}^{d' \times d'}$. Consequently, the dynamics for the secondary factor ξ_i is

$$\begin{aligned} d\xi_i &= \mu_i(\xi^p, \xi_i) dt + \sigma_i(\xi^p) dW_t^p + \varsigma_i dW_{it}, \\ \text{with } \begin{cases} \mu_i &= \kappa_i(\theta_i - \xi_i + f_i) + (\nabla f_i)^\top \mu^p + \frac{1}{2} \sum_{j,k} \Gamma_{jk}, \\ \sigma_i &= (\nabla f_i)^\top \sigma^p, \end{cases} \end{aligned} \quad (6.12)$$

where $\Gamma = (\Gamma_{jk}) = [\sigma^p(\sigma^p)^\top] \odot \nabla^2 f_i$ is the element-wise product of the covariance matrix and the Hessian matrix of f_i , $W_t^p = [W_{1t} \cdots W_{d't}]^\top$ and $\nabla = [\partial/\partial\xi_1 \cdots \partial/\partial\xi_{d'}]^\top$ is the gradient operator.

Factor decoding We keep the two factors in our numerical example as the primary ones, and append another 11 secondary factors, decoded consecutively with the objective of maximizing statistical accuracy using Algorithm 1. We choose to append 11 secondary factors because this is the minimal number of factors which ensures no static arbitrage, as seen in Table 6.D.1. Each secondary factor is modelled as a cubic function of (ξ_1, ξ_2) plus OU noises. In Figure 6.D.1, we show the scattergrams for ξ_7 and the primary factors (ξ_1, ξ_2) , and compare with those for the polynomial implied values. The residual time series is also plotted.

Metrics	Number of secondary factors											
	1	2	3	4	5	6	7	8	9	10	11	12
MAPE (%)	2.82	1.82	1.66	1.31	0.59	0.55	0.28	0.24	0.21	0.17	0.13	0.10
PDA (%)	2.81	0.89	0.77	0.70	0.60	0.59	0.52	0.38	0.31	0.28	0.23	0.21
PSAS (%)	27.51	20.88	18.25	15.33	12.35	10.92	7.65	7.45	2.09	2.01	0.00	0.00

Table 6.D.1: MAPE, PDA and PSAS metrics when appending different number of secondary factors.




Figure 6.D.1: Polynomial fitting results for ξ_7 .

Model estimation strategy We estimate the neural-SDE model for primary factors and the OU noise processes for secondary factors separately rather than jointly. Since primary factors are much more representative of the input price data, it is crucial to learn their dynamics correctly, while joint estimation with secondary factors dilute the weight of primary factors in the objective likelihood function. Precisely, with the 2 primary factors and the 11 secondary factors, we

- (i) estimate a neural-SDE model for the 2 primary factors, as we did in Section 6.4;
- (ii) calibrate a cubic function $f_i(\xi_1, \xi_2)$ for each secondary factor using the least-squares method, and compute the corresponding residual time series data $\{\varepsilon_{it_l}\}_{l=1, \dots, L+1}$;

- (iii) discretise the OU model for ε_i using the Euler-Maruyama scheme, which yields an AR(1) model with normal noises, and fit with the residual time series data.

Arbitrage-free simulation Unlike the neural-SDE model for primary factors, the model for secondary factors (6.12) does not necessarily rule out static arbitrage, particularly because its drift and diffusion do not satisfy Friedman and Pinsky’s condition outlined in (5.4). Therefore, when forward simulating factors, there needs to be additional transformations on particularly secondary factors to ensure no-arbitrage. Rather than scaling diffusion and correcting drift of (6.12), we follow a simpler strategy, considering that secondary factors are strongly dependent on primary factors and have much smaller magnitudes.

Let \mathcal{P} be the d -polytope factor state space where there is no static arbitrage. Suppose \mathcal{P}^p is the arbitrage-free d' -polytope state space for primary factors, then \mathcal{P}^p is the affine projection of $\mathcal{P} \subset \mathbb{R}^d$ onto $\mathbb{R}^{d'}$ given by $\mathcal{P}^p = \{y \in \mathbb{R}^{d'} : \exists x \in \mathcal{P}, y = [x_1 \dots x_{d'}]^\top\}$. Therefore, if we use the neural-SDE model to simulate arbitrage-free primary factors, then there must exist some secondary factors such that the full set of factors reproduces arbitrage-free call prices. Hence, we are able to simulate arbitrage-free factors by:

- (i) simulating primary factors using the arbitrage-free neural-SDE model;
- (ii) for each secondary factor, computing the polynomial term from simulated primary factors, and simulating noises from the estimated AR(1) model;
- (iii) adding up the polynomial and noise terms for each secondary factor; if needed, perturb the collection of secondary factors using the arbitrage repair algorithm proposed in Chapter 3.

We simulate factors from the estimated models, where the simulation results for the 2 primary factors have been presented in Section 6.6. In Figure 6.D.2, we compare the marginal distributions of input data, simulated data and arbitrage-repaired data for the 11 secondary factors. We see that repairing arbitrage does not significantly affect the marginal distributions of most factors.

Finally, we compute the VIX index using the simulated factor through (6.8), and compare its distribution with that of the Heston-SLV VIX. As seen in Figure 6.D.3, its marginal distribution and the joint distribution with the log-return of S look reasonably close to those of the Heston-SLV VIX. Repairing arbitrage has trivial impact on the distributions. Compared with Figure 6.6.4a, this higher-dimensional model greatly improves the VIX simulation, indicating that the inclusion of more factors leads to better option price reconstruction, thus a more accurate replication of VIX.

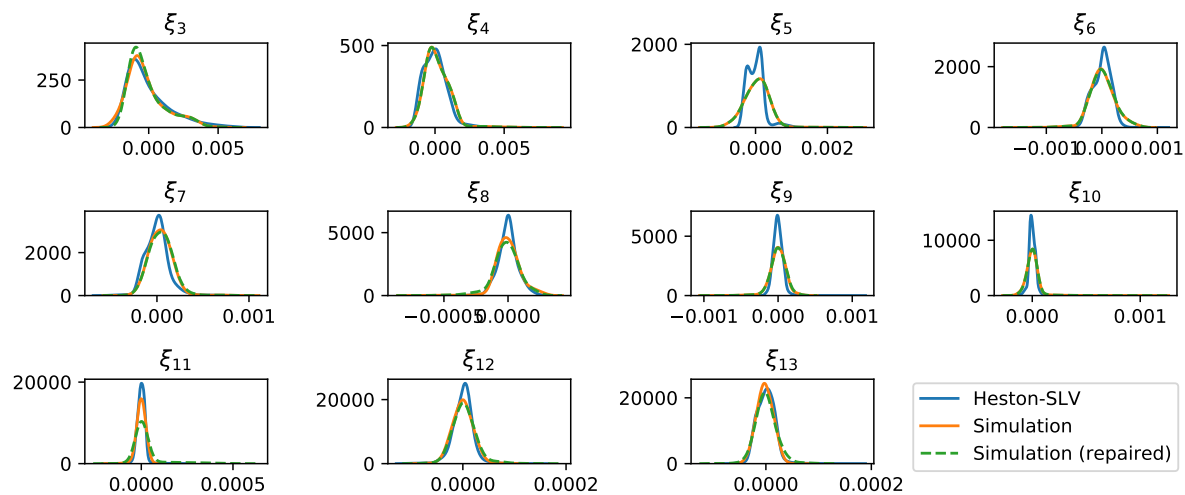
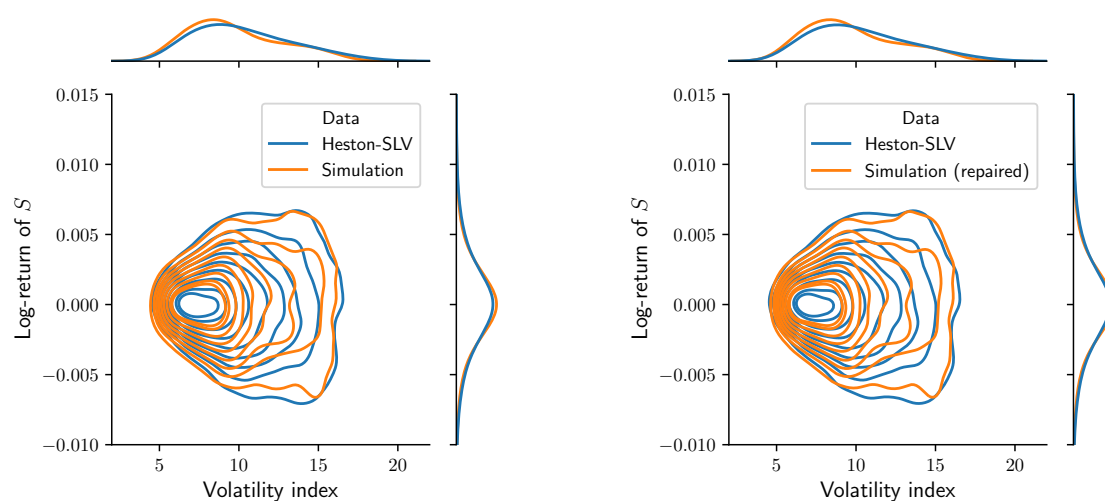


Figure 6.D.2: Marginal distributions of secondary factors.



(a) Heston-SLV VIX v.s. VIX calculated from the simulated 13 factors.

(b) Heston-SLV VIX v.s. VIX calculated from the simulated 13 factors with arbitrage repair.

Figure 6.D.3: Joint distribution of the log-return of S and the VIX-like volatility index.

Calibrating to real data: EURO STOXX 50 and DAX

We have seen decent performance of our modelling and inference approach when applied to synthetic data generated from a stochastic local volatility model (Jex, Henderson and Wang [98]), as given in Chapter 6. The focus of this chapter is to investigate its ability to capture the dynamics of real-world data, and whether practically useful conclusions can be drawn from these methods.

7.1 Option price data description and pre-processing


We study the two most traded vanilla European options listed on Eurex, namely EURO STOXX 50[®] index options and DAX[®] options. OptionMetrics' IvyDB Europe [123] offers historical daily settlement prices¹ for calls with expiries of 30, 60, 91, 122, 152, 182, 273, 365, 547, and 730 calendar days, at deltas of 0.2, 0.25, 0.3, 0.35, 0.4, 0.45, 0.5, 0.55, 0.6, 0.65, 0.7, 0.75, and 0.8. We collect daily call option prices from 2nd January, 2002 to 30th December, 2019², as well as corresponding futures curves and euro zero curves.

We first normalise call prices by discount factors and futures prices following (4.3). However, the normalised call price surfaces have fixed delta parameters, rather than fixed moneynesses, over time. Therefore, we choose the historical *median* moneynesses (for each delta-quoted option) to define our liquid option lattice \mathcal{L}_{liq} , and interpolate prices to the fixed set of moneyness parameters. We give details on the interpolation in Appendix 7.A. In addition, whenever there is static arbitrage in the reported historical option prices, we perturb the prices using the arbitrage repair algorithm, introduced in Chapter 3. In Appendix 7.A, we give statistics on the presence of arbitrage in raw and interpolated

¹These prices are calculated from implied volatilities that are interpolated using a methodology based on a kernel smoothing algorithm, according to the OptionMetrics' IvyDB Europe reference manual [123].

²By the time this chapter, as well as the associated published work [47], was written, the OptionMetrics database had yet to publish more recent data. However, we have incorporated newly released data up to 30th September, 2021 in the hedging performance analysis in Chapter 9.

option price data, and assess the impact of repairing arbitrage. We show the liquid lattice \mathcal{L}_{liq} in Figure 7.1.1a.

(a) Liquid lattice \mathcal{L}_{liq} .

(b) Historical daily prices (underlying indices and 1M ATM call options).

Figure 7.1.1: Scattergram of liquid option lattice and historical prices.

In Figure 7.1.1b, we show historical daily prices for the EURO STOXX 50 index, the DAX index, and the 1M at-the-money (ATM) normalised call options written on the two indices. Since the normalised option price has a deterministic bijective relation with the implied volatility, we see that the two options have very similar volatility dynamics in history. It is consequently reasonable to assume a common background model for the volatility processes of the two options. Therefore, we concatenate the normalised call price data of the EURO STOXX 50 index options and DAX options when decoding factors and learning factor dynamics. This enriches the dataset available for our statistical estimation problems by assuming that our market models are universal for options written on these two different underlyings. Specifically, there is only one functional form for the drift $\mu(\cdot)$ and the diffusion $\sigma(\cdot)$ for the factor dynamics, despite different option underlyings. Nevertheless, the underlying indices' models will not be restricted to be the same; as evidenced in Figure 7.1.1b, DAX index appears to have a more positive drift³. Hence, we build a universal model for the factor dynamics but separate models for each index. To better accommodate the need of training models with concatenated option price, we revisit the original SDE model specified in (4.2) and make modifications in Section 7.2.

Using $k \in \{U, D\}$ to distinguish data of EURO STOXX 50 index options and DAX options, we denote the observation times as $0 = t_1 < \dots < t_{L^k} = T^k$ and let $\delta t = t_i - t_{i-1} = T^k / (L^k - 1)$ for all $2 \leq i \leq L^k$. We split the L^k historical data samples into *training/estimation* samples $\{t_1, \dots, t_{E^k}\}$ and *testing* samples $\{t_{E^k+1}, \dots, t_{L^k}\}$. We use only the training samples for estimating a neural-SDE market model, while the testing

³Two factors contribute to the drift difference in the two indices: in addition to the fact that the two equity indices have different constituent stocks, EURO STOXX 50 is a price return index while DAX is a total return index.

samples will be used for testing performance of out-of-sample risk simulation, as will be seen in Section 8.3. In Table 7.1.1, we list the details of the data used.

Parameters	#training samples E^k	#testing samples $L^k - E^k$	#options N	δt
Values	4296 (2002-01-01–2018-12-31)	253 (2019-01-01–2019-12-31)	130	1 day ⁴

Table 7.1.1: Details of the option data used.

7.2 Revisiting the neural-SDE market models

Here we briefly re-state our market models for reminding readers the major notations and mathematical structures to be used in this chapter.

We model the joint dynamics of the stock price S and the factors ξ by a $(d + 1)$ -dimensional diffusion process solving the following time-homogeneous SDE:

$$\frac{dS_t}{S_t} = \alpha(\xi_t) dt + \gamma(\xi_t) dW_{0,t}, \quad S_0 = s_0 \in \mathbb{R}; \quad (7.1a)$$

$$d\xi_t = \mu(\xi_t) dt + \sigma(\xi_t) dW_t, \quad \xi_0 = \zeta_0 \in \mathbb{R}^d, \quad (7.1b)$$

where $W_0 \in \mathbb{R}$ and $W = [W_1 \ \cdots \ W_d]^\top \in \mathbb{R}^d$ are standard independent Brownian motions under the real-world measure \mathbb{P} . We use α_t , γ_t , μ_t and σ_t to denote $\alpha(\xi_t)$, $\gamma(\xi_t)$, $\mu(\xi_t)$ and $\sigma(\xi_t)$, respectively. We denote by $L_{\text{loc}}^p(\mathbb{R}^d)$ the space of all \mathbb{R}^d -valued, progressively measurable, and locally p -integrable (in t , \mathbb{P} -a.s.) processes, and assume that $\alpha \in L_{\text{loc}}^1(\mathbb{R})$, $\mu \in L_{\text{loc}}^1(\mathbb{R}^d)$, $\gamma \in L_{\text{loc}}^2(\mathbb{R})$ and $\sigma \in L_{\text{loc}}^2(\mathbb{R}^d)$.

Remark 7.1. We assume, in the model of ξ (7.1b), that the drift and diffusion coefficients are independent of S , in the sense that volatility surfaces (represented by ξ) are unaffected by the underlying price level. Nevertheless, it is possible, and simple, to add S as an additional argument and train the model with neural nets (as is done in Chapter 4).

Remark 7.2. For simplicity, we assume independent Brownian motions that drive the randomness in S and ξ ; one can easily add a correlation structure and estimate the resulting larger covariance matrix, albeit with a greater risk of over-fitting. In Appendix 7.C.2, we discuss the performance when estimating models for S and ξ jointly with a full $\mathbb{R}^{(d+1) \times (d+1)}$ covariance matrix. With the current model setup (7.1), in order to take into account non-negligible correlations between ξ and log-returns of S for realistic forward simulations, we randomly sample innovations from the historical residuals of the fitted models (7.1), rather than from standard Gaussians; see the discussion in Section 7.7.

⁴We use the Act/365 fixed day count convention for converting $\delta t = 1/365$ to year fraction.

7.3 Factor decoding

Suppose we observe price data $\mathbf{C}^k = (C_{lj}^k) \in \mathbb{R}^{L^k \times N}$ with $C_{lj}^k = \tilde{c}_{t_l}^k(\tau_j, m_j)$ for $k \in \{E, D\}$. Given the concatenated data $\mathbf{C} = [\mathbf{C}^E; \mathbf{C}^D] \in \mathbb{R}^{L \times N}$ with $L = L^E + L^D$, we aim to construct factor data $\mathbf{\Xi} = (\Xi_{li}) \in \mathbb{R}^{L \times d}$ with $\Xi_{li} = \xi_{i, t_l}$. In particular, we represent \mathbf{C} , with residuals $\mathbf{\Upsilon} \in \mathbb{R}^{L \times N}$, by

$$\mathbf{C} = \mathbf{1}_L \otimes \mathbf{G}_0 + \mathbf{\Xi} \mathbf{G} + \mathbf{\Upsilon}, \quad (7.2)$$

where $\mathbf{1}_L$ is an L -vector of ones, and \otimes denotes the outer product of two vectors.

Vega-weighted prices

We call $\mathbf{C} - \mathbf{\Upsilon}$ the *reconstructed prices*, and the Frobenius norm $\|\mathbf{\Upsilon}\|_F$ the *reconstruction error*. As in-the-money (ITM) and long-dated options have higher prices, the objective of minimising $\|\mathbf{\Upsilon}\|_F$ is naturally more effective in reducing relative calibration errors for ITM and long-dated option prices. To avoid overfitting ITM and long-dated option prices, practitioners often weight each option by the reciprocal of its Black–Scholes vega. In fact, using the approximation $\Delta C \approx \text{vega} \times \Delta \sigma^{\text{imp}}$, weighting prices by $1/\text{vega}$ approximately corresponds to measuring reconstruction errors in implied volatilities. To follow this approach, we introduce a constant weight matrix $\mathbf{\Lambda} = \text{diag}(\lambda_1, \dots, \lambda_N)$, and construct factors from $\mathbf{C} \mathbf{\Lambda}$ instead, as the following two optimisation problems are equivalent:

$$\min_{\mathbf{\Xi}, \mathbf{G}} \left\| (\mathbf{C} - \mathbf{1}_L \otimes \mathbf{G}_0 - \mathbf{\Xi} \mathbf{G}) \mathbf{\Lambda} \right\| \iff \min_{\hat{\mathbf{\Xi}}, \hat{\mathbf{G}}} \left\| \mathbf{C} \mathbf{\Lambda} - \mathbf{1}_L \otimes \hat{\mathbf{G}}_0 - \hat{\mathbf{\Xi}} \hat{\mathbf{G}} \right\|.$$

Now we associate the weight λ with the option vega. At time t , the Black–Scholes vega for the (τ, m) -option is

$$\mathcal{V}_t(\tau, m) = \frac{d\tilde{c}_t}{d\sigma_t^{\text{imp}}}(\sigma_t^{\text{imp}}; \tau, m) = \sqrt{\tau} \phi \left(-\frac{m}{\sigma_t^{\text{imp}} \sqrt{\tau}} + \frac{1}{2} \sigma_t^{\text{imp}} \sqrt{\tau} \right) > 0, \quad (7.3)$$

where $\phi(\cdot)$ is the density of the standard normal distribution. This leads us to set

$$\lambda_j = \frac{1}{\bar{\mathcal{V}}_j}, \text{ for each } j, \text{ where } \bar{\mathcal{V}}_j = \frac{1}{L} \sum_{l=1}^L \mathcal{V}_{t_l}(\tau_j, m_j).$$

Remark 7.3. Weighting price data by a constant matrix does not change the factor decoding algorithm (Algorithm 1 in Chapter 4), except that the coefficient and constant terms of the static arbitrage constraints (4.11) need to be revised. Specifically, we aim for a new factor representation $\mathbf{\Lambda} \mathbf{c}_t = \mathbf{G}_0 + \mathbf{G}^\top \xi_t$. Then the static arbitrage constraints $\mathbf{\Lambda} \mathbf{c}_t \geq \hat{\mathbf{b}}$ yield

$$(\mathbf{\Lambda} \mathbf{\Lambda}^{-1}) \mathbf{G}^\top \xi_t \geq \mathbf{b}^{\text{new}} := \hat{\mathbf{b}} - (\mathbf{\Lambda} \mathbf{\Lambda}^{-1}) \mathbf{G}_0.$$

Nevertheless, for notational simplicity, we write $\mathbf{C} \leftarrow \mathbf{C} \mathbf{\Lambda}$, $\mathbf{A} \leftarrow \mathbf{\Lambda} \mathbf{\Lambda}^{-1}$ and $\mathbf{b} \leftarrow \mathbf{b}^{\text{new}}$ for the rest of the chapter.

Primary and secondary factors

While decoding more factors can improve the representation of call prices, it also leads to higher-dimensional models, which may over-parametrise the dynamics of noisy factors and lead to over-fitting. To overcome these issues, we assume that some *primary* factors are modelled by the SDE (7.1) and other *secondary* factors are modelled as mean-reverting Ornstein–Uhlenbeck (OU) processes.

We now investigate how many primary and secondary factors are needed for a reasonable representation of the call price data. In Section 6.3, we define three metrics to examine the quality of factor reconstruction, which are MAPE (mean absolute percentage error), PDA (proportion of dynamic arbitrage) and PSAS (proportion of statically arbitrageable samples). As discussed in the previous section, we use vega-weighted versions of these quantities in what follows. Applying the factor decoding algorithm (Algorithm 1 in Chapter 4), we show the three metrics for a few combinations of using up to two factors in Table 7.3.1. In Appendix 7.B, we give details on how we estimate z_t (defined in (4.8)), which is required for computing PDA and decoding dynamic arbitrage factors.

Factors	MAPE	PDA	PSAS
Dynamic arb.	10.75%	2.20%	33.74%
Dynamic arb. + Statistical acc.	4.61%	1.92%	8.26%
Dynamic arb. + Static arb.	4.94%	1.99%	1.29%
Statistical acc.	5.40%	13.07%	7.13%
Statistical acc. + Dynamic arb.	5.40%	1.80%	8.13%
Statistical acc. + Static arb.	4.82%	7.83%	0.30%

Table 7.3.1: MAPE, PDA and PSAS metrics when including different combinations of factors.

Using only two factors, i.e. one statistical accuracy factor and one static arbitrage factor (last row in the table), we can represent the whole collection of call prices with reasonable accuracy. Note that the use of a static arbitrage factor is significant in reducing the number of violations of the static arbitrage constraints, as evidenced by the corresponding reduction in PSAS. Since arbitrageable data samples are invalid inputs for the model inference and need to be thrown away, it is crucial to ensure a low PSAS. The high value of PDA is of less concern for real data and will be reduced further by including more secondary factors. Hence, we will restrict our attention to this simple three-factor model (i.e. two factors ξ , in addition to the stock index price), as it also allows us to demonstrate qualitative features of the model easily. See, however, Appendix 7.C.1 for the analysis of a market model with three primary factors.

We plot the price basis functions of these two factors, denoted as G_1 and G_2 , as well as G_0 , the constant term of \tilde{c} , in Figure 7.3.1. The points in the liquid lattice (in (τ, m) coordinates) are also shown on this plot. Here the real-valued functions G_1 and G_2

are obtained by interpolating the price basis vectors \mathbf{G}_1 and \mathbf{G}_2 , and G_0 is obtained by interpolating the normalised call prices averaged over time.

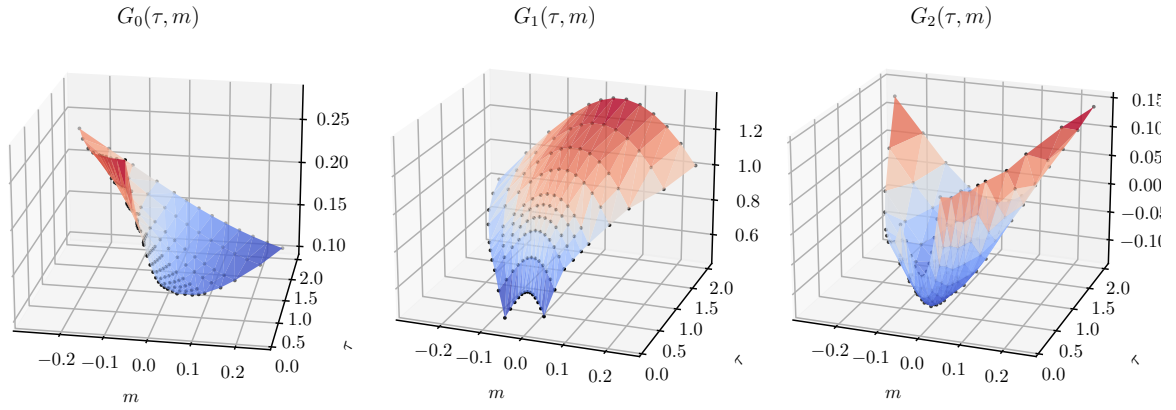


Figure 7.3.1: Price basis functions of the normalised call price surface.

Given these factors, we use the linear programming method (Caron, McDonald and Ponik [36]) to eliminate redundant constraints in the system $\mathbf{A}\mathbf{G}^\top \xi \geq \mathbf{b}$, which is the projection of the original no-arbitrage constraints, constructed in price space, to the \mathbb{R}^2 factor space. This results in only 8 constraints, which we indicate as red dashed lines in Figure 7.3.2. The convex polygonal domain bounded by these constraints (light green area) is the statically arbitrage-free zone for the factors, that is, provided the factor process remains in this region, we are guaranteed to have no static arbitrage in the reconstructed call prices on our liquid lattice \mathcal{L}_{liq} . Shown in Figure 7.3.2 is also the cloud of factor realisation decoded from the market data. We observe that it has a qualitatively similar shape to that in Figure 6.3.2 of Chapter 6 for simulated prices from a Heston SLV model.




Figure 7.3.2: Trajectory (black dots) of the primary \mathbb{R}^2 factors and the corresponding static arbitrage constraints (red dashed lines) projected to the \mathbb{R}^2 factor space.

To further reduce the factor reconstruction error, we include a few more secondary factors, whose dynamics are modelled as mean-reverting OU processes. We take the

residuals from the reconstruction of call prices with two primary factors, and decode secondary factors consecutively (using PCA) from the residuals, with the objective of maximising statistical accuracy. In Figure 7.3.3, we show how various metrics change with the inclusion of more secondary factors. In the left most plot, the factor magnitude is defined as the ratio of the min-max range of the corresponding factor over that of the first primary factor, provided unit-norm price basis vectors. It measures the relative importance of each secondary factor compared to the first primary factor, in reconstructing option prices. While MAPE and PDA decrease with inclusion of more secondary factors, PSAS is not always reduced. As an example, we choose to include 13 secondary factors resulting in MAPE = 0.48%, PDA = 0.06% and PSAS = 2.42%.




Figure 7.3.3: Price reconstruction metrics with the inclusion of more secondary factors.

Remark 7.4. PSAS initially increases with more secondary factors included and drops to zero eventually. This is because adding secondary factors leads to small perturbations of the price surface, potentially violating monotonicity and convexity and thus resulting in (small) arbitrages in reconstructed prices. In other words, the OU model for the secondary factors does not necessarily rule out static arbitrage, particularly because their drift and diffusion do not satisfy Friedman and Pinsky’s condition. This is less important when choosing the number of secondary factors, firstly as we no longer need to remove arbitrageable data points to estimate their dynamics. Moreover, the magnitudes of secondary factors are much smaller than the primary ones (left most plot in Figure 7.3.3). In addition, we could always apply the arbitrage repair algorithm (as stated in Chapter 3) to perturb our simulated secondary factors in order to restore monotonicity and convexity.

7.4 Factor dynamics estimation

The estimation of the SDE model for primary factors (7.1b) with static arbitrage constraints has been formulated as an unconstrained deep learning problem in Chapter 5. In this study, we use the network architecture

$$\phi^\theta = \mathcal{F}_2 \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{256} \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{256} \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{256}, \quad (7.4)$$

with weight and bias parameters θ , where \mathcal{F}_x is a fully connected layer, or affine transformation, with x units, and \mathcal{A}_{xxx} is an activation function. Each layer \mathcal{F} is parametric, but we omit the parameters for notational simplicity. In addition, to mitigate over-fitting problems, we train both networks with 50% sparsity, meaning that the 50% smallest weights are pruned to zero. We implement and train our model using the standard tools within the Tensorflow [2] environment. In Figure 7.4.1, we show the evolution of training losses and validation losses⁵ over epochs during the training of the model. The loss value has a rapid decline for the first a few epochs, and then gradually converges.




Figure 7.4.1: Evolution of training losses and validation losses for the model of the primary factors.

Next, we sample a few factor data points, and visualise their drift and diffusion coefficients in Figure 7.4.2. As observed in the trajectory of the factor data in Figure 7.3.2, the dataset is distributed around a lower-dimensional manifold. For samples on the periphery of the dataset, the learnt drifts tend to point inwards to the manifold, while the principal direction of their diffusions tend to align with the closest boundary. This broadly agrees with the phenomena seen in synthetic data in Chapter 6 (in particular, Figure 6.6.2).

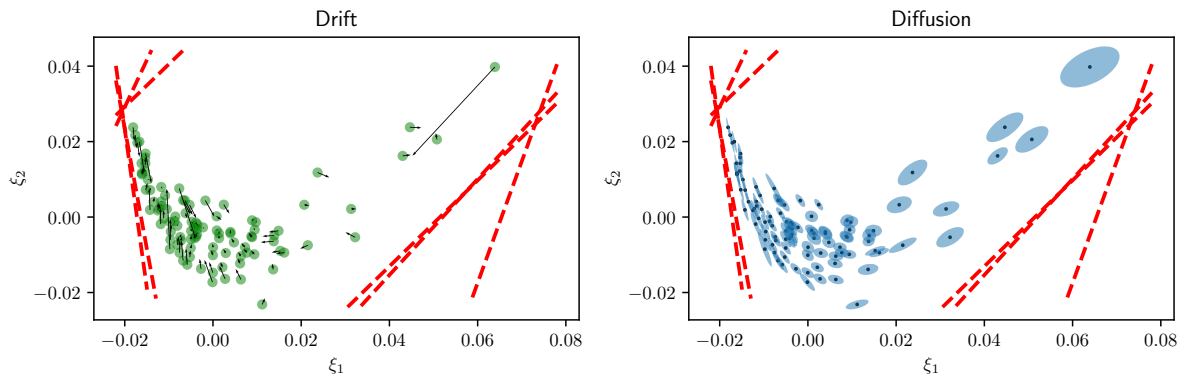


Figure 7.4.2: Drift vectors (arrows on the left plot) and diffusion matrices (ellipses representing the principal components of the diffusion on the right plot) for some randomly selected factor data points.

⁵The first 90% of the training samples are used for training and the last 10% are reserved as validation data.

For each of the secondary factors, we discretise the OU model using the Euler–Maruyama scheme, which yields an AR(1) model with normal noises, and then fit with the factor time series data using maximum likelihood estimation.

7.5 Underlying index dynamics estimation

Let $X_t = \ln S_t$, then the model for S (7.1a) implies, by Itô's lemma,

$$dX_t = r(\xi_t) dt + \gamma(\xi_t) dW_{0,t}, \text{ where } r(\xi_t) = \alpha(\xi_t) - \frac{1}{2}\gamma^2(\xi_t).$$

Since drift estimation is well known to be noisy, we assume a constant drift for X_t , i.e., $r(\xi_t) \equiv r$. With this assumption, an unbiased estimate for r is

$$\hat{r} = \frac{1}{\delta t} \left(\frac{1}{L-1} \sum_{t=1}^{L-1} \ln \left(\frac{S_{t+1}}{S_t} \right) \right) = \frac{1}{T} \ln \left(\frac{S_T}{S_0} \right). \quad (7.5)$$

Hence, for each underlying equity index, we first estimate a constant drift using historical log-returns, and then train a neural-SDE model with only its diffusion represented by a neural network. In particular, for our training samples (from 2002-01-01 to 2018-12-31), $\hat{r}^U = -0.027$, $\hat{r}^D = 0.071$. Compared with the model for the primary factors, we use a smaller network $\phi^{S,\theta} : \mathbb{R}^2 \rightarrow \mathbb{R}$ for the neural-SDE model of S :

$$\phi^{S,\theta} = \mathcal{F}_2 \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{128} \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{128} \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{128}. \quad (7.6)$$

In Figure 7.5.1, we show the evolution of training losses and validation losses over epochs during the training of the models for both EURO STOXX 50 index and DAX index. The loss value quickly drops during the first 10 epochs, and slowly declines and converges within 50 epochs.




Figure 7.5.1: Evolution of training losses and validation losses for the models of S .

To examine how well the neural network models the diffusion of S , we use a GARCH(1,1) model to estimate empirical volatilities and compare with the in-sample volatilities computed from the trained neural networks. As evidenced in Figure 7.5.2, the neural network models produce similar volatility dynamics.




Figure 7.5.2: Comparison of the volatility estimates for the option underlying indices.

7.6 Analysis of historical residuals

Historical residuals of the trained neural-SDE models (7.1) at time t_l , for $l = 1, \dots, L - 1$, are computed by

$$\hat{Z}_{t_l}^\xi = \sigma^{-1}(\xi_{t_l}) (\xi_{t_{l+1}} - \xi_{t_l} - \mu(\xi_{t_l})\delta t), \quad \hat{Z}_{t_l}^S = \gamma^{-1}(\xi_{t_l}) \left(\ln \frac{S_{t_{l+1}}}{S_{t_l}} - \hat{r}\delta t \right).$$

Our approximate likelihood corresponds to an assumption that (Z_t^S, Z_t^ξ) are independent standard normal white noise. Therefore, it is worth checking if the estimated residuals $(\hat{Z}_t^S, \hat{Z}_t^\xi)$ appear to be independent standard normals, within our training data. In Figure 7.6.1, we show the normal QQ-plots of the empirical marginal distributions for the residuals. For the primary factors, their residuals show mild fat tails. For the underlying stock indices, their residuals appear slightly left-skewed, implying higher-than-expected probabilities for large loss.




Figure 7.6.1: QQ-plots of historical in-sample model residuals.

To investigate dependence between the historical residuals for each factor, we show in Figure 7.6.2 the scatter plots of each pair of residuals. The residuals among risk factors ξ_i look fairly uncorrelated, indicating that the trained neural nets have well captured

factor covariances over time. However, there is evident negative correlation between the residuals of $\ln S$ and ξ_1 , coinciding with the commonly seen *leverage effect* [21]. In fact, we could capture the leverage effect using the neural-SDE model where a full covariance matrix between $(\ln S, \xi)$ is specified, which consequently produces uncorrelated residuals — see Appendix 7.C.2.




Figure 7.6.2: Scatter plots of historical in-sample model residuals.

7.7 Out-of-sample simulation of factors and implied volatility surface

We assess the learnt model’s ability to simulate time series data that are like the input real data. We take the trained model and simulate sample paths using a tamed⁶ Euler scheme (see Hutzenthaler, Jentzen and Kloeden [93]), which is given by

$$\xi_{t+\delta t} = \xi_t + \frac{\mu(\xi_t)}{1 + |\mu(\xi_t)|\sqrt{\delta t}}\delta t + \frac{\sigma(\xi_t)}{1 + \|\sigma(\xi_t)\|\sqrt{\delta t}}Z_t^\xi, \quad \ln \frac{S_{t+\delta t}}{S_t} = \hat{r}\delta t + \frac{\gamma(\xi_t)}{1 + |\gamma(\xi_t)|\sqrt{\delta t}}Z_t^S, \quad (7.7)$$

where the values of μ , σ and γ are approximated by the trained neural networks. Rather than generating the innovations (Z_t^S, Z_t^ξ) from normal distributions, we randomly sample them from historical residuals of the trained models. There are two major advantages of drawing innovations from historical residuals rather than normal distributions. First, historical residuals have fatter tails than normal (see Figure 7.6.1). Second, the joint historical residuals implicitly preserve the (potentially higher-order) dependence of (S, ξ) that fails to be captured by the covariance terms (see Figure 7.6.2).

We show the time series of ξ in Figure 7.7.1 (right). In the scatter plot on the left of Figure 7.7.1, we see that the dependence structure between ξ_1 and ξ_2 is well captured.

⁶We include the taming method simply to ensure stability of simulations if our neural networks were to produce unusually large values for drifts or volatilities.

In addition, the simulated factors remain within the no-arbitrage region, due to the hard constraints imposed on the drift and diffusion functions. We also simulate secondary factors from their fitted OU processes, and check the similarity in distribution, as shown in Figure 7.7.2.

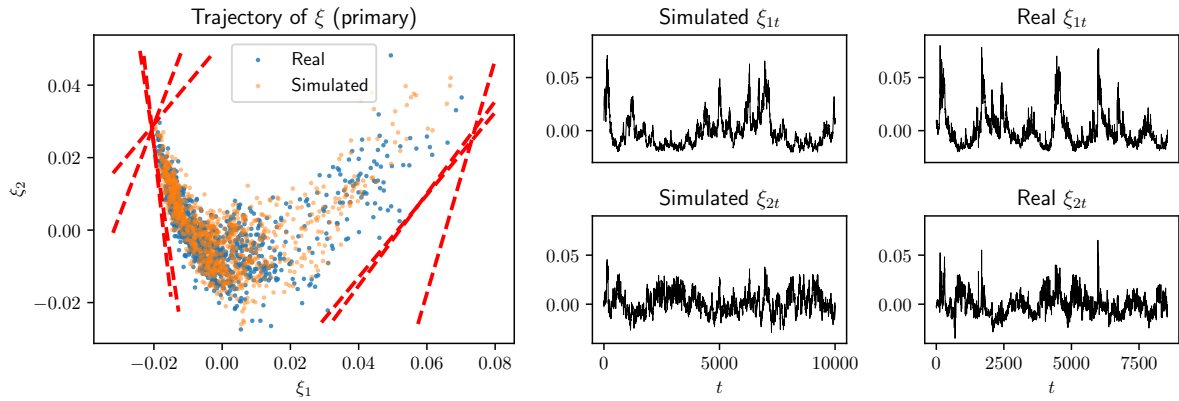


Figure 7.7.1: Simulation of primary factors from the learnt neural-SDE model, compared with the real data.

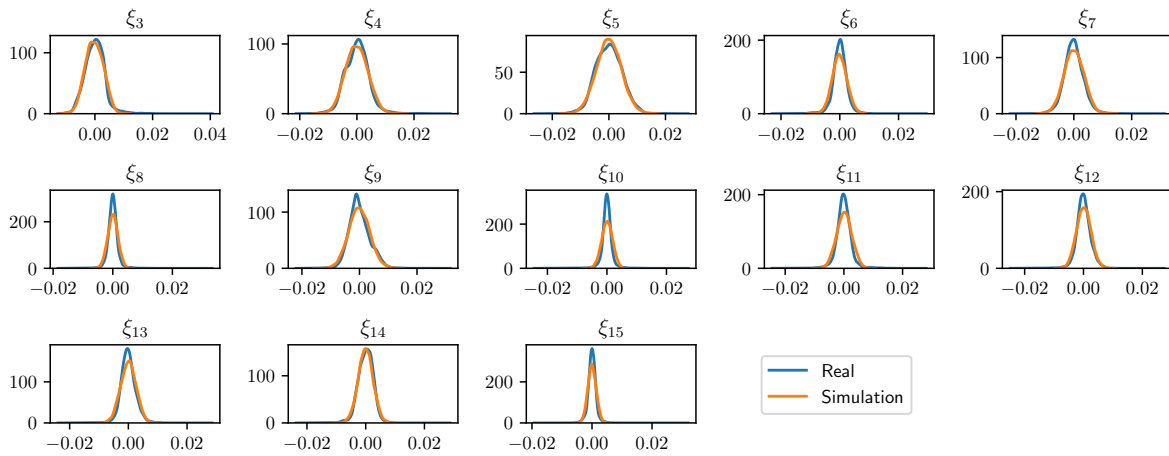


Figure 7.7.2: Distribution of OU-simulated and real secondary factors.

The learnt model is also capable of simulating a variety of realistic patterns of implied volatility surface (IVS). As demonstrated in Figure 7.7.3, we pick a few IVS patterns observed in historical data and find the closest (in the sense of minimising the ℓ^2 distance for ξ) ones in the simulated sample path.

7.8 VIX-like volatility index simulation

We follow the CBOE VIX calculation methodology [41] to compute a *volatility index* from the real and simulated option prices. Specifically, the square of VIX is a linear




Figure 7.7.3: Examples of real and simulated implied volatility surfaces. $\tau_1 = 30$ days and $\tau_{10} = 730$ days.

combination of OTM call and put option prices, which can be further written as a linear combination of call prices only, provided that put-call parity holds under no-arbitrage.

In Figure 7.8.1, we compare the real historical distribution of $(\text{VIX}, \log\text{-return of } S)$ and one simulation using the trained market model. For the simulated VIX, both its marginal distribution and the joint distribution with the log-return of S look reasonably similar to those of the real data. This demonstrates that our model is capturing the dependence structure between the volatility index and the underlying S . Looking at the simulated time series of VIX and log-return of S , we see several occurrences of volatility clustering in the return series, which always coincide with high VIX values.




Figure 7.8.1: Simulation of the VIX-like volatility index and log-return of S (EURO STOXX 50 index), compared with the real data.

Appendix

7.A Data pre-processing

Typical data vendors⁷ provide *historical* EOD (end-of-day) option prices with parameters in the format of a fixed set of representative time-to-expiries and Black–Scholes deltas, instead of traded expiries and strikes. This format is popular in industry, as it saves users great efforts in building historical implied volatility surfaces, which otherwise requires the users themselves to reconcile traded option contracts that are changing over time.

This data format is also appealing to our modelling approach, though the raw data need to be pre-processed to (1) get rid of static arbitrage, and (2) interpolate prices to a fixed set of moneyness parameters. We now explain how we pre-process the OptionMetrics dataset as described in Section 7.1.

We use the arbitrage detection method established in Chapter 3 to construct static arbitrage constraints in terms of option prices, and check the violation of these constraints by the data over time. In Figure 7.A.1, we show the histogram for the number of constructed arbitrage constraints and what percentage of them are violated over time. On average, there are about 3668 constraints per day. Historical option prices are almost free of static arbitrage after October 2008, before which there were about 8 ~ 9% of constraint violations, in the worst cases, in 2002 and 2006.




Figure 7.A.1: Statistics on the static arbitrage constraints violation for EURO STOXX 50 index option price data.

We use historical *median* moneynesses to define our liquid option lattice \mathcal{L}_{liq} (as shown in Figure 7.1.1a), and interpolate prices to the fixed set of moneyness parameters. In Figure 7.A.2, we show boxplots of moneynesses that are calculated over time for the fixed set of 13 deltas and 10 time-to-expiries. The closer to the money a delta is, the less variation we observe in the corresponding moneyness. This implies that less interpolation error is introduced to prices for close-to-the-money options, assuming a C^2 convex price

⁷For example, OptionMetrics (<https://wrds-www.wharton.upenn.edu/pages/about/data-vendors/optionmetrics/>), IVolatility (<https://www.ivolatility.com/home.j>), Nasdaq Data Link (<https://data.nasdaq.com/data/OWF-optionworks-futures-options/documentation>), Bloomberg, etc.

curve over moneyness. We use cubic splines to interpolate the call option price curve over moneyness for each expiry.

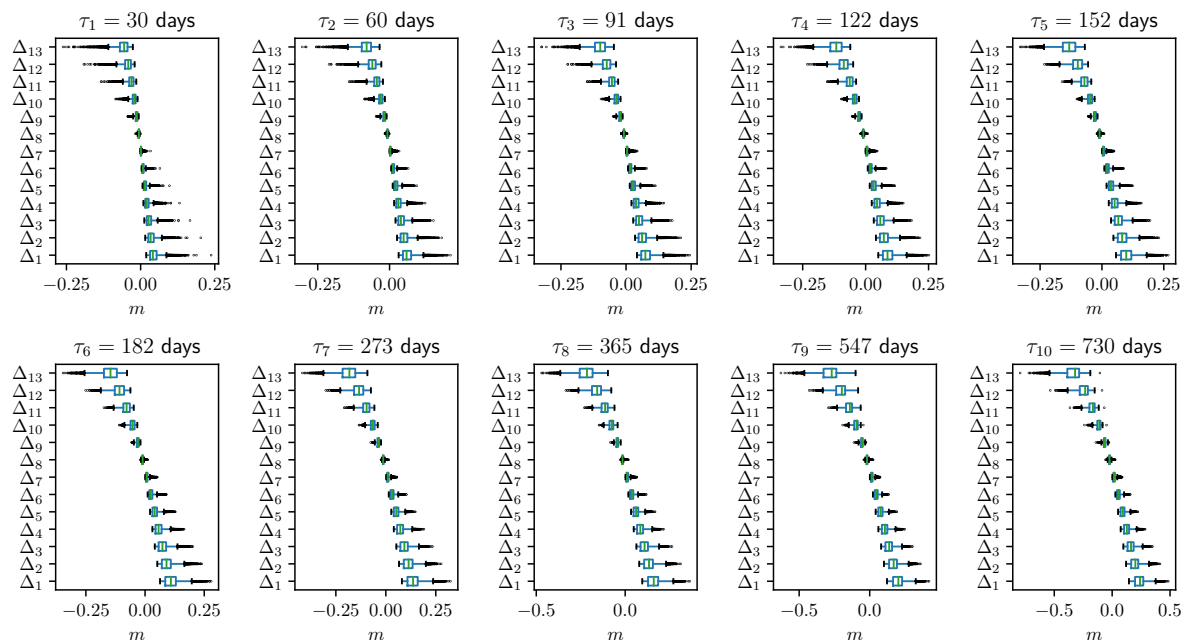


Figure 7.A.2: Boxplots of the moneynesses m corresponding to the 13 deltas for the 10 time-to-expiries. $\Delta_1 = 0.2$, $\Delta_{13} = 0.8$, and consecutive deltas have a 0.05 increment.

After interpolating option prices to the liquid lattice \mathcal{L}_{liq} for each historical date, we check the amount of static arbitrage in the interpolated prices and present the fraction of arbitrage constraint violation in Figure 7.A.3. Compared with Figure 7.A.1, the interpolation has introduced extra arbitrage for a few days, but in the worst case the introduced violation percentage is not greater than 2%. Thereafter, we perturb arbitrageable prices using the ℓ^1 -repair method described in Chapter 3 to get arbitrage-free prices.




Figure 7.A.3: Statistics on the static arbitrage constraints violation for EURO STOXX 50 index option prices interpolated to the liquid lattice \mathcal{L}_{liq} .

7.B Computing z_t for estimating dynamic arbitrage

Computing the z_t term in (4.8) requires us to estimate γ_t , the volatility of $\ln S_t$, and the partial derivatives of call prices $\partial\tilde{c}/\partial\tau$, $\partial\tilde{c}/\partial m$ and $\partial^2\tilde{c}/\partial m^2$.

To estimate γ_t , we fit an *initial* model for S with state variable $\tilde{\xi} = (S, \xi^{(0)})$, where $\xi^{(0)}$ is the factor projected from the first $d^{(0)}$ principal components of the input call prices. The initial model is also constructed and estimated by a neural-SDE model (7.1a), where the neural network architecture is the same as (7.6). We use a large $d^{(0)}$, such that the reconstruction error is small. This initial calibration gives us a candidate volatility $\hat{\gamma}_t^{(0)}$.

Next, we need to evaluate the partial derivatives $\partial\tilde{c}/\partial\tau$, $\partial\tilde{c}/\partial m$ and $\partial^2\tilde{c}/\partial m^2$ on the option lattice $\mathcal{L}_{\text{liq}} = \{(\tau_j, m_j)\}_{j=1, \dots, N}$. For each (τ, m) , we assume $\tilde{c}(\tau, m)$ behaves locally like the Black–Scholes price, such that we can approximate its partial derivatives using those of the Black–Scholes formula, and get

$$\frac{\partial\tilde{c}_t}{\partial m} = -e^m\Phi(d_2), \quad \frac{\partial^2\tilde{c}_t}{\partial m^2} = \frac{e^m}{\sigma_t^{\text{imp}}\sqrt{\tau}}\phi(d_2) + \frac{\partial\tilde{c}_t}{\partial m}, \quad \frac{\partial\tilde{c}_t}{\partial\tau} = \phi(d_1)\frac{\partial d_1}{\partial\tau} - e^m\phi(d_2)\frac{\partial d_2}{\partial\tau},$$

where

$$d_{1,2} = -\frac{m}{\sigma_t^{\text{imp}}\sqrt{\tau}} \pm \frac{1}{2}\sigma_t^{\text{imp}}\sqrt{\tau}, \quad \frac{\partial d_{1,2}}{\partial\tau} = \frac{1}{\tau} \left(\frac{m}{2\sigma_t^{\text{imp}}\sqrt{\tau}} \pm \frac{1}{4}\sigma_t^{\text{imp}}\sqrt{\tau} \right).$$

In Figure 7.B.1, we show one example of the estimated partial derivatives on \mathcal{L}_{liq} .




Figure 7.B.1: Estimated partial derivatives of the normalised call price surface (EURO STOXX 50 index option as of 2002-01-02).

7.C Alternative model configurations

Apart from the market model discussed in the main text, we present two alternative models that differ in the number of primary factors chosen and the dependence between S and ξ . We also show their corresponding VaR backtesting results in Appendix 8.C, which perform slightly worse than the model presented in the main text.

7.C.1 A three-primary-factor market model

In Table 7.C.1, we display the types of the three primary factors $\xi = [\xi_1 \ \xi_2 \ \xi_3]^\top \in \mathbb{R}^3$ we use for representing call option prices, as well as the resultant MAPE, PDA and PSAS. Compared with the two primary factors (last row in Table 7.3.1), the additional dynamic arbitrage factor reduces dynamic arbitrage by 6.1% and MAPE in price reconstruction error by 0.56%, but introduces 0.55% more static arbitrage. In Figure 7.C.1, we plot the corresponding price basis functions of these three factors.

Factors	MAPE	PDA	PSAS
Dynamic arb. + Statistical acc. + Static arb.	4.26%	1.73%	0.85%

Table 7.C.1: MAPE, PDA and PSAS metrics when including three primary factors.




Figure 7.C.1: Price basis functions of the normalised call price surface.

In Figure 7.C.2, the black dots give trajectories of (ξ_i, ξ_j) for $i, j \in \{1, 2, 3\}$. For each (i, j) , we also indicate boundaries of the polygon, that is, the 2D affine projection of the 3D arbitrage-free state space for ξ . Specifically, let \mathcal{P} be the \mathbb{R}^3 -polytope state space for the factors ξ formed by the static arbitrage constraints, then the affine projection of this polytope onto the (ξ_i, ξ_j) -space is given by

$$\pi_{i,j}(\mathcal{P}) = \{y \in \mathbb{R}^2 : \exists \xi = (\xi_k) \in \mathcal{P}, y = [\xi_i \ \xi_j]^\top\}.$$

We plot the boundaries $\partial\pi_{i,j}(\mathcal{P})$ as red dashed lines in Figure 7.C.2. These boundaries (which are, in some sense, as wide as possible) suggest how tight the arbitrage bounds are.

We take the three factors and model their dynamics using neural-SDEs. After training the models, we forward-simulate 10,000 timesteps, and compare the pairwise joint distribution of the simulated factors and that of the real data, as shown in Figure 7.C.3. The trained neural-SDE model seems to be capable of simulating long trajectory of the three factors that are similar to the real data, at least in terms of pairwise joint distribution.




Figure 7.C.2: Trajectory (black dots) of the primary \mathbb{R}^3 factors and the corresponding static arbitrage constraints (red and orange dashed lines), projected onto \mathbb{R}^2 spaces.




Figure 7.C.3: Joint distribution of pairs of the simulated primary factors, compared with that of the real data.

7.C.2 Modelling and estimating S and ξ jointly

We build and estimate an SDE model for $\tilde{\xi} = (\ln S, \xi)$ with a fully-specified diffusion matrix; in fact, the model (7.1) is a special case with a block-diagonal diffusion matrix. Given $\xi \in \mathbb{R}^d$, compared with the model (7.1), the joint model has d more \mathbb{R} -valued functions to estimate in the diffusion term, which captures dynamic correlation between $\ln S$ and each of the d factors.

For estimating the joint model, we still use the concatenated data (of EURO STOXX 50 and DAX options) for training the neural nets. We assume that the two stock index processes have the same diffusion function with respect to ξ but different constant drifts estimated by (7.5). This enables us to use the concatenated stock index data to estimate a universal volatility function for the two indices. In Figure 7.C.4, we show the loss history of training the neural nets. Both the training and validation losses decline rapidly over the first 1000 epochs and then converge gradually.

To see that the joint model has reasonably captured covariances between the underlying and the factors, we show in Figure 7.C.5 the scatter plots of all pairs of model




Figure 7.C.4: Evolution of training losses and validation losses for the joint model of $\ln S$ and ξ .

residuals. Compared with Figure 7.6.2, the joint model takes into account the leverage effect and leaves no correlation between residuals of $\ln S$ and ξ_1 .




Figure 7.C.5: Scatter plots of historical in-sample model residuals.

Application: estimating risks of option books

“凡事豫则立，不豫则废。言前定则不跲，事前定则不困，行前定则不疚，道前定则不穷。”

— 孔伋《礼记·中庸》

“In all things success depends on previous preparation, and without such previous preparation there is sure to be failure. If what is to be spoken be previously determined, there will be no stumbling. If affairs be previously determined, there will be no difficulty with them. If one’s actions have been previously determined, there will be no sorrow in connection with them. If principles of conduct have been previously determined, the practice of them will be inexhaustible.”

— Ji Kong, *The Doctrine of the Mean, The Book of Rites* (translated by James Legge)

8.1 Introduction

Managing option trade books usually requires modelling and simulating dynamics of liquid vanilla options. A desirable model, and therefore trajectories simulated from the model, should preserve statistical properties of the real data, while respecting underlying financial constraints. Previous work [46] has considered building arbitrage-free neural-SDE market models under the real-world measure, and established corresponding model estimation algorithms that take as input time series price data of finitely many liquid options. The model consists of an SDE system for states representing the underlying and a small number of factors chosen for a statistically accurate dynamic representation of the option surface and the minimisation of dynamic and static arbitrage. The coefficients of the SDE are calibrated under hard linear inequality constraints on the joint options price processes dictated by static arbitrage relationships. In terms of practical applications, [46]

only considers synthetic data simulated from a Heston stochastic local volatility model, and does not demonstrate performance on any concrete tasks. This paper validates the modelling approach in [46] using real world data, and explores the capacity of the model as a realistic option *market simulator*. In particular, we shall see that these methods yield computationally efficient and accurate models for *evaluating Value-at-Risk*, with less procyclicality and better performance than traditional filtered historical simulation [13].

Essentially, we investigate the feasibility of using a neural-SDE market model as a *generative model* for time series of option prices. Simulating synthetic data from a generative model is an important measure to ensure data privacy [20], and can be useful for boosting deep learning performance, for instance to provide an environment to train further models with reinforcement learning [29, 128]. Synthetic data generators are probabilistic models which simulate observations similar to historical data, in the sense of capturing specific features of the time series. Recent developments of generative models for financial time series focus on generative adversarial networks (GANs) [83, 122, 106, 141, 145] and variational autoencoders (VAEs) [30, 146], and typically are not driven by fundamental financial modelling principles. We highlight the work of [146] and [42], which take similar approaches to ours, using alternative representations of arbitrage conditions.

Building generative models for multiple option price time series, with the same underlying asset and different strikes and maturities, is by nature a high-dimensional problem, where complex arbitrage constraints need to be respected among different dimensions. Wiese et al. [145, 146] first transform option prices to an equivalent representation in the form of discrete local volatilities [147] with simple static arbitrage constraints; they then apply PCA to reduce the dimension and formulate a typical GAN problem. Compared with GAN approaches, our neural-SDE market models largely reduce the “black-box” nature of neural network models (see the discussion by [48]) while still retaining their computational advantages. Specifically, our neural-SDE market model gives drift and diffusion estimates of each risk factor modelled, which helps us to understand the expected return and volatility of investing in each risk factor.

Here, we verify that these neural-SDE market models produce price paths similar to historical data for European equity index markets, in that they display volatility clustering (in the returns on the underlying index, and this is replicated in an options portfolio similar to VIX) and they match marginal and joint densities for the underlying and option portfolios. We then focus on risk simulation for option portfolios, which is especially convenient because the models are built under the real-world measure. In particular, we assess how well the models evaluate risks for a comprehensive set of option portfolios which represent typical risk profiles, such as (calendar) spreads, risk reversals and strangles.

In fact, these portfolios can be viewed as statistical features of option price surfaces. We therefore verify that the neural-SDE market models are capable of simulating data similar to historical data, in the sense of capturing these highly structured while financially meaningful features.

Evaluating market risk for portfolios of nonlinear derivatives such as options is challenging from both an accuracy and speed perspective. We focus on Value-at-Risk (VaR), a risk measure promoted by J.P. Morgan RiskMetrics [112] in the mid-1990s and quickly mandated by international regulators such as the US SEC and Basel committee for disclosing portfolio risk. Mathematically, the key is to establish a loss distribution for the evaluated portfolio. As our models generate samples from the full loss distribution of the portfolio, our approach could also be used to compute estimates of other risk measures, for example expected shortfall. We focus on VaR due to its common use for regulation in financial markets, and given the existence of standard backtesting approaches, which allow us to verify the performance of our method against standard criteria.

Parametric approaches to VaR estimation, such as the delta method [112] and delta-gamma method [66, 27], give closed-form approximations of the loss distribution, which are fast, but rely heavily on distributional assumptions on the risk factors. In addition, El-Jahel et al. [63] have shown that these methods give biased VaR estimates due to poor local, linear and quadratic, approximations of the loss.

Conversely, non-parametric methods based on Monte Carlo estimation of the loss distribution are unbiased and consistent, and have flexible distributional assumptions, but can present a substantial computational burden. First, a large number of samples of the risk factors (known as “*risk scenarios*”) are needed as VaR only considers the (upper) tail of the loss distribution. Second, for each risk scenario, one needs to re-evaluate the portfolio value, which is generally not available in closed-form, so computationally costly numerical procedures are needed. If Monte Carlo simulation is used for pricing the portfolio, one ends up with a two-level (or nested) simulation scheme. Variance reduction techniques can be used in both outer and inner levels to accelerate simulations; specifically, for the outer level, where quantile estimation is made, there is a vast literature discussing the use of control variates [89, 84], antithetic variates [7], and, more popularly, importance sampling or stratified sampling [75, 76, 139]. Further non-parametric VaR estimation methods include historical simulation [31] or volatility-filtered historical simulation [12, 11, 13, 78], which sample risk scenarios based on empirical distributions of risk factors. Nevertheless, regardless of how risk scenarios are simulated, the valuation of options portfolios often requires re-evaluation of option prices based on risk factors (for multiple simulation trajectories and time points), which is computationally expensive.

In our market models, option prices are linear in the risk factors, resulting in negligible computational cost for re-evaluating an option portfolio for each simulated risk scenario. Despite the simplicity of the linear factor representation, it is sufficiently rich to achieve satisfactory statistical accuracy. As shall be detailed later, we simulate risk factors from a flexible, offline trained neural-SDE with historical innovations. Through out-of-sample backtesting, we show that our model yields statistically and economically more appealing VaR estimates for various option portfolios than traditional filtered historical simulation approaches (see, for example, Barone-Adesi, Giannopoulos and Vosper [13]).

This chapter is organised as follows. First, we explain how to simulate risk scenarios under the market model and with the filtered historical simulation method in Section 8.2. Given the real world dataset and the trained model as described in Chapter 7, we estimate VaR over various risk horizons and confidence levels for a varied set of option portfolios, and backtest the out-of-sample VaR performance in Section 8.3. Finally, Section 8.4 concludes our observations and addresses a few further extensions.

8.2 Measuring risk

As our market models are defined and trained under the real-world measure, a natural application is to simulate risk scenarios for option portfolios. We explain how this is done, and compare its performance with the filtered historical simulation method [78], which is a standard approach used widely for managing risk of futures and options by mainstream CCPs (see, for example, CME [44], Eurex [65] and LCH [109]).

Consider a portfolio Π of N call options with market value $\Pi_t = \sum_{j=1}^N w_j C_t(T_j, K_j)$ at time t , where $w_j \neq 0$ denotes the weight of the j -th call option in the portfolio. The profit-and-loss (PnL) of the portfolio at time $t + \Delta t$, for given time horizon Δt , is then

$$\text{PnL}_{\Delta t}(\Pi_t) = \Pi_{t+\Delta t} - \Pi_t = \sum_{j=1}^N w_j (C_{t+\Delta t}(T_j, K_j) - C_t(T_j, K_j)) = \sum_{j=1}^N w_j \text{PnL}_{\Delta t}(C_t(T_j, K_j)).$$

For a chosen risk measure $\varrho(\cdot)$, time t and horizon Δt , we evaluate the risk of the portfolio, $\varrho(\text{PnL}_{\Delta t}(\Pi_t))$, in broadly three steps:

- (i) Identify *risk factors* Θ and a *value function* V such that $\Pi_t \approx V(t, \Theta_t)$.
- (ii) Simulate $\Theta_{t+\Delta t}^{(m)} | \Theta_t$ based on some dynamic model for Θ_t for $m = 1, \dots, M$, and compute $\Pi_{t+\Delta t}^{(m)} \approx V(t + \Delta t, \Theta_{t+\Delta t}^{(m)})$. Each $\Pi_{t+\Delta t}^{(m)}$ is called a *risk scenario*.
- (iii) For the m -th risk scenario, calculate the *scenario PnL* as $V(t + \Delta t, \Theta_{t+\Delta t}^{(m)}) - \Pi_t$, which is assumed to be a realisation of $\text{PnL}_{\Delta t}(\Pi_t) | \Theta_t$. Therefore, the M scenario

PnLs give an estimate of the empirical distribution¹ of $\text{PnL}_{\Delta t}(\Pi_t)|\Theta_t$, which the risk measure ρ will be evaluated on.

8.2.1 Risk simulation using the market model

Our market models characterise option risk in terms of $\Theta = (S, \xi)$, for which the joint dynamics are given as the SDE system (7.1). After an offline training of the SDE model, we discretise it and simulate the factors, which are subsequently translated to arbitrage-free option prices through simple linear transformations.

More specifically, suppose we want to simulate prices for the (T, K) -call option at time $t + \Delta t$. Then the simulation consists of the following steps:

- (i) **Decode factors at time t .** We collect prices of call options on the liquid lattice \mathcal{L}_{liq} and normalise the prices using the time- t zero curve and futures curve. Given the price basis \mathbf{G}^\top , we solve for $\xi_t = (\mathbf{G}\mathbf{G}^\top)^{-1}\mathbf{G}(\mathbf{c}_t - \mathbf{G}_0^\top)$.
- (ii) **Simulate factors at time $t + \Delta t$.** Given the SDE model (7.1) for which the coefficients are estimated by neural networks, we simulate $S_{t+\Delta t}^{(m)}$ and $\xi_{t+\Delta t}^{(m)}$ using the tamed Euler scheme (7.7) iteratively over $\Delta t/\delta t$ time steps, for $m = 1, \dots, M$.
- (iii) **Compute option prices at time $t + \Delta t$.** Under the m -th simulated scenario, we compute the normalised prices for the call options on \mathcal{L}_{liq} as $\mathbf{c}_{t+\Delta t}^{(m)} = \mathbf{G}_0^\top + \mathbf{G}^\top \xi_{t+\Delta t}^{(m)}$. Let $h(\tau, m; \mathbf{c})$ be some $C^{1,2}$ function that interpolates the data $(\mathcal{L}_{\text{liq}}, \mathbf{c})$, then

$$C_{t+\Delta t}^{(m)}(T, K) = S_{t+\Delta t}^{(m)} \times h\left(\tau_t - \Delta t, m_t + \ln \frac{S_t}{S_{t+\Delta t}^{(m)}}; \mathbf{c}_{t+\Delta t}^{(m)}\right),$$

where $\tau_t = T - t$ and $m_t = \ln(K/F_t(T))$.

Remark 8.1. Our model is written not in terms of prices C but in terms of the normalised prices \tilde{c} . This transformation (4.3) involves discounting both for interest rates and dividends. While interest rates and dividends could be included as risk factors, they have only marginal impact on equity index options, especially when the longest expiry we consider is two years. We will therefore restrict our attention to the underlying risk and volatility risk. Therefore, we will evaluate risk of the revised option value $C_t = S_t \tilde{c}_t$.

¹In general, the conditional distribution of PnL is not available in closed-form, otherwise one may be able to derive an analytical formula for the risk measure without simulation.

8.2.2 Filtered historical simulation with Heston models

The historical simulation approach [12] simply uses historical changes of a collection of risk factors $\Theta = (x_1, \dots, x_d)$ to construct risk scenarios. In addition, historical changes of risk factors may be *filtered* (or *scaled*) based on current volatility, so that the simulated risk scenarios are more responsive to recent market conditions. Specifically, suppose $C_t(T, K) = V(T - t, K; \Theta_t)$ and we want to simulate prices at $t + \Delta t$. Then there are two steps:

- (i) Collect historical data of Δt -returns r_j in the j -th risk factor x_j , for $j = 1, \dots, d$, over the last M timesteps, so we have $r_j(t_m)$ for $m = 1, \dots, M$. Let $\sigma_j(t)$ be some realised volatility estimate for the j -th risk factor at t . Then we define

$$r_j^{(m)}(t + \Delta t) = r_j(t_m) \frac{\sigma_j(t)}{\sigma_j(t_m)} \quad (8.1)$$

as the risk factor return under the m -th scenario, which leads to the corresponding risk factor scenario $x_j^{(m)}(t + \Delta t)$ by applying the return² to $x_j(t)$.

- (ii) Revalue the call option $C_{t+\Delta t}^{(m)}(T, K) = V(T - t - \Delta t, K; \Theta_{t+\Delta t}^{(m)})$.

We use the parameters of calibrated Heston models [85] as risk factors. Specifically, we calibrate the Heston parameters $\Theta = (S_0, \nu_0, \theta, k, \sigma, \rho)$ as detailed in Appendix 8.A to historical vega-weighted option prices, yielding historical observations of Θ_t , which are then used to simulate historical scenarios for option portfolios.

8.3 Value-at-Risk backtesting analysis

In order to assess the scenarios produced by risk simulation engines for a comprehensive set of risk profiles, we backtest various option portfolios ranging from outright options to heavily hedged portfolios. The list of tested portfolios is given in Table 8.3.1, with detailed description in Appendix 8.B. We consider both long and short positions in each portfolio considered.

We have split the historical data into training and testing samples according to Table 7.1.1. We use only the training samples for estimating the neural-SDE market model, which will then be used to calculate risks. As a measure of risk of a portfolio Π , we consider Value-at-Risk (VaR) at a given confidence level α for a given time horizon Δt , defined as

$$\text{VaR}_{\alpha, \Delta t}(\Pi_t) = -\sup \{l \in \mathbb{R} : \mathbb{P}(\text{PnL}_{\Delta t}(\Pi_t) < l) \leq 1 - \alpha\}.$$

²We use log returns if a risk factor is positive by definition, in which case $x_j^{(m)}(t + \Delta t) = x_j(t) \exp(r_j^{(m)}(t + \Delta t))$, or absolute return otherwise, in which case $x_j^{(m)}(t + \Delta t) = x_j(t) + r_j^{(m)}(t + \Delta t)$.

Portfolio	Delta-exposed			Delta-hedged				VIX
	Outright	Delta spread	Risk reversal	Delta butterfly	Delta-hedged option	Delta-neutral strangle	Calendar spread	
Number	140	420	60	20	60	60	90	2

Table 8.3.1: Number of tested portfolios of various types.

Given M risk simulation scenarios, the VaR is estimated by

$$\widehat{\text{VaR}}_{\alpha, \Delta t}(\Pi_t) = -\sup \left\{ l \in \mathbb{R} : \frac{1}{M} \sum_{m=1}^M \mathbb{1}_{\{\text{PnL}_{\Delta t}^{(m)}(\Pi_t) < l\}} \leq 1 - \alpha \right\}.$$

For each testing sample t , we compute an ex-ante VaR forecast $\widehat{\text{VaR}}_{\alpha, \Delta t}(\Pi_t)$, and observe an ex-post realised $\text{PnL}_{\Delta t}(\Pi_t)$. Therefore, we can define the VaR breach

$$\beta(\Pi_t) = \mathbb{1}_{\{\text{PnL}_{\Delta t}(\Pi_t) < -\widehat{\text{VaR}}_{\alpha, \Delta t}(\Pi_t)\}}.$$

The coverage ratio over the $L - E$ testing periods is then computed as

$$1 - \frac{1}{L - E} \sum_{l=E+1}^L \beta(\Pi_{t_l}).$$

Prevailing VaR model backtesting methodologies mostly focus on two categories:

- (i) **Coverage tests** assess whether the empirical frequency of VaR breaches is consistent with the quantile of loss a VaR measure is intended to reflect. We use the widely cited Kupiec's proportion of failures (PF) coverage test [108], which assumes the null $H_0 : \mathbb{E}[\beta(\Pi_t) | \Theta_t] = 1 - \alpha$ and models the VaR breach events in each testing sample as independent Bernoulli trials. In addition, we also consider the Basel Committee traffic light approach [16]: based on the number of breaches experienced for the 1-day $\text{VaR}_{0.99}$ results during a 250-day testing window, the VaR measure is categorised as falling into the green (≤ 4), yellow (≥ 5 and ≤ 9) or red (≥ 10) zone.
- (ii) **Independence tests** assess whether VaR breach events appear to be independent of each other. We use Christoffersen's test [43] that considers unusually frequent consecutive breaches. Combining the test statistics of the Kupiec's PF test and the Christoffersen's independence test yields a conditional coverage (CC) mixed test [43].

We compute VaR for horizons $\Delta t = 1, 2, 5$ and 10 days, where 2-day and 5-day are typical *margin periods of risk* that European regulators require for risk-monitoring of exchange-traded and over-the-counter derivatives, respectively.

We report the backtesting coverage and independence statistics for 1-day VaRs of EURO STOXX 50 index options at 0.99 and 0.95 confidence levels in Table 8.3.2. We compute a coverage ratio for each of the tested option portfolios, and calculate its sample median and mean over the portfolios. The neural-SDE VaR model gives a higher mean coverage ratio over option portfolios, while much fewer portfolios fail the coverage and independence tests. At the 0.95 confidence level, more portfolios fail the two-sided Kupiec PF coverage test than the one-sided version, indicating that both VaR models tend to overestimate risks, but the neural-SDE approach is less prone to underestimating risks (as indicated by the one-sided test failure rate of 9.39% v.s. 18.08%)³.

	1-day VaR _{0.99}		1-day VaR _{0.95}					
	nSDE	FHS	nSDE	FHS				
Coverage ratio median	0.9921	0.9881	0.9484	0.9563				
Coverage ratio mean	0.9887	0.9742	0.9528	0.9321				
Kupiec PF (two-sided)	6.92%	25.23%	32.51%	33.69%				
Kupiec PF (one-sided)	6.92%	25.23%	9.39%	18.08%				
Christoffersen independence	0.70%	11.03%	6.10%	14.79%				
Conditional coverage	3.53%	24.88%	29.23%	34.39 %				
Basel committee traffic light	69.1%	29.7%	0.5%	62.4%	25.9%	10.8%	N.A.	N.A.

Table 8.3.2: Summary statistics on coverage and independence tests for 1-day VaRs computed by the neural-SDE (nSDE) market model and the filtered historical simulation (FHS) approach. For the four statistical tests, i.e. Kupiec PF test (two-sided), Kupiec PF test (one-sided), Christoffersen’s independence test and conditional coverage mixed test, the table lists the percentage of the 852 tested option portfolios that fail each test at the 0.05 significance level.

To diagnose VaR coverage performances over different option portfolios, we group portfolios with different primary Greek-exposures, i.e. delta or vega, and examine whether their VaR breaches are synchronised during the testing window. In Figure 8.3.1, we show the time series of 1-day nSDE-VaR_{0.99} breaches for primarily delta-exposed option portfolios, and compare with positions in the underlying index. We see a few time points when VaR breaches happen simultaneously for multiple option portfolios, which correspond to synchronised VaR breaches or close-to-breaches in holding S .

For delta-neutral and vega-exposed portfolios, we compare their VaR breaches with those of the VIX portfolio in Figure 8.3.2. There are few breaches observed, suggesting the VaR model is over conservative for these option types. We see some synchronised VaR breaches for vega-short positions in delta butterflies and the VIX. We perform a similar analysis for FHS-VaRs; as shown in Figure 8.C.2 in Appendix 8.C, FHS-VaRs underestimate risks severely for most delta-hedged portfolios.

³It is worth noting that since we reject the nulls in all the coverage and independence tests at the 0.05 significance level, a perfect VaR model is expected to fail for 5% of portfolios.




Figure 8.3.1: Times series of 1-day nSDE-VaR_{0.99} breaches (red dots) for delta-exposed option portfolios, compared with that for positions in the underlying index. (EURO STOXX 50 index option)




Figure 8.3.2: Times series of 1-day nSDE-VaR_{0.99} breaches (red dots) for delta-neutral and vega-exposed option portfolios, compared with that for positions in the VIX portfolio.

The dynamics of the nSDE-VaRs shown in Figure 8.3.1 and 8.3.2 appear to be relatively stable during the testing period, i.e. the year 2019. Nevertheless, this does not imply that the nSDE-VaRs are insensitive to changes in market conditions, as 2019 was simply an uneventful year for global stock markets. To show that the nSDE-VaRs are adaptive to market changes, we give in-sample nSDE-VaR backtesting results in Appendix 8.D.

While coverage and independence tests assess whether a VaR calculation approach achieves desirable statistical properties, regulators would also like to avoid over-*procyclical* VaR models. Procyclicality measures quantify variations in VaR values over the testing periods. Common risk models are procyclical because they estimate higher risk for the same portfolio in times of market stress and lower risk in calm markets. While procyclicality might indicate that a model is responsive to contemporary market volatility, it may also be an artefact of a poorly calibrated model. Excessive procyclicality can cause liquidity stress, where parties posting margin have to find additional liquid assets, often at times when it is most difficult for them to do so. We use the trough-to-peak ratio⁴ proposed by the Bank of England [120] to measure procyclicality. After calculating the trough-to-peak ratio for each tested portfolio, we get a distribution of the ratio over our evaluated portfolios. In Figure 8.3.3, we compare the trough-to-peak distributions computed from the neural-SDE and the FHS approach. The neural-SDE approach tends to produce less procyclical 1-day VaRs at both 0.99 and 0.95 confidence levels. In particular, for 1-day $\text{VaR}_{0.99}$, there is a substantial reduction in the number of portfolios with very low trough-to-peak ratio (corresponding to a high level of procyclicality) and an increase in the number of ratios above 0.6. Consistent results for other risk horizons and confidence levels can be seen in Appendix 8.C.




Figure 8.3.3: Distribution of trough-to-peak ratios (1-day VaRs) over different tested option portfolios.

While more backtesting results for 2-day, 5-day and 10-day VaRs and those for DAX options are reported in Appendix 8.C, here we briefly summarise the backtesting results

⁴The original paper [120] uses peak-to-trough ratio, the ratio of the maximum risk of a constant portfolio to the minimum risk over a fixed observation period. Here we take its reciprocal because for some heavily-hedged portfolio the minimum risk can be very close to 0, making the peak-to-trough ratio explode.

for 10-day VaRs of EURO STOXX 50 index options in Table 8.3.3⁵. Similar to the observations of 1-day VaRs, the neural-SDE approach leads to less portfolios that fail the one-sided coverage test, the independence test and the conditional coverage test (the nulls are rejected at the 0.05 significance level for all tests). This indicates that the neural-SDE approach has consistently superior performance for long-horizon 0.99-confidence VaR calculation. For 0.95-confidence VaR, the neural-SDE approach is more conservative than the FHS approach, and the results are more mixed than in the 1-day case.

	10-day VaR _{0.99}		10-day VaR _{0.95}	
	nSDE	FHS	nSDE	FHS
Coverage ratio median	1.0000	0.9918	0.9712	0.9637
Coverage ratio mean	0.9963	0.9800	0.9584	0.9464
Kupiec PF (two-sided)	4.10%	17.92%	51.52%	37.12%
Kupiec PF (one-sided)	4.10%	17.92%	10.07%	14.05%

Table 8.3.3: Summary statistics on coverage tests for 10-day VaRs computed by the neural-SDE (nSDE) market model and the filtered historical simulation (FHS) approach.

Computation time comparison

We have carried out the backtesting analysis on a 2.4GHz 8-core Intel Core i9-9980HK CPU with 32GB RAM, and summarised the average elapsed time for simulating one risk scenario per backtested day in Table 8.3.4.

	Simulate & revaluation		Repair arbitrage	
	nSDE	FHS (Heston)	nSDE	FHS (Heston)
Expected time (ms)	5.0	147.3	14.3	N.A.
Std. time (ms)	0.3	58.7	7.1	N.A.

Table 8.3.4: Average and standard deviation of the elapsed time for simulation one risk scenario.

In the FHS approach, evaluating option portfolios with simulated risk factor, i.e. the Heston parameters, is most time-consuming. We perform this by calling the efficiently implemented Matlab function `optByHestonNI` [115], which gives the option price by Heston model using numerical integration. In contrast, the nSDE approach evaluates option portfolios through a simple linear combination of simulated risk factors. This is the reason why the time for simulation and revaluation is reduced from 147.3 ms to 5.0 ms. The nSDE approach has an additional step of detecting and repairing arbitrage in the simulated option prices, which however is computationally cheap due to the efficient linear programming problem formulation in [45].

⁵Given we consider breaches in overlapping 10-day periods, the Christoffersen independence test is not applicable.

8.4 Conclusions and further extensions

We estimate a neural-SDE market model using historical price data of EURO STOXX 50 and DAX index options, the top two most traded equity options listed on Eurex. We assess the in-sample goodness-of-fit quality by investigating the neural network training loss convergence, the post-fitting historical residuals, and the comparison between the volatility estimates for $\ln S$ with the benchmark GARCH(1,1) estimates. Thereafter, we verify that the trained model can forward simulate long trajectories of option price (or implied volatility) data that have similar marginal and joint distributions to historical data. By backtesting the option portfolio VaR generated from the trained neural-SDE market model in an out-of-sample setting, we find that the model beats the traditional filtered historical approaches with statistically more efficient coverage, less procyclicality and less computational efforts; and this is consistently true for VaRs computed with various risk horizons and confidence levels.

While seeing the success in applying the neural-SDE market model in simulation and risk management, there remain directions to extend the current framework for a wider range of applications. For example, a simple and straightforward extension is to examine other risk measures such as expected shortfall. Second, in order to model option types other than European vanillas, it is necessary to derive arbitrage constraints among those options, or to modify the approach taken accordingly.

Appendix

8.A Heston model calibration results and historical scenarios

Heston models are constructed from the SDE

$$dS_t = r_t S_t dt + \sqrt{\nu_t} S_t dW_t^S, \quad (8.2a)$$

$$d\nu_t = k(\theta - \nu_t) dt + \sigma\sqrt{\nu_t} dW_t^\nu, \quad d\langle W_t^S, W_t^\nu \rangle = \rho dt, \quad (8.2b)$$

We re-write the variance process (8.2b) as

$$d\nu_t = k(\theta - \nu_t) dt + \eta\sqrt{\nu_t} dW_t^S + \lambda\sqrt{\nu_t} dW_t^\perp, \quad d\langle W_t^S, W_t^\perp \rangle = 0,$$

where $\eta = -\sigma\rho$, $\lambda = \sigma\sqrt{1-\rho^2}$. We characterise option price surfaces by the parameters $\Theta = (S_0, \nu_0, \theta, k, \eta, \lambda)$. On each historical date t , we calibrate parameters for a Heston model by solving

$$\hat{\Theta}_t = \arg \min_{\Theta} \sum_{j=1}^N \frac{1}{\mathcal{V}_t(\tau_j, m_j)} (c^\Theta(\tau_j, m_j) - \tilde{c}_t(\tau_j, m_j))^2,$$

where \mathcal{V} is the option vega calculated through (7.3). In Figure 8.A.1, we show the time series of the calibrated parameters and calibration errors, measured by MAPE, for historical price data of EURO STOXX 50 index options. Since ν_0 , θ , k and λ are positive by construction, we include the positivity constraint for these parameters during the calibration. While η can be positive or non-positive, the calibrated result is always positive, implying negative correlations ρ between the index and volatility over time.




Figure 8.A.1: Heston model calibrated parameters and calibration errors (measured as mean absolute percentage error) for historical price data of EURO STOXX 50 index options.

Given that all risk factors are positive, we compute their 1-day log-returns, based on which we estimate the exponentially weighted moving average (EWMA) volatilities with time decay factor 0.95. We plot the EWMA 1-day volatility estimates over the estimation window in Figure 8.A.2. The RiskMetrics Technical document [112] suggests a value of 0.94 for the decay factor when estimating daily volatility. Nevertheless, in order to make the FHS-VaR and the nSDE-VaR more comparable, we optimise the choice of the decay factor by minimising the Wasserstein distance between the resulting EWMA vol and the in-sample neural-SDE vol in S (as shown in Figure 7.5.2). This yields the decay factor 0.95. Dividing the returns by the volatility estimates, we obtain historical standardised residuals, or *historical shocks*, as shown in Figure 8.A.3. The standardised residuals appear to be roughly stationary over time and do not exhibit volatility clustering. In order to generate risk scenarios for each t in the testing window, we multiply these historical shocks by the time- t 1-day EWMA volatility, as described in (8.1), and apply these filtered returns to the time- t risk factors, as if these historical shocks actually occurred.

8.B Tested option portfolios for risk simulation

To assess performance of a risk model, we backtest its risk evaluation of a variety of representative trading strategies. In particular, defining a set of deltas $\mathcal{D} = \{0.2, 0.3, 0.4, 0.5, 0.6,$




Figure 8.A.2: EWMA 1-day volatility estimates (with time decay factor 0.95) for the Heston risk factors for EURO STOXX 50 index options.




Figure 8.A.3: Standardised residuals of the Heston risk factors for EURO STOXX 50 index options.

0.7, 0.8} and a set of time-to-expiries $\mathcal{T} = \{30, 60, 91, 122, 152, 182, 273, 365, 547, 730\}$ (in calendar days), we consider:

- (i) *Outright call* option $C(\tau, \Delta)$ with all deltas $\Delta \in \mathcal{D}$ and all time-to-expiries $\tau \in \mathcal{T}$. There are in total $7 \times 10 = 70$ outright test strategies.
- (ii) *Delta spread* $C(\tau, \Delta_1) - C(\tau, \Delta_2)$ with all distinguished delta pairs $\Delta_1 > \Delta_2 \in \mathcal{D}$ and all time-to-expiries $\tau \in \mathcal{T}$. There are in total $\binom{7}{2} \times 10 = 210$ delta spread test strategies.
- (iii) *Delta butterfly* $C(\tau, \Delta) + C(\tau, 1 - \Delta) - 2 \times C(\tau, 0.5)$ with all deltas $\Delta \in \mathcal{D} \setminus \{0.5\}$ and all time-to-expiries $\tau \in \mathcal{T}$. There are in total $3 \times 10 = 30$ delta butterfly spread test strategies.
- (iv) *Delta-hedged option* $C(\tau, \Delta) - \Delta \times S$ with $\Delta = 0.5$ and all time-to-expiries $\tau \in \mathcal{T}$. There are in total $1 \times 10 = 10$ delta-hedged option test strategies.

(v) *Delta-neutral strangle* $C(\tau, \Delta) + C(\tau, 1 - \Delta) - S$ with all deltas $\Delta \in \mathcal{D} \setminus \{0.5\}$ and all time-to-expiries $\tau \in \mathcal{T}$. There are in total $3 \times 10 = 30$ delta-neutral strangle test strategies.

Note that a typical delta-neutral strangle involves simultaneously buying a call and a put, i.e. $C(\tau, \Delta) + P(\tau, -\Delta)$. Here we use put-call parity to express such straddles in terms of only calls and the underlying, ignoring cash components.

(vi) *Risk reversal* $C(\tau, \Delta) - C(\tau, 1 - \Delta) + S$ with all deltas $\Delta \in \{0.2, 0.3, 0.4\}$ and all time-to-expiries $\tau \in \mathcal{T}$. There are in total $3 \times 10 = 30$ risk reversal test strategies. Note that a typical risk reversal consists of buying an OTM call and selling an OTM put, i.e. $C(\tau, \Delta) - P(\tau, -\Delta)$. Here we use put-call parity to express such straddles in terms of only calls and the underlying, ignoring cash components.

(vii) *Calendar spread* $C(\tau_1, \Delta) - C(\tau_2, \Delta)$ with $\Delta = 0.5$ and all distinguished time-to-expiry pairs $\tau_1 > \tau_2 \in \mathcal{T}$. There are in total $1 \times \binom{10}{2} = 45$ calendar spread test strategies.

(viii) *VIX*, the square of which is a linear combination of OTM call and put option prices, and can be further written as a linear combination of call prices only, provided that put-call parity holds under no-arbitrage; see details in [41].

In total, as we consider both long and short positions, there are $= 2 \times (70 + 210 + 30 + 10 + 30 + 30 + 45 + 1) = 852$ tested trading strategies. Among the tested trading strategies, delta butterflies, delta-hedged options, delta-neutral strangles and calendar spreads have low values of Black–Scholes delta.

8.C More VaR backtesting results

We show backtesting results, using the model in the main text, for 2-day, 5-day and 10-day VaRs for EURO STOXX 50 index option portfolios in Table 8.C.1. Similarly, we show backtesting results for 1-day, 2-day, 5-day and 10-day VaRs for DAX option portfolios in Table 8.C.2. Aligned with our observation in the main text, the nSDE-VaR model outperforms the FHS-VaR approach through the backtesting analysis in almost all metrics.

In Figure 8.C.1 and 8.C.2, we show the time series of 1-day FHS-VaR_{0.99} breaches for delta-exposed and vega-exposed option portfolios respectively, and compare with those for positions in the underlying index and the VIX portfolio. These are results for EURO STOXX 50 index options, and are supposed to provide a comparison with Figure 8.3.1 and

	2-day VaR _{0.99}		2-day VaR _{0.95}		5-day VaR _{0.99}		5-day VaR _{0.95}	
	nSDE	FHS	nSDE	FHS	nSDE	FHS	nSDE	FHS
Coverage ratio median	0.9960	0.9920	0.9602	0.9602	0.9960	0.9879	0.9677	0.9637
Coverage ratio mean	0.9950	0.9800	0.9602	0.9402	0.9944	0.9798	0.9618	0.9464
Kupiec PF (two-sided)	0.12%	13.11%	38.29%	35.71%	1.41%	12.76%	45.08%	37.12%
Kupiec PF (one-sided)	0.12%	13.11%	7.26%	15.69%	1.41%	12.76%	7.61%	14.05%

Table 8.C.1: Summary statistics on coverage tests for 2-day and 5-day VaRs computed by the neural-SDE (nSDE) market model and the filtered historical simulation (FHS) approach. These results are for EURO STOXX 50 index options.

	1-day VaR _{0.99}				1-day VaR _{0.95}			
	nSDE		FHS		nSDE	FHS		
Coverage ratio median	0.9881		0.9841		0.9484	0.9524		
Coverage ratio mean	0.9872		0.9726		0.9556	0.9314		
Kupiec PF (two-sided)	14.64%		19.56%		33.49%	34.19%		
Kupiec PF (one-sided)	14.64%		19.56%		6.21%	18.03%		
Christoffersen independence	0.47%		7.61%		4.68%	10.42%		
Conditional coverage	8.31%		19.44%		22.60%	33.49%		
Basel committee traffic light	67.3%	31.1%	0.8%	56.9%	28.7%	12.9%	N.A.	N.A.

	2-day VaR _{0.99}		2-day VaR _{0.95}		5-day VaR _{0.99}		5-day VaR _{0.95}	
	nSDE	FHS	nSDE	FHS	nSDE	FHS	nSDE	FHS
Coverage ratio median	0.9960	0.9920	0.9681	0.9482	0.9960	0.9839	0.9758	0.9637
Coverage ratio mean	0.9903	0.9811	0.9602	0.9335	0.9928	0.9778	0.9644	0.9453
Kupiec PF (two-sided)	8.31%	13.93%	54.33%	37.00%	0.00%	14.40%	58.90%	42.62%
Kupiec PF (one-sided)	8.31%	13.93%	10.42%	20.37%	0.00%	14.40%	8.78%	14.64%

	10-day VaR _{0.99}		10-day VaR _{0.95}	
	nSDE	FHS	nSDE	FHS
Coverage ratio median	1.0000	0.9794	0.9671	0.9630
Coverage ratio mean	0.9949	0.9742	0.9659	0.9367
Kupiec PF (two-sided)	0.12%	22.72%	43.44%	37.59%
Kupiec PF (one-sided)	0.12%	22.72%	7.61%	20.37%

Table 8.C.2: Summary statistics on coverage tests for 2-day, 5-day and 10-day VaRs computed by the neural-SDE (nSDE) market model and the filtered historical simulation (FHS) approach. These results are for DAX options.

8.3.2. The FHS-VaRs tend to have poor coverage performance for delta-hedged option portfolios.

In addition to the coverage and independence test results, in Figure 8.C.3 and 8.C.4, we show the comparison of more procyclicality results produced by the nSDE-VaRs and the FHS-VaRs for both index options. The nSDE-VaRs are consistently less procyclical than the FHS-VaRs in all scenarios.

Finally, we compare the VaR backtesting results for EURO STOXX 50 index options generated by the neural-SDE models with different configurations. Specifically, we have described the model with three primary factors (referred to as “nSDE-3”) in Appendix 7.C.1 and the model where a full diffusion matrix of (S, ξ) is specified (referred to as “nSDE-J”) in Appendix 7.C.2. In Table 8.C.3, we list the statistics on coverage and independence tests for 1-day, 2-day and 5-day VaRs computed by the two models. Compared




Figure 8.C.1: Times series of 1-day FHS-VaR_{0.99} breaches (red dots) for delta-exposed option portfolios, compared with that for positions in the underlying index.

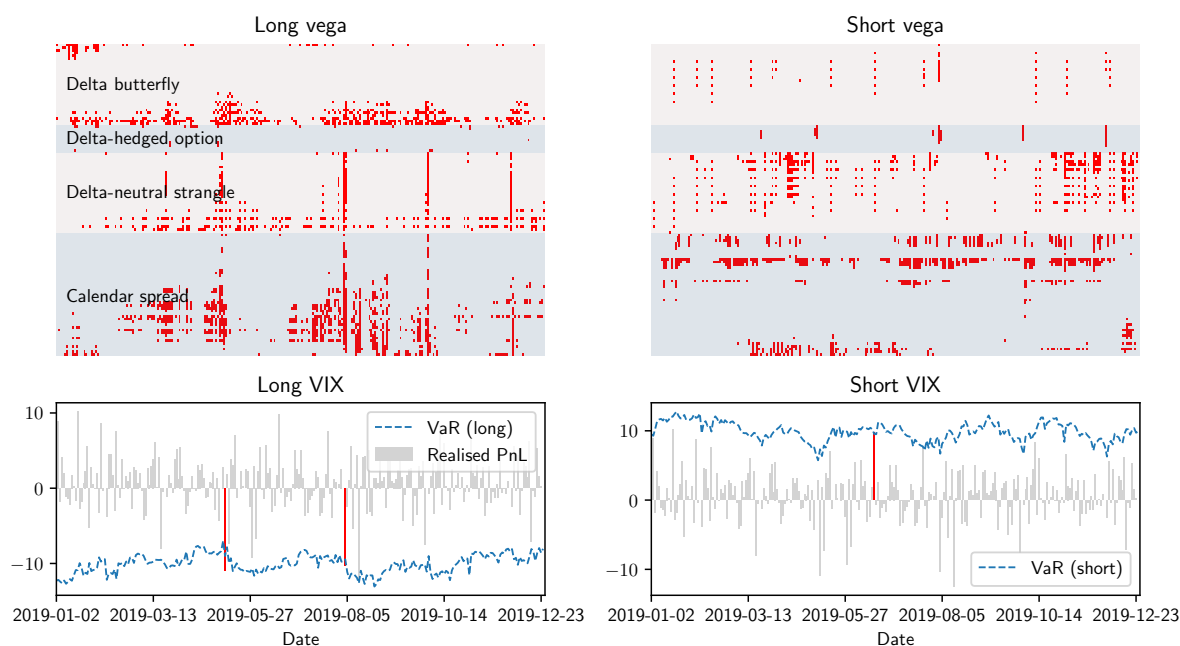


Figure 8.C.2: Times series of 1-day FHS-VaR_{0.99} breaches (red dots) for delta-neutral and vega-exposed option portfolios, compared with that for positions in the VIX portfolio.

with the original model for which the related 1-day VaR backtesting results are shown in Table 8.3.2, the 3-factor model slightly outperforms the original 2-factor model for most tests, with 17.3% more portfolios in the green zone. Nevertheless, when compared with the original model for which the related 2-day and 5-day VaR backtesting results

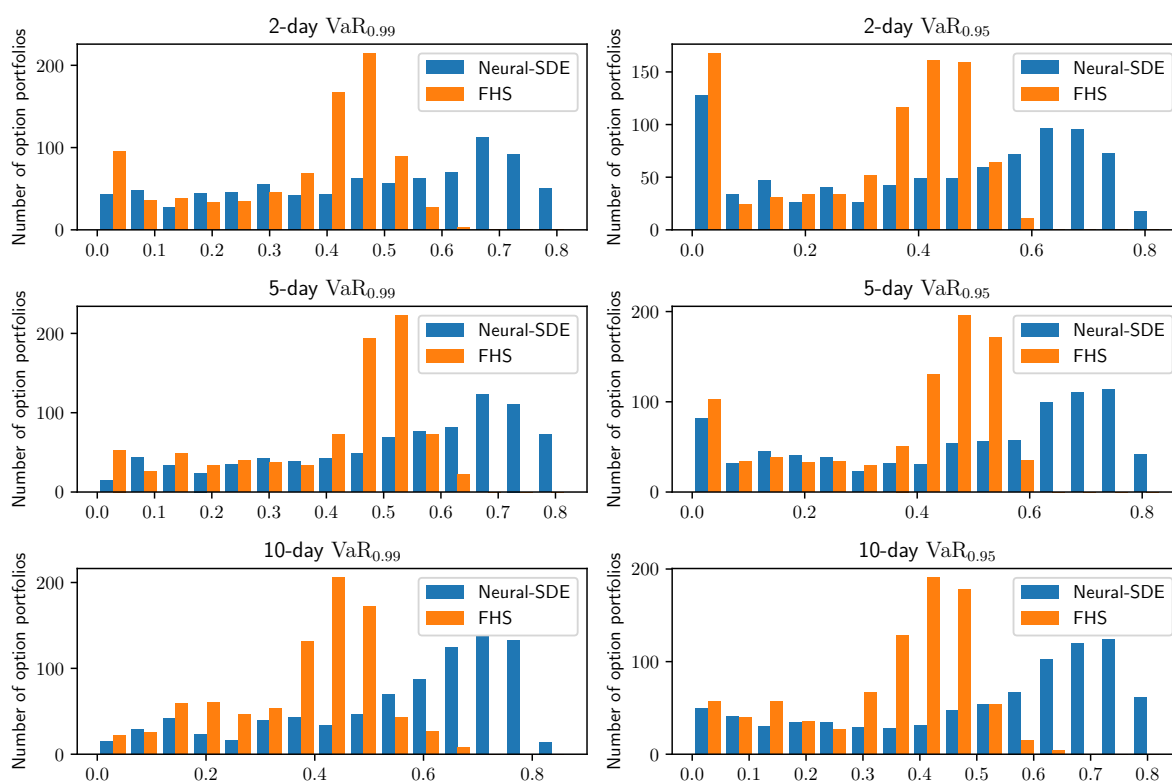


Figure 8.C.3: Distribution of trough-to-peak ratios over different tested portfolios of options on the EURO STOXX 50 index.

are shown in Table 8.C.1, neither the 3-factor model nor the original 2-factor model is significantly better than the other.

	1-day VaR _{0.99}		1-day VaR _{0.95}	
	nSDE-3	nSDE-J	nSDE-3	nSDE-J
Coverage ratio median	0.9881	0.9960	0.9643	0.9742
Coverage ratio mean	0.9891	0.9911	0.9603	0.9667
Kupiec PF (two-sided)	4.68%	11.12%	28.81%	51.41%
Kupiec PF (one-sided)	4.68%	11.12%	5.15%	6.32%

	2-day VaR _{0.99}		2-day VaR _{0.95}		5-day VaR _{0.99}		5-day VaR _{0.95}	
	nSDE-3	nSDE-J	nSDE-3	nSDE-J	nSDE-3	nSDE-J	nSDE-3	nSDE-J
Coverage ratio median	0.9960	1.0000	0.9641	0.9880	1.0000	1.0000	0.9637	0.9899
Coverage ratio mean	0.9946	0.9938	0.9634	0.9718	0.9966	0.9946	0.9694	0.9731
Kupiec PF (two-sided)	1.05%	3.04%	33.84%	65.11%	0.00%	3.75%	40.52%	75.76%
Kupiec PF (one-sided)	1.05%	3.04%	4.57%	6.91%	0.00%	3.75%	2.34%	8.08%

Table 8.C.3: Summary statistics on coverage and independence tests for 1-day, 2-day and 5-day VaRs computed by the neural-SDE market models of different configurations. These results are for EURO STOXX 50 options.

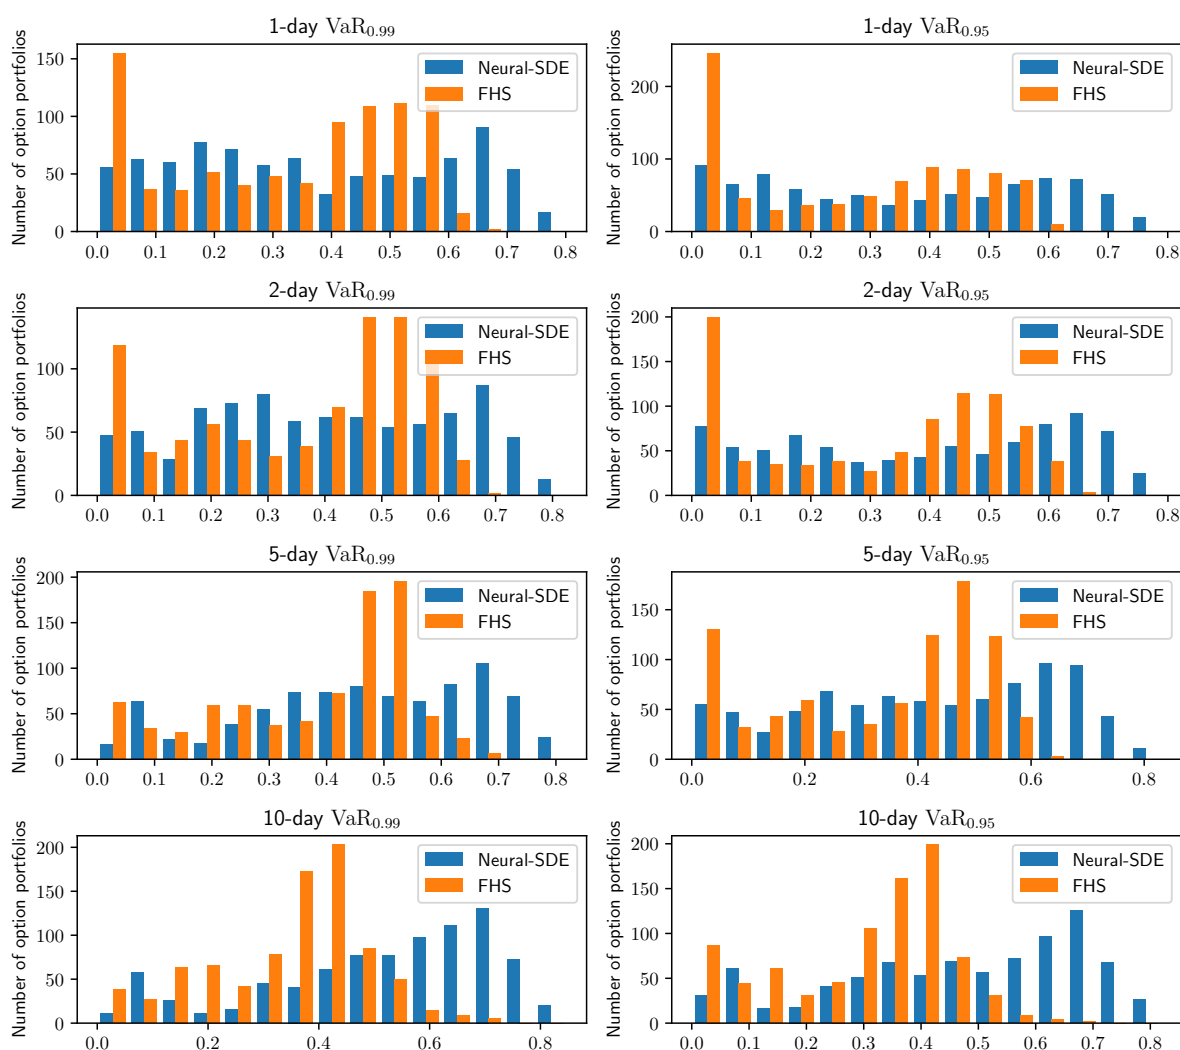


Figure 8.C.4: Distribution of trough-to-peak ratios over different tested portfolios of options on the DAX index.

8.D In-sample VaR backtesting

In addition to the out-of-sample VaR backtesting analysis, we show how in-sample nSDE-VaR looks in this section. Specifically, since we have used option data from 2002 to 2018 to train the neural-SDE model, the “in-sample” nSDE-VaRs then refer to the VaR time series from 2002 to 2018, computed using the trained neural-SDE model.

In Figure 8.D.1, we show the time series of 1-day realised PnLs, nSDE-VaR_{0.99}’s and the VaR breaches for long and short positions in EURO STOXX 50 index. The market experienced three notable stressed periods since 2002, corresponding to the three volatility clusterings in 2002 (the dot-com bubble bursting), 2008 (the global financial crisis triggered by the US subprime mortgage crisis) and 2015 (the European sovereign debt crisis). Overall, the nSDE-VaRs are reactive to changes in market conditions, and have

produced statistically sufficient risk margins over time.




Figure 8.D.1: Time series of realised 1-day PnL and nSDE-VaR_{0.99} breaches (red bars) for long and short positions in the underlying index (EURO STOXX 50).

Application: hedging risks of option books

“流水不腐，户枢不蠹，动也。”

— 吕不韦《吕氏春秋·尽数》

“Flowing waters do not stagnate and door hinges do not get mole crickets, because they move.”

— Buwei Lü, *Fullfill the Number, The Spring and Autumn Annals of Lü Buwei* (translated by John Knoblock and Jeffrey Riegel)

9.1 Introduction

Hedging is critical for managing market risks of option books; in particular, the less hedging errors there are, the smaller bid-ask spreads that market makers are capable of offering to improve liquidity. In general, to hedge an option dynamically, one starts with identifying its risk factors (e.g. the price and volatility of the underlying asset), then evaluates the sensitivities of the option value with respect to these risk factors (normally known as “Greeks”), and finally trades a certain amount of liquid instruments so that the sensitivities of the combined portfolio to *some* risk factors vanish. In order to calculate the sensitivities, or Greeks, we require a pricing model that evaluates an option’s price from its risk factors. We note that this is true even for vanilla options, so while models are typically redundant for determining the current value of vanilla options, as these are observed on the market, they are needed for finding the hedge ratios. Our neural-SDE factor-based models are extremely well suited for this task, as we elaborate now. While there is a vast literature on option pricing models and their associated hedging strategies, this chapter explores the capacity of the neural-SDE market models for producing effective hedging strategies for portfolios of vanilla European options using real world data. Though

we also discuss hedge ratios for exotics, such as barrier and lookback options, we leave empirical assessment of these for future works.

A typical option price model describes the dynamics of the underlying price by a stochastic process. The Black–Scholes model [22] assumes constant volatility (and infers a complete market), where continuously rebalanced delta-hedging using the underlying asset should completely remove risk from the option exposure. However, due to model misspecifications (for instance, the assumption of constant volatility) and the impossibility of continuous rebalancing in practice, there is residual risk for delta-hedged options. The resulting residual risk may be reduced by the addition of a gamma or vega hedge, which requires taking positions in a second option on the same underlying. The introduction of stochastic volatilities and/or jumps in models such as Heston [85], SABR [79, 80], Merton [117] and more generally exponential Lévy models [51], renders the market incomplete, making it not possible to eliminate option risk by solely hedging with the underlying asset even under continuously rebalancing. The hedging strategies associated with these models then differ¹, depending on what risk factors are included and how their dynamics are modelled.

The neural-SDE market model consists of an SDE system for the states that represent the underlying asset price and a small number of *factors*. These factors are chosen, in a data-driven manner, for a statistically accurate dynamic representation of the option surface and the minimisation of dynamic and static arbitrage. Consequently, hedging an option under the neural-SDE market models means to neutralise the sensitivities of the option’s modelled price to the underlying (i.e. delta-hedging) and the *factors*. As we shall see later, there are at least two notable differences in the derivation of hedging strategies between neural-SDE market models and traditional martingale models.

- (i) The market model leads to a straightforward computation of price sensitivities to the factors; specifically, the *normalised* call option price is linear in the factors, so that the first partial derivative with respect to a given factor is simply the factor weight. In contrast, computing sensitivities generally involves numerical integration or Monte-Carlo simulation to compute Greeks under a martingale model, due to the lack of analytical formulas for option prices.
- (ii) A model needs to be calibrated to market data before using it for hedging. The neural-SDE market models are calibrated with the *joint* objectives of (a) calibrating

¹Nevertheless, different option pricing models may yield the same hedge ratios in values (and consequently hedging performance) if they are scale-invariant; see the in-depth analysis by Alexander and Nogueira [5].

(exactly in principle) to the cross sections of option prices, which contain forward-looking information, and (b) maximising the likelihood of having observed the historical time-series of the underlying and *liquid options*. These two objectives, however, cannot be fulfilled simultaneously, in a consistent manner², when calibrating martingale models, and there is not a straightforward mechanism of factoring in historical option prices. Choosing one objective over the other can lead to very different parameter values, and therefore raises concerns of model misspecifications (see Alexander and Kaeck [4]).

We now discuss the difference between our approach and other published works on the use of neural networks for pricing and hedging options; see Ruf and Wang [129] for an overview of a large body of the literature. While most works focus on first outputting option prices and then deriving hedging strategies as sensitivities, some directly output hedging strategies from neural networks, such as Buehler et al. [29] and Ruf and Wang [130]. In constructing neural-SDE market models, the neural networks act as function approximators to the drift and diffusion coefficients of the SDE, in order to give a flexible class of models. As a result, neural-SDE market models largely reduce the “black-box” nature of neural network models (see the discussion by [48]) and enhance model interpretability while still retaining their computational advantages.

In this chapter, we will derive sensitivity-based and minimum-variance-based hedging strategies under neural-SDE market models, and examine their performance when applied to various portfolios of EURO STOXX 50 index options over usual and stressed market periods. We compare the market model hedging strategies with the Black–Scholes delta-vega hedging, along with considering the use of different hedging instruments and rebalancing frequencies. Empirical results in Section 9.4 show that hedging statistically decoded factors (e.g. using PCA) by our market models can lead to lower hedging errors than BS delta-vega hedging for various option portfolios, and is less sensitive to the choice of hedging instruments. In addition, knowledge of the correlation between factors that is modelled by the neural-SDE helps to further reduce hedging errors when short-dated options are used as the hedging instrument.

²There are attempts to estimate some parameters using time series of the underlying and calibrating the remaining parameters to option prices (for example, Bates [17] and Broadie et al. [28]). However, this approach remains inconsistent, as it still relies on daily recalibration of the model, which leads to time-varying calibrated parameters, whereas in theory the model parameters are assumed constant over time.

9.2 Option hedging problem

Suppose an option pricing model considers a set of $d + 1$ risk factors $\{S, \xi_1, \dots, \xi_d\}$. An agent sells an option with price V . To ensure that the agent's position is neutral to all $d + 1$ risk factors (at least instantaneously), it is typically necessary to hedge the sold option using $d + 1$ risky instruments. We assume that hedges are rebalanced at discrete time intervals of length Δt , called the *rebalancing horizon*. Therefore, we consider the hedged portfolio

$$\Pi(S, \xi_1, \dots, \xi_d) = V(S, \xi_1, \dots, \xi_d) - X^S S - \sum_{i=1}^d X^{C_i} C_i(S, \xi_1, \dots, \xi_d), \quad (9.1)$$

where X^S is the amount (or *hedge ratio*) held of the underlying asset with value S , and X^{C_i} is the amount of the i -th liquidly tradable call option with value C_i . We observe that V and Π are deterministic functions of the risk factors (S, ξ_1, \dots, ξ_d) , while by substituting the stochastic processes of the risk factors $(S_t, \xi_{1t}, \dots, \xi_{dt})$, we obtain processes V_t and Π_t .

In a one-period hedging problem, suppose at time t the agent holds this hedged portfolio (with all X^S, X^{C_i} constant), then closes it after Δt , and obtains a hedging error of

$$\mathcal{E}_t(\Delta t) = \Pi_t - \Pi_{t+\Delta t}. \quad (9.2)$$

For simplicity, we assume zero interest and dividend yield³. To reduce the variations of the hedging error given by (9.2), the agent finds the hedging strategy defined by $X = (X^S, X^{C_1}, \dots, X^{C_d})$ such that Π is *instantaneously* neutral to each risk factor $f \in \{S, \xi_1, \dots, \xi_d\}$. In the following sections, we introduce sensitivity-based and minimum-variance-based hedging strategies.

9.2.1 Delta-hedging

Usually an agent will not fully hedge all risk factors imposed by a model, likely due to high transaction costs and trading limits of adding extra hedges. In particular, in a delta-hedging strategy, we use the underlying asset as the only hedging instrument, and exclude the use of other options as hedging instruments. Thereafter, by Itô's lemma,

$$d\Pi = \left(\frac{\partial V}{\partial S} - X^S \right) dS + \sum_{i=1}^d [\dots] d\xi_i + [\dots] dt.$$

³The agent could invest the residual cash from the hedged portfolio (9.1) in a risk-free asset. If, for example, the risk-free asset has non-zero interest yield r_t , then the hedging error in (9.2) would be revised to $\mathcal{E}_t(\Delta t) = \Pi_t e^{\int_t^{t+\Delta t} r_s ds} - \Pi_{t+\Delta t}$ instead.

To eliminate Π 's instantaneous exposure to S , we set the coefficient of dS to zero and obtain $X^S = \partial V / \partial S$, which is the standard delta of the option. A pragmatic way of deriving this is to set $\partial \Pi / \partial S = 0$, i.e. to eliminate the (partial) sensitivity of Π to S .

When the dynamics of the other risk factors $\xi = (\xi_1, \dots, \xi_d)$ imposed by the model interact with that of the underlying S , the standard delta might not be the most efficient hedge ratio. An alternative derivation of hedging strategies has been introduced to *locally* minimise the conditional variance process of the hedging error (9.2) by Föllmer and Schweizer [70] and Schweizer [134]. The resulting hedge ratios are known as minimum variance (MV) hedge ratios. Furthermore, Bakshi et al. [10] and Poulsen et al. [125] develop a more intuitive and tractable derivation based on reducing the instantaneous covariance of a delta-hedged portfolio with the underlying to zero, i.e.

$$0 = \langle d\Pi, dS \rangle = \langle dV - X_{\text{mv}}^S dS, dS \rangle = \langle dV, dS \rangle - X_{\text{mv}}^S \langle dS, dS \rangle,$$

where we use X_{mv}^S to denote the MV delta hedge ratio, and $\langle \cdot, \cdot \rangle$ to denote the quadratic co-variation. Using Itô's lemma on V , we obtain

$$\langle dV, dS \rangle = \left\langle \frac{\partial V}{\partial S} dS + \sum_{i=1}^d \frac{\partial V}{\partial \xi_i} d\xi_i, dS \right\rangle = \frac{\partial V}{\partial S} \langle dS, dS \rangle + \sum_{i=1}^d \frac{\partial V}{\partial \xi_i} \langle d\xi_i, dS \rangle.$$

Therefore,

$$X_{\text{mv}}^S = \frac{\langle dV, dS \rangle}{\langle dS, dS \rangle} = \frac{\partial V}{\partial S} + \sum_{i=1}^d \frac{\partial V}{\partial \xi_i} \frac{\langle d\xi_i, dS \rangle}{\langle dS, dS \rangle} = X^S + \sum_{i=1}^d \frac{\partial V}{\partial \xi_i} \frac{\langle d\xi_i, dS \rangle}{\langle dS, dS \rangle}. \quad (9.3)$$

Hence, the MV delta hedge ratio X_{mv}^S is the sensitivity-based delta hedge ratio X^S with adjustments that take into account sensitivities to other risk factors, weighted by the correlation between the risk factor and the underlying price. In particular, if the underlying-factor correlations are all zero, the MV hedge ratio is identical to the sensitivity-based hedge ratio (see more discussions by Alexander and Nogueira [5]).

Remark 9.1. The prevailing presence of the leverage effect [21] in equity markets, i.e. negative correlation between stock return and volatility, has encouraged practitioners to adjust sensitivity-based deltas by either following the MV approach (9.3) or, in a model-free manner, adding other Greeks weighted by statistically regressed coefficients. For example, Hull and White [91] and Ruf and Wang [130] have demonstrated the excellent hedging performance on S&P 500 index options of the statistically adjusted Black–Scholes delta.

9.2.2 Delta-factor-hedging

Now we generalise to hedging other risk factors. Suppose we want to eliminate the exposures to S and the first d' factors (i.e. $\xi_1, \dots, \xi_{d'}$), without loss of generality, using the underlying asset and d' options $\{C_i\}_{i=1, \dots, d'}$ as the hedging instruments. In other words, we exclude the use of other options C_i for each $i > d'$ in (9.1) as hedging instruments.

We use $f \in \{S, \xi_1, \dots, \xi_{d'}\}$ to denote a risk factor. By Itô's lemma,

$$d\Pi = \sum_f \frac{\partial \Pi}{\partial f} df + [\dots] dt = \sum_f \left(\frac{\partial V}{\partial f} - \frac{\partial S}{\partial f} X^S - \sum_{i=1}^{d'} \frac{\partial C_i}{\partial f} X^{C_i} \right) df + [\dots] dt$$

Similar to the delta hedging approaches discussed previously, when hedging each risk factor f , there are two alternatives to deriving the hedge ratios, either directly using sensitivities,

$$\frac{\partial \Pi}{\partial f} = 0 \quad \Rightarrow \quad \begin{bmatrix} \partial S / \partial f \\ \partial C_1 / \partial f \\ \vdots \\ \partial C_{d'} / \partial f \end{bmatrix}^\top \begin{bmatrix} X^S \\ X^{C_1} \\ \vdots \\ X^{C_{d'}} \end{bmatrix} = \frac{\partial V}{\partial f}, \quad (9.4a)$$

or using a minimum variance criterion in terms of the risk factors,

$$\langle d\Pi, df \rangle = 0 \quad \Rightarrow \quad \begin{bmatrix} \langle dS, df \rangle \\ \langle dC_1, df \rangle \\ \vdots \\ \langle dC_{d'}, df \rangle \end{bmatrix}^\top \begin{bmatrix} X^S \\ X^{C_1} \\ \vdots \\ X^{C_{d'}} \end{bmatrix} = \langle dV, df \rangle. \quad (9.4b)$$

Either condition leads to a *linear* equation with regard to the hedging strategy X , because both the partial derivative operator and the quadratic co-variation operator are linear. Consequently, the delta-factor-hedging problem can be formulated as solving a $(d' + 1)$ -dimensional linear system of equations.

Remark 9.2. Alternatively, we can use a minimum variance criterion in terms of the hedging instruments rather than the risk factors, i.e. $\langle d\Pi, dC_i \rangle = 0$, which also leads to a linear condition with regard to the hedging strategy. We compare the difference between this approach and the two criteria above under a specific neural-SDE market model in Section 9.4.4.

9.3 Hedging with neural-SDE market models

We first briefly recall the modelling components of the neural-SDE market model from Chapter 4, and then derive corresponding hedging strategies. Specifically, suppose we

work with d primary factors $\xi = (\xi_1, \dots, \xi_d)$. Let $\tilde{\xi} = (\ln S, \xi)$, then we model its dynamics by the following SDE:

$$d\tilde{\xi}_t = \mu(\xi_t) dt + \sigma(\xi_t) dW_t, \quad (9.5)$$

where $W = [W_1 \ \dots \ W_{d+1}]^\top \in \mathbb{R}^{d+1}$ are standard independent Brownian motions under the real-world measure \mathbb{P} . We denote by $L_{\text{loc}}^p(\mathbb{R}^d)$ the space of all \mathbb{R}^d -valued, progressively measurable, and locally p -integrable (in t , \mathbb{P} -a.s.) processes, and assume that $\mu \in L_{\text{loc}}^1(\mathbb{R}^{d+1})$ and $\sigma \in L_{\text{loc}}^2(\mathbb{R}^{d+1})$.

Now we associate $\tilde{\xi}$ with option prices. We denote by $C_t(T, K)$ the market price at time t of an European call option with expiry T and strike K . In addition, to reparametrise the call price function $C_t(\cdot, \cdot)$, we introduce $\tau = T - t$ to be time-to-expiry and $m = \ln(K/S_t)$ to be moneyness. Thereafter, we assume

$$C_t(T, K) = S_t \tilde{c}_t(\tau, m), \quad (9.6a)$$

$$\tilde{c}_t(\tau, m) = G_0(\tau, m) + \sum_{i=1}^d G_i(\tau, m) \xi_{it}, \quad (9.6b)$$

where $G_i \in C^{1,2}(\mathbb{R}_{\geq 0} \times \mathbb{R})$ for $i = 0, \dots, d$ is called a *price basis function*. We call \tilde{c} the normalised call option price, and (9.6b) assumes that the normalised call price is a time-independent linear combination of the factors ξ .

Next, we derive the hedge ratios with the sensitivity-based and the MV-based approaches for a call option with maturity T^* and strike price K^* . For ease of notation, we remove the time dependence in the following derivations.

9.3.1 Delta-hedging

Let $V = C(T^*, K^*)$ be the option to hedge. Using (9.6a), in a sensitivity-based delta-hedging strategy, the hedge ratio of the underlying asset is

$$X^S = \frac{\partial V}{\partial S} = \frac{\partial}{\partial S} (S\tilde{c}) = \left(\tilde{c} - \frac{\partial \tilde{c}}{\partial m} \right) (\tau^*, m^*), \quad \text{where } \tau^* = T^* - t, m^* = \ln(K^*/S). \quad (9.7)$$

It is worth noting that the sensitivity-based delta X^S is independent of the dynamics of S and ξ in the model (9.5). In addition, if the prices of the T^* -maturity options, for a continuum of strike prices, are available, the evaluation of $\partial \tilde{c} / \partial m$ is model-free, hence X^S could be computed independently of any models. This is indeed consistent with the conclusion by Alexander and Noguera [5] that hedge ratios for European claims under scale-invariant models are model-free, provided these models have a perfect fit to market prices.

In contrast to sensitivity-based deltas, deriving MV deltas relies on the dynamics in (9.5). In particular, we assume, without loss of generality, that $\sigma = (\sigma_{ij}) \in \mathbb{R}^{(d+1) \times (d+1)}$ is a lower triangular matrix, so that we have,

$$\langle dS, dS \rangle = \sigma_{11}^2 S^2 dt, \quad \langle d\xi_i, dS \rangle = \sigma_{11} \sigma_{i+1,1} S dt.$$

In addition,

$$\frac{\partial V}{\partial \xi_i} = \frac{\partial}{\partial \xi_i} (S\tilde{c}) = S \frac{\partial}{\partial \xi_i} \left(G_0 + \sum_{j=1}^d G_j \xi_j \right) = SG_i. \quad (9.8)$$

Hence, the MV delta hedge ratio in (9.3) under the market model can be computed as

$$X_{\text{mv}}^S = X^S + \sum_{i=1}^d \frac{\sigma_{i+1,1}(\xi)}{\sigma_{11}(\xi)} G_i(\tau^*, m^*). \quad (9.9)$$

9.3.2 Delta-factor-hedging

Now we evaluate the two alternative hedging equations (9.4a) and (9.4b) under the market models given by (9.5) and (9.6). For an asset (underlying or option) U , we introduce the following notation:

$$\begin{aligned} \Delta(U) &:= \frac{\partial U}{\partial S}, & \Delta_{\text{mv}}(U) &:= \Delta(U) + \sum_{i=1}^d \Delta^i(U) \frac{\sigma_{i+1,1}}{\sigma_{11} S}, \\ \Delta^j(U) &:= \frac{\partial U}{\partial \xi_j}, & \Delta_{\text{mv}}^j(U) &:= \Delta(U) \sigma_1 \sigma_{j+1}^\top S + \sum_{i=1}^d \Delta^i(U) \sigma_{i+1} \sigma_{j+1}^\top, \end{aligned} \quad (9.10)$$

where σ_i is the i -th row of the diffusion matrix $\sigma \in \mathbb{R}^{(d+1) \times (d+1)}$. In particular, we call Δ^j (Δ_{mv}^j) the sensitivity-based (MV-based) ξ_j -*exposure*. To hedge V with the strategy $X = (X^S, X^{C_1}, \dots, X^{C_{d'}})$ following the sensitivity-based approach, we set the partial sensitivity of the hedged portfolio Π with regard to a risk factor f to zero and establish the linear equation

$$\frac{\partial \Pi}{\partial f} = 0 \Rightarrow \begin{cases} X^S + \sum_{i=1}^{d'} \Delta(C_i) X^{C_i} = \Delta(V), & \text{if } f = S; \\ \Delta^j(S) X^S + \sum_{i=1}^{d'} \Delta^j(C_i) X^{C_i} = \Delta^j(V), & \text{if } f = \xi_j, \text{ for } j = 1, \dots, d'. \end{cases} \quad (9.11)$$

And with MV-based hedging where there is zero instantaneous co-variation between the hedged portfolio Π and the risk factor f , we have

$$\langle d\Pi, df \rangle = 0 \Rightarrow \begin{cases} X^S + \sum_{i=1}^{d'} \Delta_{\text{mv}}(C_i) X^{C_i} = \Delta_{\text{mv}}(V), & \text{if } f = S; \\ \Delta_{\text{mv}}^j(S) X^S + \sum_{i=1}^{d'} \Delta_{\text{mv}}^j(C_i) X^{C_i} = \Delta_{\text{mv}}^j(V), & \text{if } f = \xi_j, \text{ for } j = 1, \dots, d'. \end{cases} \quad (9.12)$$

Both (9.11) and (9.12) indicate that, for each risk factor f , the aggregated f -exposure of all hedging instruments equals the f -exposure of the option to be hedged. In addition, the only difference between them is that the corresponding f -exposure is computed based on neutral sensitivity and minimum variance, respectively.

Remark 9.3. The sensitivity-based hedging equation does not exploit statistical information established by the SDE model (9.5), and completely relies on the factor representation of option prices given by (9.6). This means it is a simpler approach to hedging, which does not use “machine learning” techniques to a great extent.

Remark 9.4. When one risk factor f is independent of all other risk factors, its MV-based hedging condition (9.12) becomes equivalent to its sensitivity-based hedging condition (9.11). To see this, suppose S is independent of ξ , then $\sigma_{i+1,1} = 0$, which yields $\Delta_{\text{mv}}(U) = \Delta(U)$ according to (9.10); if ξ_j is independent of S and the other factors, $\Delta_{\text{mv}}^j(U) = \Delta^j(U)\sigma_{j+1,j+2}^2 \propto \Delta^j(U)$.

9.3.3 A specific model

Here we consider a simpler setting of the market model as in Chapter 7 and 8 for modelling historical price data of EURO STOXX 50 index options and DAX options and therefore constructing a realistic option market simulator for computing VaR.

In particular, we assume $d\langle \xi, S \rangle = 0$ (no correlation between S and factors ξ ; see Remark 7.2 in Section 7.2 for the discussion of this assumption) and $d = 2$ (two primary factors), which implies that we can write the diffusion matrix in (9.5) as

$$\sigma = \left[\begin{array}{c|c} \gamma & \mathbf{0} \\ \hline \mathbf{0} & \sigma_\xi \end{array} \right], \quad (9.13)$$

where $\gamma = \sigma_{11} \in \mathbb{R}$ and $\sigma_\xi \in \mathbb{R}^{d \times d}$. We focus on studying the following two problems:

- (i) In both the numerical experiments with Heston-SLV synthetic data (Chapter 6) and with real world data (Chapter 7), the first primary factor ξ_1 behaves similarly to the volatility process of the underlying. In fact, as shall be seen in Section 9.4.3, vegas of vanilla options have qualitatively similar geometry over strikes and expiries to their ξ_1 -exposures. Hence, we are interested in seeing if hedging an option's ξ_1 -exposure could achieve comparable performance with hedging its vega, while the former has computational advantages. The calculation of an option's vega relies on the model specified for the underlying and its volatility process, and, in general, is computationally costly due to lack of analytical formulas and that the model calibration could also be non-trivial. In contrast, the sensitivity-based ξ_1 -exposure,

i.e. SG_1 , is decoded from historical time series of option price data using a fast PCA-based algorithm (see Algorithm 1 in Section 4.7.1).

- (ii) Will hedging MV-based ξ_1 -exposures perform better than hedging sensitivity-based ξ_1 -exposures? Since MV-based hedging exploits the correlation between factors, captured by the SDE model (9.5), we are therefore interested in seeing whether estimating a joint-dynamics model for the factors is beneficial for hedging. In particular, with the specifications of the diffusion matrix in (9.13), we focus on evaluating the impact of modelling the correlation between ξ_1 and ξ_2 on hedging performance.

Assuming the hedging instruments available are the underlying asset S and one liquid call option C_1 , we derive their hedge ratios under the model specification (9.13) with the objective of hedging exposures to S and ξ_1 (i.e. $d' = 1$). With sensitivity-based hedging, we solve (9.11) and obtain

$$X^S = \left(\tilde{c} - \frac{\partial \tilde{c}}{\partial m} \right) (\tau^*, m^*) - X^{C_1} \left(\tilde{c} - \frac{\partial \tilde{c}}{\partial m} \right) (\tau_1, m_1), \quad X^{C_1} = \frac{G_1(\tau^*, m^*)}{G_1(\tau_1, m_1)}. \quad (9.14)$$

And with MV-based hedging, we solve (9.12) and obtain

$$X_{\text{mv}}^S = \left(\tilde{c} - \frac{\partial \tilde{c}}{\partial m} \right) (\tau^*, m^*) - X_{\text{mv}}^{C_1} \left(\tilde{c} - \frac{\partial \tilde{c}}{\partial m} \right) (\tau_1, m_1), \quad X_{\text{mv}}^{C_1} = \frac{\sum_{i=1}^d G_i(\tau^*, m^*) \sigma_{i+1} \sigma_2^T}{\sum_{i=1}^d G_i(\tau_1, m_1) \sigma_{i+1} \sigma_2^T}. \quad (9.15)$$

Here τ_1 and m_1 denote the time-to-expiry and moneyness of the hedging option C_1 .

9.3.4 Hedging strategies for more option types

The hedge ratios (9.14) and (9.15) derived in the previous section are for European vanilla call options. Nevertheless, the sensitivity-based and MV-based hedging equations (9.11) and (9.12) can be universally applied to arbitrary option types, and therefore yield corresponding hedge ratios.

To derive the hedge ratios for a generic option V , we need to evaluate $\Delta(V)$ and $\Delta^j(V)$ for which the definitions are given in (9.10). This will be straightforward if $V = h(S, C_1, \dots, C_n)$ where $\{C_1, \dots, C_n\}$ is a series of n European vanilla call options:

$$\frac{\partial V}{\partial S} = \frac{\partial h}{\partial S} + \sum_{i=1}^n \frac{\partial h}{\partial C_i} \Delta(C_i), \quad \frac{\partial V}{\partial \xi_j} = \sum_{i=1}^n \frac{\partial h}{\partial C_i} \Delta^j(C_i).$$

We discuss a few examples as follows.

European put option For a European put option with maturity T and strike price K , there is a model-free no-arbitrage relationship between the European call option with exactly the same maturity and strike price. This relation is called put-call parity and is given by $P(K) = C(K) - S + K$, where we use $P(K)$ and $C(K)$ to denote prices of the put and the call (without explicitly showing their dependence on T for ease of notation). Therefore, we have

$$\frac{\partial P(K)}{\partial S} = \Delta(C(K)) - 1, \quad \frac{\partial P(K)}{\partial \xi_j} = \Delta^j(C(K)).$$

Binary options Let $BC(K)$ denote a binary call option with strike price K . It is known that binary calls can be synthesised using an infinite number of vertical spreads of European calls:

$$BC(K) = \lim_{n \rightarrow +\infty} n \left(C(K) - C\left(K + \frac{1}{n}\right) \right) = \frac{\partial C(K)}{\partial K}.$$

Assuming that $C \in \mathcal{C}^2$ (i.e. all the second-order partial derivatives exist and are continuous),

$$\begin{aligned} \frac{\partial BC(K)}{\partial S} &= \frac{\partial \Delta(C(K))}{\partial K} = \frac{1}{K} \left(\frac{\partial \tilde{c}}{\partial m} - \frac{\partial^2 \tilde{c}}{\partial m^2} \right), \\ \frac{\partial BC(K)}{\partial \xi_j} &= \frac{\partial \Delta^j(C(K))}{\partial K} = -\frac{S}{K} \left(G_j - \frac{\partial G_j}{\partial m} \right). \end{aligned}$$

Thereafter, the corresponding partial derivatives for a binary put option $BP(K)$ can be derived easily using the model-free no-arbitrage put-call parity $BP(K) = 1 - BC(K)$.

Exotic options Carr, Ellis and Gupta [37] have developed static hedges for various exotic options such as (single and multiple) barrier options and lookback options using European vanilla puts and calls, with assumptions on the shape of the local volatility smiles⁴.

For example, a down-and-out call option⁵, denoted as $DOC(K, B)$, with strike K and barrier $B < K$ can be completely replicated by a portfolio of calls and puts

$$DOC(K, B) = C(K) - KB^{-1}P(B^2K^{-1}).$$

The replicating portfolio matches the terminal payoff (i.e. long a call $C(K)$) and the payoff along the barrier (i.e. short additional KB^{-1} puts struck at B^2K^{-1}). Using put-call parity,

⁴However, these assumptions will not generally hold for neural-SDE market models. It remains to be investigated empirically how the violation of these assumptions impacts hedging performance under neural-SDE market models.

⁵A down-and-out call option is a vanilla call option if its barrier has not been hit by its expiry date. However it becomes worthless if its barrier is hit at any time before it expires.

we can further re-write $DOC(K, B)$ as a portfolio of vanilla calls. We refer readers to [37] for how to construct vanilla call portfolios that statically replicate down-and-out, double knockouts, roll-down, ratchet, and lookback options.

In general, if an option can be statically hedged⁶ by European vanilla options, then we can derive its partial derivatives with regard to S and ξ and therefore its hedge ratios. There is vast literature on static hedging of exotic options, to name a few recent developments, [24],[103] and [104].

9.4 Empirical results

We assess the hedging performance under the market model specified in Section 9.3.3 by examining out-of-sample hedging errors where the hedge ratios are given by (9.14) and (9.15). In addition, we use the Black–Scholes (BS) delta-vega hedging approach as the benchmark, and examine if the market model yields better hedging performance than the benchmark.

9.4.1 Option data and the trained neural-SDE market model

We use the same dataset from OptionMetrics’ IvyDB Europe [123] as in Section 7.1 for training a neural-SDE model. To be specific, the training data include historical daily settlement prices for EUROSTOXX 50 index call options and DAX call options from 2nd January, 2002 to 31st December, 2018. The detailed factor decoding and SDE estimation procedures can be found in Chapter 7.

For the hedging analysis, we study EURO STOXX 50 index options, the most traded vanilla European options listed on Eurex. We choose the testing dataset in the OptionMetrics database in a different format from the training dataset (which was also the format used for the testing dataset used in the out-of-sample VaR backtesting analysis as in Section 8.3). Specifically, rather than using standardised option prices that are interpolated over a standard set of expiries and deltas, we use raw end-of-day settlement prices of *actually* traded contracts. It is worth noting that the standardised prices are interpolated from the raw prices. While the standardised prices are convenient for decoding factors and estimating statistical models, the raw prices are more accurate in reflecting options’ values, which are important for calculating hedging errors.

There is another difference in the testing dataset used. The test samples used in the VaR backtesting analysis in Section 8.3 only consist of data up to 30th December

⁶Nevertheless, the static hedging relation is usually derived with certain model assumptions, so that the replicating portfolios may only *approximate* the option’s payoff. The impact of the approximation errors on hedging performance needs to be examined with empirical evidences.

2019. By the time this chapter was written, OptionMetrics had updated its database (on 20th February, 2022) to include option prices until 30th September, 2021. We therefore include these newly released data in the out-of-sample hedging analysis for more accurate statistical results, in particular during the COVID-19 pandemic induced turmoil of early 2020.

9.4.2 Tested option portfolios and hedging error calculation

We assess the performance of using the trained neural-SDE market model for hedging a variety of representative option portfolios. In particular, we define the *naive* portfolio on some day as the equally-weighted portfolio that consists of *all* traded call option contracts within the liquid range⁷ on the day. In addition, similar to the VaR backtesting analysis in Section 8.3, we consider a variety of portfolios from outright options to heavily hedged portfolios; a detailed description is given in Appendix 8.B⁸.

Both the sensitivity-based and MV-based hedging strategies are linear with respect to the constituents of the option portfolio, because “the hedging portfolios can be interpreted as orthogonal projections of contingent claims onto the closed linear subspace of hedgeable portfolios”, as argued by Tankov [142]. Consequently, to hedge a portfolio of options written on the same underlying, we compute the hedge ratios for every option in the portfolio and then add them up.

On each trading day t in the testing data samples (between 2nd January, 2019 and 30th September, 2021), we suppose that an option portfolio with value $V(t)$ has been written. We hedge the portfolio with the underlying asset S and a single tradable call option that has approximately 7-month⁹ time-to-expiry and is closest to at-the-money. We choose this option as the hedging instrument because near ATM options are liquid and have relatively large vegas and ξ_1 -exposures, as will be seen in Section 9.4.3. We hold the hedged portfolio, denoted $\Pi(t)$, until $t + \Delta t$ and then close it, and therefore measure the Δt -hedging error over the out-of-sample testing period as

$$\hat{\mathcal{E}}(\Delta t) = \frac{\sqrt{\sum_t^{N_{\text{test}}} (\Pi(t) - \Pi(t + \Delta t))^2}}{\sqrt{\sum_t^{N_{\text{test}}} (V(t) - V(t + \Delta t))^2}} \times 100\%, \quad (9.16)$$

which is the percentage of the standard deviation of the PnL of the hedged portfolio over the standard deviation of the PnL of the unhedged position. We evaluate hedging

⁷The liquid range can be computed as the convex hull of the liquid lattice defined for the training data; see Section 7.1 for more details.

⁸These portfolios shall be constructed from options of specified time-to-expiries and moneynesses, which are in general impossible to match perfectly with traded options. Therefore, we find the options that have closest time-to-expiries and moneynesses to the desired specifications when constructing portfolios.

⁹We choose a 7-month option because, as we shall see later in Section 9.4.4, using it Black–Scholes delta-vega hedging and delta- ξ_1 hedging yield similarly good hedging performance for the naive portfolio.

strategies through the hedging error defined in (9.16) with smaller values indicating better performance¹⁰. In the following analysis, we will focus on studying hedging performance under two different rebalancing frequencies: daily ($\Delta t = 1$ trading day) and weekly ($\Delta t = 5$ trading days).

9.4.3 Hedging with the Black–Scholes model

Now we derive the hedge ratios under the benchmark Black–Scholes model. The delta and vega of a call option $U = C_t(T, K)$ under the Black–Scholes model are given by

$$\Delta_{\text{bs}}(U) = \Phi(d_1), \quad \mathcal{V}_{\text{bs}}(U) = S\sqrt{T-t}\phi(d_1), \quad \text{where } d_1 = \frac{\ln(S/K)}{\sigma\sqrt{T-t}} + \frac{1}{2}\sigma\sqrt{T-t}. \quad (9.17)$$

Therefore, neutralising the delta and vega exposures of V yields the hedge ratios

$$X_{\text{bs}}^S = \Delta_{\text{bs}}(V) - \Delta_{\text{bs}}(C_1), \quad X_{\text{bs}}^{C_1} = \frac{\mathcal{V}_{\text{bs}}(V)}{\mathcal{V}_{\text{bs}}(C_1)}. \quad (9.18)$$

Remark 9.5. There remains the choice of which volatility should be used as σ in the formulas in (9.17), as the implied volatility of the option to hedge (V) is usually different from the IV of the option used as the hedging instrument (C_1). For example, in computing the hedge ratios in (9.18), should one use V 's IV in the deltas/vegas of both V and C_1 ? In our empirical experiments that follow, we choose to use an option's own IV for computing its delta/vega. As it turns out, this has limited impact due to the robustness of the Black–Scholes model [126].

In Figure 9.4.1, we show the Black–Scholes vegas and the market model ξ_1 -exposures for a few traded call options that have selected time-to-expiries (roughly 1M, 3M, 6M, 1Y and 2Y) as of 29th September, 2021. BS vegas and ξ_1 -exposures have qualitatively similar behaviour over strikes and expiries. For a fixed strike price, they increase over expiries; for a fixed expiry, they first increase over strikes, reach a peak near the ATM strike level, and then decline.

9.4.4 Hedging error analysis for the naive portfolio

Hedging instruments of different tenors Due to liquidity concerns, we restrict our choice of C_1 , the single option used as a hedging instrument, to traded options that are closest to at-the-money. Nevertheless, we assume that we are free to choose tenors up to 1 year. We analyse the errors of hedging the naive portfolio when near ATM options of different tenors are used as the hedging instrument.

¹⁰For a different definition of hedging performance such as reductions in tail risk, then smaller hedging errors as defined in (9.16) do not necessarily indicate better hedging performance.




Figure 9.4.1: Black-Scholes vegas and market model ξ_1 -exposures of traded call options as of 29th September, 2021. Here we only show traded options that are within the liquid range and have selected time-to-expiries.

Specifically, we consider the set of 12 consecutive monthly tenors from 1M to 1Y. For each tenor, we identify the available time-to-expiry of traded options that is closest to the tenor, and then choose the closest to ATM option of the identified time-to-expiry as the hedging instrument. Thereafter, we compute the hedging errors for the naive portfolio over the out-of-sample testing period, under four hedging strategies: (1) BS delta-hedging, (2) BS delta-vega hedging, (3) market model sensitivity-based delta- ξ_1 hedging, and (4) market model MV-based delta- ξ_1 hedging.




Figure 9.4.2: Errors of hedging the naive option portfolio when near ATM options of different tenors are used as the hedging instrument.

In Figure 9.4.2, we show how hedging errors change with the use of hedging options of different tenors, for both daily and weekly rebalancing cases. Overall, the performances of both delta-vega and delta- ξ_1 (sensitivity- and MV-based) hedging are much worse when short-dated options are used as the hedging instrument. In particular, using 1M ATM options as the hedging instrument gives consistently worse errors than a simpler

delta hedging. The potential reason is that, as we have seen in Figure 9.4.1, short-dated options have relatively smaller vegas and ξ_1 -exposures, making them less effective in hedging long-dated options in the portfolio.

In the daily rebalancing case, delta- ξ_1 hedging seems to be at least as good as BS delta-vega hedging, regardless of which tenor of option is used as the hedging instrument. In addition, delta- ξ_1 hedging outperforms BS delta-vega hedging more when shorter-dated options are used as the hedging instrument. In practice, in markets where short-dated options are much more liquid than long-dated options (e.g. FX market), we would then be more comfortable with hedging ξ_1 -exposures than vegas. Overall, the performance of delta- ξ_1 hedging is *less* sensitive than that of BS delta-vega hedging to the choice of hedging option. When the hedging option's time-to-expiry is at least 7 months, the performance of delta-vega and delta- ξ_1 (sensitivity- and MV-based) hedging are similar. Comparing sensitivity-based delta- ξ_1 hedging and MV-based delta- ξ_1 hedging, taking into account the correlation between factors does reduce the hedging errors if the hedging option's time-to-expiry is less than 7 months.

When the hedged portfolios are rebalanced weekly, we observe similar behaviours of the three hedging strategies as seen in the daily rebalancing case. MV-based delta- ξ_1 hedging outperforms sensitivity-based delta- ξ_1 hedging only when the hedging option's time-to-expiry is less than 5 months.

Hedging errors over time Next, we investigate how the three hedging strategies perform over time, during the testing period. While Figure 9.4.2 shows the standard deviations of PnLs over the whole testing period, we now use a rolling window of one-year length and calculate the hedging errors within the window that rolls one day ahead each time until it has passed through the whole testing period. For each of the three hedging strategies, we choose the corresponding *optimal* hedging option that gives the smallest hedging errors in the testing period. Therefore, in the daily rebalancing case, based on the evidence in Figure 9.4.2, we choose 7M ATM options as hedging instruments for BS delta-vega hedging, and 5M ATM options as hedging instruments for both sensitivity- and MV-based delta- ξ_1 hedging. In the top plot of Figure 9.4.3, we report the one-year rolling hedging errors of the three hedging strategies that are rebalanced on a daily basis. For better illustrating the dynamic changes in the hedging errors, we provide the time series of the PnLs for the unhedged naive portfolio and for the sensitivity-based delta- ξ_1 -hedged naive portfolio in the bottom plot.

First, all the three hedging strategies yield similar dynamics of hedging errors. Hedging errors experienced a sharp rise in March 2020, when the fear of the COVID-19 pandemic severely hit global equity markets. The hedging errors maintained a higher-than-usual




Figure 9.4.3: *Top* - one-year rolling hedging errors of the three hedging strategies. *Bottom* - time series of the PnLs for the unhedged naive portfolio and for the sensitivity-based delta- ξ_1 -hedged naive portfolio (daily rebalancing).

level until the spring of 2021. Meanwhile, we could see concurrent larger fluctuations of the portfolio PnLs. When comparing across different hedging strategies, MV-based delta- ξ_1 hedging is consistently the best, and has maintained a 1% advantage in hedging errors when compared with BS delta-vega hedging. Sensitivity-based delta- ξ_1 hedging is better than BS delta-vega hedging except before March 2020.

We report the same hedging error statistics for weekly-rebalanced hedging in Figure 9.4.4. In this comparison analysis, we choose 7M ATM options as hedging options for BS delta-vega hedging, 6M ATM options as hedging options for sensitivity-based delta- ξ_1 hedging, and 5M ATM options as hedging options for MV-based delta- ξ_1 hedging. Sensitivity-based delta- ξ_1 hedging is always better than BS delta-vega hedging. However, unlike the daily rebalancing case, MV-based delta- ξ_1 hedging is worse than sensitivity-based delta- ξ_1 hedging except before March 2020.

Analysis on decoded factors Why does MV-based delta- ξ_1 hedging perform worse than sensitivity-based delta- ξ_1 hedging, in the weekly rebalancing case, after March 2020? By construction, MV-based delta- ξ_1 hedging is beneficial only when the neural-SDE model gives good adjustments to ξ_1 -exposures based on accurate *forecasts* of the correlation between ξ_1 and ξ_2 within the rebalancing horizon. When the rebalancing horizon gets larger (i.e. from one day to one week) and the instantaneous correlation between ξ_1 and ξ_2 changes drastically, the neural-SDE might give a poor model for the average correlation within the rebalancing horizon and therefore deteriorate the hedging performance.

Since the diffusion matrix of a neural-SDE is a (smooth) function of (ξ_1, ξ_2) , we now examine if (ξ_1, ξ_2) behaves differently before and after March 2020. We show the trajectory




Figure 9.4.4: *Top* - one-year rolling hedging errors of the three hedging strategies. *Bottom* - time series of the PnLs for the unhedged naive portfolio and for the sensitivity-based delta- ξ_1 -hedged naive portfolio (weekly rebalancing).

of the decoded primary factors ξ_1 and ξ_2 for the testing data in Figure 9.4.5 by black dots. The top left plot shows the factors for all the testing data (from 2nd January, 2019 to 30th September, 2021). From 24th February to 7th April in 2020, equity markets across the world suffered from a series of crashes due to panic selling following the onset of the COVID-19 pandemic. In particular in the US market, on 16th March, the sudden and dramatic drop in stock prices triggered multiple trading halts (also known as “circuit breakers”) in a single day. We divide the testing period into three sub-periods that are before, between and after the market crashes. Accordingly, we show the trajectories of (ξ_1, ξ_2) during the three sub-periods in the other subplots. During the market crashes (24th Feb – 7th April), as seen in the top right plot, the trajectory of (ξ_1, ξ_2) spreads over a wide range of the conventional area (light grey dots) where the training data usually reside. In particular, we mark on the plot the corresponding dates of the five data points that are out side of the arbitrage-free region (confined by the red dashed line and highlighted in green). Note that the 16th of March is among these dates.

The most visited areas of the factors seem to be distinct *before* and *after* the COVID-19 pandemic. Specifically, before the pandemic, as seen in the bottom left plot, factors tend to cluster around and have small variations in ξ_1 . In contrast, since the pandemic started, factors tend to spread out with much larger variations in ξ_1 , as seen in the top and bottom right plots. Furthermore, comparing with the diffusion matrices visualised by ellipses in Figure 7.4.2, factors before the pandemic have more ill-conditioned diffusion matrices (i.e. with flatter ellipses), implying that the instantaneous correlation between ξ_1 and ξ_2 is stable and close to -1 . Hence, since factors change more drastically post-pandemic, the neural-SDE model fails to give good forecasts of the average factor correlation over

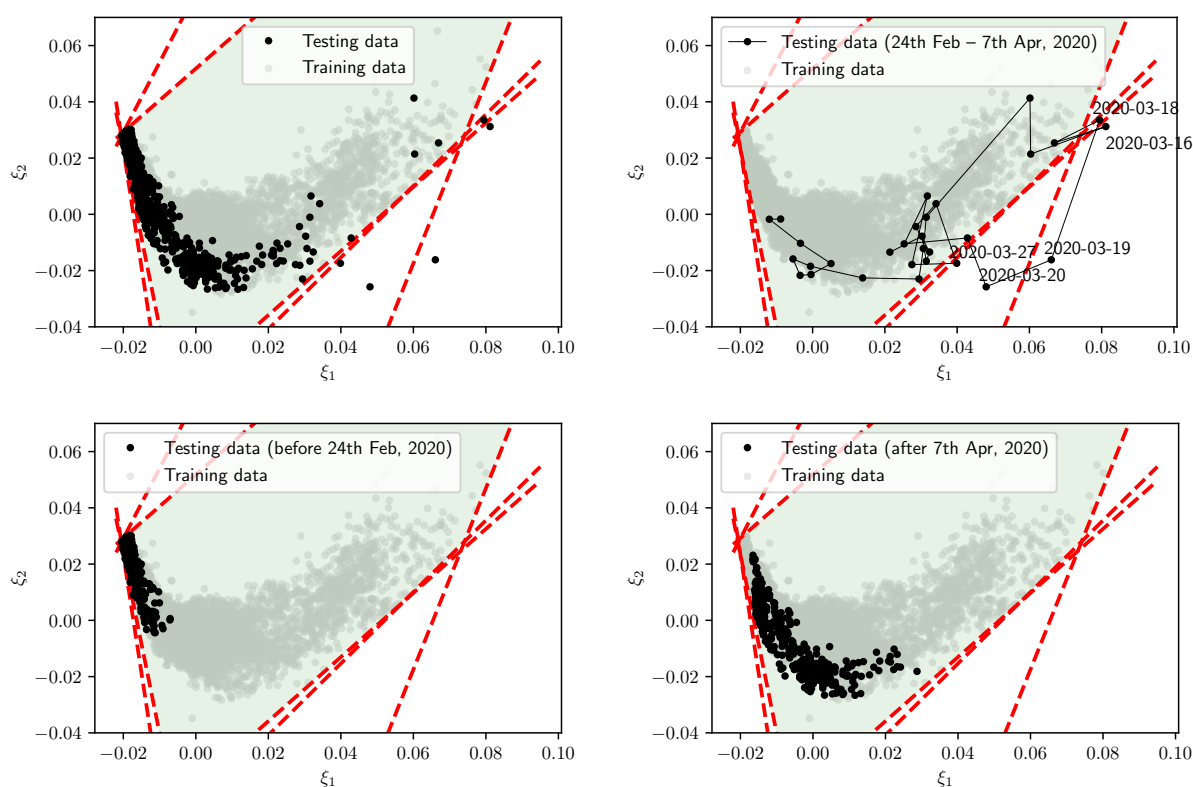


Figure 9.4.5: Scatter plots of the decoded primary factors (ξ_1, ξ_2) for the testing data (from 2nd January, 2019 to 30th September, 2021). The decoded factors for the training data are shown as light grey dots, and the corresponding static arbitrage constraints (projected to the \mathbb{R}^2 factor space) are given as red dashed lines. While the top left plot shows factors of all testing data, the other plots show factors of testing data within the specified periods.

weekly rebalancing horizons. We conjecture this to be a contributing factor why MV-based hedging becomes worse than sensitivity-based hedging after the pandemic.

To support the above arguments with more data, we evaluate the one-year rolling hedging errors (with a weekly rebalancing horizon) back to the year ending January, 2019. Note that all delta- ξ_1 hedge ratios computed before 2019 are based on in-sample estimation of price basis functions G . In Figure 9.4.6, we see that factors in 2018 cluster around similar regions to factors in 2019 and early 2020 before the pandemic. In addition, since the instantaneous correlations estimated within the regions are stable, the neural-SDE model tends to produce reasonable forecasts of the average factor correlation over weekly rebalancing horizons. The corresponding MV-based delta- ξ_1 hedge yields lower hedging errors than sensitivity-based delta- ξ_1 hedging, as seen on the right plot of Figure 9.4.6.

Hedging instruments of different moneynesses Although we have restricted the choice of hedging instruments to ATM options due to liquidity concerns, it is interesting

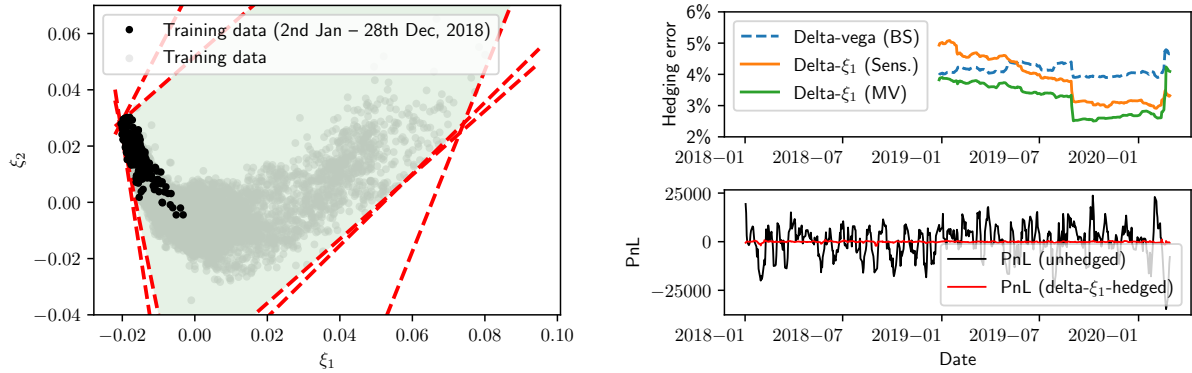


Figure 9.4.6: *Left* - trajectories of the decoded primary factors (ξ_1, ξ_2) for the price data in 2018. *Right* - one-year rolling hedging errors of the three hedging strategies and time series of the PnLs for the unhedged naive portfolio and for the sensitivity-based delta- ξ_1 -hedged naive portfolio (weekly rebalancing).

to see whether ITM and OTM options could achieve lower hedging errors. Therefore, we fix the tenors and see how options of different moneynesses perform when used as hedging instruments.

In Figure 9.4.7, we show how hedging errors change with the use of hedging options of different moneynesses, for both daily and weekly rebalancing cases. For the daily rebalancing case, we choose 7M options as hedging instruments for BS delta-vega hedging, and 5M options for both sensitivity- and MV-based delta- ξ_1 hedging. For the weekly rebalancing case, we choose 7M options as hedging instruments for BS delta-vega hedging, 6M options for sensitivity-based delta- ξ_1 hedging, and 5M options for MV-based delta- ξ_1 hedging. For all three hedging strategies, ATM options achieve the lowest hedging errors (except the MV-based delta- ξ_1 hedging in the weekly rebalancing case). This is unsurprising because both BS-vega and ξ_1 -exposure reach their peak near the ATM strike level, as seen in Figure 9.4.1; hence, ATM options are more effective in reducing a portfolio's BS-vega and ξ_1 -exposure.

MV-based direct hedging We consider an alternative MV-based hedging approach where, rather than requiring zero instantaneous co-variations between hedged portfolios and risk factors, we require a hedged portfolio to have zero instantaneous co-variations with all its hedging instruments. We call this approach MV-based *direct hedging*, and compare its performance with MV-based delta- ξ_1 hedging.

Again, assuming the hedging instruments available are the underlying asset S and one liquid call option C_1 , we derive their hedge ratios under the neural-SDE model specified in Section 9.3.3:

$$\begin{cases} \langle d\Pi, dS \rangle = 0, \\ \langle d\Pi, dC_1 \rangle = 0; \end{cases} \Rightarrow \begin{bmatrix} 1 & \Delta(C_1) \\ \Delta(C_1) & \mathcal{H}(C_1, C_1) \end{bmatrix} \begin{bmatrix} X^S \\ X^{C_1} \end{bmatrix} = \begin{bmatrix} \Delta(V) \\ \mathcal{H}(V, C_1) \end{bmatrix}, \quad (9.19)$$




Figure 9.4.7: Errors of hedging the naive option portfolio when options of different moneynesses (in deltas) are used as the hedging instrument.

where, for two options U and V ,

$$\mathcal{H}(U, V) = \Delta(U)\Delta(V) + \frac{1}{\sigma_{11}^2 S^2} \begin{bmatrix} \Delta^1(U) & \Delta^2(U) \end{bmatrix} \sigma_{\xi} \sigma_{\xi}^{\top} \begin{bmatrix} \Delta^1(V) \\ \Delta^2(V) \end{bmatrix}. \quad (9.20)$$

MV-based delta- ξ_1 hedging constructs hedge ratios by eliminating sensitivities to (S, ξ_1) with adjustments due to correlations between (S, ξ_1) and ξ_2 . Instead, MV-based direct hedging establishes hedge ratios by eliminating a combination of exposures to all risk factors (S, ξ_1, ξ_2) that is represented by the hedging option C_1 .

It turns out that *MV-based direct hedging leads to quite similar hedging errors to MV-based delta- ξ_1 hedging, regardless of which hedging option is used.* To see this, we show in Figure 9.4.8 the difference in hedging errors (errors of MV-based delta- ξ_1 hedging minus errors of MV-based direct hedging) for the naive option portfolio when options of different tenors (in months) and moneynesses (in deltas) are used as hedging instruments. The difference is close to zero in all cases, while there seems to be consistent patterns between the daily and weekly rebalancing cases.

Concluding remarks To conclude from the evidence presented in this section, choosing different options as the hedging instrument makes a substantial difference in the resultant hedging performance. In particular, the performance of BS delta-vega hedging is more sensitive to the choice of ATM hedging options of different tenors than that of market model delta- ξ_1 hedging (see Figure 9.4.2). In order to allow a fair comparison of hedging performance in the following analysis (for more option portfolios), we choose 7M ATM options as the hedging instruments for all the three hedging strategies because using them Black–Scholes delta-vega hedging and delta- ξ_1 hedging yield similarly good hedging performance for the naive portfolio (see Figure 9.4.2). This allows us to disentangle




Figure 9.4.8: Difference in hedging errors (errors of MV-based delta- ξ_1 hedging minus errors of MV-based direct hedging) for the naive option portfolio when options of different tenors (in months) and moneynesses (in deltas) are used as hedging instruments. *Left* - difference in hedging errors when ATM options of different tenors are used as hedging instruments. *Right* - difference in hedging errors when options of specified tenors (corresponding to Figure 9.4.7) and different moneynesses are used as hedging instruments.

the effect of the different payoff characteristics in a situation when the average hedging performance is comparable.

9.4.5 Hedging error analysis for more portfolios

We show the hedging errors of different categories of option portfolios in Table 9.4.1. When there are multiple portfolios within one category (for example, all except the naive and VIX portfolio), the table reports the average hedging errors over all portfolios of the same category. We have the following observations:

- (i) For outright and strangle portfolios, (sensitivity-based and MV-based delta- ξ_1) hedging consistently outperforms BS delta-vega hedging by at least 3%. Meanwhile, MV-based delta- ξ_1 hedging slightly improves sensitivity-based hedging. Compared with other portfolio categories, these two types only involve options of one direction (i.e. either all-long or all-short).
- (ii) For delta spread and delta butterfly portfolios, there is no significant difference in hedging their vegas or ξ_1 -exposures.
- (iii) Sensitivity-based delta- ξ_1 hedging achieves the lowest hedging errors for the VIX portfolio, while the slightly worse MV-based delta- ξ_1 hedging improves over BS delta-vega hedging by at least 5%. This seems to imply that ξ_1 -exposure is a better indicator of volatility risk than BS vega.

- (iv) All the given hedging strategies perform poorly for calendar spread portfolios. The best strategy is the simplest BS delta-hedging, which only manages a reduction of less than 4% of the PnL fluctuations.

Portfolio category	Daily rebalancing				Weekly rebalancing			
	Delta (BS)	Delta-vega (BS)	Delta- ξ_1 (Sens.)	Delta- ξ_1 (MV)	Delta (BS)	Delta-vega (BS)	Delta- ξ_1 (Sens.)	Delta- ξ_1 (MV)
Naive	20.19%	4.36%	4.22%	4.35%	33.28%	4.13%	4.09%	6.17%
Outright	29.00%	14.89%	11.34%	11.13%	43.27%	22.68%	16.68%	15.45%
Delta spread	12.96%	11.45%	11.69%	11.57%	21.47%	18.48%	19.47%	17.41%
Delta butterfly	62.54%	48.77%	51.84%	49.87%	72.52%	49.21%	45.58%	51.43%
Strangle	24.47%	12.05%	9.02%	8.81%	38.35%	18.30%	12.41%	12.24%
Calendar spread	96.21%	164.4%	114.2%	111.2%	96.77%	226.5%	140.1%	132.2%
VIX	81.54%	59.46%	51.12%	53.92%	87.35%	56.59%	45.20%	49.81%

Table 9.4.1: Average hedging errors of different categories of option portfolios under different hedging strategies. The lowest hedging error of each category is highlighted in **bold**.

Next, we compare the hedging performance of the three hedging strategies when applied to outright options of different expiries and moneynesses. In Figure 9.4.9, we show the heatmaps of three pairwise hedging error differences, when the hedges are rebalanced on a daily basis. In each plot, given values indicate the second listed strategy performs better, i.e. less hedging errors. The rightmost plot shows that MV-based delta- ξ_1 hedging outperforms BS delta-vega hedging for almost all combinations of tenors and moneynesses. In particular, the advantage of MV-based delta- ξ_1 hedging gets more pronounced for long-dated (18M and 24M) OTM options. The middle plot shows that, when hedging ξ_1 -exposures, it is beneficial to follow the MV-based approach and take into account the correlation between factors only when the options to hedge have long expiries (typically more than 9 months). It is also interesting to see that BS delta-vega hedging slightly outperforms delta- ξ_1 hedging for some 5M and 6M options. The reason may be that the hedging instrument we choose is a 7M near ATM option, while BS delta-vega hedging gives more efficient hedge ratios for options with specifications nearer the hedging instrument. This is consistent with our observation in Section 9.4.4 that the performance of BS delta-vega hedging is more sensitive to the choice of the hedging option than that of market model delta- ξ_1 hedging.

For completeness, we show in Figure 9.4.10 the heatmaps of the three pairwise hedging error differences when the hedges are rebalanced on a weekly basis. We have very consistent observations with the daily rebalancing case.

Revisiting the hedging of the VIX portfolio

When hedging the VIX portfolio, the best hedging strategy, i.e. sensitivity-based delta- ξ_1 hedging, only manages to reduce the PnL standard deviations by half (see Table 9.4.1).




Figure 9.4.9: Comparison of hedging errors for positions of outright options with various tenors (i.e. time-to-maturities) and moneynesses (measured in deltas). The hedging positions are rebalanced on a daily basis.




Figure 9.4.10: Comparison of hedging errors for positions of outright options with various tenors (i.e. time-to-maturities) and moneynesses (measured in deltas). The hedging positions are rebalanced on a weekly basis.

This is likely due to the fact that the VIX portfolio consists of approximately 1M options [41] while the hedging option used in all the hedging strategies have 7M tenor. Hence, to find the optimal hedging instrument to use for the VIX portfolio, we perform a similar hedging instrument choice analysis as for the naive portfolio in Section 9.4.4.

In Figure 9.4.11, we show how hedging errors change with the use of hedging options of different tenors, for both daily and weekly rebalancing. In both cases, MV-based delta- ξ_1 hedging is at least as good as BS delta-vega hedging, regardless of which tenor of option is used as the hedge. In practice, it might be more sensible to hedge the VIX with 1M options, which are usually more liquid; and when 1M options are used as hedges, MV-based delta- ξ_1 hedging performs the best and significantly improves sensitivity-based delta- ξ_1 hedging. Similar to hedging the naive portfolio, the performance of delta- ξ_1 hedging is *less* sensitive than that of BS delta-vega hedging to the choice of the hedging option for the VIX portfolio.

We report the one-year rolling hedging errors of the three hedging strategies using 1M ATM options as the hedge (that are rebalanced on a daily basis and on a weekly basis




Figure 9.4.11: Errors of hedging the VIX option portfolio when near ATM options of different tenors are used as the hedge.

respectively) in Figure 9.4.12 and 9.4.13. For better illustrating the dynamic changes in the hedging errors, we append the time series of the PnLs for the unhedged VIX portfolio and for the MV-based delta- ξ_1 -hedged VIX portfolio at the bottom. MV-based delta- ξ_1 hedging consistently improves sensitivity-based delta- ξ_1 hedging over the course of the testing period, where the largest improvement amounts to about 20%. MV-based delta- ξ_1 hedging is also consistently better than BS delta-vega hedging most of time, while the advantage is small, especially in the daily rebalancing case.




Figure 9.4.12: *Top* — one-year rolling hedging errors of the three hedging strategies. *Bottom* — time series of the PnLs for the unhedged naive portfolio and for the sensitivity-based delta- ξ_1 -hedged VIX portfolio (daily rebalancing).




Figure 9.4.13: *Top* — one-year rolling hedging errors of the three hedging strategies. *Bottom* — time series of the PnLs for the unhedged naive portfolio and for the sensitivity-based delta- ξ_1 -hedged VIX portfolio (weekly rebalancing).

9.5 Conclusion

We derive a generic system of linear equations of solving sensitivity-based and MV-based hedging strategies under neural-SDE market models. In particular, we focus on a specific two-primary-factor market model and examine the performance of its derived single-option hedging strategies, called delta- ξ_1 hedging, when applied to various portfolios of EURO STOXX 50 index options over usual and stressed market periods. We compare delta- ξ_1 hedging with Black–Scholes delta-vega hedging, along with considering the use of different hedging instruments and rebalancing frequencies. The empirical results address the following two questions:

- (i) *Does market model delta- ξ_1 hedging outperform Black–Scholes delta-vega hedging?*

We have shown that the hedging performance of a strategy highly depends on the choice of hedging instruments, and therefore we answer this question in two aspects:

- (a) BS delta-vega hedging seems to be more sensitive to the choice of the hedging instrument than delta- ξ_1 hedging, where we give evidence when hedging the naive portfolio and the VIX portfolio. Traders might favour less sensitive hedging strategies as they will have more flexibility in choosing the hedging instrument and can better accommodate liquidity and other practical restrictions.
- (b) When optimal hedges are chosen, delta- ξ_1 hedging is shown to outperform BS delta-vega hedging consistently over time and at different rebalancing frequencies, for the naive and the VIX portfolios. In particular, delta- ξ_1 hedging leads to the largest improvements for hedging long-dated and OTM options.

However, there is no significant difference in their hedging performance when hedging delta spread and delta butterfly portfolios.

- (ii) *In a market model, does the dependence between factors that is estimated by the neural-SDE improve delta- ξ_1 hedging? In other words, is MV-based delta- ξ_1 hedging better than sensitivity-based delta- ξ_1 hedging?* There is not a definite answer to this question, as it depends on which hedging instrument is used. From the evidence in the study of the naive and VIX portfolios, MV-based delta- ξ_1 hedging can improve on the sensitivity-based counterpart when short-dated options are used as the hedge. Specifically, it is worth highlighting that when hedging the VIX using 1M options, MV-based delta- ξ_1 hedging can improve sensitivity-based delta- ξ_1 hedging by as much as 20%.

The hedging performance of “PCA-factor” hedging is impressively good: though ξ_1 is a factor decoded statistically from historical option prices using a simple, model-free PCA-based method, hedging its exposure is at least as good as BS vega hedging, which is a commonly acknowledged robust model-based hedging strategy. Meanwhile, since MV-based delta- ξ_1 hedging is not necessarily better than the model-free sensitivity-based counterpart, in some cases it may not be necessary to establish and estimate a model for factor dynamics at all.

Conclusions and extensions

“客亦知夫水与月乎？逝者如斯，而未尝往也；盈虚者如彼，而卒莫消长也。盖将自其变者而观之，则天地曾不能以一瞬；自其不变者而观之，则物与我皆无尽也，而又何羨乎？且夫天地之间，物各有主，苟非吾之所有，虽一毫而莫取。惟江之清风，与山间之明月，耳得之而为声，目遇之而成色，取之无禁，用之不竭。是造物者之无尽藏也，而吾与子之所共适。”

— 苏轼《赤壁赋》

“Have you considered the water and the moon? Water flows away but is never lost; the moon waxes and wanes, but neither increases nor diminishes. If you look at its changing aspect, the universe passes in the twinkling of an eye; but if you look at its changeless aspect, all creatures including ourselves are imperishable. What reason have you to envy other things? Besides, everything in this universe has its owner; and if it does not belong to me not a tiny speck can I take. The sole exceptions are the cool breeze on the river, the bright moon over the hills. These serve as music to our ears, as colour to our eyes; these we can take freely and enjoy forever; these are inexhaustible treasures supplied by the Creator, and things in which we can delight together.”

— Shi Su, *Ode on the Red Cliff* (translated by Xianyi Dai and Naidie Dai)

We have constructed nonparametric factor-based and arbitrage-free market models for a stock and finitely many European call options on the stock with various expiries and monynesses, and establish an efficient inference method for estimating the models by leveraging deep learning algorithms. In particular, our models (i) permit no arbitrage, (ii) allow, in principal, exact cross-sectional calibration, and (iii) reflect stylised facts observed from market price dynamics. Importantly, it is practically convenient to estimate these models, given that observations are discrete time series of prices for a large but finite collection of options.

We choose the market model approach over the martingale approach mainly because

martingale models usually impose restrictive dynamics for options and may give complex expressions for option prices (which make calibration computationally infeasible). However, unlike martingale models, market models are not arbitrage-free by construction, and ensuring no-arbitrage for option market models is difficult and has always been a central topic in the relevant literature. To address this challenge, in Chapter 4, we develop a linear factor representation of option prices and construct models given by a finite system of SDEs for the factors and the stock price. Thereafter, we derive an HJM-type drift condition on the factor SDEs which guarantees freedom from dynamic arbitrage, and the state space of the market factor processes where the models are free from static arbitrage. The static arbitrage constraints can be constructed efficiently using the method proposed in Chapter 2. To calibrate our models, we represent the drift and diffusion functions by neural networks, referred to as *neural-SDE*. By leveraging deep learning algorithms, we train the neural networks by maximising the likelihood of observing the factor paths, subject to the derived arbitrage constraints. In particular, we need to tackle the problem of calibrating a neural-SDE with a convex polytope state space. We propose a novel hard constraint approach that modifies the network to respect sufficient conditions on the drift and diffusion to restrain the process within the polytope in Chapter 5. Thereafter, in Chapter 6 and 7, we validate the modelling using both synthetic data from a known Heston-SLV model and real world data. Finally, we explore the capacity of the model in two important applications about managing risks of option books, namely the Value-at-Risk (VaR) calculation and hedging in Chapters 8 and 9, respectively. In terms of calculating VaR, we show that our models are more computationally efficient and accurate for evaluating risks of option portfolios, with better coverage performance and less procyclicality than standard filtered historical simulation approaches. In the hedging application, we derive partial derivatives and minimum variance hedging ratios under neural-SDE market models, and compare the hedging performance with standard Greeks-based hedging under the Black–Scholes model, demonstrating favourable performance of our approach.

We believe that the work presented in this thesis has offered a new, and practically workable, way of modelling dynamics of option prices (or implied volatility surfaces). It is a mix of simple ideas: (1) market models for dynamics of the directly observable, factor-based models for objects living in low-dimensional manifolds and (2) neural network approximator for (possibly) high-dimensional and nonlinear functions. Our major contributions are then to combine these parts and complete a mathematically solid methodological framework and demonstrate its effectiveness in real world applications. There are still many interesting open problems derived from the results of the thesis, for which we highlight a few as follows.

- *Other option types.* Nowadays American options, Asian options and some first generation exotic options like barriers become more and more liquid and standard. There is a growing interest of modelling these options. While our modelling framework is not restricted to a specific option type, identifying arbitrage constraints for these options remains a challenge.
- *Optimal neural network training configurations.* The presented numerical results are not necessarily optimal. To enhance the goodness-of-fit, stability and robustness of the trained model, it may be worth tweaking the training processes by, for example, cross validation, batch normalisation, and adding regularisation terms in the loss function.
- *Other applications.* We have demonstrated the superior performance of our model in risk management. As a generic dynamic model of option prices under the real world measure, it would be interesting to see how it applies to market simulation and hedging of exotic options.

Bibliography

- [1] QuantLib, a free/open-source library for quantitative finance. <https://www.quantlib.org/>. Accessed: 2021-04-13.
- [2] M. Abadi, A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, A. Davis, J. Dean, M. Devin, S. Ghemawat, I. Goodfellow, A. Harp, G. Irving, M. Isard, Y. Jia, R. Jozefowicz, L. Kaiser, M. Kudlur, J. Levenberg, D. Mané, R. Monga, S. Moore, D. Murray, C. Olah, M. Schuster, J. Shlens, B. Steiner, I. Sutskever, K. Talwar, P. Tucker, V. Vanhoucke, V. Vasudevan, F. Viégas, O. Vinyals, P. Warden, M. Wattenberg, M. Wicke, Y. Yu, and X. Zheng. TensorFlow: Large-scale machine learning on heterogeneous systems, 2015. Software available from tensorflow.org.
- [3] Y. Ait-Sahalia and J. Duarte. Nonparametric option pricing under shape restrictions. *Journal of Econometrics*, 116(1-2):9–47, 2003.
- [4] C. Alexander and A. Kaeck. Does model fit matter for hedging? Evidence from FTSE 100 options. *Journal of Futures Markets*, 32(7):609–638, 2012.
- [5] C. Alexander and L. M. Nogueira. Model-free hedge ratios and scale-invariant models. *Journal of Banking & Finance*, 31(6):1839–1861, 2007.
- [6] P. C. Andreou, C. Charalambous, and S. H. Martzoukos. Generalized parameter functions for option pricing. *Journal of Banking & Finance*, 34(3):633–646, 2010.
- [7] A. N. Avramidis and J. R. Wilson. Correlation-induction techniques for estimating quantiles in simulation experiments. *Operations Research*, 46(4):574–591, 1998.
- [8] K. A. Babbar. *Aspects of Stochastic Implied Volatility in Financial Markets*. PhD thesis, Imperial College London, 2001.
- [9] A. Bain, M. Mariapragassam, and C. Reisinger. Calibration of local-stochastic and

-
- path-dependent volatility models to vanilla and no-touch options. *The Journal of Computational Finance*, 24(4), 2021.
- [10] G. Bakshi, C. Cao, and Z. Chen. Empirical performance of alternative option pricing models. *The Journal of Finance*, 52(5):2003–2049, 1997.
- [11] G. Barone-Adesi and K. Giannopoulos. Non-parametric VaR techniques – myths and realities. *Economic Notes*, 30(2):167–181, 2001.
- [12] G. Barone-Adesi, K. Giannopoulos, and L. Vosper. VaR without correlations for nonlinear portfolios. *Journal of Futures Markets*, 1997.
- [13] G. Barone-Adesi, K. Giannopoulos, and L. Vosper. Backtesting derivative portfolios with filtered historical simulation (FHS). *European Financial Management*, 8(1):31–58, 2002.
- [14] A. R. Barron. Approximation and estimation bounds for artificial neural networks. *Machine Learning*, 14(1):115–133, 1994.
- [15] E. Barucci, U. Cherubini, and L. Landi. Neural networks for contingent claim pricing via the Galerkin method. In *Computational Approaches to Economic Problems*, pages 127–142, 1997.
- [16] Basel Committee on Banking Supervision. Supervisory framework for the use of “backtesting” in conjunction with the internal models approach to market risk capital requirements. 1996. <https://www.bis.org/publ/bcbs22.htm>, Accessed: 2021-11-04.
- [17] D. S. Bates. Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics*, 94(1-2):181–238, 2000.
- [18] C. Bayer, B. Horvath, A. Muguruza, B. Stemper, and M. Tomas. On deep calibration of (rough) stochastic volatility models, 2019. arXiv:1908.08806.
- [19] C. Beck, S. Becker, P. Cheridito, A. Jentzen, and A. D. Neufeld. Deep splitting method for parabolic PDEs. *SIAM Journal of Scientific Computing*, 43:A3135–A3154, 2021.
- [20] S. M. Bellovin, P. K. Dutta, and N. Reiter. Privacy and synthetic datasets. *Stanford Technology Law Review*, 22(1), 2019.
- [21] F. Black. Studies of stock market volatility changes. *1976 Proceedings of the American statistical association business and economic statistics section*, pages 177–181, 1976.
- [22] F. Black and M. Scholes. The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3):637–54, 1973.

-
- [23] P. M. Blacque-Florentin and B. Missaoui. Nonparametric and arbitrage-free construction of call surfaces using ℓ^1 -recovery, 2015. arXiv:1506.06997.
- [24] S. Boyarchenko and S. Levendorskiĭ. Static and semistatic hedging as contrarian or conformist bets. *Mathematical Finance*, 30(3):921–960, 2020.
- [25] S. Boyd and L. Vandenberghe. *Convex Optimization*. Cambridge University Press, New York, NY, USA, 2004.
- [26] D. T. Breeden and R. H. Litzenberger. Prices of state-contingent claims implicit in option prices. *The Journal of Business*, 51(4):621–51, 1978.
- [27] M. Britten-Jones and S. M. Schaefer. Non-linear value-at-risk. *Review of Finance*, 2(2):161–187, 1999.
- [28] M. Broadie, M. Chernov, and M. Johannes. Model specification and risk premia: evidence from futures options. *The Journal of Finance*, 62(3):1453–1490, 2007.
- [29] H. Buehler, L. Gonon, J. Teichmann, and B. Wood. Deep hedging. *Quantitative Finance*, 19(8):1271–1291, 2019.
- [30] H. Buehler, B. Horvath, T. Lyons, I. P. Arribas, and B. Wood. A data-driven market simulator for small data environments, 2020. arXiv:2006.14498.
- [31] J. Butler and B. Schachter. Estimating value-at-risk with a precision measure by combining kernel estimation with historical simulation. *Review of Derivatives Research*, 1:371–390, 1997.
- [32] E. Candes, M. Rudelson, T. Tao, and R. Vershynin. Error correction via linear programming. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pages 668–681, 2005.
- [33] R. Carmona. *HJM: A Unified Approach to Dynamic Models for Fixed Income, Credit and Equity Markets*, pages 1–50. Paris–Princeton Lectures on Mathematical Finance 2004. Springer, Berlin Heidelberg, 2007.
- [34] R. Carmona and S. Nadtochiy. Local volatility dynamic models. *Finance and Stochastics*, 13(1):1–48, 2009.
- [35] R. Carmona and S. Nadtochiy. Tangent Lévy market models. *Finance and Stochastics*, 16(1):63–104, 2012.
- [36] R. J. Caron, J. F. McDonald, and C. M. Ponik. A degenerate extreme point strategy for the classification of linear constraints as redundant or necessary. *Journal of Optimization Theory and Applications*, 62(2):225–237, 1989.

-
- [37] P. Carr, K. Ellis, and V. Gupta. Static hedging of exotic options. In *Quantitative Analysis In Financial Markets: Collected Papers of the New York University Mathematical Finance Seminar*, pages 152–176. World Scientific, 1999.
- [38] P. Carr, H. Geman, D. Madan, and M. Yor. Stochastic volatility for Lévy processes. *Mathematical Finance*, 13(3):345–382, 2003.
- [39] P. Carr and D. B. Madan. A note on sufficient conditions for no arbitrage. *Finance Research Letters*, 2(3):125–130, 2005.
- [40] C. Cartis, J. Fiala, B. Marteau, and L. Roberts. Improving the flexibility and robustness of model-based derivative-free optimization solvers. *ACM Transactions on Mathematical Software*, 45(3):Article 32, 2019.
- [41] CBOE. White paper: CBOE volatility index, 2019. <https://cdn.cboe.com/resources/vix/vixwhite.pdf>, Accessed: 2021-04-08.
- [42] M. Chataigner, S. Crépey, and M. Dixon. Deep local volatility. *Risks*, 8(3), 2020.
- [43] P. Christoffersen. Evaluating interval forecasts. *International Economic Review*, 39(4):841–862, 1998.
- [44] CME Group. Futures and options margin model. <https://www.cmegroup.com/clearing/risk-management/futures-and-options-margin-model.html>. Accessed: 2022-01-18.
- [45] S. N. Cohen, C. Reisinger, and S. Wang. Detecting and repairing arbitrage in traded option prices. *Applied Mathematical Finance*, 27(5):345–373, 2020.
- [46] S. N. Cohen, C. Reisinger, and S. Wang. Arbitrage-free neural-SDE market models, 2021. arXiv:2105.11053.
- [47] S. N. Cohen, C. Reisinger, and S. Wang. Estimating risks of option books using neural-SDE market models, 2022. arXiv:2202.07148.
- [48] S. N. Cohen, D. Snow, and L. Szpruch. Black-box model risk in finance, 2021. arXiv:2102.04757.
- [49] R. Cont and J. da Fonseca. Dynamics of implied volatility surfaces. *Quantitative Finance*, 2(1):45–60, 2002.
- [50] R. Cont, J. da Fonseca, and V. Durrleman. Stochastic models of implied volatility surfaces. *Economic Notes*, 31(2):361–377, 2002.
- [51] R. Cont and P. Tankov. *Financial modelling with jump processes*. Chapman and Hall/CRC, 2004.

-
- [52] L. Cousot. Conditions on option prices for absence of arbitrage and exact calibration. *Journal of Banking & Finance*, 31(11):3377–3397, 2007.
- [53] C. Cuchiero, W. Khosrawi, and J. Teichmann. A generative adversarial network approach to calibration of local stochastic volatility models. *Risks*, 8(4), 2020.
- [54] T. Daglish, J. Hull, and W. Suo. Volatility surfaces: theory, rules of thumb, and empirical evidence. *Quantitative Finance*, 7(5):507–524, 2007.
- [55] M. Davis and J. Oblój. Market completion using options. *Advances in Mathematics of Finance*, 83:49–60, 2008.
- [56] M. H. A. Davis and D. G. Hobson. The range of traded option prices. *Mathematical Finance*, 17(1):1–14, 2007.
- [57] F. Delbaen and W. Schachermayer. A general version of the fundamental theorem of asset pricing. *Mathematische Annalen*, 300(1):463–520, 1994.
- [58] E. Derman and I. Kani. Riding on a smile. *Risks*, 7, 1994.
- [59] E. Derman and I. Kani. Stochastic implied trees: Arbitrage pricing with stochastic term and strike structure of volatility. *International Journal of Theoretical and Applied Finance*, 01(01):61–110, 1998.
- [60] D. L. Donoho and M. Elad. Optimally sparse representation in general (nonorthogonal) dictionaries via ℓ^1 minimization. *Proceedings of the National Academy of Sciences*, 100(5):2197–2202, 2003.
- [61] C. Dugas, Y. Bengio, F. Bélisle, C. Dadeau, and R. Garcia. Incorporating functional knowledge in neural networks. *Journal of Machine Learning Research*, 10:1239–1262, 2009.
- [62] B. Dupire. Pricing with a smile. *Risk Magazine*, 7:18–20, 1994.
- [63] L. El-Jahel, W. Perraudin, and P. Sellin. Value at risk for derivatives. *The Journal of Derivatives*, 6(3):7–26, 1999.
- [64] R. Eldan and O. Shamir. The power of depth for feedforward neural networks. In V. Feldman, A. Rakhlin, and O. Shamir, editors, *29th Annual Conference on Learning Theory*, volume 49 of *Proceedings of Machine Learning Research*, pages 907–940, Columbia University, New York, New York, USA, 23–26 Jun 2016. PMLR.
- [65] Eurex Clearing. Eurex clearing prisma: portfolio-based risk management. https://www.eurexchange.com/resource/blob/32818/7bcf119060b658ad4e487f588744140d/data/brochure_eurex_clearing_prisma.pdf. Accessed: 2022-01-18.

-
- [66] W. Fallon. *Calculating value-at-risk*. Wharton School, University of Pennsylvania Philadelphia, 1996.
- [67] M. Fengler. Arbitrage-free smoothing of the implied volatility surface. *Quantitative Finance*, 9:417–428, 2009.
- [68] M. Fengler. Option data and modeling BSM implied volatility. *Handbook of Computational Finance*, pages 117–142, 2012.
- [69] M. R. Fengler and L.-Y. Hin. Semi-nonparametric estimation of the call-option price surface under strike and time-to-expiry no-arbitrage constraints. *Journal of Econometrics*, 184(2):242–261, 2015.
- [70] H. Föllmer and M. Schweizer. Hedging of contingent claims under incomplete information. In M. Davis and R. Elliott, editors, *Applied Stochastic Analysis*, pages 389–414. Stochastics Monographs Gordon & Breach, London, 1990.
- [71] A. Friedman and M. A. Pinsky. Asymptotic stability and spiraling properties for solutions of stochastic equations. *Transactions of the American Mathematical Society*, 186:331–358, 1973.
- [72] J. Gatheral and A. Jacquier. Arbitrage-free SVI volatility surfaces. *Quantitative Finance*, 14(1):59–71, 2014.
- [73] S. Gerhold and I. C. Gülüm. Consistency of option prices under bid–ask spreads. *Mathematical Finance*, 30(2):377–402, 2020.
- [74] P. Gierjatowicz, M. Sabate-Vidales, D. Šiška, Ł. Szpruch, and Ž. Žurič. Robust pricing and hedging via neural SDEs, 2020. arXiv:2007.04154.
- [75] P. Glasserman, P. Heidelberger, and P. Shahabuddin. Variance reduction techniques for estimating value-at-risk. *Management Science*, 46(10):1349–1364, 2000.
- [76] P. Glasserman, P. Heidelberger, and P. Shahabuddin. Portfolio value-at-risk with heavy-tailed risk factors. *Mathematical Finance*, 12(3):239–269, 2002.
- [77] R. Gribonval and M. Nielsen. Sparse representations in unions of bases. *IEEE Transactions on Information Theory*, 49(12):3320–3325, 2003.
- [78] P. Gurrola-Perez and D. Murphy. Filtered historical simulation Value-at-Risk models and their competitors. *Bank of England Working Paper*, (525), 2015.
- [79] P. S. Hagan, D. Kumar, A. S. Lesniewski, and D. E. Woodward. Managing smile risk. *Wilmott Magazine*, 1:84–108, 2002.
- [80] P. S. Hagan, D. Kumar, A. S. Lesniewski, and D. E. Woodward. Arbitrage-free SABR. *Wilmott*, 69:60–75, 2014.

-
- [81] J. M. Harrison and D. Kreps. Martingales and arbitrage in multiperiod securities markets. *Journal of Economic Theory*, 20(3):381–408, 1979.
- [82] D. Heath, R. Jarrow, and A. Morton. Bond pricing and the term structure of interest rates: A new methodology for contingent claims valuation. *Econometrica*, 60(1):77–105, 1992.
- [83] P. Henry-Labordere. Generative models for financial data. 2019. Available at SSRN 3408007.
- [84] T. C. Hesterberg and B. L. Nelson. Control variates for probability and quantile estimation. *Management Science*, 44(9):1295–1312, 1998.
- [85] S. L. Heston. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies*, 6:327–343, 1993.
- [86] K. Hornik, M. Stinchcombe, and H. White. Multilayer feedforward networks are universal approximators. *Neural Networks*, 2(5):359–366, 1989.
- [87] K. Hornik, M. Stinchcombe, and H. White. Universal approximation of an unknown mapping and its derivatives using multilayer feedforward networks. *Neural Networks*, 3(11):551–560, 1990.
- [88] B. Horvath, A. Muguruza, and M. Tomas. Deep learning volatility: a deep neural network perspective on pricing and calibration in (rough) volatility models. *Quantitative Finance*, 21(1):11–27, 2021.
- [89] J. C. Hsu and B. L. Nelson. Control variates for quantile estimation. *Management Science*, 36(7):8350–851, 1990.
- [90] P. Huber, J. Wiley, and W. InterScience. *Robust statistics*. Wiley New York, 1981.
- [91] J. Hull and A. White. Optimal delta hedging for options. *Journal of Banking & Finance*, 82:180–190, 2017.
- [92] J. M. Hutchinson, A. W. Lo, and T. Poggio. A nonparametric approach to pricing and hedging derivative securities via learning networks. *The Journal of Finance*, 49(3):851–889, 1994.
- [93] M. Hutzenthaler, A. Jentzen, and P. E. Kloeden. Strong convergence of an explicit numerical method for SDEs with nonglobally Lipschitz continuous coefficients. *The Annals of Applied Probability*, 22(4):1611–1641, 2012.
- [94] A. Itkin. Deep learning calibration of option pricing models: some pitfalls and solutions, 2019. arXiv:1906.03507.

-
- [95] A. Ivanovas. *Option data, missing tails, and the intraday variation of implied moments*. PhD thesis, University of St. Gallen, 2015.
- [96] P. J. Schonbucher. A market model for stochastic implied volatility. *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 357:2071–2092, 1999.
- [97] J. Jacod and P. Protter. Risk-neutral compatibility with option prices. *Finance and Stochastics*, 14(2):285–315, 2010.
- [98] M. Jex, R. Henderson, and D. Wang. Pricing exotics under the smile. *Risk*, pages 72–75, November 1999.
- [99] N. Kahale. An arbitrage-free interpolation of volatilities. *Risk Magazine*, 17:102–106, 2004.
- [100] J. Kallsen and P. Krühner. On a Heath–Jarrow–Morton approach for stock options. *Finance and Stochastics*, 19:583–615, 2013.
- [101] I. Karatzas and C. Kardaras. The numéraire portfolio in semimartingale financial models. *Finance and Stochastics*, 11:447–493, 2007.
- [102] H. G. Kellerer. Markov-Komposition und eine Anwendung auf Martingale. *Mathematische Annalen*, 198:99–122, 1972.
- [103] K.-K. Kim and D.-Y. Lim. Static replication of barrier-type options via integral equations. *Quantitative Finance*, 21(2):281–294, 2021.
- [104] J. L. Kirkby and S. Deng. Static hedging and pricing of exotic options with payoff frames. *Mathematical Finance*, 29(2):612–658, 2019.
- [105] P. N. Kolm and G. Ritter. Dynamic replication and hedging: A reinforcement learning approach. *The Journal of Financial Data Science*, 1(1):159–171, 2019.
- [106] A. Koshiyama, N. Firoozye, and P. Treleaven. Generative adversarial networks for financial trading strategies fine-tuning and combination. *Quantitative Finance*, 21(5):797–813, 2021.
- [107] D. Kreps. Arbitrage and equilibrium in economies with infinitely many commodities. *Journal of Mathematical Economics*, 8(1):15–35, 1981.
- [108] P. H. Kupiec. Techniques for verifying the accuracy of risk measurement models. *Journal of Derivatives*, 3(2):73–84, 1995.
- [109] LCH. Risk management - LTD: margin models & governance. <https://www.lch.com/risk-management/risk-management-ltd>. Accessed: 2022-01-18.

-
- [110] F. Le Floch and C. W. Oosterlee. Model-free stochastic collocation for an arbitrage-free implied volatility, part II. *Risks*, 7(1), 2019.
- [111] H. Lim. Improved methods for implied volatility surface and implied distributions. *SSRN preprint 3561100*, 2020.
- [112] J. Longestaey and M. Spencer. RiskMetrics™—technical document. *Morgan Guaranty Trust Company of New York: New York*, 51:54, 1996.
- [113] T. J. Lyons. Derivates as tradable assets. RISK 10th Anniversary Global Summit, 1997.
- [114] L. V. Martin S. Andersen, J Dahl. CVXOPT: A python package for convex optimization. [Version 1.2.5; available at cvxopt.org].
- [115] MathWorks. optByHestonNI. <https://uk.mathworks.com/help/fininst/optbyhestonni.html>. Accessed: 2022-01-25.
- [116] P. Meier. *Essays on pricing kernel estimation, option data filtering and risk-neutral density tail estimation*. PhD thesis, University of St. Gallen, 2015.
- [117] R. C. Merton. Option pricing when underlying stock returns are discontinuous. *Journal of financial economics*, 3(1-2):125–144, 1976.
- [118] H. N. Mhaskar and T. Poggio. Deep vs. shallow networks: An approximation theory perspective. *Analysis and Applications*, 14(06):829–848, 2016.
- [119] T. Motzkin, H. Raiffa, G. L. Thompson, and R. M. Thrall. The double description method. In *Contributions to the Theory of Games II*. Princeton University Press, 1953.
- [120] D. Murphy, M. Vasios, and N. Vause. An investigation into the procyclicality of risk-based initial margin models. *Bank of England Financial Stability Paper*, (29), 2014.
- [121] B. K. Natarajan. Sparse approximate solutions to linear systems. *SIAM Journal on Computing*, 24(2):227–234, 1995.
- [122] H. Ni, L. Szpruch, M. Wiese, S. Liao, and B. Xiao. Conditional sig-wasserstein gans for time series generation, 2020. arXiv:2006.05421.
- [123] OptionMetrics. IvyDB Europe Reference Manual, May 2020. Version 2.4.
- [124] V. Piterbarg. Markovian projection method for volatility calibration. SSRN preprint 906473, 2006.
- [125] R. Poulsen, K. R. Schenk-Hoppé, and C.-O. Ewald. Risk minimization in stochastic

-
- volatility models: model risk and empirical performance. *Quantitative Finance*, 9(6):693–704, 2009.
- [126] R. Rebonato. *Volatility and correlation: the perfect hedger and the fox*. John Wiley & Sons, 2005.
- [127] D. Reiswich and U. Wystup. FX volatility smile construction. CPQF Working Paper Series 20, Frankfurt School of Finance and Management, Centre for Practical Quantitative Finance (CPQF), 2009. Available at <https://EconPapers.repec.org/RePEc:zbw:cpqfwp:20>; accessed 2021-05-20.
- [128] G. Ritter and P. N. Kolm. Dynamic replication and hedging: A reinforcement learning approach. *Journal of Financial Data Science*, 1(1):159–171, 2019.
- [129] J. Ruf and W. Wang. Neural networks for option pricing and hedging: a literature review. *Journal of Computational Finance*, 24(1):1–46, 2020.
- [130] J. Ruf and W. Wang. Hedging with linear regressions and neural networks. *Journal of Business & Economic Statistics*, pages 1–13, 2021.
- [131] M. Sabate-Vidales, D. Siska, and L. Szpruch. Unbiased deep solvers for parametric PDEs, 2021. arXiv:1810.05094.
- [132] I. Safran and O. Shamir. Depth-width tradeoffs in approximating natural functions with neural networks. In D. Precup and Y. W. Teh, editors, *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 2979–2987. PMLR, 06–11 Aug 2017.
- [133] D. C. Schwarz. Market completion with derivative securities. *Finance and Stochastics*, 21(1):263–284, 2017.
- [134] M. Schweizer. Option hedging for semimartingales. *Stochastic processes and their Applications*, 37(2):339–363, 1991.
- [135] M. Schweizer and J. Wissel. Arbitrage-free market models for option prices: the multi-strike case. *Finance and Stochastics*, 12(4):469–505, 2008.
- [136] M. Schweizer and J. Wissel. Term structures of implied volatilities: Absence of arbitrage and existence results. *Mathematical Finance*, 18(1):77–114, 2008.
- [137] M. Shaked and J. Shanthikumar. *Stochastic Orders*. Springer Series in Statistics. Springer New York, 2007.
- [138] S. Stoikov. The micro-price: a high-frequency estimator of future prices. *Quantitative Finance*, 18:1–8, 2018.

-
- [139] L. Sun and L. J. Hong. Asymptotic representations for importance-sampling estimators of value-at-risk and conditional value-at-risk. *Operations Research Letters*, 38(4):246–251, 2010.
- [140] L. Szpruch and X. Zhang. V-integrability, asymptotic stability and comparison theorem of explicit numerical schemes for SDEs. *Mathematics of Computations*, 87:755–783, 2018.
- [141] S. Takahashi, Y. Chen, and K. Tanaka-Ishii. Modeling financial time-series with generative adversarial networks. *Physica A*, 527(121261):797–813, 2019.
- [142] P. Tankov. Pricing and hedging in exponential lévy models: review of recent results. *Paris-Princeton Lectures on Mathematical Finance 2010*, pages 319–359, 2011.
- [143] A. Ventzel'. On equations of the theory of conditional Markov processes. *Theory of Probability and its Applications*, 10(2):357–360, 1965.
- [144] Y. Wang, H. Yin, and L. Qi. No-arbitrage interpolation of the option price function and its reformulation. *Journal of Optimization Theory and Applications*, 120(3):627–649, 2004.
- [145] M. Wiese, L. Bai, B. Wood, and H. Buehler. Deep hedging: Learning to simulate equity option markets, 2019. arXiv:1911.01700.
- [146] M. Wiese, B. Wood, A. Pachoud, R. Korn, H. Buehler, P. Murray, and L. Bai. Multi-asset spot and option market simulation, 2021. arXiv:2112.06823.
- [147] J. S. Wissel. *Arbitrage-free market models for liquid options*. PhD thesis, ETH Zürich, 2008.
- [148] U. Wystup. *FX options and structured products*. John Wiley & Sons, Ltd, 2 edition, 2017.
- [149] D. Yarotsky. Error bounds for approximations with deep ReLU networks. *Neural Networks*, 94:103–114, 2017.