

Arbitrage-Free Neural-SDE Market Models

Samuel N. Cohen, Christoph Reisinger & Sheng Wang

To cite this article: Samuel N. Cohen, Christoph Reisinger & Sheng Wang (2023) Arbitrage-Free Neural-SDE Market Models, *Applied Mathematical Finance*, 30:1, 1-46, DOI: [10.1080/1350486X.2023.2257217](https://doi.org/10.1080/1350486X.2023.2257217)

To link to this article: <https://doi.org/10.1080/1350486X.2023.2257217>



© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.



Published online: 28 Sep 2023.



Submit your article to this journal [↗](#)



Article views: 793



View related articles [↗](#)



View Crossmark data [↗](#)



Citing articles: 1 View citing articles [↗](#)

Arbitrage-Free Neural-SDE Market Models

Samuel N. Cohen, Christoph Reisinger and Sheng Wang

Mathematical Institute, University of Oxford, Oxford, UK

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 paper develops a nonparametric model for the European options book respecting underlying financial constraints and 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 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 data generated from a Heston stochastic local volatility model.

ARTICLE HISTORY

Received 30 June 2023
Accepted 28 August 2023

KEYWORDS



Market models; no-arbitrage;
European options; neural
networks; neural SDE;
constrained diffusions

1. Introduction

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 $T \in \mathcal{T}$ 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 stylized 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
- (iii) reflect stylized facts observed from market price dynamics.

CONTACT Samuel N. Cohen  samuel.cohen@maths.ox.ac.uk  Mathematical Institute, University of Oxford, Oxford, OX2 6GG, UK

© 2023 The Author(s). Published by Informa UK Limited, trading as Taylor & Francis Group.
This is an Open Access article distributed under the terms of the Creative Commons Attribution License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided the original work is properly cited. The terms on which this article has been published allow the posting of the Accepted Manuscript in a repository by the author(s) or with their consent.

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-nets where the resulting behaviour (when used in an SDE) will satisfy a family of linear (in)equalities. In this paper we will develop methods to solve this challenge, and demonstrate their effectiveness when building a financial model. The goal of this paper is to provide the core details of the approach and its performance in an idealized (simulated) setting, where it can be compared with a known ground-truth. The application of this methodology to risk estimation can be found in Cohen, Reisinger, and Wang (2022a), and to hedging in Cohen, Reisinger, and Wang (2022b).

1.1. 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 and Itô calculus to derive option prices, written as discounted conditional expectation of options' payoffs under \mathbb{Q} . For example, $C_t(T, K) = D_t(T) \cdot \mathbb{E}^{\mathbb{Q}}[(S_T - K)^+ | \mathcal{F}_t]$, where $D_t(T)$ is the time t price of the zero-coupon bond expiring at time T . Prestigious examples of martingale models are the Black–Scholes (BS) (Black and Scholes 1973) model, the local volatility (LV) model of Dupire (1994) and Derman and Kani (1994), stochastic volatility (SV) models such as Heston (1993) and Hagan et al. (2002, 2014), and stochastic local volatility (SLV) models (see, for example, Jex, Henderson, and Wang (1999)). A data-driven approach to learn an over-parametrized martingale model with SDE coefficients expressed as neural-nets is taken by Gierjatowicz et al. (2020).

The martingale approach guarantees no-arbitrage by the First Fundamental Theorem of Asset Pricing (FFTAP) (Harrison and Kreps 1979). Therefore, the martingale approach eliminates arbitrage by construction, and the LV model and SV models can calibrate exactly 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;
- (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. This idea originates from the framework of Heath, Jarrow, and Morton (1992) 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. In other words, the state space of these processes is heavily constrained.

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 (2005) and Davis and Hobson (2007)).

Conveniently, the model-free arbitrage restrictions correspond to *static* arbitrage constraints, that is, to restrictions on the state-space of the price process. In previous work (Cohen, Reisinger, and Wang 2020) (drawing on (Cousot 2007)), we have reduced these constraints, for a collection of arbitrary strikes and maturities, to an efficient set of linear inequalities.

In this paper, 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.

1.2. Factor-Based Arbitrage-Free Market Models

To allow practical calibration, we will focus on constructing *market models* for *finitely* many options that may have an *arbitrarily-shaped* lattice.

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 1.² This causes problems for methods based on a rectangular lattice. One might augment market data using arbitrage-free interpolation techniques (see e.g., Fengler 2009; Kahale 2004), 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 stylized liquidity profile of the options market.

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. Unlike the neat static arbitrage constraints for options on a rectangular lattice, e.g., see Carr and Madan (2005), 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 (2007) and Cousot (2007), our previous work (Cohen, Reisinger, and Wang 2020) produces an efficient construction of static arbitrage constraints for large-scale practical problems. Second, we aim both to guarantee the statistical accuracy of our model (by training with historical data) and to minimize 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




Figure 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.

from which our training data are sampled). This involves compatibility restrictions on the model coefficients, analogous to the Heath, Jarrow, and Morton (1992) drift restrictions.

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 and Fonseca 2002; Cont, Fonseca, and Durrleman 2002) for the use of Karhunen–Loève decomposition on implied volatility (IV) surfaces to extract low-dimensional factor models. Rather than decomposing IVs, we decode factors directly from option prices, preserving the linearity of the static arbitrage constraints in the factor space. This results in a convex polytope state space for the factors.

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. To summarize, from the no-arbitrage conditions, we derive a class of low-dimensional market models with drift restrictions and a convex polytope 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.

1.2.1. Comparison with Codebook Models

Most work on arbitrage-free market models for option prices has sought a convenient parametrization, or *codebook* in the language of Carmona (2007), 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}). Schonbucher (1999) proposes a market model in terms of the BS *implied volatility* σ_{imp} for a single option, i.e., the case $\mathcal{K} = \{K\}$ and $\mathcal{T} = \{T\}$. The model is free from static arbitrage if $S > 0$ and $\sigma_{\text{imp}} \geq 0$. However, as discussed by Schweizer and Wissel (2008a), 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 Gatheral and Jacquier (2014)). For the term structure case, $\mathcal{K} = \{K\}$ and $\mathcal{T} = [0, +\infty)$, Schonbucher (1999) and Schweizer and Wissel (2008b) 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 $T \in \mathcal{T}$. Jacod and Protter (2010) 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 $\mathcal{T} = \{T\}$, Schweizer and Wissel (2008a) introduce a new parametrization of option prices called *local implied volatilities* $X(K)$ and *price level* Y , see Definition 4.1 and 4.4 of Schweizer and Wissel (2008a). 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 $\mathcal{T} = [0, +\infty)$ has been considered by Derman and Kani (1998), Carmona and Nadtochiy (2009, 2012) and Kallsen and Krühner (2013). 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 parametrizations make the static arbitrage constraints naturally hold.

The non-trivial construction of codebook processes, i.e., a PDE for local volatility (Carmona and Nadtochiy 2009) and Fourier transforms for Lévy process (Carmona and Nadtochiy 2012; Kallsen and Krühner 2013), 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 (2008). Wissel combines ideas from the LV model of Dupire with the market model of Schweizer and Wissel (2008a) and uses *local implied volatilities* and *price level* as a codebook for parametrizing statically arbitrage-free option prices on a *rectangular* lattice.³ 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. Model Inference and Constrained Neural SDE

For our models, inference consists of two independent steps: factor decoding and SDE model calibration.

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 (2020)

and Gierjatowicz et al. (2020). By leveraging deep learning algorithms, we train the neural networks by maximizing 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.

No-arbitrage conditions are embedded as part of model inference. Specifically, static arbitrage constraints are characterized by a convex polytope 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 (1973). Consequently, the neural network that is used to parametrize the drift and diffusion functions needs to be constrained.⁴ We achieve this by developing appropriate transformations for the output of the neural network.

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 stylized 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.

1.4. Key Contributions

We summarize the key contributions of this paper as follows.

- (i) We construct a family of factor-based market models (Section 2), 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 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 Section 3.

2. Arbitrage-Free Market Models

We assume a frictionless market in continuous-time, with a distant fixed finite horizon T^* . We allow non-zero interest and dividends,⁵ 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^* .

2.1. 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.

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 parametrization is given by the *forward log-moneyness* m_t , which is

$$m_t = M(K; F_t(T)) := \ln \left(\frac{K}{F_t(T)} \right). \quad (1)$$

The function $M(\cdot)$ is referred to as the moneyness function. In Figure 2, 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 grey dotted area) and was never traded after its time-to-expiry became shorter than 0.4 years.

We assume the range of liquid options in the (τ, m) -coordinates is fixed over time, and develop arbitrage-free models on 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 3.

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 $\mathcal{L}_{\text{liq}} = \{(\tau_i, m_j^i) \in \mathcal{R}_{\text{liq}} : \underline{m}(\tau_i) = m_1^i < \dots < m_{n_i}^i = \bar{m}(\tau_i), \underline{\tau} = \tau_1 < \dots < \tau_p = \bar{\tau}, 1 \leq j \leq n_i, 1 \leq i \leq p\}$. 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,




Figure 2. 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.




Figure 3. An example of liquid range and lattice.

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 $[1\text{-day}, 1\text{-week}, 1\text{-month}, \dots] \times [10\text{-delta}, 15\text{-delta}, \dots]$ and smoothly interpolate them for pricing a specific contract.

2.2. 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 $\bar{W} = [W_0 \ W_1 \ \dots \ W_d]^\top$, 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 $\mathcal{F}_t^0 := \Sigma(\{\bar{W}_s : 0 \leq s \leq t\}) \vee \mathcal{F}_0$ (here we use $\Sigma(\cdot)$ to denote the sigma-algebra generator). Taking Ω to be the space of paths of \bar{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:

$$\begin{cases} dS_t = (\alpha(\tilde{\xi}_t) - q_t) S_t dt + \gamma(\tilde{\xi}_t) S_t dW_{0,t}, & S_0 = s_0 \in \mathbb{R}; \\ d\tilde{\xi}_t = \mu(\tilde{\xi}_t) dt + \sigma(\tilde{\xi}_t) dW_t, & \tilde{\xi}_0 = \xi_0 \in \mathbb{R}^d, \end{cases} \quad (2)$$

where $\tilde{\xi}_t = (S_t, \xi_t)$. We use $\alpha_t, \gamma_t, \mu_t$ and σ_t to denote $\alpha(\tilde{\xi}_t), \gamma(\tilde{\xi}_t), \mu(\tilde{\xi}_t)$ and $\sigma(\tilde{\xi}_t)$, respectively. 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. Denote by $C_t(T, K)$ the market price at time t of a 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 parametrized 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 $t \in [0, T^*]$:

$$\hat{c}_t(T, K) = \frac{C_t(T, K)}{D_t(T)F_t(T)}, \quad \tilde{c}_t(T - t, M(K; F_t(T))) = \hat{c}_t(T, K), \quad (3)$$

where \hat{c}_t and \tilde{c}_t are the *normalized* prices but with different arguments. The parametric surface we shall model is the normalized 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 normalized option price $\tilde{c}_t(\tau, m)$ is determined by the latent factor ξ through a time-independent linear transformation

$$\tilde{c}_t(\tau, m) = g(t, \tau, m, \xi_t) := G_0(\tau, m) + \sum_{i=1}^d G_i(\tau, m) \xi_{it}, \quad (4)$$

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 normalized call price $\tilde{c}(\tau, m)$ and the underlying price S , it does relate the two variables through the moneyness m .

Combining (2), (3) and (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 (2), we argue that a linear transformation in (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 2.1: 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). As this involves all liquid option prices, it determines the law of S under \mathbb{Q} in the future (through the Breeden–Litzenberger formula (Breeden and Litzenberger 1978)); 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 (2008), Schwarz (2017) and Wissel (2008).

2.2.1. Convergence to Terminal Payoff

To ensure prices of call options converge to correct payoffs at expiry, independently of the behaviour of ξ , by examining (4) it is clearly sufficient (and generally necessary) that

$$G_0(0, m) = (1 - e^m)^+, \text{ and } G_i(0, m) = 0, \quad \forall i = 1, \dots, d. \quad (5)$$

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

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

2.2.2. 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 $d_{1,2}(\tau, m) = -\frac{m}{\gamma\sqrt{\tau}} \pm \frac{1}{2}\gamma\sqrt{\tau}$,

$$\tilde{c}_t(\tau, m) = G(\tau, m) := \Phi(d_1(\tau, m)) - e^m \Phi(d_2(\tau, m)). \quad (6)$$

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 (3) indicates that

$$\begin{aligned}\tilde{c}_t(\tau, m) &= \frac{C_t(t + \tau, e^m F_t(t + \tau))}{D_t(t + \tau) F_t(t + \tau)} \\ &= \Phi\left(-\frac{m}{\sigma_t^{\text{imp}} \sqrt{\tau}} + \frac{1}{2} \sigma_t^{\text{imp}} \sqrt{\tau}\right) - e^m \Phi\left(-\frac{m}{\sigma_t^{\text{imp}} \sqrt{\tau}} - \frac{1}{2} \sigma_t^{\text{imp}} \sqrt{\tau}\right),\end{aligned}\quad (7)$$

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 normalized 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 2007), where the process for implied volatility depends on T, K, F and t only through its dependence on τ and m .

2.3. 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). After the landmark work of Harrison and Kreps (1979), there are various versions of the FFTAP and extensions of the no-arbitrage concept (e.g., no free lunch (Kreps 1981), no free lunch with vanishing risk (Delbaen and Schachermayer 1994), no unbounded profit with bounded risk (Karatzas and Kardaras 2007)). In this article, we work with a simplified version of FFTAP as follows. 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 1, by analysing 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

$$\begin{cases} \mathbf{G}^\top \sigma_t \psi_t = \mathbf{G}^\top \mu_t - z_t, & \text{where} \\ z_t = \left(-\frac{\partial}{\partial \tau} - \frac{1}{2} \gamma_t^2 \frac{\partial}{\partial m} + \frac{1}{2} \gamma_t^2 \frac{\partial^2}{\partial m^2}\right) \begin{bmatrix} \tilde{c}_t(\tau_1, m_1) \\ \vdots \\ \tilde{c}_t(\tau_N, m_N) \end{bmatrix}. \end{cases}\quad (8)$$

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 (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 (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 (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 2.2: 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.

2.4. 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.

2.4.1. 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 $\check{c}(T, k) \in C^{1,2}(\mathbb{R}_{\geq 0}^2)$.

Carr et al. (2003) and Carr and Madan (2005) establish the equivalence of no static arbitrage and the existence of a Markov Martingale Measure (MMM) \mathbb{Q} , under which

$$\check{c}(T, k) = \mathbb{E}^{\mathbb{Q}}[(M_T - k)^+ | \mathcal{F}_0]. \quad (9)$$

For an arbitrary but fixed T , using (Breedon and Litzenberger 1978) analysis, the marginal measure $\mathbb{Q}_T = \mathbb{Q}(\cdot | \mathcal{F}_T)$ exists if $-1 \leq \partial \check{c} / \partial k \leq 0$ and $\partial^2 \check{c} / \partial k^2 \geq 0$. Furthermore, if a family of marginal measures $\{\mathbb{Q}_T\}_{T \in [0, T^*)}$ on $(\mathbb{R}, \mathcal{B}(\mathbb{R}))$ exist, which are non-decreasing in convex order⁶ (NCDO), then there exists a Markov martingale measure with these marginals, by Kellerer's theorem (Kellerer 1972). The convex order can be characterized in terms of call price functions:

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

Hence, $\{\mathbb{Q}_T\}_{T \in [0, T^*)}$ is NDCO if $\partial \check{c} / \partial T \geq 0$.

Therefore, 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

$$\mathcal{S}(D) = \left\{ s(x, y) \in C^{1,2}(D) : 0 \leq s \leq 1, \frac{\partial s}{\partial x} \geq 0, -1 \leq \frac{\partial s}{\partial y} \leq 0, \frac{\partial^2 s}{\partial y^2} \geq 0 \right\}, \quad (10)$$

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.,

$$\exists s \in \mathcal{S}(\check{\mathcal{R}}_{\text{liq}}), \text{ s.t. } \forall (\tau, k) \in \check{\mathcal{L}}_{\text{liq}}, \quad s(\tau, k) = \check{c}(\tau, k). \quad (11)$$

- (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 interpolating function then belongs to $\mathcal{S}(\check{\mathcal{R}}_{\text{liq}})$. For example, Fengler and Hin (2015) 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.

2.4.2. 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 $c_j = \check{c}(\tau_j, k_j) = \check{c}(\tau_j, m_j = \ln k_j)$.

Cousot (2007) finds practically verifiable conditions for (11), which we have simplified in Cohen, Reisinger, and Wang (2020). We can write these constraints in the form $\mathbf{A}\mathbf{c} \geq \hat{\mathbf{b}}$, where $\mathbf{c} = [c_1 \ \dots \ c_N]^\top \in \mathbb{R}^N$ is the vector of normalized call prices, $\mathbf{A} = (A_{ij}) \in \mathbb{R}^{R \times N}$ and $\hat{\mathbf{b}} = (\hat{b}_j) \in \mathbb{R}^R$ are a known constant matrix and vector. Here R is the number of static arbitrage constraints. The construction of \mathbf{A} and $\hat{\mathbf{b}}$ depends on the choice of \mathcal{L}_{liq} , and is given in Cohen, Reisinger, and Wang (2020). We assume for simplicity that \mathcal{L}_{liq} is fixed over time.

Using our factor representation (4) of option prices, we have $\mathbf{c}_t = \mathbf{G}_0 + \mathbf{G}^\top \xi_t$, with \mathbf{G} as in (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

$$\mathbf{A}\mathbf{G}^\top \xi_t \geq \mathbf{b} := \hat{\mathbf{b}} - \mathbf{A}\mathbf{G}_0^\top, \text{ for all } t \in [0, T^* - \bar{\tau}]. \quad (12)$$

The dynamics of the diffusion process ξ are characterized 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 3.1.

2.5. Model Inference

Given an observable normalized 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 normalized 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 $d \ll N \ll L$.

2.5.1. 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

$$\mathbf{C} = \mathbf{1}_{L+1} \otimes \mathbf{G}_0 + \mathbf{\Xi} \mathbf{G} + \mathbf{\Upsilon}, \quad (13)$$

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 (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 (12) 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 2.3: 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 (8), denoted as $\mathbf{Z} \in \mathbb{R}^{(L+1) \times N}$; constraint matrices \mathbf{A}, \mathbf{b} from (12).

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 $\forall l = 1, \dots, L+1$,

$$\mathbf{G}_i = i\text{-th principal component of } \mathbf{Z}, \quad \Xi_{li} = \langle (\mathbf{R}_0)_l, \mathbf{G}_i \rangle,$$
 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 $\Xi_k = [\Xi_{1k} \cdots \Xi_{L+1,k}]$;
- 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 $\forall l = 1, \dots, L+1$,

$$\mathbf{G}_i = (i - d^{\text{da}})\text{-th principal component of } \mathbf{R}_{d^{\text{da}}}, \quad \Xi_{li} = \langle (\mathbf{R}_{d^{\text{da}}})_l, \mathbf{G}_i \rangle,$$
 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 $\Xi_k = [\Xi_{1k} \cdots \Xi_{L+1,k}]$;
- 8 Compute the covariance of the residual data $\Sigma_i = \mathbf{R}_i^\top \mathbf{R}_i$;
- 9 Eigen-decompose $\Sigma_i = \mathbf{Q}_i \mathbf{\Lambda}_i \mathbf{Q}_i^{-1}$ (where \mathbf{Q}_i is a unitary matrix);
- 10 For some $\mathbf{w} \in \mathbb{R}^d$ such that $\|\mathbf{w}\|_2 = 1$, define

$$\mathbf{q}_i(\mathbf{w}) = \mathbf{Q}_i \mathbf{w}, \quad \mathbf{s}_i(\mathbf{w}) = \mathbf{R}_i \mathbf{q}_i(\mathbf{w}).$$
 Then solve the following minimization problem:

$$\mathbf{w}^* = \underset{\|\mathbf{w}\|_2=1}{\operatorname{argmin}} \left(- \sum_{l=1}^{L+1} \mathbb{1} \left\{ \langle \mathbf{A} \mathbf{q}_i(\mathbf{w}), \mathbf{s}_i(\mathbf{w}) \rangle \geq \mathbf{b} - \mathbf{A} \left(\sum_{k=1}^{i-1} \Xi_{lk} \mathbf{G}_k \right) \right\} + \lambda \|\mathbf{R}_i - \mathbf{s}_i(\mathbf{w}) \otimes \mathbf{q}_i(\mathbf{w})\|_2 \right), \quad (14)$$
 where we require $\lambda < \min(1, \frac{1}{\|\mathbf{R}_i\|})$ 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 (14) 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 an example in the numerical experiments in Section 4.
- (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.

2.5.2. 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})

$$dS_t = \left(\alpha(\tilde{\xi}_t) - q_t \right) S_t dt + \gamma(\tilde{\xi}_t) S_t dW_{0,t}, \quad (15a)$$

$$d\tilde{\xi}_t = \mu(\tilde{\xi}_t) dt + \sigma(\tilde{\xi}_t) dW_t. \quad (15b)$$

To ensure reconstructed prices do not generate static arbitrage, μ, σ in (15b) should be chosen to guarantee that $\tilde{\xi}$ takes values in the (deterministic) polytope given by (12).

The SDE system (15) can be divided into two subsystems and estimated independently, as the constraints are only needed in (15b), and do not depend on S . In the next section, we will discuss an estimation method for the model of $\tilde{\xi}$ (15b), which is generic enough for estimating the model of S (15a).

3. Neural SDE Constrained by a Polytope

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

$$dY_t = \mu(Y_t) dt + \sigma(Y_t) dW_t, \quad (16)$$

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 (16), 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.

3.1. 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 (16) as

$$1 - \mathbb{P}(Y_t \in \mathcal{P} \quad \text{for all } t > 0 \mid Y_0 \in \mathcal{P}) = \mathbb{P}(\exists t > 0 \text{ s.t. } Y_t \in \mathcal{P}^c \mid Y_0 \in \mathcal{P}) = 0. \quad (17)$$

Friedman and Pinsky (1973) give a set of sufficient conditions for (17) 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 $v_i = \partial \rho_k / \partial y_i$ on $\partial \mathcal{P}_k$. Defining $a := \sigma \sigma^\top$, to ensure that \mathcal{P}_k^c is nonattainable we verify

$$\sum_{i,j=1}^d a_{ij} v_i v_j = 0 \quad \text{on } \partial \mathcal{P}_k, \quad (18a)$$

$$\sum_{i=1}^d \mu_i v_i + \frac{1}{2} \sum_{i,j=1}^d a_{ij} \frac{\partial \rho_k}{\partial y_i \partial y_j} \geq 0 \quad \text{on } \partial \mathcal{P}_k, \quad (18b)$$

assuming that ρ_k has second-order derivatives in a neighbourhood of $\partial \mathcal{P}_k$. This is formally written as Theorem 3.1.

Theorem 3.1 (Friedman and Pinsky (1973)): *Suppose μ and σ satisfy Lipschitz and linear growth conditions, and let Y_t be any solution to (16) with initial data $Y_0 \in \mathcal{P}$. If (18) holds for all k , then \mathcal{P}^c is nonattainable by Y_t , i.e., $\mathbb{P}(\exists t > 0 \text{ s.t. } Y_t \in \mathcal{P}^c \mid Y_0 \in \mathcal{P}) = 0$.*

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 (18) as, for all k ,

$$\mathbf{v}_k^\top a(y) \mathbf{v}_k = 0 \quad \text{if } \mathbf{v}_k^\top y = b_k, \quad (\text{no diffusion over the boundary}); \quad (19a)$$

$$\mathbf{v}_k^\top \mu(y) \geq 0 \quad \text{if } \mathbf{v}_k^\top y = b_k, \quad (\text{inward-pointing drift at the boundary}). \quad (19b)$$

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} .

3.2. 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 (16) is formulated as the constrained optimization problem:

$$\max_{\mu, \sigma} p(y_{0:T}; \mu, \sigma), \quad \text{subject to (19a) and (19b), for all } k. \quad (20)$$

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 (20) 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

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

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

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

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

$$\begin{aligned} \ln p(y_{0:T} | y_0; \mu, \sigma) \\ \propto -\frac{1}{2} \sum_{i=0}^{L-1} \left[\ln |a(i)| + \frac{1}{\Delta t} \|y_{t_{i+1}} - y_{t_i}\|_{a(i)}^2 + \|\mu(i)\|_{a(i)}^2 \Delta t - 2(\mu(i), y_{t_{i+1}} - y_{t_i})_{a(i)} \right]. \end{aligned} \tag{21}$$

3.3. Deep Learning a Function Approximation

Given the approximated likelihood function (21), the optimization (20) 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 (19). 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 (19) holds for any functions $\hat{\mu}(\cdot)$ and $\hat{\sigma}(\cdot)$, but impose essentially no further restrictions. We then proceed by modelling the functions

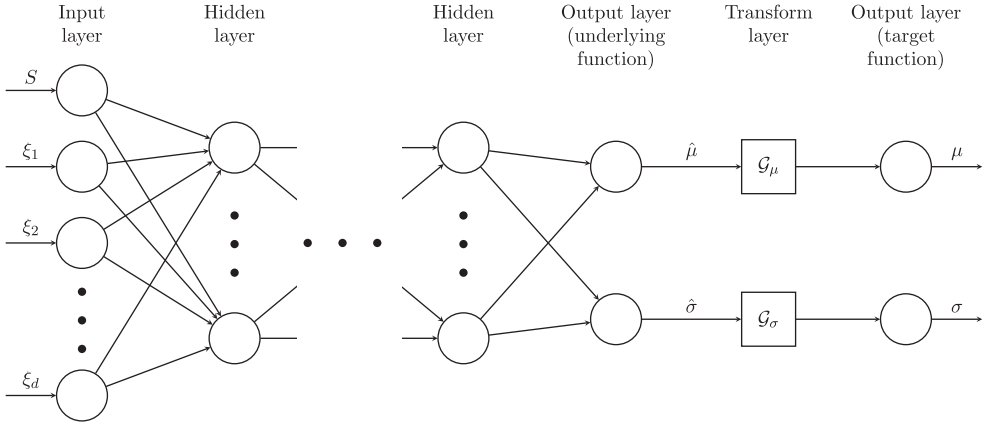


Figure 4. Constrained neural network.

$\hat{\mu}(\cdot)$ and $\hat{\sigma}(\cdot)$ using a standard neural network. In Figure 4, we show an illustration of the proposed constrained neural network model.

3.3.1. 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 $\rho_k \downarrow 0$. Accordingly, we construct an operator which is able to act when near multiple boundaries, and should be unaffected by boundaries distant from y . Specifically, we define

$$\sigma(y) = \mathcal{G}_\sigma[\hat{\sigma}](y) := (\mathbf{P}(y))^\top \hat{\sigma}(y), \quad (22)$$

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

Proposition 3.2: *Given the operator \mathcal{G}_σ defined by (22) and (24), 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,

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

Algorithm 2: Constructing diffusion shrinking transformation matrix

Input : A non-empty d -polytope $\mathcal{P} = \{x \in \mathbb{R}^d : \mathbf{V}x \geq \mathbf{b}, \mathbf{V} \in \mathbb{R}^{R \times d}, \mathbf{b} \in \mathbb{R}^R\}$; 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 normalized 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 3.1.

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

$$\mathbf{q}_k = \frac{1}{\sqrt{1 - \sum_{j=1}^{k-1} \langle \mathbf{q}_j, \mathbf{v}_{(k)} \rangle^2}} \left(\mathbf{v}_{(k)} - \sum_{j=1}^{k-1} \langle \mathbf{q}_j, \mathbf{v}_{(k)} \rangle \mathbf{q}_j \right), \quad (23)$$

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 */*

11 Define

$$\mathbf{P}(y) = \text{diag}(\sqrt{\varepsilon_{(1)}}, \dots, \sqrt{\varepsilon_{(d)}}) \times \mathbf{Q}; \quad (24)$$

/ Unless the directions $\{\mathbf{v}_{(k)}\}_k$ are orthogonal, shrinking the diffusion along \mathbf{q}_i 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 $\sqrt{\varepsilon_{(k)}}$. To mitigate this compounded shrinking issue, we propose an alternative method to compute adaptive shrinking scales in Appendix 2. */*




Figure 5. 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 $\mathbf{P}(y)$.




Figure 6. An example of ρ^* -interior points for calculating drift correction directions.

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,

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

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

3.3.2. 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 an unconstrained locally (Lipschitz) continuous function $\hat{\mu}$, corrected by some \mathcal{P} -inward-pointed d -vectors, i.e.,

$$\mu(y) = \mathcal{G}_\mu[\hat{\mu}](y) := \hat{\mu}(y) + \sum_k \lambda_k(y) (\zeta_k - y), \quad (25)$$

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 3 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 6.

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

$$\lambda_k(y) = \left(\frac{\mathbf{v}_k^\top \hat{\mu} - \varepsilon_k^\mu(y)}{\mathbf{v}_k^\top (\zeta_k - y)} \right)^+, \quad (26)$$

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 3.3: *Given the operator \mathcal{G}_μ defined by (25) and (26), 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

$$\mathbf{v}_k^\top (\zeta - y) = (\mathbf{v}_k^\top \zeta - b_k) - (\mathbf{v}_k^\top y - b_k) \geq \rho^* - \rho^* = 0. \quad (27)$$

Consequently, by the definition of λ_k in (26),

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

If $y \in \mathcal{P}^*$, 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 (27) and the definition of λ_k in (26). It follows that, when $\rho_k(y) = 0$, we have $\mathbf{v}_k^\top \mu \geq -\varepsilon_k^\mu(y) = 0$. ■

3.3.3. 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 (20) to the following unconstrained optimization problem:

$$\begin{aligned} & \min_{\hat{\mu}, \hat{\sigma}} J[\hat{\mu}, \hat{\sigma}] \\ &= \sum_{i=0}^{L-1} \left[\ln |a(i)| + \frac{1}{\Delta t} \|y_{t_{i+1}} - y_{t_i}\|_{a(i)}^2 + \|\mu(i)\|_{a(i)}^2 \Delta t - 2(\mu(i), y_{t_{i+1}} - y_{t_i})_{a(i)} \right] \\ &+ \lambda \mathcal{R}(\hat{\mu}, \hat{\sigma}), \end{aligned} \quad (28)$$

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 parametrizing $\hat{\mu}$, $\hat{\sigma}$ 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

$$\hat{\sigma} = \begin{bmatrix} \exp(\phi_1^\theta) & 0 & \cdots & 0 \\ \phi_2^\theta & \exp(\phi_3^\theta) & \cdots & 0 \\ \vdots & \vdots & \ddots & \vdots \\ \phi_{\frac{1}{2}d(d-1)+1}^\theta & \phi_{\frac{1}{2}d(d-1)+2}^\theta & \cdots & \exp\left(\phi_{\frac{1}{2}d(d+1)}^\theta\right) \end{bmatrix}, \quad \hat{\mu} = \begin{bmatrix} \phi_{\frac{1}{2}d(d+1)+1}^\theta \\ \phi_{\frac{1}{2}d(d+1)+2}^\theta \\ \vdots \\ \phi_{\frac{1}{2}d(d+3)}^\theta \end{bmatrix}. \quad (29)$$

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

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

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

4. Numerical Results – Estimating the Market Model

We validate our modelling and inference approach with numerical experiments using data generated from a stochastic local volatility model (Jex, Henderson, and Wang 1999), 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 shall explore in future work.

4.1. Input Simulation Data

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

$$\begin{aligned} dS_u &= (r_u - q_u) S_u du + \mathcal{L}_t(u, S_u) \sqrt{v_t} S_u dW_u^S, \\ dv_u &= \kappa(\theta - v_u) du + \sigma \sqrt{v_u} dW_u^v, \\ d\langle W_u^S, W_u^v \rangle &= \rho du, \quad u \in (t, T^*). \end{aligned} \quad (30)$$

This model is popular in industry for its realistic dynamic properties through the stochastic volatility component $\sqrt{v_t}$, and its ability to exactly reprice vanilla options through fitting




Figure 7. The leverage function $\mathcal{L}(t, S)$ of the Heston-SLV model used for simulation.

Table 1. Parameter values.

Parameter	Heston-SLV model								Simulation		
	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	10,000	0.0001	46

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 optimization of the fit over the Heston parameters, followed by a calibration of the leverage function using Markovian projections (see Piterbarg 2006 for a review of early approaches and the references in Bain, Mariapragassam, and Reisinger (2021) 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 (Quantlib 2021). The Heston parameters are listed in Table 1, and we plot the leverage function \mathcal{L} in Figure 7.

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 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 (30) on $[0, T^*]$,⁷ and, for each t , option values at (S_t, ν_t) are produced using the calibrated Heston parameters and the initial portion, i.e., $\mathcal{L}_t(u, \cdot) = \mathcal{L}_0(u - t, \cdot)$, $u \in (t, T^*)$.⁸

This step is only required to generate the simulated input data, but not when real market data are used. The SDE (30) 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 (30) 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.⁹




Figure 8. Partial derivatives of the normalized call price surface approximated by interpolation at $t = 0$.

4.2. Implementation Details of the Model Estimation Strategy

Approximation of z_t in (8): We need $\mathbf{Z} = (z_t)_t$, as inputs for Algorithm 1 to decode our factors. However, neither γ_t (see (2) and thereafter) nor the partial derivatives of call option prices, as seen in (8), are readily available for computing z_t .

To estimate γ_t , we fit an *initial model*¹⁰ 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. 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 functions,¹¹ from which we can approximate the required derivatives. In Figure 8, 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 .

Decorrelation and normalization 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 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 normalize the decorrelated factor data $\Xi^{(1)}$ as $\Xi^{(2)} = \Xi^{(1)}\mathbf{\Lambda}^{-1}$, where

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

Here λ_0 is a positive constant; we choose $\lambda_0 = 10$ in our numerical example so that the min-max ranges of all normalized factor data are $1/\lambda_0 = 0.1$. Consequently, the factor representation of call price data as described in (13) can be re-written as

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

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

Parametrization 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 (29), 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 (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 (8) implies¹² that $\mu_t = \Lambda^{-2} \mathbf{G} z_t$. In the following tests, we first perform a regression on the data $\{(\tilde{\xi}_{t_l}, \Lambda^{-2} \mathbf{G} z_{t_l})\}_l$ to obtain a baseline drift model $\tilde{\xi} \mapsto \bar{\mu}(\tilde{\xi})$.

4.3. Factor Construction and No-Arbitrage Boundaries

We apply Algorithm 1 to decode factors from the normalized 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:

- Mean absolute percentage error (MAPE):

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

- Proportion of dynamic arbitrage (PDA):

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

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

- Proportion of statically arbitrageable samples (PSAS):

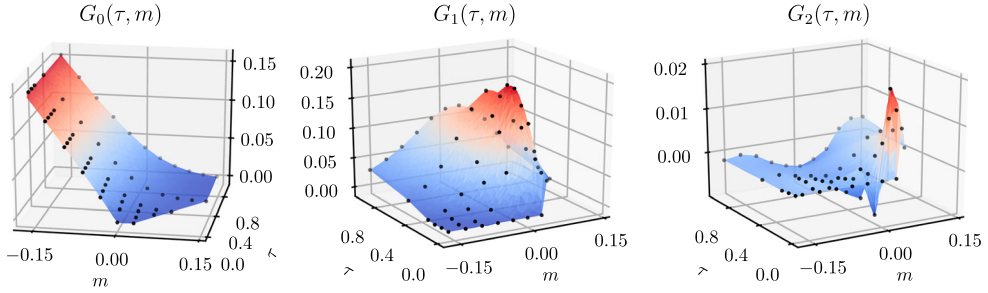

$$\text{PSAS} = 1 - \frac{\sum_{l=0}^L \mathbf{1}_{\{\mathbf{A} \mathbf{G}^\top \xi_{t_l} \geq \mathbf{b}\}}}{L+1}.$$

We show the three metrics for a few combinations of factors in Table 2. Using only two factors, i.e., one dynamic arbitrage factor and one static arbitrage factor¹³ (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.

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 9. The points in the liquid lattice (in (τ, m) coordinates)¹⁴ 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

Table 2. MAPE, PDA and PSAS metrics when including different combinations of factors.

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%

**Figure 9.** Price basis functions of the normalized call price surface.**Figure 10.** 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.

G_0 is obtained by interpolating the normalized 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 (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 (5) before interpolation.

Given these factors, we use the linear programming method (Caron, McDonald, and Ponik 1989) 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 10. 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}_{lip} .




Figure 11. Evolution of training losses and validation losses.

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 region.¹⁵

4.4. Neural Network Training Results

We apply the method in Section 3 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 architecture¹⁶ that is a composition of fully-connected layers and activation functions in the following orders:

$$\phi^\theta = \mathcal{F}_{d+1} \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{256} \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{256} \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{256}, \quad (31)$$

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 :

$$\phi^{S,\theta} = \mathcal{F}_{d+1} \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{128} \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{128} \circ \mathcal{A}_{\text{ReLU}} \circ \mathcal{F}_{128}. \quad (32)$$

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 (Abadi et al. 2015) environment.

In Figure 11, we show the evolution of training losses and validation losses¹⁷ 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.




Figure 12. Estimated drift and diffusion coefficients for S . (a) Scattergrams of $\phi_{\sigma}^{S,\mu}(\tilde{\xi}_t)$ against S , ξ_1 and ξ_2 . The estimated drift $\phi_{\sigma}^{S,\mu}(\tilde{\xi}_t)$ has values around 0, the ground-truth drift. (b) Scattergram of $\phi_{\sigma}^{S,\theta}(\tilde{\xi}_t)$ against σ_t^S .

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.

4.5. In-Sample Test

We assess how well the learnt model recovers the ground truth Heston-SLV model (30). Specifically, we study $\sigma_t^S = \mathcal{L}(t, S_t) \sqrt{v_t} S_t$ (the diffusion coefficient of S), $\mu_t = \kappa(\theta - v_t)$ (the drift coefficient of v) and $\sigma_t = \eta \sigma \sqrt{v_t}$ (the diffusion coefficient of v), 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 $\{\tilde{\xi}_t = (S_t, \xi_t)\}_{t=0,\dots,L}$, we can compute $\{\phi_{\mu}^{S,\theta}(\tilde{\xi}_t)\}_l$ and $\{\phi_{\sigma}^{S,\theta}(\tilde{\xi}_t)\}_l$.

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 12(a), 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 $\{\phi_{\sigma}^{S,\theta}(\tilde{\xi}_t)\}_l$ approximate $\{\sigma_t^S\}_l$, which is generated from the ground truth Heston-SLV model. As we see in Figure 12(b), 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 $\frac{1}{L+1} \sum_l |\phi_{\sigma}^{S,\theta}(\tilde{\xi}_t) - \sigma_t^S| / \sigma_t^S$.

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 4.6; here




Figure 13. The linear relationship between ξ_1 and ν .

we present in-sample performance to avoid any potential biases introduced by simulation methods.

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 13. Using this, an approximated ground-truth model for ξ_1 can be used to benchmark the learnt neural network model.

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 (30) to get

$$\begin{aligned} d\xi_{1t} &= \mu^{\xi_1}(\xi_{1t})dt + \sigma^{\xi_1}(\xi_{1t})dW_t^\nu, \\ \mu^{\xi_1}(x) &= \frac{\kappa(\theta - \beta_0 - \beta_1 x)}{\beta_1}, \quad \sigma^{\xi_1}(x) = \frac{\eta\sigma\sqrt{\beta_0 + \beta_1 x}}{\beta_1}. \end{aligned} \quad (33)$$

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 (33). We plot and compare these data against $\{\xi_{1t_l}\}_l$ in Figure 14. 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 behaviour in the diffusion of ν has been particularly well captured, as has the decreasing drift, despite the presence of noise.

Market price of risk: As discussed in Section 2.3, the no-arbitrage HJM-type drift restriction implies an over-determined linear system (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 (8) using the ordinary least squares method,

$$\hat{\psi}_t = \left(\sigma_t^\top \mathbf{G} \mathbf{G}^\top \sigma_t \right)^{-1} \sigma_t^\top \mathbf{G} \left(\mathbf{G}^\top \mu_t - z_t \right). \quad (34)$$

The violation to the drift restriction can be measured by

$$\chi_t = \left\| \mathbf{G}^\top \mu_t - z_t - \mathbf{G}^\top \sigma_t \hat{\psi}_t \right\|_2. \quad (35)$$




Figure 14. Comparison of the estimated drift and diffusion functions and the (approximated) ground-truth drift and diffusion functions for ξ_1 .




Figure 15. Statistics of the estimated market price of risk.

We estimate the market price of risk ψ_t from the learnt drift and diffusion coefficients using (34) for all in-sample data points. In Figure 15, we show the heatmaps of the violation of the HJM-type drift restriction χ , as defined in (35), 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 4.1) may lead to a non-vanishing market price of risk in the underlying ground truth model. We also note, comparing with Figures 10, 16 and 18, that the most common values of ξ_1 are slightly negative, coinciding with the region where the market price of risk is low.

4.6. 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 tamed¹⁸ Euler scheme (see Hutzenthaler, Jentzen, and Kloeden 2012

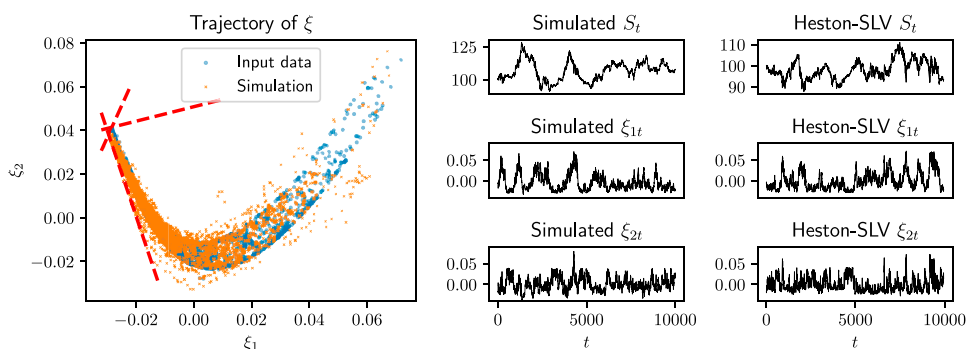


Figure 16. Simulation of S and ξ from the learnt neural network model, compared with the Heston-SLV model generated data.

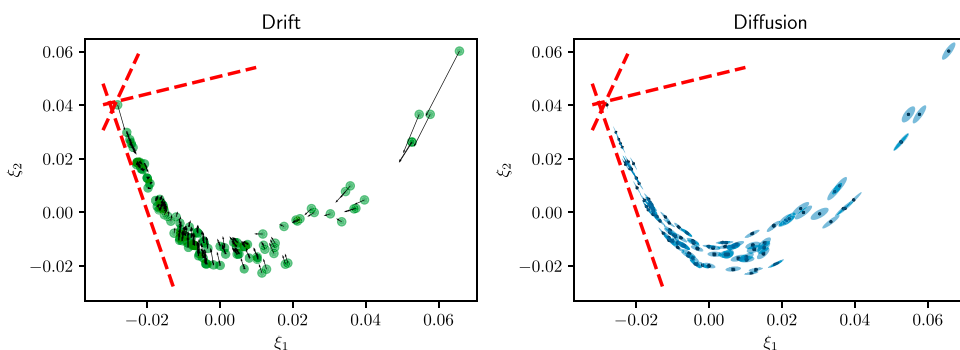


Figure 17. Drift vectors (arrows on the left plot) and diffusion matrices (ellipses representing the principal components of the diffusion on the right plot) for some randomly selected factor data points.




Figure 18. 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).

or Szpruch and Zhang 2018), which is given by

$$\begin{aligned} S_{t+\Delta t} &= S_t + \frac{\mu^S(\tilde{\xi}_t)}{1 + |\mu^S(\tilde{\xi}_t)|\sqrt{\Delta t}} \Delta t + \frac{\sigma^S(\tilde{\xi}_t)}{1 + |\sigma^S(\tilde{\xi}_t)|\sqrt{\Delta t}} (W_{0,t+\Delta t} - W_{0,t}), \\ \xi_{t+\Delta t} &= \xi_t + \frac{\mu(\tilde{\xi}_t)}{1 + |\mu(\tilde{\xi}_t)|\sqrt{\Delta t}} \Delta t + \frac{\sigma(\tilde{\xi}_t)}{1 + \|\sigma(\tilde{\xi}_t)\|\sqrt{\Delta t}} (W_{t+\Delta t} - W_t). \end{aligned} \quad (36)$$

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).

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 10,000 steps for $\tilde{\xi} = (S, \xi)$ in Figure 16, right.¹⁹ In the scatter plot on the left of Figure 16, 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 17.

In Figure 18, 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 4, including the simulated paths of implied volatilities for a variety of option specifications.

VIX simulation: In addition, we follow the CBOE VIX calculation methodology (CBOE 2019) to compute a *volatility index* from the simulated option prices. Specifically, 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 $VIX_t := \mathbf{h}^\top \mathbf{c}_t$ 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

$$VIX_t^s = \mathbf{h}^\top \mathbf{c}_t^s = \mathbf{h}^\top \mathbf{G}_0 + \mathbf{h}^\top \mathbf{G}^\top \xi_t^s. \quad (37)$$

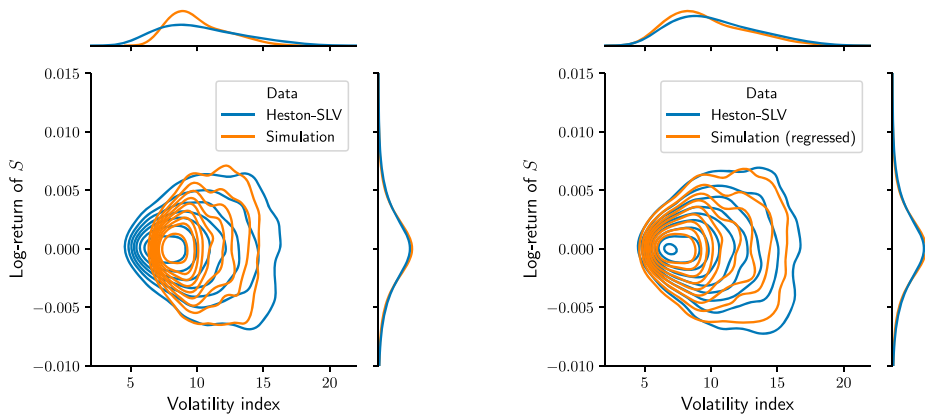
In Figure 19(a), we plot the marginal and joint distributions of the log-return of S and the VIX index, for both the input data and the simulation data. While the distributions of the log-return of S are very close, our simulation gives a VIX of much lower kurtosis. This error is principally due to our 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 minimize the impact of the factor reconstruction error. Consider fitting the linear regression model

$$VIX_t = \beta_0 + \boldsymbol{\beta}^\top \xi_t + \varepsilon_t, \quad (38)$$

Using this relationship, we can compute VIX directly from the simulated factors ξ_t^s using

$$VIX_t^r = \hat{\beta}_0 + \hat{\boldsymbol{\beta}}^\top \xi_t^s. \quad (39)$$



(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 19. Joint distribution of the log-return of S and the VIX-like volatility index. (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 20. Simulated time series of VIX^r , together with log-returns of the underlying stock.

We plot the distributions for this factor-regression based VIX in Figure 19(b). 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 20. We see several occurrences of volatility clustering in the return series, which always coincide with high VIX values.

Remark 4.1: The nontrivial improvement in VIX simulation from (37) to (39) indicates that the decoded factors and their corresponding price basis do not optimize 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 minimizing

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.)

4.7. 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) of normalized call prices in terms of latent factors. In fact, we see that ξ_1 has a strong linear relationship with ν (Figure 13), where most of the residuals from this linearity has been captured by ξ_2 . In addition, as seen in Figure 16, 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.

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 (35) 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 (35)).

Notes

1. Although we refer to the asset as ‘stock’ throughout the paper, our methods are equally applicable to other asset classes such as currencies or commodities.
2. Unsurprisingly, 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 2009).
3. The 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.
4. Chataigner, Crépey, and Dixon (2020) 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 et al. 2009), while the soft constraint approach introduces penalty terms favouring the satisfaction of the constraints (see Itkin 2019). 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.
5. When applying our methods to other asset classes, dividends of stocks are comparable to foreign currency interest for FX, or convenience yield for commodities.
6. The convex order is given by $\mathbb{Q}_{T_1} \geq_{cx} \mathbb{Q}_{T_2}$ if $\int_{\mathbb{R}} \phi d\mathbb{Q}_{T_1} \geq \int_{\mathbb{R}} \phi d\mathbb{Q}_{T_2}$ for each convex function $\phi : \mathbb{R} \rightarrow \mathbb{R}$. We say $\{\mathbb{Q}_T\}_{T \in [0, T^*)}$ is non-decreasing in convex order if $\mathbb{Q}_{T_1} \geq_{cx} \mathbb{Q}_{T_2}$ whenever $T_1 \geq T_2$.
7. Since 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 16.
8. This is consistent with the re-calibration approach in financial practice, which lends justification to our assumption of a stationary model for each fixed maturity.

9. In 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.
10. We use the neural SDE estimation method described in Section 3, where the neural network architecture is defined in (32).
11. One 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 .
12. The normalization of factors results in $\mathbf{G}\mathbf{G}^\top = \mathbf{\Lambda}^2$, rather than an identity matrix.
13. The minimization problem (14) in the factor decoding algorithm has a discrete objective function, where gradient-based optimization methods do not apply. We use Py-BOBYQA (Cartis et al. 2019), a derivative-free optimization solver, to find its global optimum heuristically.
14. These lattice points are chosen to mimic a typical liquid range of EURUSD options traded at CME, as seen in Figure 1.
15. We are solving a supervised learning problem with the input-label data as $(x_l, y_l) = (\tilde{\xi}_l, \xi_{l+1} - \xi_l)$ 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.
16. The robustness of the estimated neural network is assessed via sensitivity analysis in Appendix 6.
17. The first 90% of the dataset is used for training and the last 10% is reserved as validation data.
18. We 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.
19. There is no pathwise similarity here, as our learned factor model is based on a larger number of Brownian motions than the ground truth model.
20. If 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 (2007).

Acknowledgements

We thank Florian Huchede, Director of Quantitative Risk Management, and other colleagues at CME Group for providing valuable data access, suggestions from the business perspective, and continued support.

Disclosure statement

No potential conflict of interest was reported by the author(s).

Funding

This publication is based on work supported by the EPSRC Centre for Doctoral Training in Industrially Focused Mathematical Modelling [grant number EP/L015803/1] in collaboration with CME Group. Samuel Cohen and Christoph Reisinger acknowledge the support of the Oxford-Man Institute for Quantitative Finance, and Samuel Cohen also acknowledges the support of the Alan Turing Institute under the Engineering and Physical Sciences Research Council [grant number EP/N510129/1].

References

- Abadi, M., A. Agarwal, P. Barham, E. Brevdo, Z. Chen, C. Citro, G. S. Corrado, et al. 2015. “TensorFlow: Large-Scale Machine Learning on Heterogeneous Systems.” <https://tensorflow.org/>.

- Bain, A., M. Mariapragassam, and C. Reisinger. 2021. "Calibration of Local-Stochastic and Path-Dependent Volatility Models to Vanilla and No-Touch Options." *The Journal of Computational Finance* 24 (4): 115–161. <https://doi.org/10.21314/JCF.2020.400>.
- Black, F., and M. Scholes. 1973. "The Pricing of Options and Corporate Liabilities." *Journal of Political Economy* 81 (3): 637–54. <https://doi.org/10.1086/260062>.
- Breeden, D. T., and R. H. Litzenberger. 1978. "Prices of State-Contingent Claims Implicit in Option Prices." *The Journal of Business* 51 (4): 621–51. <https://ssrn.com/abstract=2642349>.
- Carmona, R. 2007. "HJM: A Unified Approach to Dynamic Models for Fixed Income, Credit and Equity Markets." Paris–Princeton Lectures on Mathematical Finance 2004, 1–50. Berlin, Heidelberg: Springer.
- Carmona, R., and S. Nadtochiy. 2009. "Local Volatility Dynamic Models." *Finance and Stochastics* 13 (1): 1–48. <https://doi.org/10.1007/s00780-008-0078-4>.
- Carmona, R., and S. Nadtochiy. 2012. "Tangent Lévy Market Models." *Finance and Stochastics* 16 (1): 63–104. <https://doi.org/10.1007/s00780-011-0158-8>.
- Caron, R. J., J. F. McDonald, and C. M. Ponik. 1989. "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. <https://doi.org/10.1007/BF00941055>.
- Carr, P., H. Geman, D. B. Madan, and M. Yor. 2003. "Stochastic Volatility for Lévy Processes." *Mathematical Finance* 13 (3): 345–382. <https://doi.org/10.1111/1467-9965.00020>.
- Carr, P., and D. B. Madan. 2005. "A Note on Sufficient Conditions for No Arbitrage." *Finance Research Letters* 2 (3): 125–130. <https://doi.org/10.1016/j.frl.2005.04.005>.
- Cartis, C., J. Fiala, B. Marteau, and L. Roberts. 2019. "Improving the Flexibility and Robustness of Model-Based Derivative-Free Optimization Solvers." *ACM Transactions on Mathematical Software* 45 (3): 32. <https://doi.org/10.1145/3338517>.
- CBOE. 2019. "White Paper: CBOE Volatility Index." April 2021. <https://cdn.cboe.com/resources/vix/vixwhite.pdf>.
- Chataigner, M., S. Crépey, and M. Dixon. 2020. "Deep Local Volatility." *Risks* 8 (3): 1. <https://doi.org/10.3390/risks8030082>.
- Choudhary, V., S. Jaimungal, and M. Bergeron. 2023. "FuNVol: A Multi-Asset Implied Volatility Market Simulator Using Functional Principal Components and Neural SDEs." ArXiv Preprint ArXiv:2303.00859.
- Cohen, S. N., C. Reisinger, and S. Wang. 2020. "Detecting and Repairing Arbitrage in Traded Option Prices." *Applied Mathematical Finance* 27 (5): 345–373. <https://doi.org/10.1080/1350486X.2020.1846573>.
- Cohen, S. N., C. Reisinger, and S. Wang. 2022a. "Estimating Risks of European Option Books Using Neural Stochastic Differential Equation Market Models." *Journal of Computational Finance* 26 (3): 33–72.
- Cohen, S. N., C. Reisinger, and S. Wang. 2022b. "Hedging Option Books Using Neural-SDE Market Models." *Applied Mathematical Finance* 29 (5): 366–401. <https://doi.org/10.1080/1350486X.2023.2221448>.
- Cont, R., and J. da Fonseca. 2002. "Dynamics of Implied Volatility Surfaces." *Quantitative Finance* 2 (1): 45–60. <https://doi.org/10.1088/1469-7688/2/1/304>.
- Cont, R., J. da Fonseca, and V. Durrleman. 2002. "Stochastic Models of Implied Volatility Surfaces." *Economic Notes* 31 (2): 361–377. <https://doi.org/10.1111/1468-0300.00090>.
- Cont, R., and M. Vuletić. 2022. "Simulation of Arbitrage-Free Implied Volatility Surfaces." Available at SSRN: <https://ssrn.com/abstract=4299363>.
- Cousot, L. 2007. "Conditions on Option Prices for Absence of Arbitrage and Exact Calibration." *Journal of Banking & Finance* 31 (11): 3377–3397. <https://doi.org/10.1016/j.jbankfin.2007.04.006>.
- Cuchiero, C., W. Khosrawi, and J. Teichmann. 2020. "A Generative Adversarial Network Approach to Calibration of Local Stochastic Volatility Models." *Risks* 8 (4). <https://doi.org/10.3390/risks8040101>.
- Daglish, T., J. Hull, and W. Suo. 2007. "Volatility Surfaces: Theory, Rules of Thumb, and Empirical Evidence." *Quantitative Finance* 7 (5): 507–524. <https://doi.org/10.1080/14697680601087883>.

- Davis, M. H. A., and D. G. Hobson. 2007. "The Range of Traded Option Prices." *Mathematical Finance* 17 (1): 1–14. <https://doi.org/10.1111/j.1467-9965.2007.00291.x>.
- Davis, M., and J. Oblój. 2008. "Market Completion Using Options." *Advances in Mathematics of Finance* 83: 49–60. <https://doi.org/10.4064/bc83-0-4>.
- Delbaen, F., and W. Schachermayer. 1994. "A General Version of the Fundamental Theorem of Asset Pricing." *Mathematische Annalen* 300 (1): 463–520. <https://doi.org/10.1007/BF01450498>.
- Derman, E., and I. Kani. 1994. "Riding on a Smile." *Risk* 7 (2): 32–39.
- Derman, E., and I. Kani. 1998. "Stochastic Implied Trees: Arbitrage Pricing with Stochastic Term and Strike Structure of Volatility." *International Journal of Theoretical and Applied Finance* 1 (1): 61–110. <https://doi.org/10.1142/S0219024998000059>.
- Dugas, C., Y. Bengio, F. Bélisle, C. Dadeau, and R. Garcia. 2009. "Incorporating Functional Knowledge in Neural Networks." *Journal of Machine Learning Research* 10 (42): 1239–1262.
- Dupire, B. 1994. "Pricing with a Smile." *Risk Magazine* 7: 18–20.
- Fengler, M. 2009. "Arbitrage-Free Smoothing of the Implied Volatility Surface." *Quantitative Finance* 9 (4): 417–428. <https://doi.org/10.1080/14697680802595585>.
- Fengler, M. R., and L.-Y. Hin. 2015. "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. <https://doi.org/10.1016/j.jeconom.2014.09.003>.
- Friedman, A., and M. A. Pinsky. 1973. "Asymptotic Stability and Spiraling Properties for Solutions of Stochastic Equations." *Transactions of the American Mathematical Society* 186: 331–358. <https://doi.org/10.1090/S0002-9947-1973-0329031-5>.
- Gatheral, J., and A. Jacquier. 2014. "Arbitrage-Free SVI Volatility Surfaces." *Quantitative Finance* 14 (1): 59–71. <https://doi.org/10.1080/14697688.2013.819986>.
- Gierjatowicz, P., M. Sabate-Vidales, D. Šiška, Ł. Szpruch, and Ž. Žurič. 2020. "Robust Pricing and Hedging via Neural SDEs." ArXiv:2007.04154.
- Hagan, P. S., D. Kumar, A. S. Lesniewski, and D. E. Woodward. 2002. "Managing Smile Risk." *The Best of Wilmott* 1: 249–296.
- Hagan, P. S., D. Kumar, A. S. Lesniewski, and D. E. Woodward. 2014. "Arbitrage-Free SABR." *Wilmott* 69: 60–75. <https://doi.org/10.1002/wilm.10290>.
- Harrison, J. M., and D. Kreps. 1979. "Martingales and Arbitrage in Multiperiod Securities Markets." *Journal of Economic Theory* 20 (3): 381–408. [https://doi.org/10.1016/0022-0531\(79\)90043-7](https://doi.org/10.1016/0022-0531(79)90043-7).
- Heath, D., R. Jarrow, and A. Morton. 1992. "Bond Pricing and the Term Structure of Interest Rates: A New Methodology for Contingent Claims Valuation." *Econometrica* 60 (1): 77–105. <https://doi.org/10.2307/2951677>.
- Heston, S. L. 1993. "A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options." *Review of Financial Studies* 6 (2): 327–343. <https://doi.org/10.1093/rfs/6.2.327>.
- Hutzenthaler, M., A. Jentzen, and P. E. Kloeden. 2012. "Strong Convergence of An Explicit Numerical Method for SDEs with Nonglobally Lipschitz Continuous Coefficients." *The Annals of Applied Probability* 22 (4): 1611–1641. <https://doi.org/10.1214/11-AAP803>.
- Itkin, A. 2019. "Deep Learning Calibration of Option Pricing Models: Some Pitfalls and Solutions." ArXiv:1906.03507.
- Jacod, J., and P. Protter. 2010. "Risk-Neutral Compatibility with Option Prices." *Finance and Stochastics* 14 (2): 285–315. <https://doi.org/10.1007/s00780-009-0109-9>.
- Jex, M., R. Henderson, and D. Wang. 1999. "Pricing Exotics under the Smile." *Risk* 72–75.
- Kahale, N. 2004. "An Arbitrage-Free Interpolation of Volatilities." *Risk Magazine* 17:102–106.
- Kallsen, J., and P. Krühner. 2013. "On a Heath–Jarrow–Morton Approach for Stock Options." *Finance and Stochastics* 19: 583–615. <https://doi.org/10.1007/s00780-015-0263-1>.
- Karatzas, I., and C. Kardaras. 2007. "The Numéraire Portfolio in Semimartingale Financial Models." *Finance and Stochastics* 11 (4): 447–493. <https://doi.org/10.1007/s00780-007-0047-3>.
- Kellerer, H. G. 1972. "Markov-Komposition Und Eine Anwendung Auf Martingale." *Mathematische Annalen* 198 (3): 99–122. <http://eudml.org/doc/162296>. <https://doi.org/10.1007/BF01432281>.
- Kreps, D. 1981. "Arbitrage and Equilibrium in Economies with Infinitely Many Commodities." *Journal of Mathematical Economics* 8 (1): 15–35. [https://doi.org/10.1016/0304-4068\(81\)90010-0](https://doi.org/10.1016/0304-4068(81)90010-0).

- Motzkin, T. S., H. Raiffa, G. L. Thompson, and R. M. Thrall. 1953. “The Double Description Method.” In *Contributions to the Theory of Games II*. Princeton University Press.
- Piterbarg, V. 2006. “Markovian Projection Method for Volatility Calibration.” Available at SSRN: <https://ssrn.com/abstract=906473>.
- Quantlib. 2021. “QuantLib, a Free/Open-Source Library for Quantitative Finance.” Accessed April 13, 2021. <https://www.quantlib.org/>.
- Reiswich, D., and U. Wystup. 2009. “FX Volatility Smile Construction.” CPQF Working Paper Series 20. Frankfurt School of Finance and Management, Centre for Practical Quantitative Finance (CPQF). Accessed May 20, 2021. <https://EconPapers.repec.org/RePEc:zbw:cpqfw:20>.
- Schonbucher, P. J. 1999. “A Market Model for Stochastic Implied Volatility.” *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 357 (1758): 2071–2092. <https://doi.org/10.1098/rsta.1999.0418>.
- Schwarz, D. C. 2017. “Market Completion with Derivative Securities.” *Finance and Stochastics* 21 (1): 263–284. <https://doi.org/10.1007/s00780-016-0317-z>.
- Schweizer, M., and J. Wissel. 2008a. “Arbitrage-Free Market Models for Option Prices: The Multi-Strike Case.” *Finance and Stochastics* 12 (4): 469–505. <https://doi.org/10.1007/s00780-008-0068-6>.
- Schweizer, M., and J. Wissel. 2008b. “Term Structures of Implied Volatilities: Absence of Arbitrage and Existence Results.” *Mathematical Finance* 18 (1): 77–114. <https://doi.org/10.1111/j.1467-9965.2007.00323.x>.
- Szpruch, Ł., and X. Zhang. 2018. “V-Integrability, Asymptotic Stability and Comparison Theorem of Explicit Numerical Schemes for SDEs.” *Mathematics of Computations* 87:755–783. <https://doi.org/10.1090/mcom/2018-87-310>.
- Ventzel, A. D. 1965. “On Equations of the Theory of Conditional Markov Processes.” *Theory of Probability and Its Applications* 10 (2): 357–360. <https://doi.org/10.1137/1110045>.
- Wissel, J. S. 2008. “Arbitrage-Free Market Models for Liquid Options.” PhD diss., ETH Zürich.

Appendices

Appendix 1: 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 (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 (3) and the latent factor representation (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

$$\begin{aligned}
 d\bar{c}_t &= dg(t, \xi_t) \\
 &= \left[\frac{\partial g}{\partial t}(t, \xi_t) + \sum_{i=1}^d \mu_{i,t} \frac{\partial g}{\partial y_i}(t, \xi_t) + \frac{1}{2} \sum_{i=1}^d \sum_{k=1}^d \sum_{j=1}^d \sigma_{ij,t} \sigma_{kj,t} \frac{\partial^2 g}{\partial y_i \partial y_k}(t, \xi_t) \right] dt \\
 &\quad + \sum_{i=1}^d \sum_{j=1}^d \sigma_{ij,t} \frac{\partial g}{\partial y_i}(t, \xi_t) dW_{j,t} \\
 &=: \tilde{\eta}_t dt + \tilde{\nu}_t dW_t.
 \end{aligned} \tag{A1}$$

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

$$dm_t = dM(F_t) = \left[(\alpha_t - r_t)F_t \frac{\partial M}{\partial f}(F_t) + \frac{1}{2} \gamma_t^2 F_t^2 \frac{\partial^2 M}{\partial f^2}(F_t) \right] dt + \gamma_t F_t \frac{\partial M}{\partial f}(F_t) dW_{0,t}. \quad (\text{A2})$$

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 (A2) and $\tau_t = T - t$. Using a generalized Itô–Wentzell formula (Ventzel' 1965),

$$\begin{aligned} d\hat{c}_t(T, K) &= d\tilde{c}_t(T - t, m_t) \\ &= \left\{ \tilde{\eta}_t(T - t, m_t) - \frac{\partial g}{\partial \tau}(t, T - t, m_t) + \frac{1}{2} \left[\gamma_t F_t \frac{\partial M}{\partial f}(F_t) \right]^2 \frac{\partial^2 g}{\partial m^2}(t, T - t, m_t) \right. \\ &\quad \left. + \left[(\alpha_t - r_t)F_t \frac{\partial M}{\partial f}(F_t) + \frac{1}{2} \gamma_t^2 F_t^2 \frac{\partial^2 M}{\partial f^2}(F_t) \right] \frac{\partial g}{\partial m}(t, T - t, m_t) \right\} dt \\ &\quad + \gamma_t F_t \frac{\partial M}{\partial f}(F_t) \frac{\partial g}{\partial m}(t, T - t, m_t) dW_{0,t} + \tilde{v}_t(T - t, m_t) dW_t \\ &=: \eta_t(T, K) dt + v_t(T, K) d\bar{W}_t. \end{aligned} \quad (\text{A3})$$

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

$$d\bar{c}_t = dh(F_t, \hat{c}_t) = [\eta_t F_t + (\alpha_t - r_t) \bar{c}_t + v_{0,t} \gamma_t F_t] dt + (v_t F_t + \gamma_t \bar{c}_t \mathbf{e}_1) d\bar{W}_t, \quad (\text{A4})$$

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

A.1. 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 $\tilde{W}_t = [\tilde{W}_{0,t} \ \cdots \ \tilde{W}_{d,t}]^\top$ have dynamics given by

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

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's²⁰ 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

$$\frac{d\mathbb{Q}}{d\mathbb{P}} \Big|_{\mathcal{F}_t} = \exp \left(- \int_0^t \varphi_s^\top d\bar{W}_s - \frac{1}{2} \int_0^t \varphi_s^\top \varphi_s ds \right). \quad (\text{A5})$$

Under \mathbb{Q} the process \tilde{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 (A5) from the market price of risk process φ , we see that no-arbitrage requires the existence of an appropriate φ . Letting

$$\psi_{0,t} = \frac{\alpha_t - r_t}{\gamma_t}, \quad (\text{A6})$$

under the measure \mathbb{Q} , the discounted stock price $\tilde{S}_t = S_t \exp(\int_0^t (q_s - r_s) ds)$ and the forward price are \mathbb{Q} -local martingales, i.e., $d\tilde{S}_t = \gamma_t \tilde{S}_t d\tilde{W}_{0,t}$ and $dF_t = \gamma_t F_t d\tilde{W}_{0,t}$. Since we have assumed that $\gamma S \in L_{\text{loc}}^2(\mathbb{R}^{d+1})$, it follows that \tilde{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 (A4) and (A6), we have, under \mathbb{Q} ,

$$d\bar{c}_t = (\eta_t + v_{0,t}\gamma_t - v_t \cdot \varphi_t)F_t dt + (v_t F_t + \gamma_t \bar{c}_t \mathbf{e}_1) d\tilde{W}_t. \quad (\text{A7})$$

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

$$v_t(T, K) \cdot \varphi_t(T, K) = \eta_t(T, K) + v_{0,t}(T, K)\gamma_t. \quad (\text{A8})$$

Given the parametrization in terms of moneyness m_t and the market factor representation (4), we can simplify the drift restriction (A8). 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

$$\begin{aligned} \eta_t(T, K) = & \sum_{i=1}^d \left\{ \left[\mu_{i,t} + \xi_{i,t} \left(-\frac{\partial}{\partial \tau} + \left(r_t - \alpha_t + \frac{1}{2}\gamma_t^2 \right) \frac{\partial}{\partial m} + \frac{1}{2}\gamma_t^2 \frac{\partial^2}{\partial m^2} \right) \right] G_i(\tau_t, m_t) \right\} \\ & + \left[-\frac{\partial}{\partial \tau} + \left(r_t - \alpha_t + \frac{1}{2}\gamma_t^2 \right) \frac{\partial}{\partial m} + \frac{1}{2}\gamma_t^2 \frac{\partial^2}{\partial m^2} \right] G_0(\tau_t, m_t), \end{aligned} \quad (\text{A9})$$

$$v_{j,t}(T, K) = \begin{cases} -\sum_{i=1}^d \xi_{i,t} \gamma_t \frac{\partial G_i(\tau_t, m_t)}{\partial m} - \gamma_t \frac{\partial G_0(\tau_t, m_t)}{\partial m}, & j = 0; \\ \sum_{i=1}^d \sigma_{ij,t} G_i(\tau_t, m_t), & 1 \leq j \leq d. \end{cases} \quad (\text{A10})$$

Substituting η_t and v_t into the drift restriction (A8) yields

$$\begin{aligned} & \sum_{j=1}^d \left(\sum_{i=1}^d \sigma_{ij,t} G_i(\tau_t, m_t) \right) \psi_{j,t} \\ &= \sum_{i=1}^d \left\{ \mu_{i,t} G_i(\tau_t, m_t) + \left[-\frac{\partial}{\partial \tau} - \frac{1}{2}\gamma_t^2 \frac{\partial}{\partial m} + \frac{1}{2}\gamma_t^2 \frac{\partial^2}{\partial m^2} \right] G_i(\tau_t, m_t) \xi_{i,t} \right\} \\ &= \sum_{i=1}^d \mu_{i,t} G_i(\tau_t, m_t) + \left(-\frac{\partial}{\partial \tau} - \frac{1}{2}\gamma_t^2 \frac{\partial}{\partial m} + \frac{1}{2}\gamma_t^2 \frac{\partial^2}{\partial m^2} \right) \tilde{c}_t(\tau_t, m_t). \end{aligned} \quad (\text{A11})$$

It follows that a necessary and sufficient condition for our models to be dynamic arbitrage-free is that (A11) admits solutions with sufficient integrability.

Appendix 2: 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 $\mathbf{P}(y)$ (24) in Algorithm 2 is as follows:

Similarly to 2, we define

$$\mathbf{P}(y) = \text{diag}(\sqrt{\epsilon_1}, \dots, \sqrt{\epsilon_d}) \times \mathbf{Q},$$

where $\{\sqrt{\epsilon_k}\}_k$ is a set of adaptive shrinking scales. The novelty here is that we compute these scales recursively, together with the Gram–Schmidt process in (23).

Let $\epsilon_1 = \varepsilon_{(1)}$. For each $k \geq 2$, suppose $\{\epsilon_j\}_{j \leq k-1}$ 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 $\sqrt{\epsilon_j}$ for all j , \mathbf{v}_k will shrink to

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

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

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

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.

Appendix 3: 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 et al. 1953) 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**
-

Appendix 4: 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 A1 and Figure A2.

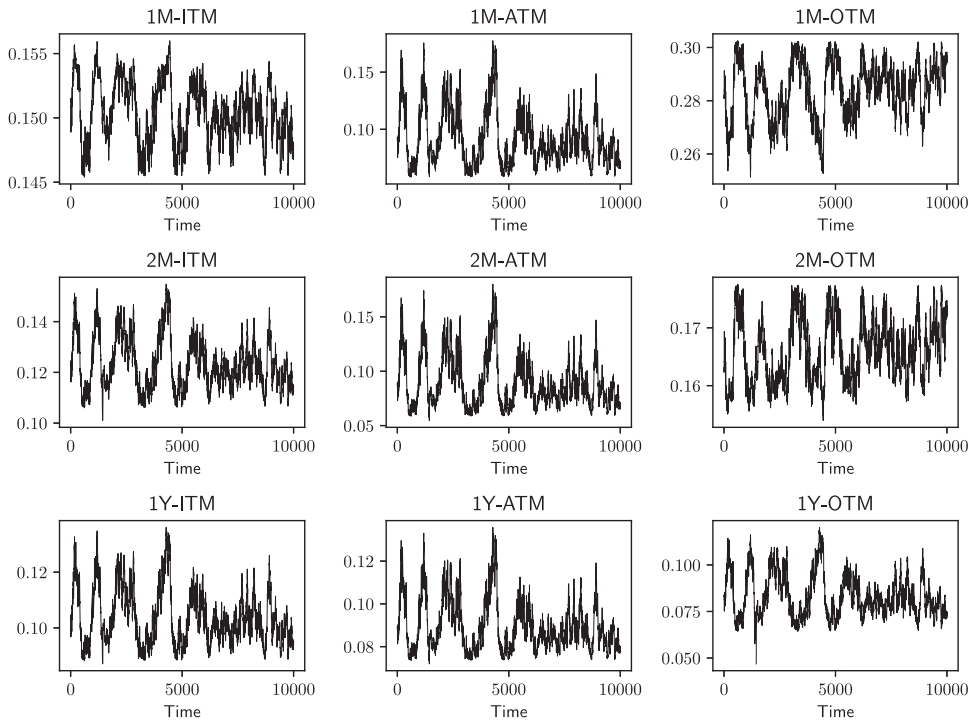


Figure A1. Simulated evolution of implied volatilities for multiple expiries and moneynesses. ATM: $m = 0$; ITM: $m = -0.0943$; OTM: $m = 0.1133$.




Figure A2. Simulated implied volatility smiles for 1M, 2M, 6M and 1Y options on multiple days.

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

A variant of our approach was more recently studied in Cont and Vuletić (2022), who consider linear factor models for the implied volatility surface instead of the option price surface. As the static no-arbitrage constraints are then nonlinear in the factors even in a linear model, a penalization approach is used to attach low probability to scenarios with static arbitrage. In contrast, Choudhary, Jaimungal, and Bergeron (2023) use functional principal components of the volatility surface in conjunction with neural SDEs and find that generated scenarios are largely free of static arbitrage without explicitly enforcing it.




Figure A3. Evolution of training losses and validation losses (without imposing drift correction and diffusion shrinking).

Appendix 5: 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 (19) are satisfied. These conditions are sufficient for constraining the factor process within the state space where there is no static arbitrage.

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 A3. We find that, compared with Figure 11,

- (1) The losses converge faster. This is probably because that diffusion shrinking and drift correction constrain the gradients of the loss function when using backpropagation.
- (2) 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.

We investigate how well the learnt model can simulate out-of-sample data. As shown in Figure A4, the simulated factors can easily break the static arbitrage constraints (the dashed red lines), and generate very unrealistic values. Compared with the simulation results shown in Figure 16 (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.

Appendix 6: 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 4.4 the *benchmark* network. The values of the hyperparameters for the benchmark network for ξ are listed in Table 1. To carry out sensitivity analysis, we

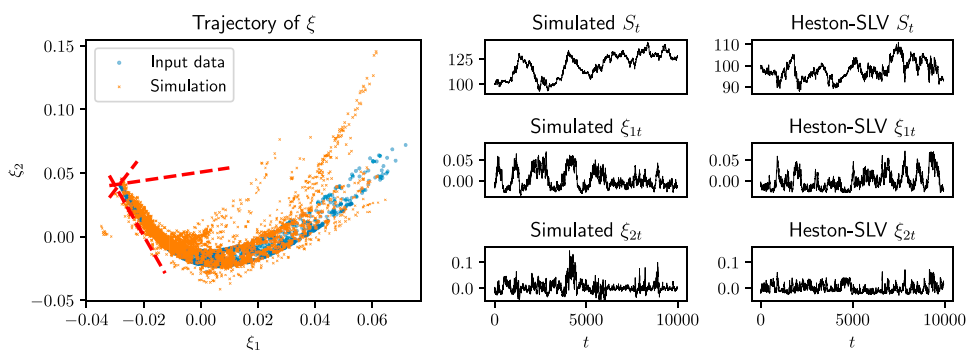


Figure A4. Simulation of S and ξ from the learnt neural network model (without imposing drift correction and diffusion shrinking).

Table A1. Hyperparameters used for the benchmark model.

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

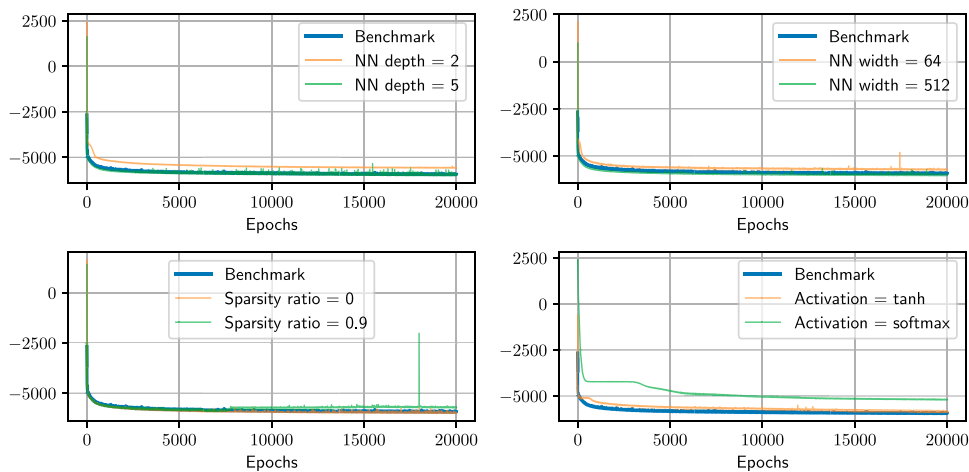


Figure A5. Evolution of training losses under different neural network architectures.

will train a collection of modified networks by varying one hyperparameter from the benchmark network (Table A1).

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 A5. We have shown in Figure 18 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 A2.

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:

Table A2. Wasserstein distances ($\times 10^{-4}$) between the empirical distributions of simulated data and the input data.

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

- 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 .