

HEC MONTRÉAL

Parametric Implied Volatility Surface Model for VIX Options

par Anthony Thibeault

Geneviève Gauthier

HEC Montréal

Directrice de recherche

Sciences de la gestion (Spécialisation Ingénierie financière)

Mémoire présenté en vue de l'obtention du grade de maîtrise ès sciences en gestion (M. Sc.)

Décembre 2024

© Anthony Thibeault 2024

Parametric Implied Volatility Surface Model for VIX Options

Anthony Thibeault

December 18, 2024

Abstract

This paper proposes a new parametric model for the daily implied volatility surface of VIX options. The model diverges from traditional methods by using a parametric structure that directly describes the daily surface of VIX options. It leverages interpretable regression factors to capture changes in maturity and moneyness dimensions. The proposed model contributes in three specific ways: an extrapolation capability for illiquid contracts through stable factors, the construction of a twice continuously differentiable surface to extract the VIX risk-neutral density function for pricing purposes, and a pricing framework for a wide range of VIX derivatives which is coherent with quoted options on that day. Compared to the benchmark Heston model, the proposed model demonstrates superior performance in fitting implied volatilities, particularly for deep out-of-the-money and shorter maturity options. The model's stability makes it well-suited to limit arbitrage opportunities.

Keywords: VIX Options, Implied volatility surfaces, Incomplete Markets, Derivatives pricing, Factor models.

Acknowledgements

I want to express my gratitude to my professor Geneviève Gauthier for her unwavering support throughout this project. Her availability and invaluable advice have been truly priceless. Geneviève has passed on to me her passion for research, encouraging me to step out of my comfort zone and strive for excellence. Thanks to her exemplary work ethic, I have developed a wide range of tools that will greatly benefit me in both my academic and professional journey. I am honored to consider Geneviève as a mentor whose impact on my academic journey at HEC Montreal has been the most significant.

I also want to thank my parents, Bruno and Chantal, not only for supporting me throughout this project but for my entire academic journey. Thank you for being there for me, encouraging me, and instilling in me great life values and ethics that have enabled me to accomplish this step in my journey.

Finally, I want to thank my girlfriend and her family for their support. I am eternally grateful to my girlfriend, Emmy, for being by my side through my highs and lows. I want to thank her even though her understanding of what I am doing does not correlate with the level of support she has provided.

1 Introduction

In the decades since their introduction to financial markets, derivatives have become an integral component of financial markets. In today's financial environment, volatility trading has significantly grown in recent years, playing a crucial role in risk management and portfolio diversification. A notable development in this field occurred in 1993 when the Chicago Board Options Exchange (CBOE) introduced the VIX, its primary volatility index. This index approximates expected market volatility implied by the S&P 500 index options over a one-month horizon. The CBOE subsequently expanded its offerings, launching VIX futures in 2004 and options in 2006. These introductions have gained popularity among market participants, providing new opportunities to manage volatility risk. This thesis aims to estimate the implied volatility surface of VIX options to enable the pricing of other volatility derivatives by applying the Breeden and Litzenberger (1978) formula. By doing so, it also becomes possible to retrieve the risk-neutral density function of the VIX for a given horizon, thereby facilitating the pricing of these derivatives using the risk-neutral density.

The financial crisis has led to extensive research into stochastic volatility models to price VIX derivatives, highlighting the growing importance of volatility risk. Two primary approaches have emerged for the valuation of these instruments. The first approach views the VIX as a portfolio of S&P 500 options. This structural approach involves modeling the return of the S&P 500 in a risk-neutral measure to calculate the index options price. The option prices are used to calculate the VIX dynamics under the real measure, which are then used to price VIX options. However, this method involves multiple complex steps, making the model difficult to estimate and potentially leading to significant differences between theoretical VIX option prices and their observed market prices. This method includes the works of Lian and Zhu (2013), Pacati et al. (2018), and Bardgett et al. (2019). The second approach treats the VIX independently of the S&P 500, considering it as having its own dynamic, and includes the works by Mencia and Sentana (2013), Park (2016), and Yuan (2022). This method often requires a complex stochastic model to capture the VIX dynamics. These models have numerous parameters and latent variables, such as volatility and jumps, that must be estimated using robust statistical filtering techniques.

Our method diverges from these two approaches by taking a different route to model option

prices. We build on recent literature based on the work of Francois et al. (2022), which suggests using the implied volatility (IV) surface not as a model output but instead as an input. We propose an extension of the authors' parametric IV surface for VIX options. To our knowledge, we are the first to propose this type of model for VIX options. Traditional pricing models for derivatives propose a coherent internal dynamic model that allows the creation of option prices free of arbitrage given a set of parameters. However, this does not guarantee that the generated option prices reproduce the observable prices on the surface. Our model takes a different approach by constructing a daily surface that best explains the daily options via regression. Instead of capturing a dynamic, we focus on the daily informational content of the options through interpretable regression factors.

This paper has three significant contributions to the field of derivative pricing on the VIX. First, the paper contributes by identifying and assigning interpretable factors to the implied volatility surface. These factors are designed to reflect changes in both the maturity and moneyness dimensions. They capture the curvature and convexity related to time-to-maturity in the maturity dimension. In the moneyness dimension, they capture the slope and extension of the smile, both of which recreate the convexity of the implied volatility smile. Additionally, the smirk is captured as an additional factor on the implied volatility surface. Second, the factors are designed to be asymptotically stable, allowing for reliable and easy extrapolation beyond the quoted moneyness levels and maturities. This is particularly useful for pricing illiquid options, like over-the-counter transactions, which often involve maturities and moneyness that differ from those of publicly quoted options. Third, these factors are set up to ensure that the surface is twice continuously differentiable, allowing for the extraction of the risk-neutral density function for each time horizon. This makes it possible to price a wide range of derivatives on the VIX. The IV surface specification outlined in this paper has successfully passed the static arbitrage detection test by using vertical and butterfly spreads. These spreads provide necessary conditions in our case for preventing static arbitrage.

The paper is organized as follows. Section 2 describes the data. Section 3 outlines the model specification and evaluates its fitting performance against its benchmark. Section 4 examines the presence of arbitrage opportunities in the smoothed IV surface. Section 5 discusses the applications of the IV surface model in derivatives pricing. Finally, Section 6 concludes the paper.

2 Data Description

The dataset extracted from the OptionMetrics database consists of daily quotes for European call and put options on the volatility index (VIX), traded on the CBOE exchange, spanning from January 6, 2006, to December 31, 2022. Each quote contains bid and ask prices, which are used to calculate mid-prices as option prices. Additionally, the dataset includes each option's strike price and expiration date. The maturity is calculated based on a 365 trading days basis and is annualized.

The risk-free rate is calculated daily from the Discount Bond Database¹, using the method explained in Filipovic et al. (2022). This method estimates the values of zero-coupon Treasury bonds by minimizing the difference between observed bond prices and their theoretical values, given by the discount curve. The bond value is calculated using kernel ridge regression. This machine learning approach provides a closed-form solution that is represented as a linear combination of kernel functions. The dataset includes daily estimates of annualized, continuously compounded zero-coupon yields for a vast range of daily maturities. We retrieve rates matching the option's time-to-maturity to determine the risk-free rate associated with each quoted option. This ensures that each option is paired with its appropriate risk-free rate. For days that rates are not available but quoted options are, we take the closest previous trading day in which rates were available.

2.1 Data Cleaning

We use multiple criteria to filter options from our dataset. We exclude :

1. options with a bid price of zero
2. options with a dollar spread greater or equal to \$2
3. options with a time-to-maturity of less than 7 days

Additionally, we excluded data from the year 2006 as it marks the introduction of VIX options data. Therefore, our analysis focuses on data from 2007 onwards. The original dataset comprised 2,288,789 daily observations. After applying our filtration criteria, the dataset was reduced to 1,817,775 observations, representing approximately 79.42% of the original dataset. Figure 1 shows how the daily filtration of the original dataset evolved, while Table 1 presents the number of options per time-

¹Data retrieved from Discount Bond Data: <https://www.discount-bond-data.org/>

to-maturity buckets. Table 1 demonstrates that most of the options that were removed had a bid price of 0\$ which is also confirmed by Figure 1. The majority of the discarded contracts, regardless of the criteria, had a short time-to-maturity. Figure 1 also supports this finding over time, showing that the majority of removed observations belonged to contracts with a bid price of 0\$. On rare occasions, many observations were removed due to their large spread causing downward spikes in the data retained. A comprehensive analysis of these specific days is presented in Table 9 in Appendix A.1. Starting from 2016, more observations with a maturity of less than 7 days were removed since they began to be traded more frequently.

	$0 < \tau \leq \frac{30}{365}$	$\frac{30}{365} < \tau \leq \frac{90}{365}$	$\frac{90}{365} < \tau \leq \frac{180}{365}$	$\frac{180}{365} < \tau \leq 1$	$\tau > 1$	All
<u>Excluded observations</u>						
2006 Observations	4,290	6,576	3,525	5,890	0	20,281
Bid price = 0\$	192,512	71,368	44,440	7,741	61	316,122
Spread \geq 2\$	36,555	9,006	4,357	1,772	16	51,706
Time-to-maturity \leq 6/365	82,825	0	0	0	0	82,825
Observations retained	461,214	515,218	694,952	144,566	1,825	1,817,775
Total Observations	777,396	602,168	747,274	159,969	1,902	2,288,709

Summary of the data cleaning procedure for VIX options based on their time-to-maturity. The following categories of observations have been excluded: (i) Observations from the year 2006, (ii) observations with a bid price of 0\$, (iii) observations with a spread equal to or greater than 2\$, and (iv) observations with a time-to-maturity of 6 days or less. Observations retained represents the observations that are kept in the sample after the data cleaning procedure, and total observations represents the total number of raw observations belonging to each bucket. For clarity, τ is reported in days on a 365-day scale.

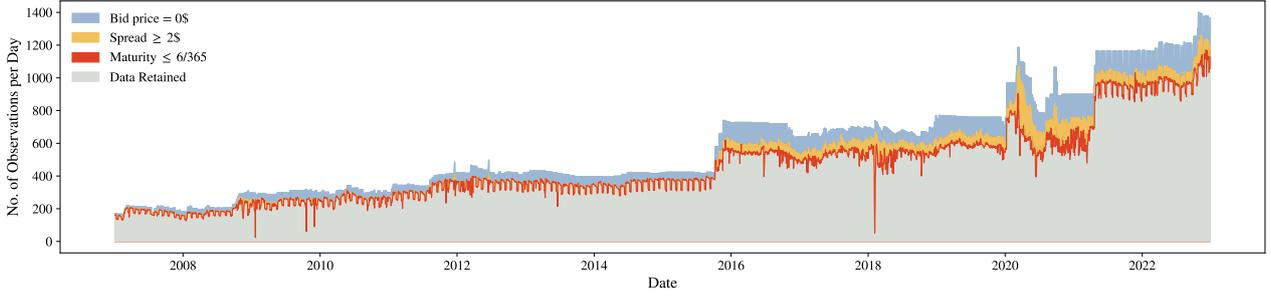
Table 1: Filtered data by time-to-maturity

2.2 Forward Price

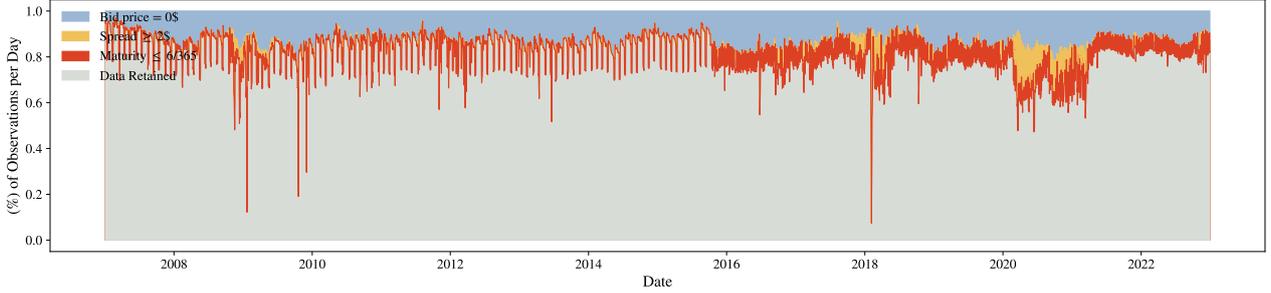
The OptionMetrics forward price for the S&P 500 is computed as $F_{t,\tau} = Se^{(r_{t,\tau}-q_t)\tau}$ where $r_{t,\tau}$ is the time- t continuously compounded risk-free rate (annualized) for time-to-maturity τ , and q_t is the annualized dividend yield. As the VIX is not an asset that can be directly traded, one cannot use the no-arbitrage principle to establish a simple relationship between VIX forward and the VIX, as one can with stock forward and stock prices. This makes pricing the VIX forward using the cost of carry method not suitable. CBOE (2023) computes the VIX forward price F using put-call parity. As outlined in the CBOE white paper,

$$F_{t,\tau} = K + \frac{1}{D_{t,\tau}} (C_t(K, \tau) - P_t(K, \tau)) \quad (1)$$

Panel A: Number of Daily Observations



Panel B: Proportion of Daily Observations



Each surface represents the number of options excluded based on the following criteria: (i) Bid price = 0\$, (ii) spread $\geq 2\%$, and (iii) time-to-maturity $\leq 6/365$. The grey surface represents the data retained after the data-cleaning procedure.

Figure 1: OptionMetrics data from January 2007 to December 2023

where $D_{t,\tau} = e^{-r_{t,\tau}\tau}$ is the discount factor, $r_{t,\tau}$ represents the risk-free rate obtained from the Discount Bond Database, τ is the annualized time-to-maturity of the option, K is the strike price, and $C_t(K, \tau)$ and $P_t(K, \tau)$ represent the call and put option prices, which are the midpoints between bid and ask prices from OptionMetrics. To compute Equation (1), we need the pair $C_t(K, \tau)$ and $P_t(K, \tau)$. Among the 1,817,775 observations, the majority of $C_t(K, \tau)$ and $P_t(K, \tau)$ have their pair. However, a subset of 248,613 observations lacked a put-call parity counterpart. These counterparts were removed during the data cleaning process, as outlined in Section 2.1.

Furthermore, for a given time-to-maturity, many strike prices are traded. Let $K_{t,\tau} = \{K_1, K_2, \dots, K_n\}$ be the set of available strike prices for that particular time-to-maturity. Each of these strike prices induces its own implied forward price through the put-call parity. To get a unique forward price for a respective date and maturity, we utilize a two-step optimization procedure that involves calculating the forward price that minimizes the put-call parity violation across all strike prices associated with the same date and expiration. Let

$$\epsilon_{t,\tau}^K(F) = K + \frac{1}{D_{t,\tau}} (C_t(K, \tau) - P_t(K, \tau)) - F \quad (2)$$

be a measure of the put-call parity violation. The forward price minimizes the sum of squared put-call parity violation:

$$F_{t,\tau} = \underset{F}{\operatorname{argmin}} \sum_{i=1}^n (\epsilon_{t,\tau}^{K_i}(F))^2 \quad (3)$$

In the two-step optimization, the initial pass computes Equation (3) using all available pairs of put and call options. The residuals of Equation (2) are calculated for each contract from which we compute a z-score :

$$z\text{-score}_{\epsilon_{t,\tau}^K} = \frac{\epsilon_{t,\tau}^K(F_{t,\tau}) - \overline{\epsilon_{t,\tau}}}{S_{\epsilon_{t,\tau}}} \quad (4)$$

where

$$\overline{\epsilon_{t,\tau}} = \frac{1}{n} \sum_{i=1}^n \epsilon_{t,\tau}^{K_i}(F_{t,\tau}) \quad (5)$$

is the residual mean and

$$S_{\epsilon_{t,\tau}} = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (\epsilon_{t,\tau}^{K_i}(F_{t,\tau}) - \overline{\epsilon_{t,\tau}})^2} \quad (6)$$

is the standard deviation.

If a z-score is greater than an absolute threshold of 3, it suggests a significant deviation from the rest of the sub-sample and is considered to violate the put-call parity. Contracts associated with residuals having absolute z-scores greater than 3 are identified as outliers and removed from our data sample. In the second optimization pass, we compute the forward value again using Equation (3), but this time, we only consider the contracts that were retained from the first pass.

During the first pass, observations without a counterpart are excluded from the calculation and reintroduced after the second pass. Each observation is attributed its forward value from the second pass, matching its date and expiration. 1,949 observations were discarded because no forward value couldn't be attributed to them. This occurs when all observations for a specific date and expiration were lacking their counterpart and no implied forward could be calculated. Another 2,754 observations were discarded due to the fact that their absolute z-score was higher than 3. Table 2 summarizes the exclusions made across various maturities. It can be observed that the majority of the removed observations are associated with a short maturity. Figure 8 in the Appendix A.2 displays the daily residual mean for both passes of the forward optimization.

	$\frac{6}{365} < \tau \leq \frac{30}{365}$	$\frac{30}{365} < \tau \leq \frac{90}{365}$	$\frac{90}{365} < \tau \leq \frac{180}{365}$	$\frac{180}{365} < \tau \leq 1$	$\tau > 1$	All
<u>Excluded observations</u>						
Missing Counterpart	1,470	368	36	74	1	1,949
Z-score >3	1,406	774	484	90	0	2,754
Proportion (%)	0.62	0.22	0.07	0.11	0.05	0.26
Total Observations	461,214	515,218	694,952	144,566	1,825	1,817,775

Summary of the exclusion process used in the two-pass forward optimization for VIX options across different time-to-maturity buckets. Observations with an absolute z-score greater than 3 are excluded in the second pass and the dataset. Observations that cannot be attributed a forward price are also excluded. The price of the forward is calculated to minimize the put-call parity violation using the equation referred to as Equation (3). For each time-to-maturity bucket, the proportion (%) of excluded observations is calculated by dividing the sum of excluded observations by the total observations.

Table 2: Descriptive statistics of exclusions in two-pass forward optimization

2.3 Moneyness

The moneyness associated with each option quote is defined as stated in Francois et al. (2022) :

$$M = \frac{1}{\sqrt{\tau}} \log\left(\frac{F_{t,\tau}}{K}\right), \quad \text{for } \tau > 0. \quad (7)$$

where $F_{t,\tau}$ represents the forward price estimated from the two-pass optimization process. Due to the volatility scaling property, the range of traded strike prices widens as time-to-maturity increases, resulting in a broader range of strike prices as time-to-maturity increases. Therefore, scaling the moneyness by $\sqrt{\frac{1}{\tau}}$ creates a consistent range of moneyness across various maturities. The value of M being zero indicates at-the-money (ATM) options, while $M < 0$ represents out-of-the-money (OTM) calls, and $M > 0$ signifies OTM puts.

2.4 Data Sample

The sample used to compute put-call parity violation requires both ITM and OTM options. For the rest of this study, our data sample only includes out-of-the-money options. (i.e., puts with $M \geq 0$ and calls with $M < 0$). The sample consists of 792,318 OTM option quotes over a period of 4,028 days, where the implied volatility $\sigma_t(M, \tau)$ of each option is computed by inverting the Black (1976) equation. The put and call price formulas utilize the implied forward from Equation (3) and the moneyness from Equation (7). The moneyness can be inverted to retrieve the actual strike price

through $K = F_{t,\tau} e^{-M\sqrt{\tau}}$. The call and put prices can be expressed as:

$$C_t(M, \tau) = e^{-r\tau} F_{t,\tau} \left(\Phi(d_1(M, \tau)) - e^{-M\sqrt{\tau}} \Phi(d_2(M, \tau)) \right), \quad (8)$$

$$P_t(M, \tau) = e^{-r\tau} F_{t,\tau} \left(e^{-M\sqrt{\tau}} \Phi(-d_2(M, \tau)) - \Phi(-d_1(M, \tau)) \right). \quad (9)$$

where $d_1(M, \tau)$ and $d_2(M, \tau)$ are

$$d_1(M, \tau) = \frac{M}{\sigma_t(M, \tau)} + \frac{1}{2} \sigma_t(M, \tau) \sqrt{\tau},$$

$$d_2(M, \tau) = \frac{M}{\sigma_t(M, \tau)} - \frac{1}{2} \sigma_t(M, \tau) \sqrt{\tau},$$

and $\Phi(\cdot)$ is the standard normal cumulative distribution function.

Descriptive statistics for option implied volatilities (IV) across different moneyness and time-to-maturity are presented in Table 3. The implied volatility surfaces display a significant asymmetry in its smile. Table 3 shows a smaller number of options with short maturities (≤ 30 days) than those with long maturities (>30 days).

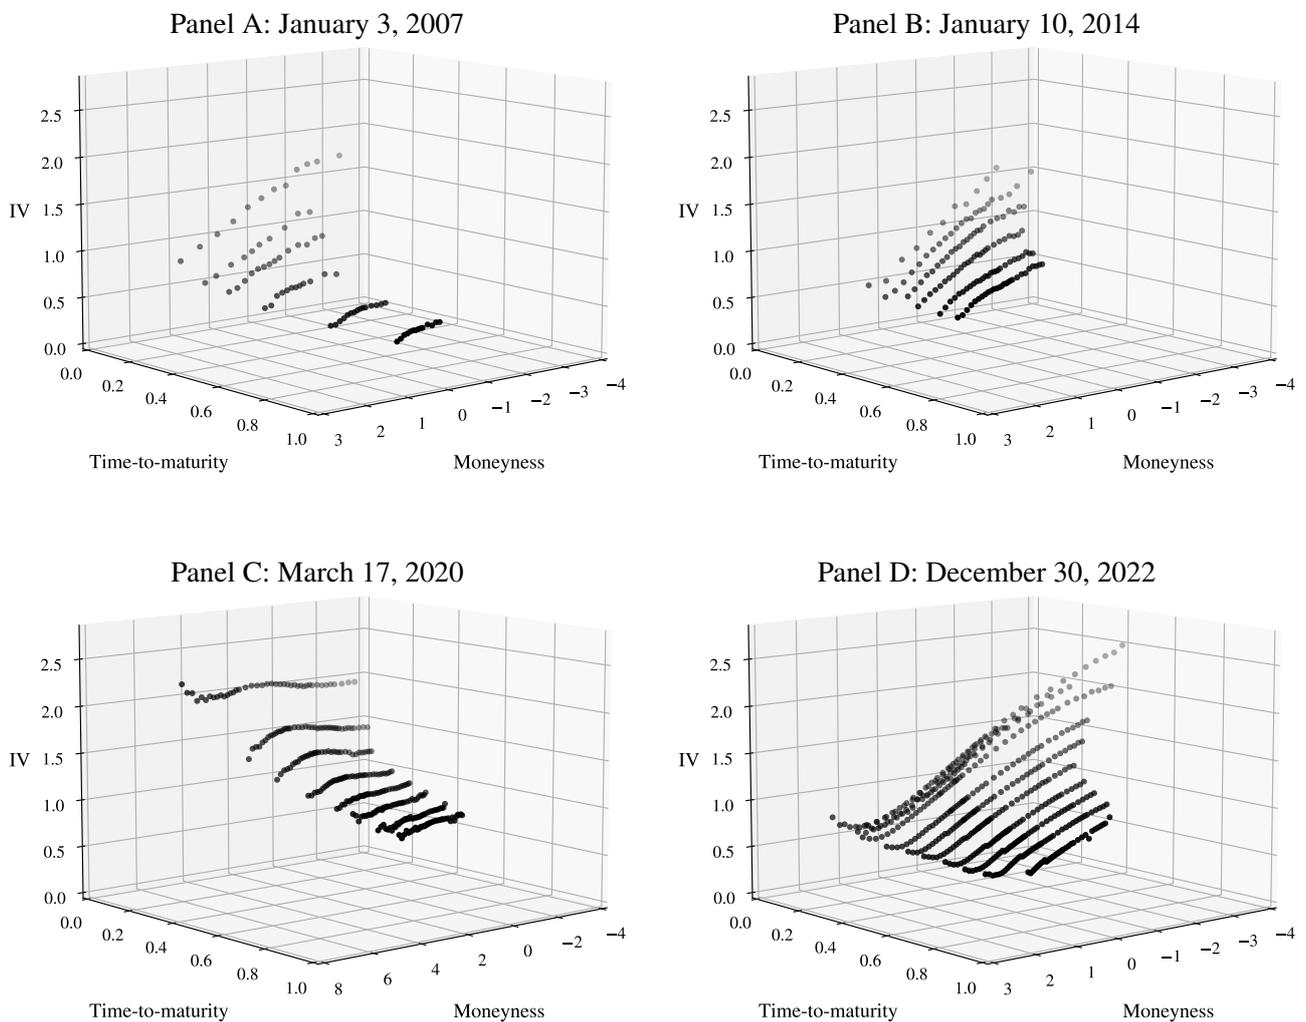
	Calls			Puts		All
	$M \leq -2$	$-2 < M \leq -1$	$-1 < M \leq 0$	$0 < M \leq 1$	$M > 1$	
Average IV(%)	152.53	105.25	84.01	67.03	88.85	96.64
Std. dev. IV(%)	39.80	23.87	20.76	18.65	28.76	38.30
No. of contracts	134,501	189,116	251,876	195,917	20,908	792,318
	$\frac{6}{365} < \tau \leq \frac{30}{365}$	$\frac{30}{365} < \tau \leq \frac{90}{365}$	$\frac{90}{365} < \tau \leq \frac{180}{365}$	$\frac{180}{365} < \tau \leq 1$	$\tau > 1$	All
Average IV(%)	134.79	103.37	78.03	69.05	41.79	96.64
Std. dev. IV(%)	46.59	31.31	19.52	16.54	3.50	38.30
No. of contracts	170,960	224,609	326,989	68,873	887	792,318

Summary of the implied volatility (IV) daily data for VIX options across various time to maturity and moneyness buckets. The data spans from January 3, 2007, to December 30, 2022. M denotes the moneyness, as defined in Equation (7).

Table 3: Descriptive statistics of the VIX options data

Figure 2 displays the surface of all option IV values considered on four different days, which shows specific features of the VIX IV surface. Upon comparing the first day of the sample with subsequent days, it is observed that the number of quoted options has significantly increased over time. This increase in the number of quoted options has primarily resulted in much lower strike intervals between the quotes of a given maturity and additional traded maturities. In Panel B, the surface shows the slope of the moneyness curve decreases as time-to-maturity increases. This indicates a decrease in the level of smile as maturity increases. Panel C presents the implied volatility surface during the

COVID-19 crisis, illustrating a decrease in the IV as the option's time-to-maturity increases. During times of crisis, the moneyness slope is inverted. Toward the end of the sampling period, Panel D displays the smirk of the volatility where OTM calls tend to have a higher IV than ATM options. The implied volatility surface of the VIX, unlike that of the S&P 500, exhibits an inverted shape due to the predominance of call options in its most traded contracts. This inversion intensifies the smirk effect on the call side, contrasting with the traditional skew in the S&P 500. As a result, the VIX usually has more liquidity for OTM calls, which is contrary to what is seen in the S&P 500 implied volatility surface.



The four selected days' retained option quotes are displayed in panels, exhibiting their associated set of implied volatilities. The sample includes four significant dates. The first day of the sample is January 3, 2007. January 10, 2014, is a day with low volatility. March 17, 2020, marks the beginning of a crisis, specifically the COVID-19. The last day in the sample is December 30, 2022. The moneyness is defined in Equation (7).

Figure 2: Observed Implied Volatility (IV) surfaces for different dates

3 Parametric Model Specification and Performance

We choose a static representation of the IV surface using six factors that capture the main empirical characteristics. We consider three requirements : (i) factor interpretability, (ii) twice differentiable and continuous, and (iii) extrapolation ability.

Our factors are crafted to match the observed patterns on IV surfaces using moneyness and time-to-maturity as functions. Each factor plays a specific role and contributes to accurately calibrating the observed IV surfaces with clear financial interpretation. The factors have functional expressions that are twice continuously differentiable, which limits the presence of arbitrage opportunities. More precisely, Breeden and Litzenberger (1978) demonstrated the risk-neutral probability density function at time $t + \tau$ given the time t information satisfies

$$\varphi_{t+\tau}(v_t, \tau) = \exp(r_{t,\tau}, \tau) \frac{\partial^2 C_t(K, \tau)}{\partial K^2} \Big|_{K=v_t} \quad (10)$$

where, in our case, v_t represents the VIX level at time t . For this density to exist, the function must be twice continuously differentiable.²

In the following sections, we introduce our IV model and demonstrate its superiority over the Heston model benchmark. We conduct a screening process that identifies the presence of theoretical arbitrage opportunities. This process is applied to both the raw data and the collection of IV surfaces calibrated by our model. Our IV model, with its asymptotically stable factors, has significant practical implications. It enhances various applications, including option pricing and risk management, by allowing for extrapolation beyond observed moneyness levels and maturities.

²The price of a VIX call option with a strike price of K and time-to-maturity τ is

$$C_t(K, \tau) = \int_K^\infty e^{-r_{t,\tau}\tau} (v_t - K) \varphi_{t+\tau}(v_t, \tau) dv_t$$

3.1 Factorized IV Approach

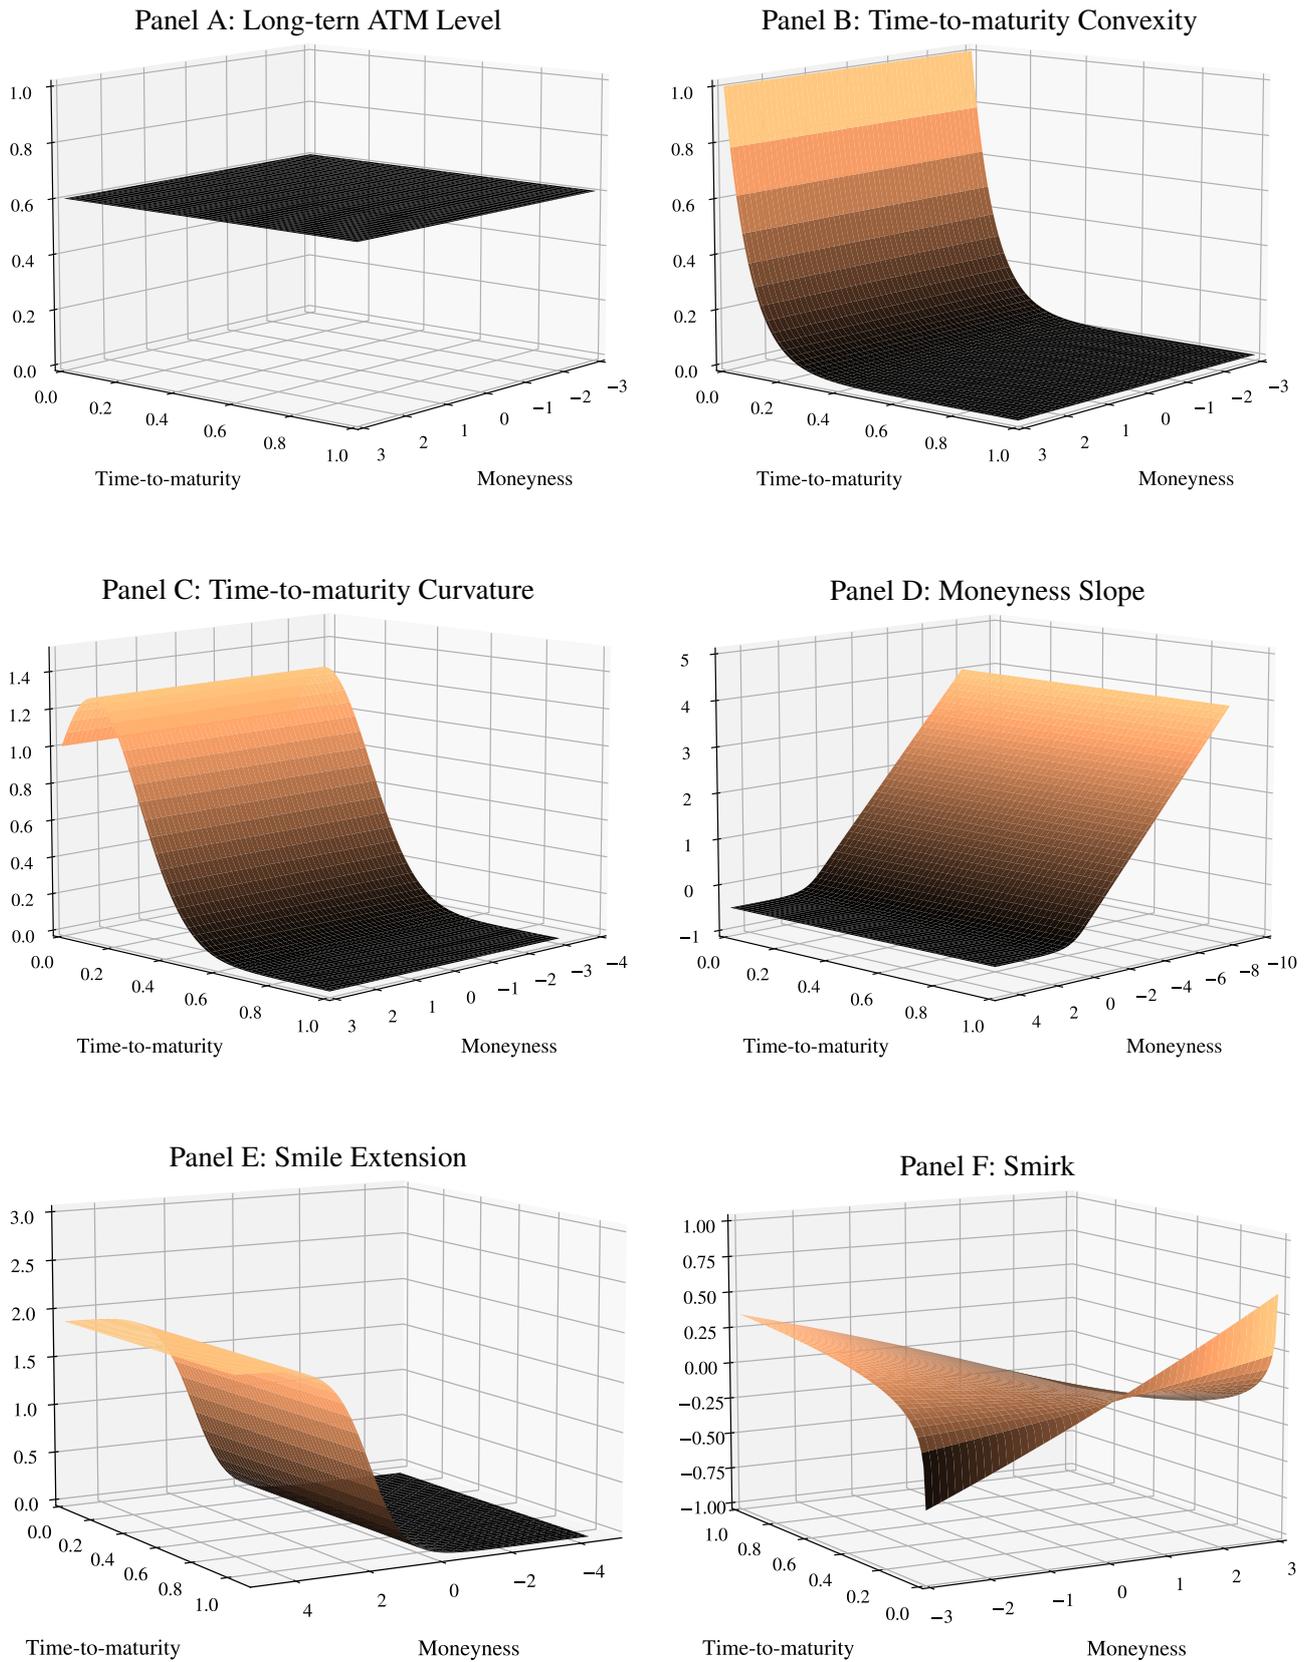
The IV $\sigma(M, \tau)$ observed on a given day for an option with moneyness M defined in Equation (7) and time-to-maturity τ is modeled as

$$\begin{aligned} \sigma(M, \tau) = & \underbrace{\beta_1}_{\text{Long-term ATM level}} + \beta_2 \underbrace{e^{-\frac{\tau}{T_{Conv}}}}_{\text{Time-to-maturity convexity}} + \beta_3 \underbrace{e^{\frac{\tau}{T_{Hump}} \left(1 - \frac{\tau}{T_{Hump}}\right)}}_{\text{Time-to-maturity curvature}} \quad (11) \\ & + \beta_4 \underbrace{\left(-M \mathbb{1}_{M \leq 0} - \frac{e^{2M} - 1}{e^{2M} + 1} \mathbb{1}_{M > 0}\right)}_{\text{Moneyness slope}} + \beta_5 \underbrace{\left(\frac{e^{2(M-1)} - 2}{e^{2(M-1)} + 2} - \frac{e^{-2} - 2}{e^{-2} + 2}\right)}_{\text{Smile extension}} + \beta_6 \underbrace{\left(1 - e^{\frac{M}{M_{Max}}}\right) \log\left(\frac{\tau}{T_{Conv}} + \alpha\right)}_{\text{Smirk}} \end{aligned}$$

The values of M_{Max} , T_{Conv} , T_{Hump} and ϵ are fixed and are used as scaling factors. These values are selected based on empirical observations. The scaling factor M_{Max} corresponds to the maximum absolute moneyness represented by the model. Even though the sample's furthest absolute moneyness for options is around 16, the model is set to $M_{Max} = 20$ to allow for extrapolation beyond the observed moneyness. The parameter $T_{Conv} = 1/12$ (1 month) denotes the point in the IV term structure where there is a rapid shift in convexity concerning time-to-maturity. Furthermore, the parameter $T_{Hump} = 3/12$ (3 months) determines the peak position and declining speed of the hump-shaped component of the IV term structure concerning time-to-maturity. Both of these values were optimized by selecting the value that minimizes the average root mean square error (RMSE) over the sample. The parameter α is set to the smallest maturity allowed in the model, which is $\frac{1}{365}$. This is done despite the fact that options with a maturity of less than $7/365$ are excluded from the dataset. The primary purpose of this approach is to enable the model to extrapolate for options that have a maturity period of less than seven days.

Equation (11) represents the six-factor model designed to capture the key characteristics of the IV surface. We denote this model as the Parametric VIX Implied Volatility Surface (PVIVS) model. Figure 3 displays each factor function in relation to time-to-maturity and moneyness. The coefficient β_1 is interpreted as the long-term ATM implied volatility because $\lim_{\tau \rightarrow \infty} \sigma(0, \tau) = \beta_1$.³ The coefficient β_2 is used to measure the convexity of ATM IV in relation to the time-to-maturity. Since the shape of IV is more accentuated for options with shorter maturities, the convexity adjustment is amplified for options with maturities less than 1 month by associating β_2 to a nonlinear function of $\frac{\tau}{T_{Conv}}$. The

³The proofs for the limits, the function's continuity, and the continuity of the first two derivatives are presented in the Appendix B.1



Model factors with respect of time-to-maturity and moneyness. Panel A: Long-term ATM level, Panel B: Moneyness Slope, Panel C: Time-to-maturity Convexity, Panel D: Time-to-maturity Curvature, Panel E: Smile Extension, Panel F: Smirk.

Figure 3: Model Factors

coefficient β_3 capture the time-to-maturity curvature of the IV. This curvature is characterized by a "hump" shape. The hump shape typically peaks between 0-3 months of maturity and decays as maturity passes 3 months. This curvature correction is achieved by relating β_3 with the scaling factor T_{Hump} . While β_2 affects the short term component of the time-to-maturity, β_3 contributes on the medium part of the IV time-to-maturity as it starts out at one and decay overtime at zero. The expression $\lim_{\tau \rightarrow 0} \sigma(0, \tau) = \beta_1 + \beta_2 + \beta_3$ can be interpreted as the ATM Short Term IV. Furthermore, the time-to-maturity slope of the ATM IV can be characterized as the difference between the short-term and long-term ATM implied volatilities : $\lim_{\tau \rightarrow 0} \sigma(0, \tau) - \lim_{\tau \rightarrow \infty} \sigma(0, \tau) = \beta_2 + \beta_3$. The fourth factor captures the moneyness slope for call and near-the-money put options. Our dataset shows there is a nearly linear increase in call IV as moneyness decrease. Conversely, put IV initially decreases as moneyness increases, but then starts rising for deeper OTM puts, as illustrated in Panel D of Figure 2. The coefficient β_5 is responsible for measuring the moneyness slope, particularly for deep out-of-the-money put options. As observed in Figure 2, put contracts are exhibiting increased activity over time. This trend extends the previous moneyness slope, evolving it into a more traditional smile-shaped surface. The smile extension factor adjusts the observed increase in IV for deep OTM put options. The last coefficient β_6 is responsible for capturing the tilt in the smile for deep OTM puts, also known as the IV smirk. The smirk factor diminishes as the time-to-maturity lengthens.

3.2 Daily Estimation

To estimate the daily set of coefficients $\beta_t = (\beta_{t,1}, \beta_{t,2}, \beta_{t,3}, \beta_{t,4}, \beta_{t,5}, \beta_{t,6})$, a two-pass optimization approach is used. This approach minimizes the sum of the squared fitting errors between the model and observed IVs since they are in the same order of magnitude across moneyness and maturity. In the first pass, all observations are taken into account to estimate β_t . In the second pass, we calculate a z-score of the daily fitting errors from the first pass and exclude any with a z-score greater than 3. The details of this exclusion process are presented in Figure 9 in Appendix B.2. The model's performance and stability are improved by excluding temporarily a small number of outliers, corresponding to 9,657 observations out of 792,318 observations. It results in reduced daily variations in the coefficients.

More precisely, for each $\sigma\{(M_i, \tau_i)\}_{i=1}^n$ representing the pairs of moneyness and time-to-maturity

available on a specific day, we compute the z-score:

$$z\text{-score}_{\epsilon_{t,M_i,\tau_i}} = \frac{\epsilon_{t,M_i,\tau_i} - \bar{\epsilon}_t}{S_t} \quad (12)$$

where

$$\epsilon_{t,M_i,\tau_i} = \sigma_t^{Obs}(M_i, \tau_i) - \sigma_t^{PVIVS}(M_i, \tau_i, \hat{\beta}_t) \quad (13)$$

is the fitting error on the implied volatility,

$$\bar{\epsilon}_t = \frac{1}{n} \sum_{i=1}^n \epsilon_{t,M_i,\tau_i} \quad (14)$$

is the daily mean, and

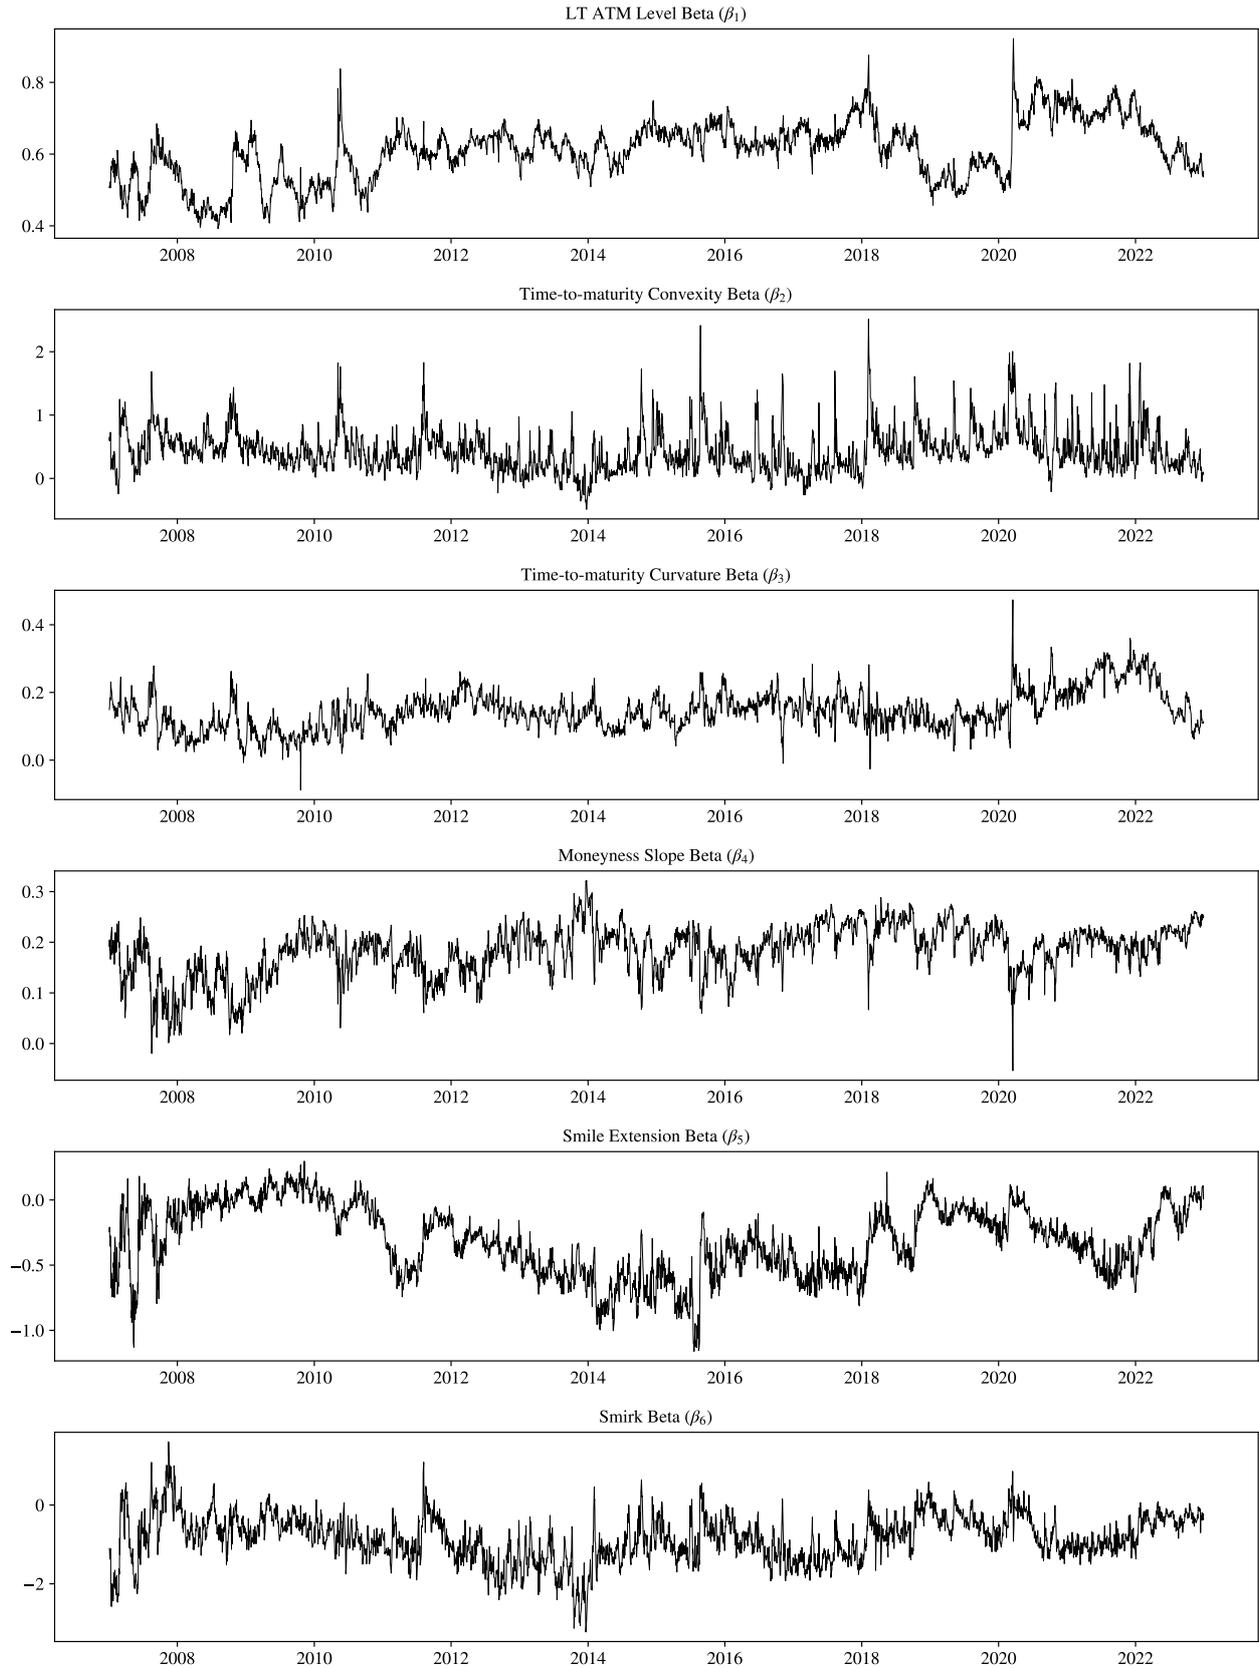
$$S_t = \sqrt{\frac{1}{n-1} \sum_{i=1}^n (\epsilon_{t,M_i,\tau_i} - \bar{\epsilon}_t)^2} \quad (15)$$

is the daily standard deviation. The Equations (14) and (15) are calculated using all IV observations on a specific day. We then recalibrate the model by minimizing the sum of the squared fitting errors.

The time series of the estimated coefficients obtained through the second pass are shown in Figure 4.⁴ The results between the two passes are similar, but the second pass improves coefficient variability. The ATM IV level typically fluctuates between 40% and 70%, but spikes during periods of uncertainty such as the financial crisis of 2008, flash crash of 2010, "Volmageddon"⁵ of 2018, and COVID-19 crisis of 2020. The time-to-maturity convexity coefficient β_2 exhibits increased volatility in the latter half of the sample period, characterized by a higher frequency of spikes. This increased volatility is particularly clear after 2020, a period during which reflects market participants expectations of central bank decisions and economic outcomes in turbulent market conditions. The coefficient of the time-to-maturity curvature appears to remain stable and consistent throughout most of the sample period. During the COVID-19 crisis, its stability was temporarily disrupted, exhibiting large variations on both sides, after which it returned to its previous normal level. The slope coefficient of moneyness is typically positive; however, certain volatility events can cause sudden drops (e.g., the flash crash of 2010, the "Volmageddon" in 2018) and even temporary inversions (e.g., the COVID-19 crisis of 2020). The first part of the dataset shows a lot of variation in the coefficient for smile extension, which is due to the lack of put options. For most of the time, this coefficient stays

⁴The results of the first and second pass processes are presented in the Appendix B.2.

⁵"Volmageddon" refers to the event on February 5, 2018, when a sudden spike in market volatility led to over a 90% loss in value of short volatility exchange-traded products (ETPs).



Daily estimation of the coefficients $\beta_t = (\beta_{t,1}, \beta_{t,2}, \beta_{t,3}, \beta_{t,4}, \beta_{t,5}, \beta_{t,6})$ for the IV model. Each of the factor coefficient time series are illustrated above capturing respectively the long-term ATM level, the moneyness slope, the time-to-maturity convexity, the time-to-maturity curvature, the smile extension and the smirk from January 1, 2007 to December 30, 2022.

Figure 4: Daily coefficients calibration of the IV model

negative. Towards the end of the dataset, there is a clear rise in the value of β_5 , indicating that there are more put options active on the market. The smirk coefficient shows decreased stability in the beginning of the dataset, with significant fluctuations around 2013 and early 2014. However, in the later part of the dataset, there is a notable increase in the coefficient's stability.

3.3 Benchmarking

To evaluate the calibration performance of the PVIVS model, we use a stochastic volatility model (SV) to calibrate the implied volatility surface, comparing the results with the parametric model proposed by Heston (1993). As introduced in Heston (1993), the stochastic volatility framework has gained widespread acceptance in derivative valuation models, particularly for the S&P 500. The benchmark model fundamentally differs from the PVIVS model, as it belongs to a distinct family of models characterized by temporal dynamics in the absence of arbitrage. For a fair comparison, the estimation of the benchmark is performed with daily calibration of the parameters. The risk-neutral dynamic of the Heston model is

$$dF_{t,\tau} = \sqrt{V_t} F_{t,\tau} dW_t^Q \quad (16)$$

$$dV_t = \kappa(\theta - V_t)dt + \sigma_V \sqrt{V_t} dB_t^Q \quad (17)$$

where $d\langle W, B \rangle_t = \rho dt$ capture part of the dependence between the underlying forward price $F_{t,\tau}$ and its instantaneous variance V_t . The parameters are explained in detail in the Appendix B.3. With Heston (1993) framework, the European options prices are available in a quasi-analytical form and can be transformed into implied volatilities by inverting the Black (1976) formula. The model's parameters, including its instantaneous variance, are calibrated daily. A two-pass optimization approach, similar to that used in the PVIVS model, is adopted for daily calibration. The daily calibration aims to determine the best combination of parameters each day to best fit the IV surface. This approach ensures a consistent and fair comparison of the fitting performance with the proposed PVIVS model. In the first pass, all available observations are used to estimate the Heston model parameters. In the second pass, z-scores of the daily fitting errors from the first pass are computed. Observations with absolute z-scores exceeding 3, as well as negative prices, are temporarily set aside. Negative prices that persist after the second pass are set to zero. An analysis of these negative prices is presented in

Appendix B.3.

Rather than calibrating directly on the sum of squared residuals of implied volatilities, the calibration procedure instead minimizes the sum of squared residuals of call prices, scaled by the inverse of their vega. The objective function is

$$\hat{\Theta}_t = \underset{\Theta_t}{\operatorname{argmin}} \sum_{i=1}^n \left(\frac{C_t^{Obs}(M_i, \tau_i) - C_t^{Heston}(M_i, \tau_i, \Theta_t)}{\zeta_t^{Obs}(M_i, \tau_i)} \right)^2 \quad (18)$$

where $C_t^{Obs}(M_i, \tau_i)$ and $C_t^{Heston}(M_i, \tau_i, \Theta_t)$ represent the observed and model call prices for the i -th option, respectively, $\zeta_t^{Obs}(M_i, \tau_i)$ is the vega of the i -th option, Θ_t denotes the set of model parameters and is equal to $\Theta_t = \{V_t, \theta_t, \kappa_t, \rho_t, \sigma_t\}$. With the optimal parameter set $\hat{\Theta}_t$, model prices are retrieved. Put prices are computed using forward put-call parity and the implied volatilities are estimated from the model-derived prices.

$$P_t^{Heston}(K, \tau) = C_t^{Heston}(K, \tau) + D_{t,\tau}(K - F_{t,\tau}), \quad (19)$$

Equation (18) is not a fully equivalent to directly calibrating on implied volatilities. By calibrating on call prices rather than implied volatilities, this approach enhances computational efficiency. It provides only an approximation of IV calibration by normalizing residuals using the vega corresponding to each call price.

3.4 Calibration Performance

To assess the calibration performance of each model over the entire sample, we compute the daily RMSE :

$$RMSE_t = \sqrt{\frac{1}{n} \sum_{i=1}^n (\sigma_t^{Model}(M_i, \tau_i) - \sigma_t^{Obs}(M_i, \tau_i))^2}, \quad (20)$$

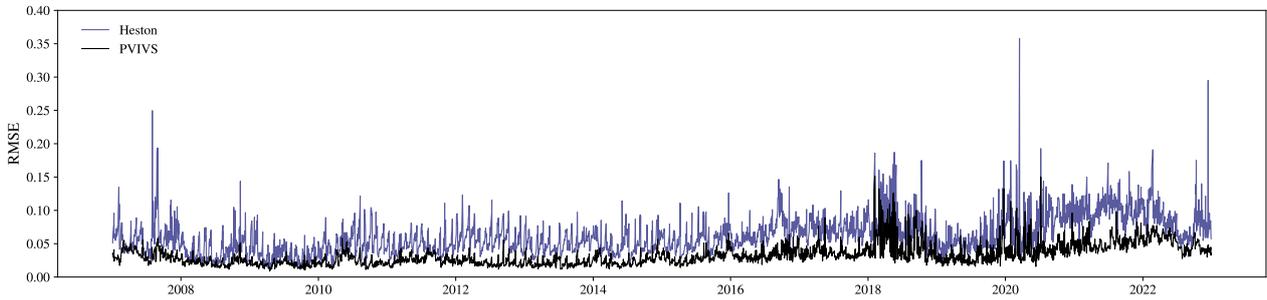
where $\sigma_t^{Obs}(M_i, \tau_i)$ represents the observed IV for the i th option contract on a specific day, and $\sigma_t^{Model}(M_i, \tau_i)$ is the model option IV with respect to its moneyness M_i and its time-to-maturity τ_i .

We also compute the daily absolute relative pricing error (ARPE):

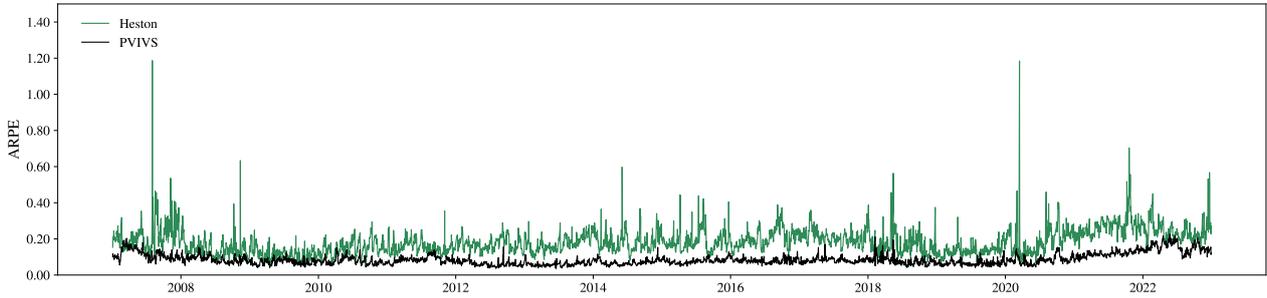
$$ARPE_t = \frac{1}{n} \sum_{i=1}^n \frac{|O_t^{Model}(M_i, \tau_i) - O_t^{Obs}(M_i, \tau_i)|}{O_t^{Obs}(M_i, \tau_i)}, \quad (21)$$

where $O_t^{Obs}(M_i, \tau_i)$ denotes the observed price for the i th option contract with respect to its moneyness M_i and its time-to-maturity τ_i , while $O_t^{Model}(M_i, \tau_i)$ represents the model option price on a specific day. Panel A of Figure 5 exhibits the daily RMSE for the PVIVS model alongside the benchmark. Panel B of Figure 5 presents the daily ARPE for both the PVIVS and Heston model. The PVIVS model demonstrates consistent performance throughout the sample period, characterized by smaller errors in the estimated implied volatility (IV) and pricing, leading to fewer fluctuations in both the RMSE and ARPE metrics. This indicates that the model is robust and stable across diverse market conditions.

Panel A: RMSE of IV



Panel B: ARPE



Daily RMSE and ARPE. Panel A presents the root mean square error (RMSE) on the IV for our model and its benchmark. The Panel B presents the absolute relative pricing error (ARPE) between the Heston model and the PVIVS model.

Figure 5: Performance metrics comparison

Table 4 displays the average RMSE across the whole sample over different subregions of the IV surface. The average RMSE for the PVIVS model across the entire surface exhibits superior performance compared to the benchmark. This is consistent across all buckets of moneyness M and time-to-maturity τ . The average RMSEs corresponding to the PVIVS model are observed to be relatively uniform across all categories of moneyness and time-to-maturity. In contrast, the Heston model demonstrates a decreased fit quality for options with short-term maturity and deep OTM call options. Table 5 illustrates the average ARPE over the entire sample across various subregions of

the IV surface. The results indicate that the PVIVS model outperforms the benchmark in terms of average ARPE across the full surface. This superior performance is consistently observed across different buckets of moneyness M and time-to-maturity τ .

	Calls			Puts		All
	$M \leq -2$	$-2 < M \leq -1$	$-1 < M \leq 0$	$0 < M \leq 1$	$M > 1$	
PVIVS	0.0572	0.0316	0.0244	0.0356	0.1172	0.0407
Heston	0.1083	0.0579	0.0536	0.0780	0.1487	0.0761
No. of contracts	134,501	189,116	251,876	195,917	20,908	792,318
	$\frac{6}{365} < \tau \leq \frac{30}{365}$	$\frac{30}{365} < \tau \leq \frac{90}{365}$	$\frac{90}{365} < \tau \leq \frac{180}{365}$	$\frac{180}{365} < \tau \leq 1$	$\tau > 1$	All
PVIVS	0.0680	0.0351	0.0215	0.0272	0.0537	0.0407
Heston	0.1240	0.0632	0.0520	0.0643	0.0552	0.0761
No. of contracts	170,960	224,609	326,989	68,873	887	792,318

Summary of the average RMSE (ARMSE) for both PVIVS and Heston models across various time to maturity and moneyness buckets. The data spans from January 3, 2007, to December 30, 2022.

Table 4: Average RMSE (ARMSE) of Calibrated IV Surfaces

	Calls			Puts		All
	$M \leq -2$	$-2 < M \leq -1$	$-1 < M \leq 0$	$0 < M \leq 1$	$M > 1$	
PVIVS	0.0276	0.0438	0.0617	0.0619	0.0387	0.0511
Heston	0.0473	0.0925	0.1395	0.1405	0.0735	0.1111
No. of contracts	134,501	189,116	251,876	195,917	20,908	792,318
	$\frac{6}{365} < \tau \leq \frac{30}{365}$	$\frac{30}{365} < \tau \leq \frac{90}{365}$	$\frac{90}{365} < \tau \leq \frac{180}{365}$	$\frac{180}{365} < \tau \leq 1$	$\tau > 1$	All
PVIVS	0.0400	0.0473	0.0499	0.0942	0.2443	0.0511
Heston	0.0698	0.1001	0.1174	0.2128	0.2289	0.1111
No. of contracts	170,960	224,609	326,989	68,873	887	792,318

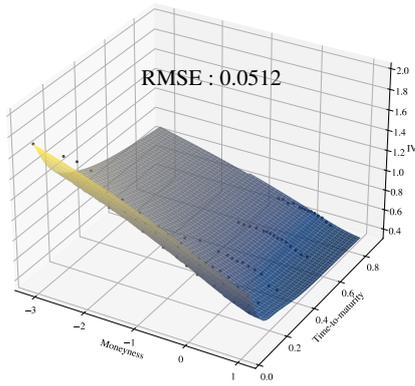
Summary of the average ARPE (AARPE) for both PVIVS and Heston models across various time to maturity and moneyness buckets. The data spans from January 3, 2007, to December 30, 2022.

Table 5: Average ARPE (AARPE) of Calibrated IV Surfaces

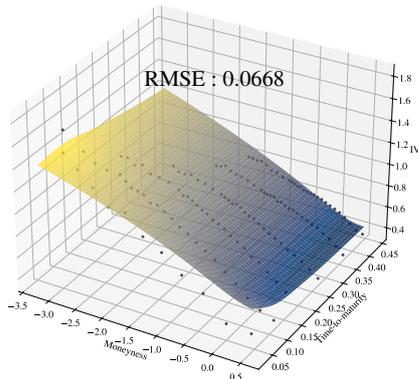
Figure 6 displays the fitting of the PVIVS model and the benchmark to the IV data points for the four dates selected in Figure 2. This demonstrates the smooth extrapolation of the IV surface using the PVIVS model. Specifically, the shape of the IV surface remains consistent with the limited implied volatilities observed in regions of far time-to-maturity and moneyness. Accurately representing the surface shape with deeper OTM options becomes challenging as the benchmark approaches its limits, leading to pricing difficulties for deep OTM options. However, the PVIVS model remains stable in areas with extreme moneyness.

Heston

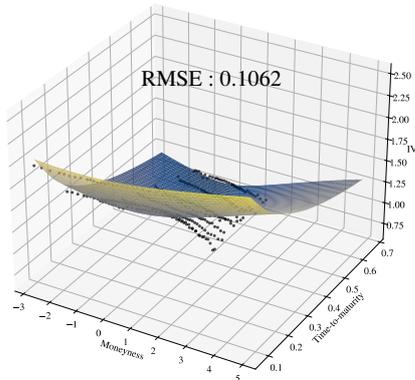
Panel A: January 3, 2007



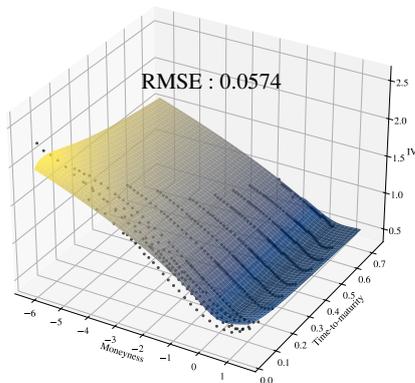
Panel B: January 10, 2014



Panel C: March 17, 2020

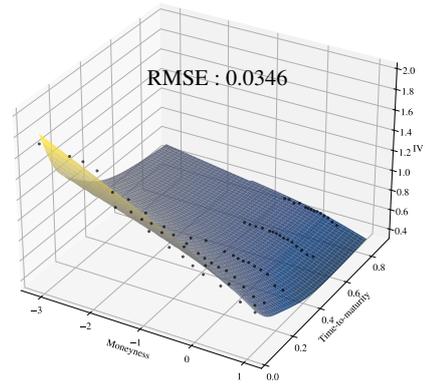


Panel D: December 30, 2022

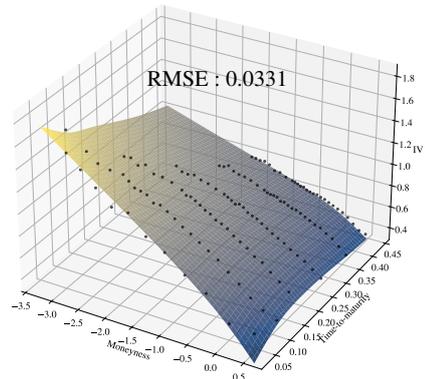


PVIVS

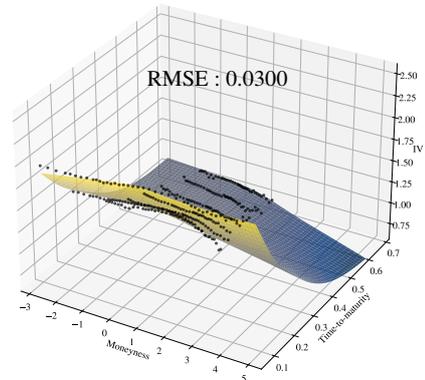
Panel E: January 3, 2007



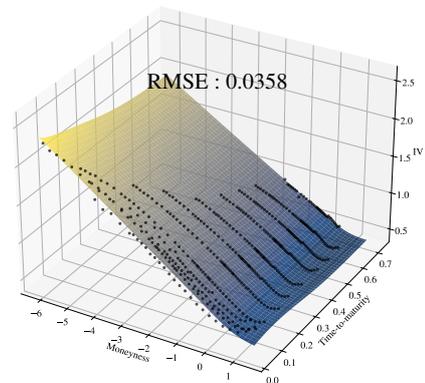
Panel F: January 10, 2014



Panel G: March 17, 2020



Panel H: December 30, 2022



PVIVS model's surfaces fitted on four distinct days with its benchmark. SV parameters are calibrated daily. The first day of the sample is January 3, 2007. January 10, 2014, is a day with low volatility. March 17, 2020, marks the beginning of a crisis, specifically the COVID-19. The last day in the sample is December 30, 2022.

Figure 6: Model Factors
21

4 Arbitrage Opportunities

We examine prices that violate the no-arbitrage constraints on both the calibrated surfaces and sample observations. The goal is to either produce prices that are inconsistent with the no-arbitrage condition or rectify any arbitrage opportunities present in the data. To determine if the PVIVS model generates prices that adhere to the no-arbitrage principle, we follow a similar method used by Carr and Madan (2005). The authors established certain conditions that prevent the existence of static arbitrage opportunities within a given set of call option prices for multiple strikes and maturities on a single underlying asset. They showed that the absence of arbitrage violations in the prices of a group of vertical spreads, butterfly spreads, and calendar spreads is a sufficient condition for the entire set of option prices to be free from arbitrage. In our case, we exclude the test on calendar spread arbitrage. Unlike calendar spread arbitrage, vertical and butterfly spreads have their cash flows on the same day at maturity, making detecting true static arbitrage opportunities possible. The term structure of these instruments can occasionally result in situations where calendar spreads appear to offer credit opportunities. However, these apparent opportunities do not necessarily represent real static arbitrage. Instead, they may be attributed to the varying forward levels of the volatility index across different maturities. Furthermore, the uncertainty surrounding the payoff in calendar spreads of volatility options is a significant consideration. The profit from calendar spreads is not riskless and, therefore, cannot be treated as a static arbitrage opportunity. We view the absence of arbitrage violations in the prices of all vertical spreads and butterfly spreads as a necessary condition, but not a sufficient one, to eliminate all static arbitrages from a group of option price quotes across different strikes and maturities.

For each trading day t , every time-to-maturity τ is characterized by a set of all available strike prices. These strike prices are arranged in ascending order such that, $K_0^{(\tau)} < K_1^{(\tau)} < \dots < K_{n(\tau)}^{(\tau)}$. The term $n(\tau)$ denotes the total number of strike prices associated with the specific maturity τ . The vertical spread, butterfly spread and calendar spread are defined as in Carr and Madan (2005). For the value of the vertical spread $VS_i^{(\tau)}$, $i = 1, \dots, n(\tau)$ can be define as :

$$VS(K_i^{(\tau)}, \tau) = \frac{C(K_{i-1}^{(\tau)}, \tau) - C(K_i^{(\tau)}, \tau)}{K_i^{(\tau)} - K_{i-1}^{(\tau)}} \quad (22)$$

$C(K, \tau)$ denotes the call price for a given time-to-maturity τ and strike price K . The value of $VS_i^{(\tau)}$ needs to be bounded between 0 and 1 to prevent vertical spread arbitrage. The butterfly spread value $BS_i^{(\tau)}, i = 1, \dots, n(\tau) - 1$ is defined as

$$BS(K_i^{(\tau)}, \tau) = C(K_{i-1}^{(\tau)}, \tau) - \frac{K_{i+1}^{(\tau)} - K_{i-1}^{(\tau)}}{K_{i+1}^{(\tau)} - K_i^{(\tau)}} C(K_i^{(\tau)}, \tau) + \frac{K_i^{(\tau)} - K_{i-1}^{(\tau)}}{K_{i+1}^{(\tau)} - K_i^{(\tau)}} C(K_{i+1}^{(\tau)}, \tau) \quad (23)$$

and needs to be positive to ensure there's no butterfly spread arbitrage opportunity.

For observed data, we incorporate both mid-prices and the bid-ask prices of the observed data. Furthermore, both bid and ask prices provide a better approach to determining the spread when working with option quotes. This approach aligns more closely with realistic conditions, as using mid prices could potentially identify unattainable arbitrage opportunities due to the limitations of trading within the bid-ask range. Specifically, we use the ask price when buying a call option and the bid price when selling a call option. For calibrated surfaces, we utilize the prices generated by the PIVS model with the same characteristics as the observed quotes to ensure a fair and practical assessment of potential arbitrage opportunities. In addition, we only use call option prices following the methodology outlined in Carr and Madan (2005). We convert the put options from the PIVS model into calls using the put-call parity. The observed call options with matching dates, expiration, and strike prices to those in the sample during the second pass, excluding outliers identified in the first pass, are retained.

The total number of vertical spread arbitrage opportunities identified in the entire dataset and on fitted surfaces is presented in Table 6. These numbers represent the total across all dates in the dataset.

		Observed mid prices		Observed bid-ask prices		Fitted surfaces		Total checks
		Arbitrage detected	% Arbitrage detected	Arbitrage detected	% Arbitrage detected	Arbitrage detected	% Arbitrage detected	
<hr/>								
6/365 < τ ≤ 30/365	$M \leq -2$	1,739	3.14	1	0.00	0	0.00	55,317
	$-2 < M \leq -1$	622	1.78	0	0.00	0	0.00	34,883
	$-1 < M \leq 0$	229	0.63	0	0.00	0	0.00	36,066
	$0 < M \leq 1$	106	0.35	0	0.00	0	0.00	29,961
	$M > 1$	11	0.17	0	0.00	0	0.00	6,553
30/365 < τ ≤ 90/365	$M \leq -2$	178	0.37	0	0.00	0	0.00	48,412
	$-2 < M \leq -1$	132	0.24	1	0.00	0	0.00	54,438
	$-1 < M \leq 0$	74	0.12	0	0.00	0	0.00	62,894
	$0 < M \leq 1$	20	0.04	0	0.00	0	0.00	46,546
	$M > 1$	0	0.00	0	0.00	0	0.00	3,788
90/365 < τ ≤ 180/365	$M \leq -2$	31	0.11	0	0.00	0	0.00	27,066
	$-2 < M \leq -1$	205	0.25	0	0.00	0	0.00	83,582
	$-1 < M \leq 0$	18	0.01	1	0.00	0	0.00	126,167
	$0 < M \leq 1$	6	0.01	0	0.00	0	0.00	77,254
	$M > 1$	0	0.00	0	0.00	0	0.00	1,588
180/365 < τ ≤ 1	$M \leq -2$	35	0.94	0	0.00	0	0.00	3,706
	$-2 < M \leq -1$	97	0.60	0	0.00	0	0.00	16,213
	$-1 < M \leq 0$	31	0.12	0	0.00	0	0.00	26,197
	$0 < M \leq 1$	4	0.02	0	0.00	0	0.00	20,398
	$M > 1$	0	0.00	0	0.00	0	0.00	235
$\tau > 1$	$M \leq -2$	0	0.00	0	0.00	0	0.00	0
	$-2 < M \leq -1$	0	0.00	0	0.00	0	0.00	0
	$-1 < M \leq 0$	0	0.00	0	0.00	0	0.00	551
	$0 < M \leq 1$	0	0.00	0	0.00	0	0.00	273
	$M > 1$	0	0.00	0	0.00	0	0.00	0

Summary statistics on violations per time-to-maturity and moneyness buckets of no-arbitrage constraints on vertical spreads based on Carr and Madan (2005) methodology. The numbers and proportions of violations, aggregated across all sample dates, are reported for both the data sample and the fitted surfaces of the PVIVS model.

Table 6: Detected vertical arbitrage opportunities

The cumulative count of butterfly spread arbitrage opportunities detected throughout the dataset and on the fitted surfaces is also shown in Table 7. The table represents the aggregate across all dates included in the dataset.

		Observed mid prices		Observed bid-ask prices		Fitted surfaces		Total checks
		Arbitrage detected	% Arbitrage detected	Arbitrage detected	% Arbitrage detected	Arbitrage detected	% Arbitrage detected	
6/365 < τ ≤ 30/365								
	M ≤ -2	12,742	26.90	1	0.00	0	0.00	47,367
	-2 < M ≤ -1	8,697	25.10	1	0.00	0	0.00	34,656
	-1 < M ≤ 0	5,806	16.10	0	0.00	0	0.00	36,064
	0 < M ≤ 1	3,501	11.69	1	0.00	0	0.00	29,960
	M > 1	1,013	15.46	1	0.01	0	0.00	6,553
30/365 < τ ≤ 90/365								
	M ≤ -2	10,169	24.75	0	0.00	0	0.00	41,080
	-2 < M ≤ -1	11,629	21.84	2	0.00	0	0.00	53,253
	-1 < M ≤ 0	8,883	14.13	1	0.00	0	0.00	62,880
	0 < M ≤ 1	3,868	8.31	0	0.00	5	0.01	46,546
	M > 1	706	18.64	0	0.00	0	0.00	3,788
90/365 < τ ≤ 180/365								
	M ≤ -2	5,024	23.71	0	0.00	0	0.00	21,190
	-2 < M ≤ -1	16,578	21.14	0	0.00	0	0.00	78,408
	-1 < M ≤ 0	20,036	15.92	1	0.00	0	0.00	125,885
	0 < M ≤ 1	7,085	9.17	0	0.00	272	0.35	77,254
	M > 1	353	22.23	0	0.00	0	0.00	1,588
180/365 < τ ≤ 1								
	M ≤ -2	669	25.74	0	0.00	0	0.00	2,599
	-2 < M ≤ -1	3,077	19.62	0	0.00	0	0.00	15,684
	-1 < M ≤ 0	4,066	15.82	0	0.00	0	0.00	25,709
	0 < M ≤ 1	3,139	15.39	0	0.00	731	3.58	20,398
	M > 1	69	29.36	0	0.00	0	0.00	235
$\tau > 1$								
	M ≤ -2	0	0.00	0	0.00	0	0.00	0
	-2 < M ≤ -1	0	0.00	0	0.00	0	0.00	0
	-1 < M ≤ 0	74	15.16	0	0.00	0	0.00	488
	0 < M ≤ 1	86	31.50	0	0.00	63	23.08	273
	M > 1	0	0.00	0	0.00	0	0.00	0

Summary statistics on violations per time-to-maturity and moneyness buckets of no-arbitrage constraints on butterfly spreads based on Carr and Madan (2005) methodology. The numbers and proportions of violations, aggregated across all dates of the sample, are reported for both the data sample and the fitted surfaces of the PVIVS model. The static arbitrage opportunities on the mid prices all fall within the bid-ask range and therefore remain theoretical arbitrage. The fitted surfaces also admit static arbitrage but to a lesser extent than the mid prices.

Table 7: Detected butterfly arbitrage opportunities

The 1071 butterfly spread opportunities detected on the fitted surface are primarily attributable to converting put options into call options using the put-call parity with a moneyness close to zero. The estimated forward price and the calculated mid price from the bid and ask prices influence the moneyness of the options, thereby affecting their prices. The summary statistics of these detected opportunities are presented in the Table 8. The observations indicate a close proximity to a moneyness of 0, with an average time to maturity exceeding 6 months. It can be inferred that this occurrence impacts at-the-money put options with a longer time-to-maturity.

	Moneyiness	Time-to-maturity
Min	0.0000	0.2274
Mean	0.0272	0.6273
Max	0.0802	1.2493

Summary statistics of identified butterfly arbitrage opportunities for the fitted surfaces of the PVIVS model. The statistics on violations are presented based on their moneyiness and time-to-maturity and are aggregated across all sample dates.

Table 8: Summary of detected butterfly arbitrage opportunities

The arbitrage opportunities identified on the fitted surfaces are comparable to those that could be based on the bid-ask prices. This suggests that both the PVIVS model and the observed data, which fall within the actual trading range, limit the possibility of static arbitrage opportunities.

5 Model Applications

Having a complete surface with implied volatilities available over a wide range of moneyiness and maturity has several practical applications. This section presents one of the applications, which is its use in derivatives pricing.

The PVIVS model can be used in the pricing of complex financial derivatives. The model has the flexibility of directly interpolate or extrapolate IVs, and extracting density of the VIX risk-neutral density. These techniques are necessary to price a broad choice of derivatives. Given the extensive range of strike prices available on the VIX, deep OTM contracts frequently exhibit a lack of liquidity. Pricing these options using mark-to-market procedures can be complex. As depicted in Figure 2, the measure of moneyiness can reach extreme values, becoming increasingly relevant in the current time period. For such contracts, the pricing process can be simplified using the PVIVS model, as the IV is solely a function of the option's moneyiness and its time-to-maturity.

One important use of option pricing across a range of strike prices involves calculating the risk-neutral density of the underlying asset. Breeden and Litzenberger (1978) established that the risk-neutral probability density function, which reflects the price of the underlying asset, follows Equation (10). Appendix C.1 provides the proof demonstrating that the integral of that $\int_0^\infty g_\tau(K)dK = 1$. This condition ensures that the total probability across all possible values of K is equal to 1, which is a fundamental property of probability density functions.

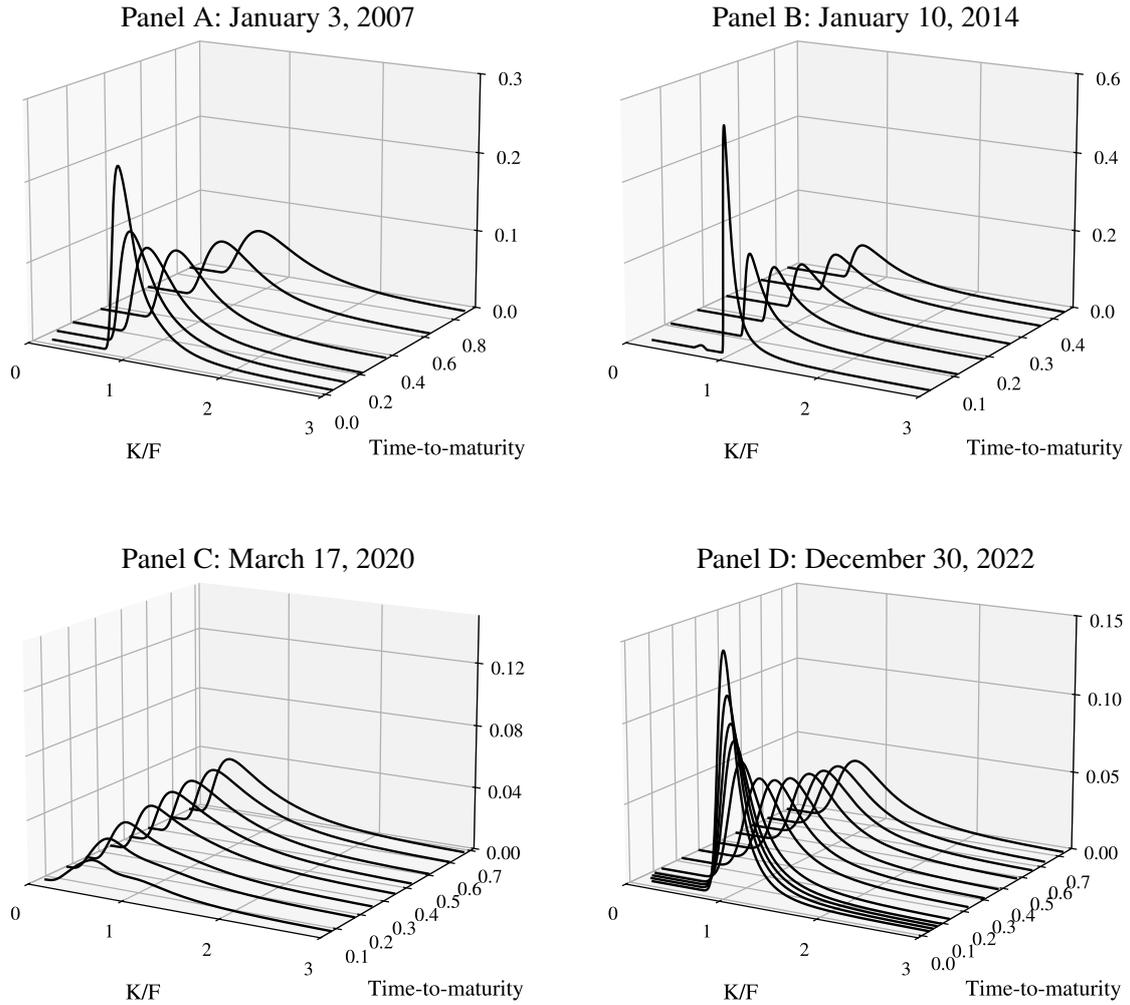
In the framework of the PVIVS model, the risk-neutral density function for a maturity τ is :

$$e^{r\tau} \frac{\partial^2 c}{\partial K^2} = \frac{F_{t,\tau}}{K^2} \varphi(d_1) \left(\frac{1}{\sqrt{\tau}} \frac{\partial d_1}{\partial M} - \frac{1}{\sqrt{\tau}} d_1 \frac{\partial d_1}{\partial M} \frac{\partial \sigma}{\partial M} + \frac{1}{\sqrt{\tau}} \frac{\partial^2 \sigma}{\partial M^2} \right) \quad (24)$$

where $\varphi(z) = \sqrt{\frac{1}{2\pi}} \exp\left(-\frac{z^2}{2}\right)$ is the density function of a standard normal random variable and

$$\begin{aligned} \frac{\partial d_1}{\partial M} &= \frac{1}{\sigma} - \left(\frac{M}{\sigma^2} - \frac{1}{2\sqrt{\tau}} \right) \frac{\partial \sigma}{\partial M}, \\ \frac{\partial \sigma}{\partial M} &= -\beta_4 \mathbb{1}_{M \leq 0} + \beta_4 \left(\left(\frac{e^{2M} - 1}{e^{2M} + 1} \right)^2 - 1 \right) \mathbb{1}_{M > 0} + \beta_5 \left(\frac{2e^{2M-2}}{e^{2M-2} + 2} - \frac{2e^{2M-2}(e^{2M-2} - 2)}{(e^{2M-2} + 2)^2} \right) \\ &\quad - \beta_6 \left(\frac{e^{\frac{M}{M_{Max}}} \log\left(\frac{\tau}{T_{Conv}} + \epsilon\right)}{M_{Max}} \right) \\ \frac{\partial^2 \sigma}{\partial M^2} &= \beta_4 8e^{2M} \frac{e^{2M} - 1}{(e^{2M} - 1)^3} \mathbb{1}_{M > 0} + \beta_5 \left((e^{2M-2} - 2) \left(\frac{8e^{4M-4}}{(e^{2M-2} + 2)^3} - \frac{4e^{2M-2}}{(e^{2M-2} + 2)^2} \right) + \frac{4e^{2M-2}}{e^{2M-2} + 2} - \frac{8e^{4M-4}}{(e^{2M-2} + 2)^2} \right) \\ &\quad - \beta_6 \left(\frac{e^{\frac{M}{M_{Max}}} \log\left(\frac{\tau}{T_{Conv}} + \epsilon\right)}{M_{Max}^2} \right) \end{aligned} \quad (25)$$

For the identical days depicted in Figure 2, we illustrate a series of risk-neutral densities $g_\tau(\nu)$, each associated with a distinct maturity τ in Figure 7. Volatility tends to rise as time-to-maturity increases. However, during the COVID-19 crisis, the volatility decreases with time-to-maturity. This effect is more pronounced for short-term maturities. When examining the implied densities from the PVIVS model across four distinct dates, a positive skewness was observed. This skewness is especially noticeable in the March 2020 crisis.



Risk-neutral densities implied by the PVIVS model with the function: $g_{\tau}(v)$. The first day of the sample is January 3, 2007. January 10, 2014, is a day with low volatility. March 17, 2020, marks the beginning of COVID-19. The last day in the sample is December 30, 2022. Instead of using the moneyness defined in Equation (7), the densities are presented with the moneyness K/F .

Figure 7: Risk-neutral densities surfaces for different dates

6 Conclusion

This paper presents a new parametric model for the implied volatility surface for VIX options. The model's design emphasizes computational efficiency and ease of estimation. The model captures the essential characteristics of VIX option surfaces while maintaining simplicity and interpretability by incorporating key factors such as the long-term level, time-to-maturity convexity, time-to-maturity curvature, moneyness slope, smile extension, and smirk. These factors are tailored to reflect empirical characteristics observed in the VIX IV surfaces.

The PVIVS model can generate IV surfaces that are twice continuously differentiable, ensuring consistency with no-arbitrage principles. Moreover, the factors are constructed to exhibit stable asymptotic properties, making the model suitable for extrapolating beyond the observed ranges of maturities and moneyness levels. This ability to extrapolate is particularly beneficial for pricing illiquid derivatives, where quoted prices are often unavailable.

The model's performance is evaluated against the Heston (1993) model, using historical VIX option data from the OptionMetrics database. Results show that the proposed model outperforms its benchmark, particularly in its ability to handle deep out-of-the-money options and shorter maturity options. This calibration performance makes the model particularly suited for market environments where robust extrapolation is required. Additionally, the model ensures a smooth surface fit and limits static arbitrage opportunities. The constructed surface passes arbitrage detection tests using vertical and butterfly spreads, which are used to identify static arbitrage opportunities. Furthermore, the applicability of the PVIVS model extends beyond VIX European option pricing to potential applications in other derivative instruments requiring risk-neutral density estimations.

Further work could expand on the development of a dynamic model, create hedging strategies specifically tailored for VIX options, and enable the pricing of more complex products, such as forward start variance swaps. Additionally, future research could also explore the use of risk factors based on the PVIVS model's coefficients to enhance both pricing and hedging strategies.

References

- Augustin, P., I.-H. Cheng, and L. Van den Bergen. 2021. Volmageddon and the failure of short volatility products. *Financial Analysts Journal* 77:35–51.
- Bardgett, C., E. Gourier, and M. Leippold. 2019. Inferring volatility dynamics and risk premia from the SP 500 and VIX markets. *Journal of Financial Economics* 131:593–618.
- Black, F. 1976. The pricing of commodity contracts. *Journal of Financial Economics* 3 (1-2):167–179.
- Breedon, D. T., and R. H. Litzenberger. 1978. Prices of State-Contingent Claims Implicit in Option Prices. *Journal of Business* 51:621–651.
- Carr, P., and D. Madan. 2005. A note on sufficient conditions for no arbitrage. *Finance Research Letters* 2:125–130.
- CBOE. 2023. Volatility Index Methodology: Cboe Volatility Index. *White Paper* .
- Filipovic, D., M. Pelger, and Y. Ye. 2022. Stripping the Discount Curve - a Robust Machine Learning Approach. *Swiss Finance Institute Research Paper No. 22-24*.
- Francois, P., R. Galarneau-Vincent, G. Gauthier, and F. Godin. 2022. Venturing into Uncharted Territory: An Extensible Parametric Implied Volatility Surface Model. *Journal of Futures Markets* 42:1912–1940.
- Heston, S. 1993. A closed-form solution for options with stochastic volatility with applications to bond and currency options. *Review of Financial Studies* 6:327.
- Lian, G.-H., and S.-P. Zhu. 2013. Pricing VIX options with stochastic volatility and random jumps. *Decisions in Economics and Finance* 36:71–88.
- Mencia, J., and E. Sentana. 2013. Valuation of VIX Derivatives. *Journal of Financial Economics* 108:367–391.
- Pacati, C., G. Pompa, and R. Renò. 2018. Smiling Twice: The Heston ++ Model. *Journal of Banking and Finance* 96:185–206.

Park, Y.-H. 2016. The Effects of Asymmetric Volatility and Jumps on the Pricing of VIX Derivatives. *Journal of Econometrics* 192:313–328.

Yuan, P. 2022. Time-Varying Skew in VIX Derivatives Pricing. *Management Science* 68:7761–7791.

A Data Cleaning

A.1 Cleaning Criteria

Some exclusions occurred on three specific days in 2009: January 21, October 20, and December 2. These exclusions coincided with periods of high market volatility, largely due to the ongoing recovery from the 2008 financial crisis. Although there were no clear, specific events on these days that could have triggered such high uncertainty in the VIX options, the market conditions at the time were characterized by heightened uncertainty. In addition to the instances in 2009, there were also significant exclusions on February 5 and 6, 2018, during the "Volmageddon" event. This period of extreme market volatility, driven by short volatility strategies and the rapid growth of Exchange Traded Products (ETPs), resulted in a one-day loss of over 90% in the value of short-volatility ETPs. Augustin et al. (2021) outlined that this event was fueled by a combination of factors, including the need for ETPs to rebalance their hedges and leverage ratios in a highly volatile and concentrated market. The rebalancing caused a feedback loop, which drove futures prices up and resulted in major losses for investors who had bet on low volatility.

There are 5 days in our sample that contain an important number of options with large spreads. These specific days are typically marked by a lack of market liquidity, affecting a substantial number of contracts and consequently resulting in an illiquid options surface. This exclusion led to noticeable decreases in the retained data. Table 9, along with the count of removed observations due to their large spread.

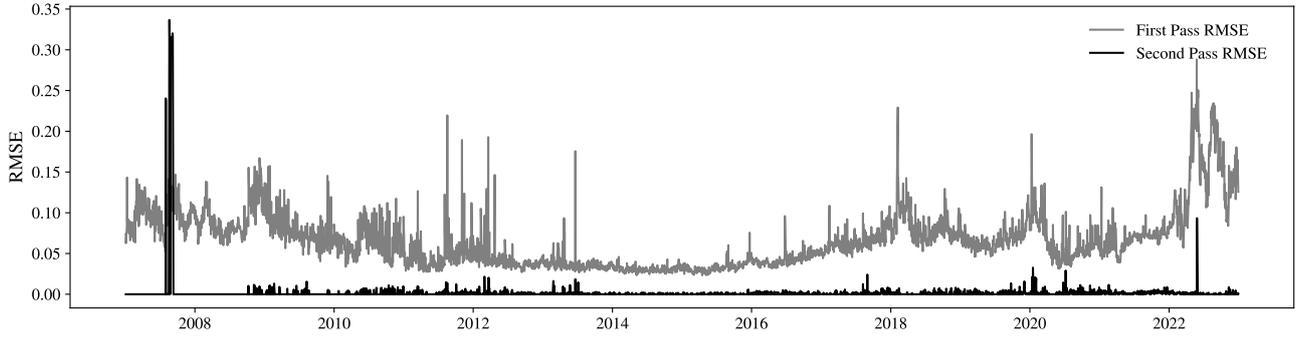
Date	Removed Observations	Removed Proportion(%)	Total Observations
2009-01-21	121	61.42	197
2009-10-20	140	43.75	320
2009-12-02	144	47.37	304
2018-02-05	451	65.56	688
2018-02-06	473	63.92	740

Summary of the days when a significant number of observations were excluded due to their large spread. The columns represent the following: Date: The specific day on which the large exclusions occurred. Removed Observations: The number of observations that were removed on the given day due to their large spread. Proportion Removed(%): The percentage of the total observations that were removed on the given day. Total Observations: The total number of observations on the given day.

Table 9: Days of data exclusion and removed observations due to large spread

A.2 Forward Estimation

Figure 8 represents the RMSE of Equation (3). The grey line represents the first pass, which incorporates outliers. The black line represents the second pass that excludes these outliers. Outliers present during the first pass can often skew the RMSE of the forward estimation. The spikes observed around 2007 and 2022 in the second pass RMSE are attributed to the concentration of values retained from the first pass. The values estimated using OLS method can deviate significantly from the median implied forward, which increases the residuals for many implied forward values on those days.



RMSE for both passes of the forward optimization from Equation (3). The gray line represents the first pass, which includes outliers that skew the mean on certain days. The black line represents the second pass, which excludes outliers.

Figure 8: Daily RMSE for two-pass forward optimization

B Model Specification

B.1 Proof of Theorem for PVIVS Surface Function

Theorem 1. Let $\sigma(M, \tau) : \mathbb{R} \times \mathbb{R}^+ \rightarrow \mathbb{R}^+$ be the PVIVS surface function defined as:

$$\begin{aligned} \sigma(M, \tau) = & \beta_1 + \beta_2 e^{\frac{-\tau}{T_{Conv}}} + \beta_3 e^{\frac{\tau}{T_{Hump}}(1 - \frac{\tau}{T_{Hump}})} + \beta_4 \left(-M \mathbb{1}_{M \leq 0} - \frac{e^{2M} - 1}{e^{2M} + 1} \mathbb{1}_{M > 0} \right) \\ & + \beta_5 \left(\frac{e^{2(M-1)} - 2}{e^{2(M-1)} + 2} - \frac{e^{-2} - 2}{e^{-2} + 2} \right) + \beta_6 \left(1 - e^{\frac{M}{M_{Max}}} \right) \log \left(\frac{\tau}{T_{Conv}} + \alpha \right) \end{aligned}$$

where $M \in \mathbb{R}$ denotes the moneyness, $\tau \in \mathbb{R}^+$ the time to maturity, and $\beta_i, T_{Conv}, T_{Hump}, M_{Max}, \alpha \in \mathbb{R}$ are the model's parameters. Then:

1. The limiting behavior of $\sigma(M, \tau)$ as M approaches zero is given by:

$$\lim_{M \rightarrow 0} \sigma(M, \tau) = \beta_1 + \beta_2 e^{-\tau/T_{Conv}} + \beta_3 e^{\tau/T_{Hump}(1 - \tau/T_{Hump})}$$

2. For at-the-money options ($M = 0$), the limiting behavior of $\sigma(0, \tau)$ as τ approaches zero and

infinity is given by:

$$\lim_{\tau \rightarrow 0^+} \sigma(0, \tau) = \beta_1 + \beta_2 + \beta_3$$

$$\lim_{\tau \rightarrow \infty} \sigma(0, \tau) = \beta_1$$

3. $\sigma(M, \tau)$ is continuous for all $M \in \mathbb{R}$, including at $M = 0$.
4. $\sigma(M, \tau)$ is twice continuously differentiable with respect to M for all M .

Proof. Part A: The asymptotic behavior of each term in the PVIVS surface function as $M \rightarrow 0$ are analyzed:

1. The first three terms, β_1 , $\beta_2 e^{-\tau/T_{\text{Conv}}}$, and $\beta_3 e^{\tau/T_{\text{Hump}}(1-\tau/T_{\text{Hump}})}$, are independent of M and thus remain invariant as $M \rightarrow 0$.
2. For the moneyness slope term:

$$\begin{aligned} \lim_{M \rightarrow 0} \beta_4 \left(-M \mathbb{1}_{M \leq 0} - \frac{e^{2M} - 1}{e^{2M} + 1} \mathbb{1}_{M > 0} \right) &= \beta_4 \left(0 - \lim_{M \rightarrow 0^+} \frac{e^{2M} - 1}{e^{2M} + 1} \right) \\ &= \beta_4 \left(0 - \frac{0}{2} \right) = 0 \end{aligned}$$

3. For the smile extension term:

$$\begin{aligned} \lim_{M \rightarrow 0} \beta_5 \left(\frac{e^{2(M-1)} - 2}{e^{2(M-1)} + 2} - \frac{e^{-2} - 2}{e^{-2} + 2} \right) \\ = \beta_5 \left(\frac{e^{-2} - 2}{e^{-2} + 2} - \frac{e^{-2} - 2}{e^{-2} + 2} \right) = 0 \end{aligned}$$

4. For the smirk term:

$$\begin{aligned} \lim_{M \rightarrow 0} \beta_6 \left(1 - e^{M/M_{\text{Max}}} \right) \log \left(\frac{\tau}{T_{\text{Conv}}} + \alpha \right) \\ = \beta_6 (1 - 1) \log \left(\frac{\tau}{T_{\text{Conv}}} + \alpha \right) = 0 \end{aligned}$$

Therefore, by the sum rule of limits, it is concluded that

$$\lim_{M \rightarrow 0} \sigma(M, \tau) = \beta_1 + \beta_2 e^{-\tau/T_{\text{Conv}}} + \beta_3 e^{\tau/T_{\text{Hump}}(1-\tau/T_{\text{Hump}})}$$

This limit represents the at-the-money (ATM) implied volatility, as $M = 0$ corresponds to the ATM option. \square

Proof. Part B: When $M = 0$, the function simplifies to:

$$\sigma(0, \tau) = \beta_1 + \beta_2 e^{-\tau/T_{\text{Conv}}} + \beta_3 e^{\tau/T_{\text{Hump}}(1-\tau/T_{\text{Hump}})}$$

For the limits as τ approaches zero and infinity:

1. As $\tau \rightarrow 0^+$:

$$\begin{aligned}\lim_{\tau \rightarrow 0^+} \sigma(0, \tau) &= \beta_1 + \beta_2 \lim_{\tau \rightarrow 0^+} e^{-\tau/T_{\text{Conv}}} + \beta_3 \lim_{\tau \rightarrow 0^+} e^{\tau/T_{\text{Hump}}(1-\tau/T_{\text{Hump}})} \\ &= \beta_1 + \beta_2 \cdot 1 + \beta_3 \cdot 1 \\ &= \beta_1 + \beta_2 + \beta_3\end{aligned}$$

2. As $\tau \rightarrow \infty$:

$$\begin{aligned}\lim_{\tau \rightarrow \infty} \sigma(0, \tau) &= \beta_1 + \beta_2 \lim_{\tau \rightarrow \infty} e^{-\tau/T_{\text{Conv}}} + \beta_3 \lim_{\tau \rightarrow \infty} e^{\tau/T_{\text{Hump}}(1-\tau/T_{\text{Hump}})} \\ &= \beta_1 + \beta_2 \cdot 0 + \beta_3 \cdot 0 \\ &= \beta_1\end{aligned}$$

For the second term, $\lim_{\tau \rightarrow \infty} e^{-\tau/T_{\text{Conv}}} = 0$ since T_{Conv} is positive. The third term can also be expressed as

$$\frac{\tau}{T_{\text{Hump}}} \left(1 - \frac{\tau}{T_{\text{Hump}}} \right) = \frac{\tau}{T_{\text{Hump}}} - \left(\frac{\tau}{T_{\text{Hump}}} \right)^2$$

As $\tau \rightarrow \infty$, this expression approaches $-\infty$. Consequently, $e^{\tau/T_{\text{Hump}}(1-\tau/T_{\text{Hump}})} \rightarrow 0$.

□

Proof. Part C: For each term of the PVIVS surface function $\sigma(M, \tau)$,

1. The first three terms are continuous functions of τ and independent of M , so they are continuous for all M .
2. The moneyness slope term, $\beta_4(-M\mathbb{1}_{M \leq 0} - \frac{e^{2M}-1}{e^{2M}+1}\mathbb{1}_{M > 0})$, is continuous at $M = 0$ since

$$\begin{aligned}\lim_{M \rightarrow 0^-} (-M) &= 0 \\ \lim_{M \rightarrow 0^+} \left(-\frac{e^{2M}-1}{e^{2M}+1} \right) &= 0\end{aligned}$$

And both parts are continuous on their respective domains $(-\infty, 0]$ and $[0, \infty)$.

3. The smile extension term, $\beta_5(\frac{e^{2(M-1)}-2}{e^{2(M-1)}+2} - \frac{e^{-2}-2}{e^{-2}+2})$, is a composition of continuous functions and thus continuous for all $M \in \mathbb{R}$.
4. The smirk term, $\beta_6(1 - e^{\frac{M}{M_{\text{Max}}}}) \log(\frac{\tau}{T_{\text{Conv}}} + \alpha)$, is also a composition of continuous functions and thus continuous for all $M \in \mathbb{R}$.

As the sum of continuous functions is continuous, $\sigma(M, \tau)$ is continuous for all $M \in \mathbb{R}$, including at $M = 0$. □

Proof. Part D: Taking the two derivatives of $\sigma(M, \tau)$ with respect to M for each term,

1. The first three terms are independent of M , so their first and second derivatives with respect to M are zero and thus continuous for all $M \in \mathbb{R}$.

2. For the moneyness slope term, $f(M) = \beta_4(-M\mathbb{1}_{M \leq 0} - \frac{e^{2M}-1}{e^{2M}+1}\mathbb{1}_{M > 0})$: The first derivative is

$$f'(M) = \begin{cases} -\beta_4 & \text{if } M < 0 \\ \beta_4 \left(\left(\frac{e^{2M}-1}{e^{2M}+1} \right)^2 - 1 \right) & \text{if } M > 0 \end{cases}$$

At $M = 0$, both left and right limits of $f'(M)$ equal $-\beta_4$, therefore $f'(M)$ is continuous.

$$f'(0) = \begin{cases} -\beta_4 & \text{if } M < 0 \\ -\beta_4 & \text{if } M > 0 \end{cases}$$

The second derivative is

$$f''(M) = \begin{cases} 0 & \text{if } M < 0 \\ \beta_4 \frac{8e^{2M}(e^{2M}-1)}{(e^{2M}+1)^3} & \text{if } M > 0 \end{cases}$$

At $M = 0$, both left and right limits of $f''(M)$ equal 0. Therefore, $f''(M)$ is continuous.

$$f''(0) = \begin{cases} 0 & \text{if } M < 0 \\ 0 & \text{if } M > 0 \end{cases}$$

3. For the smile extension term, $g(M) = \beta_5 \left(\frac{e^{2(M-1)}-2}{e^{2(M-1)}+2} - \frac{e^{-2}-2}{e^{-2}+2} \right)$, is a composition of infinitely differentiable functions. Therefore, its first and second derivatives are continuous for all $M \in \mathbb{R}$.

4. For the smirk term, $h(M) = \beta_6 \left(1 - e^{\frac{M}{M_{\text{Max}}}} \right) \log \left(\frac{\tau}{T_{\text{Conv}}} + \alpha \right)$, is also a composition of infinitely differentiable functions. Its first and second derivatives are continuous for all $M \in \mathbb{R}$.

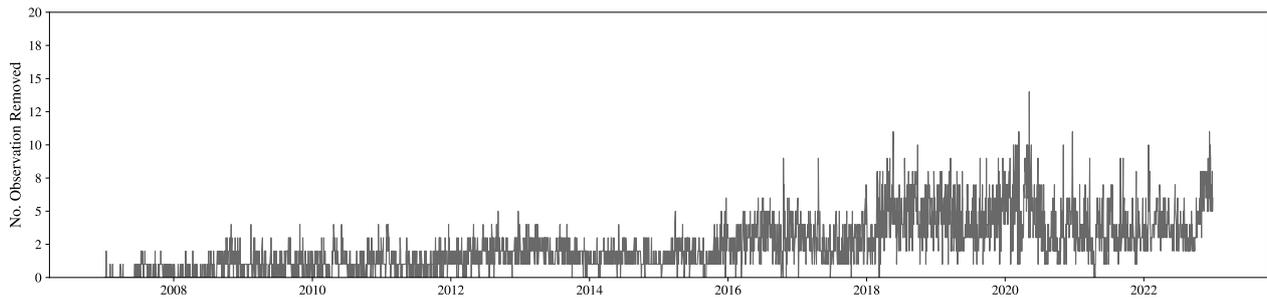
As the sum of functions with continuous first and second derivatives is itself a function with continuous first and second derivatives, the first two derivatives of the implied volatility surface function $\sigma(M, \tau)$ with respect to M are continuous for all $M \in \mathbb{R}$. \square

B.2 Least Squares Optimization

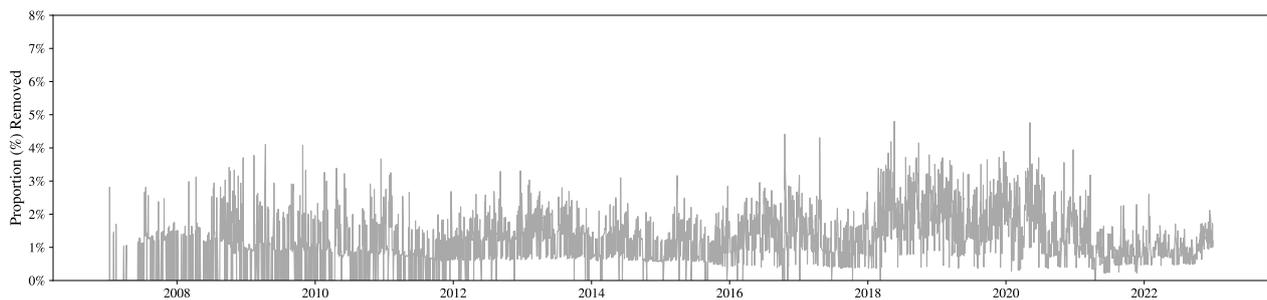
The two-pass least squares optimization reduces coefficient variability within the model. This is achieved by excluding a minor proportion of outliers, which never exceeds 5% of the total observations on a given day. This exclusion procedure enhances the model's robustness against outliers, thereby improving its overall performance. Figure 9 shows the number and proportion of outliers for each day in the sample.

Furthermore, the two-pass optimization reduces the variability of the coefficients in the model. Figure 10 displays the coefficients of Model (11) for its initial pass in red and the second pass in black. Since the model is a function of two variables: (i) moneyness and (ii) time-to-maturity, the second pass shows smaller variation in the coefficients capturing the moneyness, while the impact on the coefficients for time-to-maturity was less pronounced.

Panel A: Daily observations excluded for fitting errors with $|z\text{-score}| > 3$



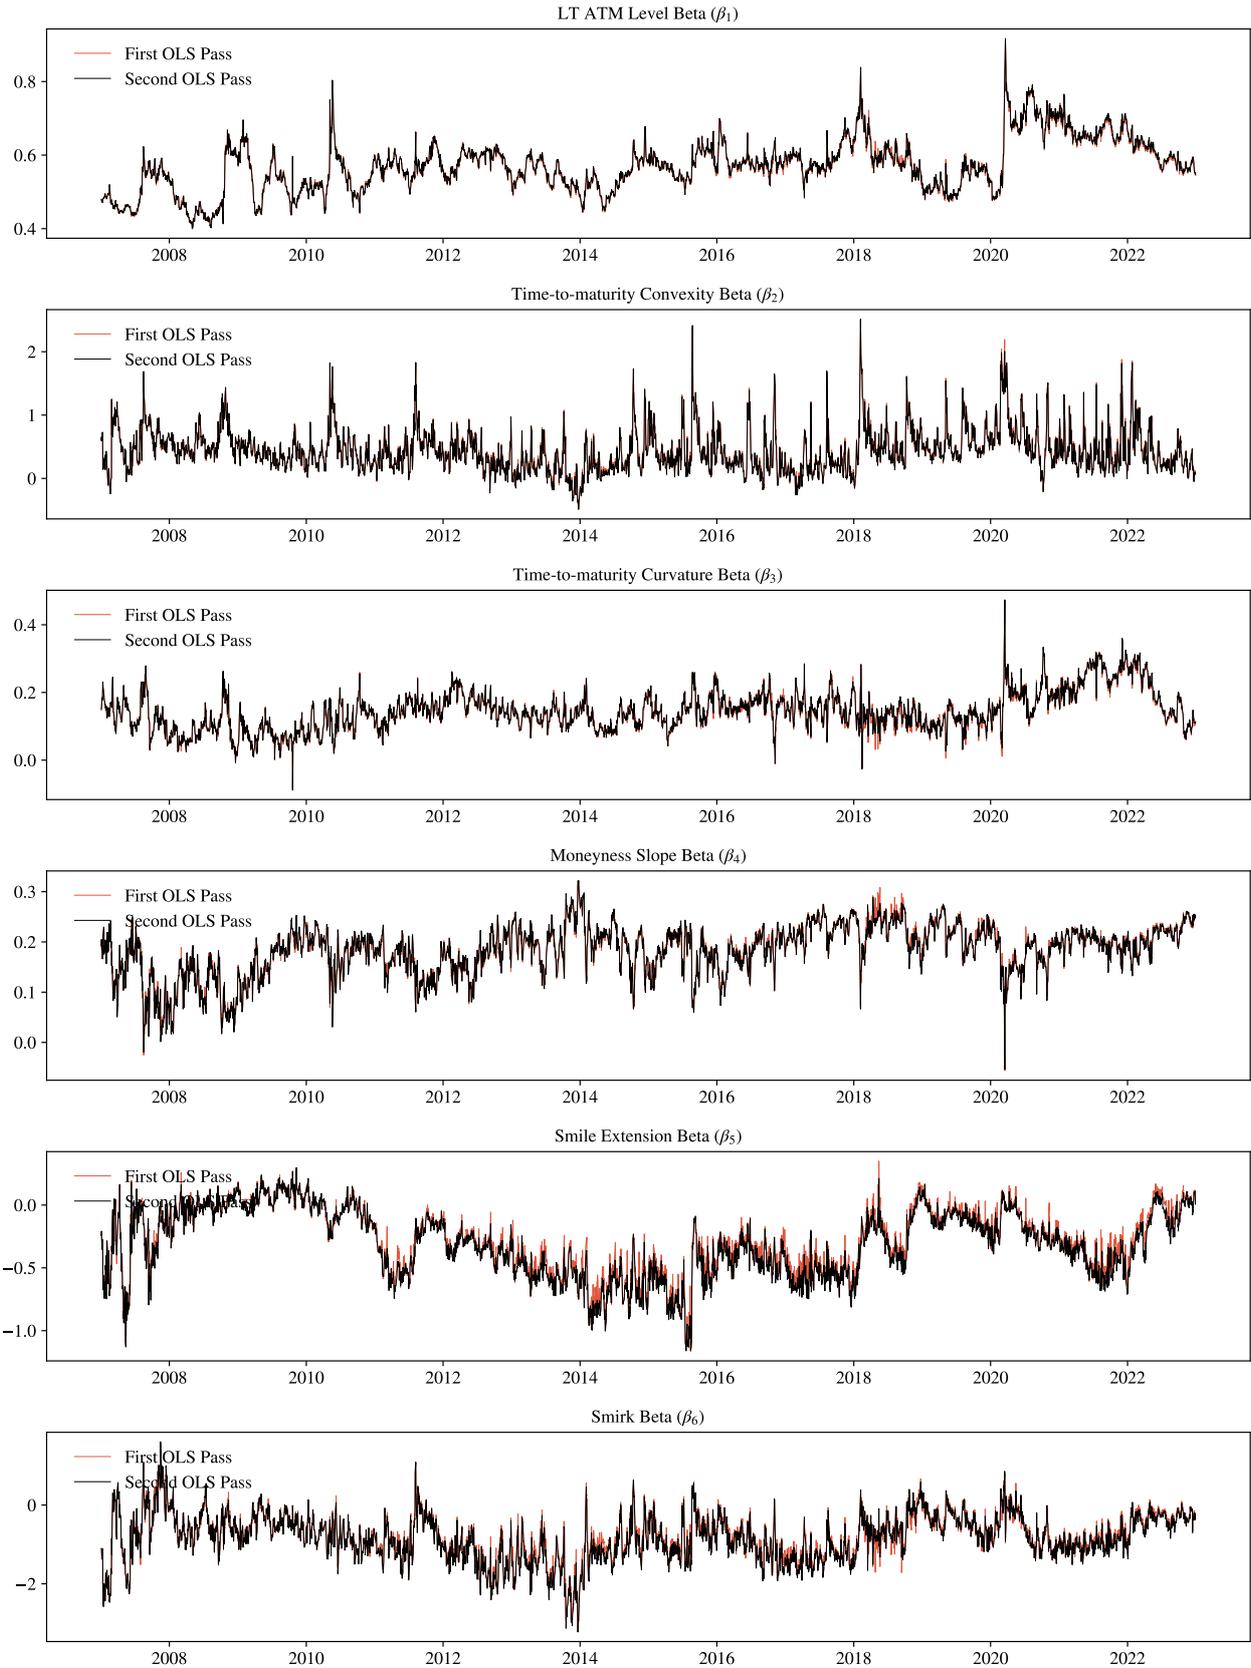
Panel B: Proportion (%) of daily observations excluded for fitting errors with $|z\text{-score}| > 3$



Exclusion of observations during the two-pass least squares optimization. The observations were excluded if their fitting error had an absolute z-score greater than 3. Panel A shows the number of daily observations that were removed. Panel B illustrates the percentage of daily observations removed each day, calculated by dividing the number of excluded observations ($|z\text{-score}| > 3$) by the total number of observations on that specific date.

Figure 9: Daily observations removed from the two-pass OLS optimization

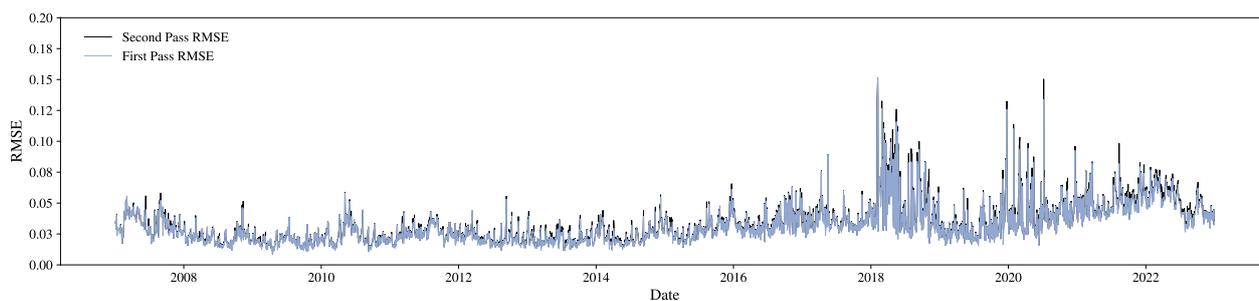
Panel A of Figure (11) indicates that the second pass has a small impact on the RMSE of the calibration. However, it enhances the overall stability of the model by smoothing the moneyness coefficients. Panel B illustrates an improvement in pricing accuracy, as the second pass reduces pricing errors and exhibits lower susceptibility to large variations compared to the first pass.



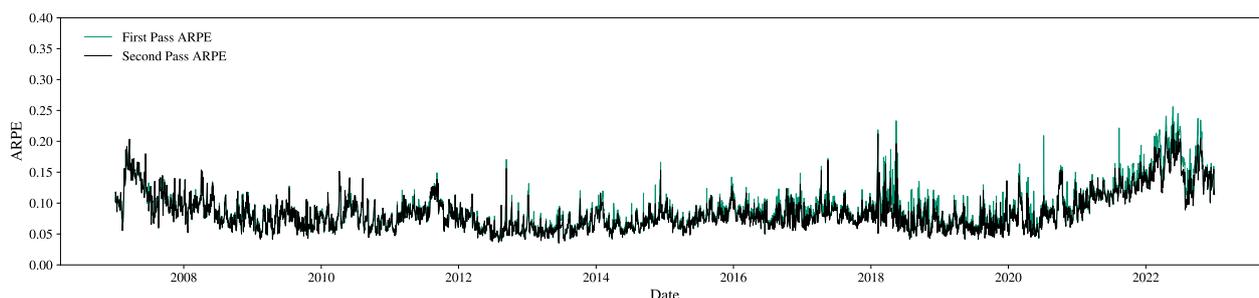
Daily estimation of the coefficients $\beta_t = (\beta_{t,1}, \beta_{t,2}, \beta_{t,3}, \beta_{t,4}, \beta_{t,5}, \beta_{t,6})$ for the IV model. Each of the factor coefficient time series are illustrated above capturing respectively the long-term ATM level, the moneyness slope, the time-to-maturity convexity, the time-to-maturity curvature, the smile extension and the smirk from January 1, 2007 to December 30, 2022.

Figure 10: Daily coefficients calibration of the IV model

Panel A: RMSE Two-pass Comparison



Panel B: ARPE Two-pass Comparison



RMSE and APRE for both pass are reported daily for the PVIVS model. The ARPE is the absolute relative pricing error, and the RMSE is the root mean square error.

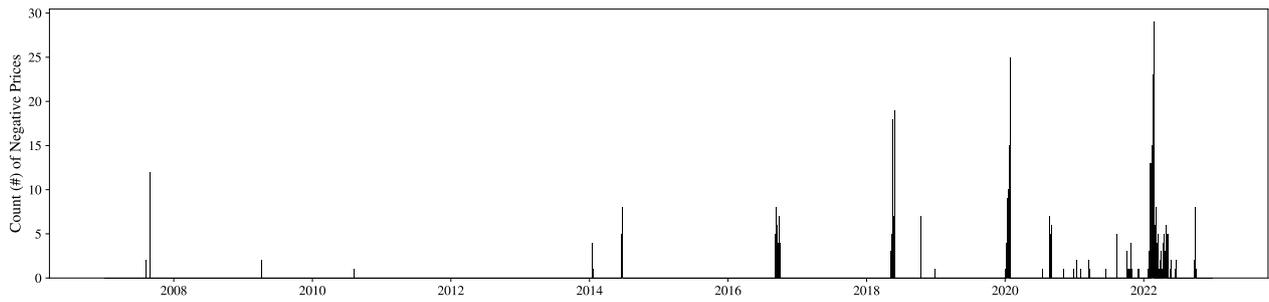
Figure 11: Model Performance : First and Second OLS Pass

B.3 Benchmark Model

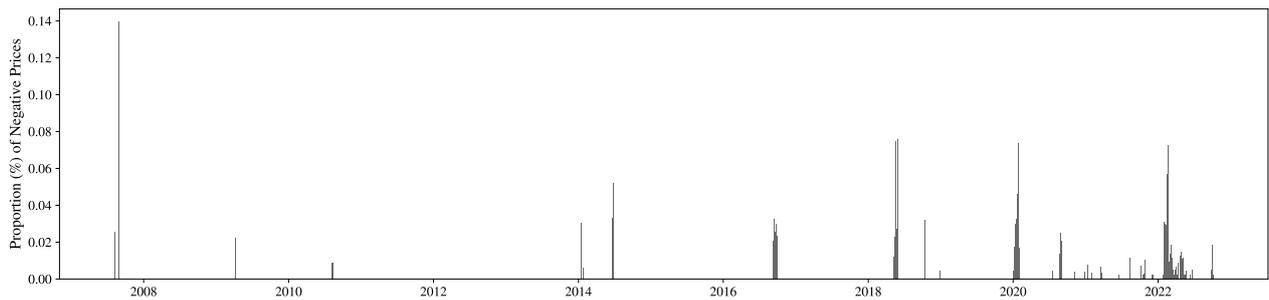
In this affine diffusion model, several parameters are used to price derivatives. The mean-reversion rate κ , indicates the speed at which the variance reverts to its long-term mean, denoted by θ . The parameter σ_v is the volatility of the variance process. The correlation between the asset price and its volatility is measured by ρ , indicating how much the asset price and its volatility move in relation to each other. The initial value of the volatility process is represented by V_0 . These parameters are typically estimated from market data. Each of the 5 parameters undergoes daily calibration using the `lsqnonlin` optimizer. For each daily calibration, the parameters from the previous day are used as the initial guess.

Once calibrated, the Heston model generates negative prices for deep out-of-the-money (OTM) put options. While no call options exhibited negative prices, 852 put options had negative prices. Panel A of Figure (12) displays the count of put options with negative prices. Panel B presents the proportion of negative prices relative to the total daily observations.

Panel A: Daily observations from Heston model with negative pricing



Panel B: Proportion (%) of daily observations from Heston model with negative pricing

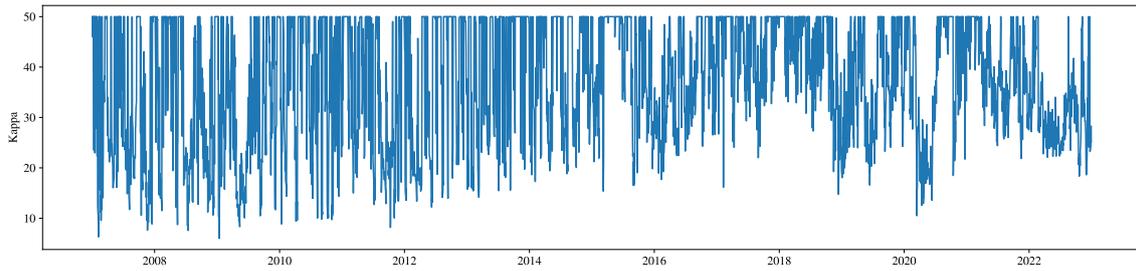


Daily observations from the Heston model identified instances of negative pricing. Panel A displays the number of daily observations with negative prices. Panel B illustrates the percentage of daily observations with negative pricing by dividing the number of negative prices by the total number of observations on that specific date.

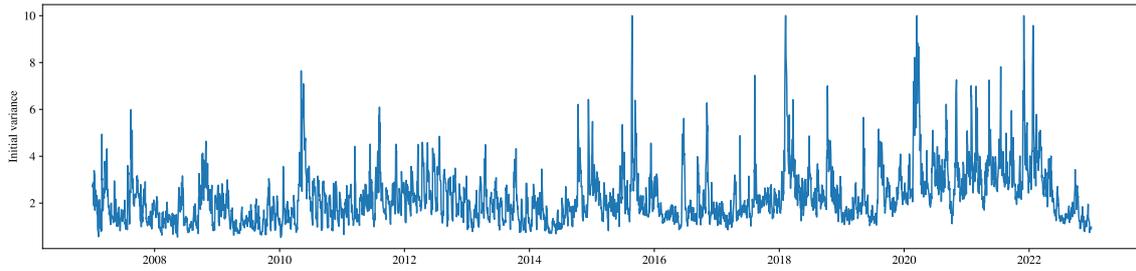
Figure 12: Negative prices from the Heston model

Figure (13) presents the time series of the Heston model parameters calibrated from the second pass. Compared to the PVIVS model, the parameters exhibit greater volatility and occasionally reach extreme values. Notably, the mean reversion speed parameter (κ) demonstrates significant instability.

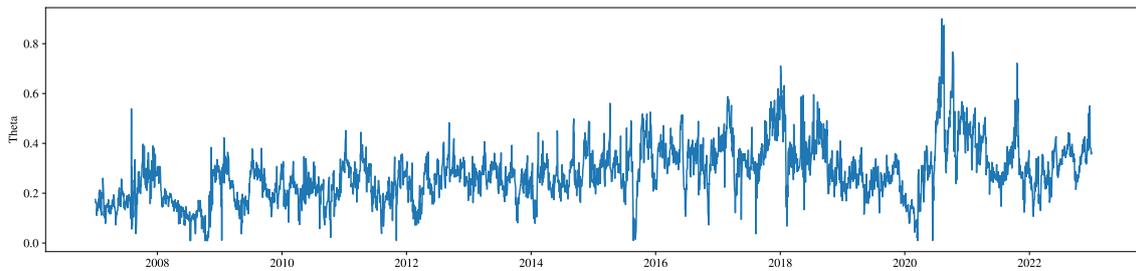
Panel A: Kappa (κ)



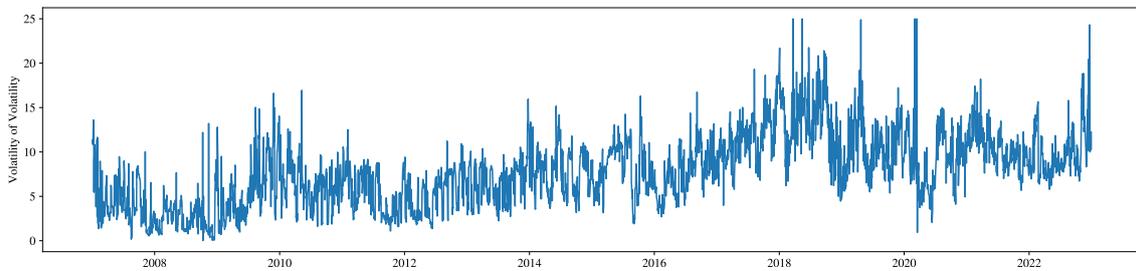
Panel B: Initial Variance (V_t)



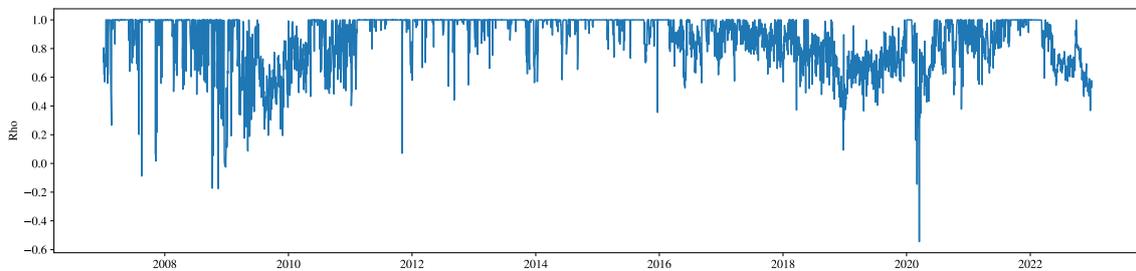
Panel C: Theta (θ)



Panel D: Volatility of Volatility (σ)



Panel E: Rho (ρ)



Daily calibrated parameters of the Heston model from January 2007 to December 30, 2022. Panel A is the mean reversion speed (κ), Panel B shows the initial variance (V_t), Panel C illustrates the long-term variance level (θ), Panel D represents the volatility of volatility (σ), and Panel E displays the correlation coefficient (ρ).

Figure 13: Time series of daily calibrated parameters for the Heston model

C Model Applications

C.1 Risk Neutral Density Function

Note that the derivations of the following partial derivatives: $\frac{\partial c}{\partial K}$, $\frac{\partial^2 c}{\partial K^2}$, $\frac{\partial d_1}{\partial M}$, and $\frac{\partial d_2}{\partial M}$, are identical to those presented in the work of Francois et al. (2022). The probability density function integrates to one since

$$\begin{aligned}\int_0^\infty e^{r\tau} \frac{\partial^2 c}{\partial K^2} dK &= e^{r\tau} \left(\lim_{K \rightarrow \infty} \frac{\partial c}{\partial K} - \lim_{K \rightarrow 0} \frac{\partial c}{\partial K} \right) \\ &= \lim_{M \rightarrow \infty} \left(\Phi(d_2) + e^{M\sqrt{\tau}} \varphi(d_1) \frac{\partial \sigma}{\partial M} \right) - \lim_{M \rightarrow -\infty} \left(\Phi(d_2) + e^{M\sqrt{\tau}} \varphi(d_1) \frac{\partial \sigma}{\partial M} \right) \\ &= \lim_{M \rightarrow \infty} \Phi(d_2) + \lim_{M \rightarrow \infty} e^{M\sqrt{\tau}} \varphi(d_1) \frac{\partial \sigma}{\partial M} = 1.\end{aligned}$$