Learning the Stochastic Discount Factor via Nonparametric Option Portfolios *

Emanuele Luzzi
USI Lugano and SFI
emanuele.luzzi@usi.ch

Paul Schneider
USI Lugano and SFI
paul.schneider@usi.ch

Rohan Sen
USI Lugano
rohan.sen@usi.ch

Abstract

We estimate the stochastic discount factor (SDF) by recovering the Sharpe-optimal nonlinear claim through a trading strategy in delta-hedged option portfolios. Our nonparametric approach leverages the classical duality between the minimum-variance SDF and the maximum Sharpe ratio portfolio, and comes with finite-sample performance guarantees, as well as a formal testing framework for the monotonicity and convexity of the SDF. We perform an empirical study in the S&P 500 market and find heterogeneous shapes across different states of the world as measured by the price of volatility and the maturities of options. While SDF implied by monthly options are monotonically decreasing, their convexity/concavity is less pronounced. Ultra-short ODTE options, on the contrary, exhibit a pronounced U-shape in higher-volatility states. Our empirical results are robust across various models of the information set.

JEL Classification: C14, C58, G11, G12

Keywords: Stochastic Discount Factor, Machine Learning, Portfolio Choice, Options

*We are grateful for the helpful comments received at the 17th Annual Meeting of the Society for Financial Econometrics (SoFiE) in 2025, the 7th QMUL Economics and Finance Workshop for PhD and Post-doctoral Students in 2025, and the HKUST IAS-SBM Joint Workshop on Financial Econometrics in the Big Data Era. We thank Loriano Mancini, Patrick Gagliardini, Damir Filipović, Peter Reinhard Hansen, Roberto Renò, Nicola Fusari, and Yingying Li for useful comments. We are also grateful to Darius Nik Nejad for valuable discussions. We gratefully acknowledge support from the SNF grant 215528 *Large-scale kernel methods in financial economics*.

1 Introduction

The fundamental challenge in asset pricing revolves around understanding investors' risk preferences and how these preferences shape market dynamics. However, these are unobservable, and strong assumptions are necessary to derive information about them. At the heart of this endeavor lies a persistent empirical conundrum: the stochastic discount factor puzzle. Standard economic preference theory predicts that the stochastic discount factor (SDF) should decrease monotonically with market returns, reflecting risk-averse behavior, and it should be convex, reflecting risk prudence. However, empirical estimates consistently reveal concavities and non-monotonic regions where the SDF increases with returns, contradicting these most fundamental notions of decision making. The verification of the very presence of this puzzle requires researchers to devise the most general and powerful econometric techniques possible.

In this paper, we estimate the empirical SDF projected onto the return space without imposing any parametric constraints or structural assumptions, thereby allowing the data to reveal a pure view of the structure of investors' risk preferences. We transform the SDF estimation problem into a tradable option portfolio allocation problem, where the trading strategy directly replicates the stochastic discount factor through the optimal nonlinear claim that maximizes the Sharpe ratio. This transformation allows us to estimate the SDF from option market data directly.

Our key innovation is twofold. Firstly, we leverage the duality between the Hansen-Jagannathan minimum-variance discount factor and the maximum Sharpe ratio portfolio to work directly with a delta-hedged option trading strategy to identify the SDF. The estimation of the SDF is thus implicit and revealed as a byproduct of the optimal allocation function of the strategy. Secondly, we avoid the problematic separate estimation of risk-neutral and physical (Lebesgue) densities to form the ratio ($d\mathbb{Q}/d\mathbb{P}$), which is common in traditional approaches. Instead, we estimate the SDF directly without the need to make the density ratio explicit at any point in the procedure, merely assuming it to be a twice-differentiable function of the market gross return.

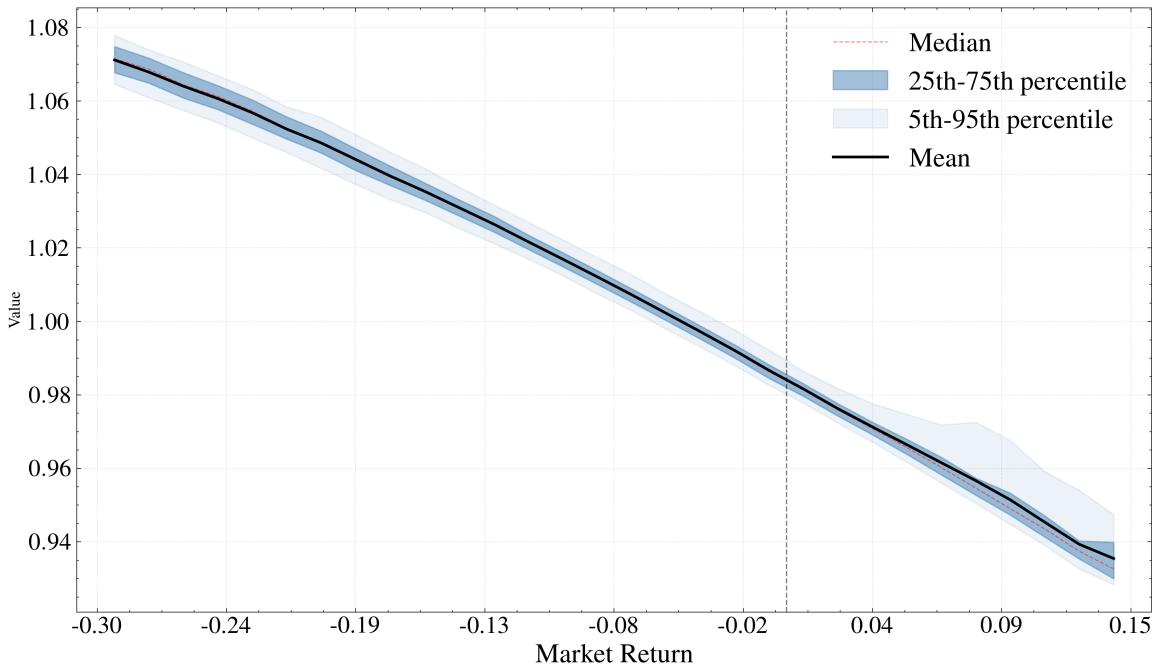

Figure 1.1: Distribution of the estimated stochastic discount factor as a function of S&P 500 market returns using monthly out-of-the-money European put and call options for the period 2000-2022. The figure displays the median (black solid line), mean (red dashed line), 25th–75th percentile ribbon (darker shaded region), and 5th–95th percentile ribbon (lighter shaded region) of SDF values across different return levels.

Our empirical analysis on S&P 500 options reveals near-linearity with a slight concave curvature in the estimated SDF, with higher values during negative market returns and minimum values for positive returns. Importantly, we find no support for the U-shaped patterns commonly reported in the literature, with the exception of ultra-short-term maturity options. Moreover, a formal test of the economic null hypothesis that the SDF is monotonically decreasing and convex reveals substantial state-dependence contingent on market volatility as measured by the VIX index, and option maturity. For monthly options, monotonicity cannot be rejected across volatility states, consistent with robust risk aversion in these markets. However, the volatile microstructure of ultra-short-term options precludes the same conclusion for 0DTEs, where monotonicity is frequently rejected at low-to-medium VIX levels. Regarding curvature, neither convexity nor concavity can be statistically rejected for options at any maturity, indicating near-linear SDF profiles. Interpreting these results through the lens of traditional preference theory (Eeckhoudt and Schlesinger, 2006),

we conclude that risk aversion, as manifested through monotonicity, represents a statistically robust feature of monthly index option markets, whereas higher-order risk preferences such as downside risk aversion (prudence) lack statistical significance in both monthly and ultra-short maturity options.

At the core of the empirical SDF puzzle lies a joint hypothesis problem: when we test an SDF and reject it, are we rejecting the economic model of risk preferences, or are we rejecting the statistical model used for the estimation of the pricing equation? Most existing approaches rely on parametric assumptions about either the utility function of investors or the return-generating process, making it difficult to disentangle these effects. Misspecification in either component can lead to the appearance of anomalies that may not reflect market behavior. The mainstream literature on the empirical SDF leverages instead the ratio between the distribution of risk-neutral and physical probabilities (Jackwerth (2015)), which trades off stability for accuracy. Our approach leads to a model-free framework that replicates the empirical stochastic discount factor implied by the market without structural assumptions or the need to estimate the distribution of state prices ($d\mathbb{Q}$) and physical probabilities ($d\mathbb{P}$) separately. To do so, we transform the estimation problem into an allocation problem, seeing it through the lens of a tradable strategy yielding the maximum Sharpe ratio. In fact, by leveraging the duality between the Hansen and Jagannathan (1991) minimum-variance discount factor and the maximum Sharpe ratio portfolio, we reframe the SDF estimation as an optimal portfolio choice problem in the options and underlying markets, where the complex, non-linear relationships between asset characteristics can be captured without parametric assumptions.

The methodological insight of our paper lies in developing a framework that estimates the SDF through optimal portfolio allocations, requiring only that the SDF, when expressed as a function of the terminal stock price, be twice differentiable. This minimal assumption allows us to leverage the Carr and Madan (2001) replication formula to decompose the problem into a portfolio of options whose allocations are determined by the first and second derivatives of the SDF. These derivatives carry direct economic content. Monotonicity reflects risk aversion: the SDF should decrease with market returns as investors value consumption more in adverse states. Convexity governs how rapidly marginal utility

declines across states. Violations of these properties are economically informative, potentially revealing heterogeneous beliefs, incomplete markets, or state-dependent preferences that representative agent models abstract away. Our nonparametric approach imposes no restrictions on the SDF's shape beyond smoothness, allowing such structures to emerge directly from option prices (Figure 1.1).

Beyond estimation, our methodology delivers a testing framework for the properties of the recovered SDF. The structure of our estimator enables formal hypothesis testing of shape restrictions – monotonicity and convexity – that carry economic implications. Monotonicity violations signal departures from standard risk aversion, while convexity reflects how rapidly marginal utility declines across states, relating to downside risk aversion and risk prudence. Our nonparametric test assesses these properties without imposing distributional assumptions. It emerges as a byproduct of our estimation approach, enabling us to assess SDFs that reflect more complex preference structures formally. The nonparametric nature of our approach, together with its direct link to trading strategies, grants an unprecedented integrated view into the S&P 500 index and index options market. In particular, we find no evidence for the convexity of the SDF, which has been a robust trait emphasized by many parametric studies, and which lies at the heart of explanations for the pronounced negative sign of the variance risk premium from downside risk aversion. As such, our approach underscores the importance of unifying statistical inference with implementable trading strategies in empirical asset pricing, and motivates applying this methodology to other asset classes and market settings where the interplay between risk preferences and market prices remains contested.

Related literature. Our work bridges different strands of literature, leveraging options and portfolio choice to estimate the SDF nonparametrically. The classic duality between the minimum-variance SDF and maximum Sharpe ratio portfolio (Hansen and Richard (1987)) provides the theoretical foundation for transforming the SDF estimation into an option trading strategy. While seminal work from Carr and Wu (2008) demonstrates how synthesized variance swap portfolios may be used to reveal the variance risk premium, we ought to extend this insight by showing how a particular portfolio of options recovers the empirical SDF. Building on Brandt et al. (2009), who parameterize portfolio weights as

functions of asset characteristics, we develop a fully nonparametric framework where the optimal allocations are learned directly from option characteristics without functional form restrictions. These allocations correspond to the first and second derivatives following from the Carr and Madan (2001) decomposition. Our approach fundamentally differs from the extensive literature on recovering the SDF through the $d\mathbb{Q}/d\mathbb{P}$ densities ratio. The standard paradigm for SDF estimation extracts risk-neutral densities from option prices and physical densities from historical returns separately, then forms their ratio. While Breeden and Litzenberger (1978) provided the theoretical foundation and subsequent work has refined the risk-neutral density extraction Aït-Sahalia and Duarte (2003); Qu and Zhang (2025), the complete density-ratio approach to SDF estimation (Aït-Sahalia and Lo (2000)) creates critical methodological issues through this two-step procedure. In particular, this separation creates two problems: compounding estimation errors and temporal mismatches between forward-looking options and backward-looking returns. These methodological issues manifest in the "SDF puzzle" documented by Jackwerth (2015), i.e., the puzzling U-shaped patterns in estimated SDFs that contradict basic economic theory. Bondarenko (2014) shows how one shade of this puzzle, the anomaly in overly expensive put contracts, struggles to be explained by conventional models. More recently, Schreindorfer and Sichert (2025) argues that most macrofinance models hardly explain conditional risk prices. A parallel strand of research investigates SDF bounds. Building on variance bounds, Snow (1991), Backus et al. (2011), Chabi-Yo (2007), Almeida and Garcia (2017) extend these restrictions by incorporating higher-order moments, conditioning information, and option-implied distributions to test whether candidate SDFs are admissible. While these approaches provide valuable frameworks to test pricing models, our method directly estimates the empirical SDF through optimal portfolio allocation, transforming the estimation problem into a tradable option strategy that replicates the SDF payoff without requiring bound-based inference or parametric restrictions. Nevertheless, theoretical explanations invoke preference-based theories, including state-dependent preferences (Grith et al. (2016)), heterogeneous beliefs (Bakshi et al. (2010)), and differential awareness of tail risks. However, each requires strong assumptions about investor behavior or market structure. A complementary strand of research directly estimates heterogeneous investor preferences through demand systems

(Koijen and Yogo (2019)), revealing how differences in risk appetites across investor types shape asset prices. Crescini et al. (2025) extract heterogeneous investors' subjective expected returns from option holdings and prices, documenting that beliefs vary wildly across investor types. While our approach shares the spirit of inferring risk preferences directly from market data, we focus on the aggregate market-implied SDF rather than modeling investor-level demand heterogeneity. Recent attempts to resolve the SDF puzzle remain constrained by the density separation problem. Linn et al. (2017) propose an information-consistent nonparametric estimator and argue that existing methods are flawed. Song and Xiu (2016) condition on volatility factors using VIX options and uncover similar U-shaped puzzles in conditional densities. Barone-Adesi et al. (2020) addresses the puzzle by incorporating option trading activity into a Bayesian framework to adjust the physical measure. Christoffersen et al. (2013) parametrically specify the SDF as a function of GARCH-driven variance. The recovery literature offers partial solutions: Ross (2015)'s theoretical recovery requires restrictive assumptions about transition-independent SDFs, and Jackwerth and Menner (2020) demonstrates its empirical issues. Schneider and Trojani (2019)'s model-free approach, while avoiding parametric assumptions, requires projecting onto specific moment portfolios constructed from variance swaps rather than estimating the SDF directly.

By transforming SDF estimation into an option portfolio allocation problem, we circumvent these issues entirely. Rather than estimating densities to form their ratio, we learn the SDF directly through optimal allocations for an option trading strategy, eliminating both the compounding errors and temporal mismatches that affect traditional approaches. Nonetheless, our use of machine learning for portfolio construction is connected to Kozak (2020) and the work of Filipović et al. (2022) and Boudabsa and Filipović (2022). However, we tailor these methods specifically to the nature of our problem. To ensure robustness, we augment our conditioning set beyond standard moneyness and VIX and include credit spreads, trading volume, and macroeconomic factors along the lines of Nagel and Singleton (2011). Testing for alternative maturities reveals that shorter-term options also yield results that mirror our baseline findings. A particular case corresponds to ultra-short zero-day-to-expiry options (0DTEs), for which a non-monotonic U-shaped SDF is recovered in line with Almeida et al. (2024), who document that the SDF of 0DTEs exhibits a fun-

damental reversal where investors consistently require compensation for upside risk. This distinctive SDF shape aligns with Bandi et al. (2023) evidence of instantaneous return and variance risk premia that deviate from longer-maturity patterns and requires tilting the conditional distribution of the return process over ultra-short horizons, suggesting unique risk preference structures. Notably, Chong and Todorov (2024) derive SDF restrictions on the agreement between equity and options markets regarding spot volatility and the continuous component of small price moves through tradable volatility arbitrage portfolios combining static 0DTEs positions with dynamic hedging, finding no evidence of violations of these restrictions empirically.

Our approach addresses longstanding challenges in the estimation of risk preferences in both the options and asset pricing literatures, while resembling a tradable option strategy. In doing so, our framework bridges portfolio choice, option-based risk-premia, and nonparametric recovery, offering a tractable, implementable solution for the empirical SDF.

The remainder of the paper is organized as follows. Section 2 describes the theoretical background connecting the Hansen-Jagannathan bound, the SDF puzzle, and the replication of non-linear payoffs through options. Section 3 introduces our methodology in detail, showing the cautious embedding of nonlinear relationships between characteristics in the estimation. Section 4 presents our data. Section 5 illustrates our findings and the resulting properties of the estimated SDF, and explains the output of our testing framework. We conclude in Section 6. The appendix contains further implementation details and the related mathematical framework.

2 Financial-economic background

In this section, we outline the financial-economic rationale of our approach and how it contrasts with existing frameworks. The classical Hansen and Jagannathan (1991) minimum-second moment SDF recovery takes asset pricing returns as given, and yields the optimal SDF as linear in the assets it prices. This minimization problem comes with an equivalent portfolio mean-variance utility dual maximization problem that yields the same solution. The innovation of our approach is to start from the solution of this dual mean-variance

problem and maximize it over the set of claims $h(S_T/S_t)$, where h is twice-differentiable almost everywhere, and S_T/S_t is the forward gross return associated with the S&P 500 from time t to $T \geq t$. From Hansen and Richard (1987) and Hansen and Jagannathan (1991), $h(S_T/S_t)$ can then be associated with the minimum-variance maximum Sharpe ratio SDF in an economy that accommodates any (twice-differentiable and possibly nonlinear) derivative claim on S_T/S_t.

To appreciate the setting of our investigation, consider an economy on the complete probability space $(\Omega, \mathcal{F}, \{\mathcal{F}_t\}_{t\geq 0}, \mathbb{P})$, where Ω represents the state space, $\mathcal{F}$ is the σ-algebra of measurable events, $\{\mathcal{F}_t\}_{t\geq 0}$ denotes the information filtration, and $\mathbb{P}$ is the physical probability measure. By the fundamental theorem of asset pricing, the absence of arbitrage implies the existence of an equivalent forward measure $\mathbb{Q} \sim \mathbb{P}$ under which discounted asset prices are martingales. Without loss of generality, we set the risk-free rate to zero throughout our theoretical exposition and work with forward prices.

A fundamental result in option pricing is that, in a no-arbitrage economy, we may replicate any twice-differentiable claim h in discrete time via the Carr and Madan (2001) spanning formula, using a portfolio of options and a hedge component. The first and second derivatives of the claim give the allocations of the strategy components. Let forward price $F_{t,T} = \mathbb{E}_t^{\mathbb{Q}}[F_{T,T}]$ and set $S_t := F_{t,T}$ and $S_T := F_{T,T}$. Carr and Madan (2001)'s formula rests on Taylor's formula with remainder,

$$\begin{aligned} h(S_{t+1}) = h(S_t) + h'(S_t)(S_{t+1} - S_t) \\ + \int_0^{S_t} h''(K)(K - S_{t+1})^+ dK + \int_{S_t}^{\infty} h''(K)(S_{t+1} - K)^+ dK. \end{aligned} \tag{2.1}$$

Denote by $R_{t+1} := S_{t+1}/S_t$. Then a discrete and scaled approximation of (2.1) in a real options market with discretely quoted strikes,

$$\begin{aligned} h(R_{t+1}) - \underbrace{h(1)}_{\text{Cash}} \approx \underbrace{h'(1)(R_{t+1} - 1)}_{\text{Hedge component}} \\ + \underbrace{\sum_{i=1}^{N_K(t)} h''(m_{it}) \left(\mathbb{1}_{m_{it}>1}(m_{it} - R_{t+1}) + \mathbb{1}_{m_{it}\leq 1}(R_{t+1} - m_{it})\right) w_{it}}_{\text{Option Portfolio}} \end{aligned} \tag{2.2}$$

where moneyness $m_{it} := \frac{K_i}{S_t}$, with $K_1 < \cdots < K_{N_K(t)}$ the ordered strike prices, and $N_K(t)$ the number of strike prices (and thus contracts) in period t. Furthermore, the weights w_{it} correspond to the discretization of the Carr-Madan formula, see (2.1), as a consequence of which we use an approximation symbol in (2.2).

To connect options trading strategies to the SDF, let $\tilde{M}_t$ be defined as the Radon-Nikodým derivative,

$$\tilde{M}_t := \left.\frac{d\mathbb{Q}}{d\mathbb{P}}\right|_{\mathcal{F}_t}. \tag{2.3}$$

This is the object of our interest, of which we assume the second moment to exist. It is, however, unobservable in general and may depend on many unobservable factors. For practical purposes, we therefore work with the projection,

$$M_{t+1} := \mathbb{E}^{\mathbb{P}}[\tilde{M}_{t+1}|R_{t+1}, \mathcal{I}_t], \tag{2.4}$$

where $\sigma(R_{t+1}) \subseteq \mathcal{F}_{t+1}$ denotes the information generated by R_{t+1} and $\mathcal{I}_t \subseteq \mathcal{F}_t$ represents the variables modulating the information known at t. Bakshi et al. (2010) make a similar modeling choice, since pricing claims depending on R_{t+1} is invariantly achieved by either $\tilde{M}_{t+1}$ or M_{t+1}, due to the law of iterated expectations. For simplicity and ease of notation, we refer to M_{t+1} as the SDF henceforth. Denoting by $\mathbb{E}_t^{\mathbb{M}}[\cdot] := \mathbb{E}^{\mathbb{M}}[\cdot|\mathcal{F}_t]$ for $\mathbb{M} \in \{\mathbb{P}, \mathbb{Q}\}$, the SDF M_{t+1} is then $\mathcal{F}_{t+1}$-measurable with finite second moments.

Empirically, we model below the information set $\mathcal{I}_t$ through different conditioning variables v_t on $\mathcal{V}$. The SDF can then be expressed as a function $M\colon \mathbb{R}^+ \times \mathcal{V} \to \mathbb{R}$ of the underlying asset: $M_{t+1} = M(R_{t+1}, v_t)$, which we assume to be twice differentiable with respect to R_{t+1}. Below, we will establish a simple relation between the SDF M and the delta-hedged option trading strategy h. For $h \in \mathcal{H}$, a function space to be discussed in detail below, we define the payoff

$$\mathcal{R}_{t+1}(h) := h(R_{t+1}, v_t) - \mathbb{E}_t^{\mathbb{Q}}[h(R_{t+1}, v_t)], \tag{2.5}$$

Motivated by the possibility of trading any twice-differentiable payoff through 2.2, we further

define the unconditional mean-variance loss function,

$$\mathcal{E}(h) := -\mathbb{E}^{\mathbb{P}}[\mathcal{R}_{t+1}(h)] + \frac{1}{2}\mathbb{V}^{\mathbb{P}}[\mathcal{R}_{t+1}(h)]. \tag{2.6}$$

From Hansen and Richard (1987) and Hansen and Jagannathan (1991), minimizing $\mathcal{E}(h)$ over the space of twice-differentiable functions yields the negative of the maximum-Sharpe-ratio SDF that prices all twice-differentiable derivatives. In other words, this framework allows us to effectively trade the SDF via Equation 2.2. For economy of notation, henceforth, we shall use the following term for the options portfolio part from (2.1)

$$R^{O}_{i,t+1} := \mathbb{1}_{m_{it}>1}(m_{it} - R_{t+1}) + \mathbb{1}_{m_{it}\leq 1}(R_{t+1} - m_{it}), \quad i = 1, \ldots, N_K(t).$$

Furthermore, the weights w_{it} refer to the *composite quadrature rule* that serves to approximate the Carr-Madan formula via discretization. Following from (2.2) and (2.5), denote by $\mathcal{R}^{\pi}_{t+1}(h)$ the Carr-Madan replicating option trading strategy at time $t+1$,

$$\mathcal{R}^{\pi}_{t+1}(h) = h'(1; v_t)\,(R_{t+1} - 1) + \sum_{i=1}^{N_K(t)} h''(m_{it}; v_t) R^{O}_{i,t+1} w_{it}, \tag{2.7}$$

and

$$\mathcal{E}^{\pi}(h) := -\mathbb{E}^{\mathbb{P}}[\mathcal{R}^{\pi}_{t+1}(h)] + \frac{1}{2}\mathbb{V}^{\mathbb{P}}[\mathcal{R}^{\pi}_{t+1}(h)]. \tag{2.8}$$

However, minimizing $\mathcal{E}^{\pi}(h)$ over the space of twice differentiable functions appears an insurmountable obstacle at first glance. Firstly, it is not obvious how to obtain the allocations needed for the replication of $\mathcal{R}^{\pi}_{t+1}(h)$. Secondly, its infinite-dimensional nature leaves it unclear whether there exists a solution and whether it is computable. The next section outlines how we tackle both problems.

3 Methodology

The duality investigated in Section 2 shows that it is possible to link the mean-variance frontier problem to the replication of the optimal non-linear trade, and thus the maximal Sharpe ratio attainable by delta-hedged option trading strategies. To derive the replication,

we leverage the spanning formula from Carr and Madan (2001). Importantly, the optimal strategy yielding the maximum Sharpe claim goes *short* the SDF (Schneider, 2015, Result 2.1), we take this into account by defining

$$M_{t+1} = M(R_{t+1}, v_t) := 1 - h(R_{t+1}, v_t). \tag{3.1}$$

This definition owes to the fact that our trading strategy h does not identify the cash component in (2.2) (as the forward price of a constant is the constant, so that cash and the price of cash cancel). Furthermore, we will also introduce regularization, which will shrink h to zero, so that M_{t+1} is mean-one and represents a valid minimum-variance (forward) stochastic discount factor.

With Definition 3.1, it is thus sufficient for us to focus on h. The Carr-Madan formula now gives a decomposition of h in terms of a hedge component and a portfolio of options, with the allocations being the first and second derivatives of h, respectively. Conducting a parametrization of the problem (for example, considering a linear approach) would be too restrictive, since the space of twice-differentiable functions is much larger. We therefore develop a nonparametric methodology based on supervised kernel learning to solve the mean-variance frontier dual problem.

3.1 Mean-variance optimization

Motivated by retrieving the SDF via an optimal nonlinear options trading strategy, we perform a regularized mean-variance optimization using the loss function defined in (2.8) in a space of admissible nonlinear claims strategy given by shorting the SDF, i.e., h as defined above. Precisely, we "learn" the function h that gives rise to the optimal trading strategy in the mean-variance sense. We thus perform mean-variance optimization in a supervised kernel learning setup that learns the ideal portfolio weights h' and h'', which correspond to the delta-hedge component and the OTM options portfolio component, respectively, in Equation 2.7.

Towards this end, we denote by $\mathcal{H}$ a possible hypothesis space that contains at least twice continuously differentiable functions.[1] As a natural choice for this endeavor, we propose to

[1]Second-order continuous differentiability is not strictly necessary, but is necessary for using the Carr-

specify $\mathcal{H}$ as the *reproducing kernel Hilbert space (RKHS)* defined using a *Sobolev kernel*. This particular nonparametric modeling choice has several benefits for this problem. We defer the discussion on RKHS and the specification of $\mathcal{H}$ to the Appendix A. For now, we point out certain facts related to RKHS that would enable us to write a regularized version of the loss function in (2.8) as a bounded, linear, real-valued function defined on this $\mathcal{H}$.

Any RKHS $\mathcal{H}$ has an associated implicit *feature map* $\phi\colon \mathcal{X} \to \mathcal{H}$ for nonlinear data transformation into the possibly infinite-dimensional space $\mathcal{H}$ of functions, and an explicit *kernel function* $k(x, x') = \langle \phi(x), \phi(x') \rangle_{\mathcal{H}}$ that captures the "similarity" of transformed data. Such an embedding allows for representing the function evaluation as an inner product as $f(x) = \langle f, \phi(x) \rangle_{\mathcal{H}}$ for any function $f \in \mathcal{H}$. Adopting a standard procedure from kernel-based supervised learning methods, and as anticipated by the SDF definition (3.1), we consider the following regularized population mean-variance optimization function,

$$h_\lambda = \underset{h \in \mathcal{H}}{\operatorname{argmin}}\ J(h) := \mathcal{E}^\pi(h) + \frac{\lambda}{2}\|h\|_{\mathcal{H}}^2 = -\mathbb{E}[\mathcal{R}_{t+1}^\pi(h)] + \frac{1}{2}\mathbb{V}[\mathcal{R}_{t+1}^\pi(h)] + \frac{\lambda}{2}\|h\|_{\mathcal{H}}^2, \tag{3.2}$$

where we use the convention $\mathbb{E}[\cdot] \equiv \mathbb{E}^{\mathbb{P}}[\cdot]$, $\mathbb{V}[\cdot] \equiv \mathbb{V}^{\mathbb{P}}[\cdot]$. The empirical counterpart of population Problem 3.2 is given by

$$\widehat{h}_\lambda = \underset{h \in \mathcal{H}}{\operatorname{argmin}}\ \widehat{J}(h) := -\widehat{\mathbb{E}}[\mathcal{R}_{t+1}^\pi(h)] + \frac{1}{2}\widehat{\mathbb{V}}[\mathcal{R}_{t+1}^\pi(h)] + \frac{\lambda}{2}\|h\|_{\mathcal{H}}^2. \tag{3.3}$$

Here, $\widehat{\mathbb{E}}[\cdot]$ and $\widehat{\mathbb{V}}[\cdot]$ refer to the sample mean and variance respectively. In this setting, our samples comprise the family of moneyness (for the option contracts) and any exogenous conditioning variable(s) v_t that we may use, for each month $t = 0, \ldots, T-1$. One of the most important aspects of optimizing Problem 3.3 is that we can derive a finite-dimensional matrix formulation of it. This is one of the several key benefits of adopting the RKHS framework that allows for such, and is facilitated by the *representer theorem* below.

Theorem 3.1 (Representer theorem)**.** *The optimal solution to Problem 3.3 has the form*

$$\widehat{h}_\lambda = \sum_{t=0}^{T-1} \widehat{c}_{1,t}\, \phi'(1; v_t) + \sum_{t=0}^{T-1} \sum_{i=1}^{N_K(t)} \widehat{c}_{2,it}\, \phi''(m_{it}; v_t). \tag{3.4}$$

Madan spanning formula.

Remark 3.2. *The representer theorem establishes that the hypothesis space for the optimal solution of the optimization problem at hand is defined by the first and second derivatives of the samples' feature map. This pivotal result allows the transformation of an infinite-dimensional optimization problem into a finite-dimensional one. It demonstrates the effectiveness of the RKHS framework, enabling us to recover the nonlinear claim nonparametrically while still employing conventional quadratic optimization in finite dimensions, as detailed in* (3.5).

3.2 Sample estimator

As noted in Remark 3.2, the representer theorem ensures that any minimizer of Problem 3.3 admits the finite expansion in (3.4). Consequently, evaluating $\widehat{h}_\lambda$ amounts to determining the coefficients in (3.4), since once the hypothesis space $\mathcal{H}$ is fixed, the kernel and its derivatives are known explicitly. Hence, the sample estimator $\widehat{h}_\lambda \in \mathcal{H}$ can be parameterized by the coefficients obtained from Problem 3.5. This yields a computationally tractable procedure even in a nonparametric regime: the RKHS allows us to work implicitly over a function space while reducing the task to a finite-dimensional problem amenable to standard linear-algebraic methods. With this motivation, we present below the practical optimization formulation and defer the derivation to Appendix B.

$$\widehat{\boldsymbol{c}} = \underset{\boldsymbol{c}\in\mathbb{R}^N}{\operatorname{argmin}} \; \frac{1}{2}\boldsymbol{c}^\top \left(\boldsymbol{K}\boldsymbol{\Sigma}\boldsymbol{K} + \lambda \boldsymbol{K}\right)\boldsymbol{c} - \boldsymbol{c}^\top \left(\boldsymbol{K}\bar{\boldsymbol{a}}\right). \tag{3.5}$$

Remark 3.3. $\boldsymbol{K}$ *is the block-kernel matrix formed by evaluating in aggregate* $k(\cdot,\cdot)$ *and its derivatives following from the representer theorem,* $\bar{\boldsymbol{a}}$ *is the stacked vector of averaged returns,* $\boldsymbol{\Sigma}$ *is the covariance matrix of returns, refer to Appendix B for details. However, the important thing to notice here is that Problem 3.5 is a tractable program that can be solved in closed form.*

Following Sen (2025), see also Filipović and Schneider (2025), we can provide the following statistical properties of $\widehat{h}_\lambda$.

Theorem 3.4 (Statistical properties)**.** *The regularized sample estimator* $\widehat{h}_\lambda$ *satisfies,*

1. *Asymptotic consistency:* $\widehat{h}_\lambda \xrightarrow{a.s.} h_\lambda$ *as* $T \to \infty$.

2. *Finite-sample properties: for any* $\delta \in (0,1)$*, it holds with sampling probability at least* $(1-\delta)$*,*

$$\|\widehat{h}_\lambda - h_\lambda\|_{\mathcal{H}} \leq C_{FS}(\delta, \|h_\lambda\|_{\mathcal{H}})\lambda^{-1}T^{-1/2},$$

for some coefficient $C_{FS}(\delta, \|h_\lambda\|_{\mathcal{H}})$*.*

With the statistical properties of the SDF estimator well established, the next section introduces the data used for the empirical study, as well as the validation procedure for the regularization parameter λ.

4 Data

Here, we introduce the option market data and outline our empirical estimation procedure.

4.1 Options

We focus on monthly OTM European put and call options on the S&P 500. Option and underlying data are pulled from the CBOE OptionMetrics from January 2000 to December 2022, and the VIX is taken from CRSP. Consistent with the literature on index-option returns (e.g., Hu and Liu (2022)), we examine monthly contracts maturing 25–33 days after the day following the third Friday of each month. The information obtained includes details on the contracts (e.g., maturity, strikes, etc.), as well as option-specific characteristics (Greeks, volume, open interest, etc.). The dataset includes 24,011 monthly contracts. It takes the shape of an unbalanced panel, since the number of strikes, and thus moneyness, varies at each date. Options and underlying are selected and evaluated every month, representing an investor trading monthly contracts and holding them until expiration. This is in line with the information needed for the estimation of the empirical SDF across states, for a fixed horizon. Conventionally, the expiration for a given month is given by the third Friday, which sets the starting date for the holding period. The following are filtered out of the dataset: ITF contracts, negative bids, negative bid-ask spreads, zero open interest contracts, zero volume contracts, redundant contracts, in which case, the one with the highest volume is kept, and extreme outliers on option-specific characteristics.

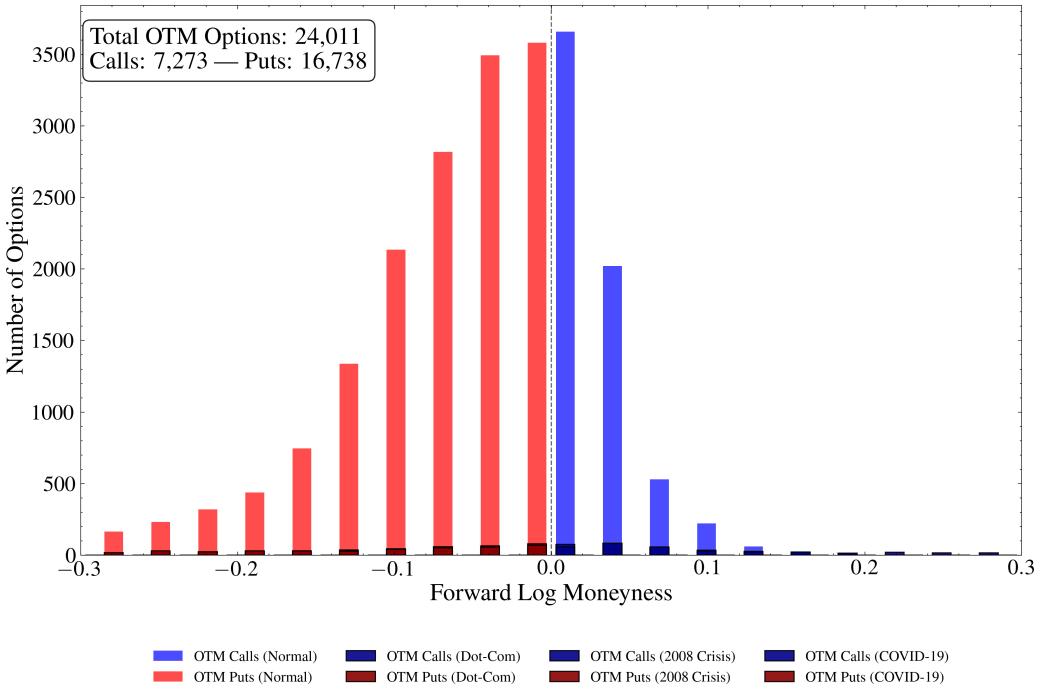

Figure 4.1: The figure shows the distribution of S&P 500 monthly out-of-the-money options used for our baseline study for the period 2000-2022. In red OTM European put options, in blue OTM European call options. The shaded areas refer to volumes emitted during the Dot-Com bubble, 08-Crisis and COVID periods. The moneyness range is capped at [-0.3, 0.3].

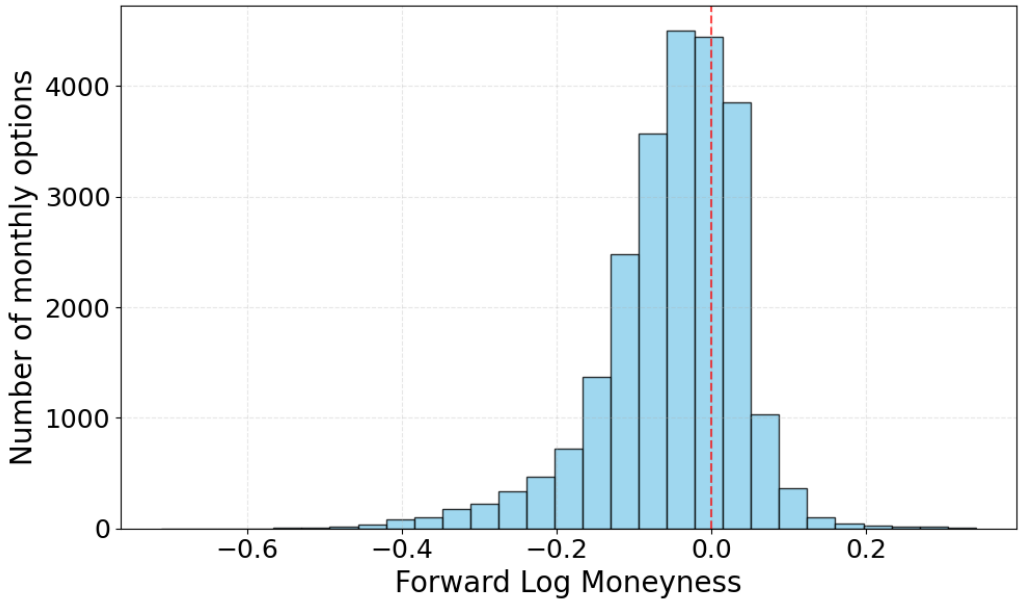

Figure 4.2: Total distribution of S&P 500 monthly out-of-the-money options for the period 2000-2022.

4.2 Validation

Our nonparametric estimation approach recovers the optimal allocation function $\widehat{h}_\lambda$ through a supervised kernel learning framework. As developed in Section 3, rather than imposing a parametric structure on the SDF, we solve the regularized mean-variance optimization problem in Equation 3.3, which converges asymptotically to the population optimizer h_λ, see Problem 3.2; this yields a trading strategy that replicates the minimum-variance SDF, up to regularization. The estimation hinges on a single regularization hyperparameter $\lambda > 0$ that controls the complexity of the learned function $\widehat{h}_\lambda$ within the reproducing kernel Hilbert space $\mathcal{H}$, analogous to the penalty parameter in ridge regression. The resulting hyperparameter selection problem is inherently non-convex, a fundamental feature of kernel-based learning, which we address through time series cross-validation.

We partition the dataset into disjoint training and validation sets with $T_{\text{val}} < T_{\text{train}}$, where the regularization parameter λ is selected out-of-sample through the validation split, despite conducting the overall empirical analysis in-sample. To preserve temporal structure and avoid look-ahead bias, we implement time series cross-validation through a time-aware

k-fold procedure (typically $k \in \{5, 10\}$). Training and validation folds are constructed sequentially, respecting chronological ordering and grouped by date (see Algorithm B.1 in Appendix B), mimicking an investor who optimally allocates based solely on historical information.

Let $\widehat{\boldsymbol{c}}_{\text{IS}} \in \mathbb{R}^N$ denote the optimal coefficient vector (that parameterizes the optimal solution $\widehat{h}_{\lambda,\text{IS}}$) estimated on in-sample training data, where N represents the combined dimension of hedge and option allocation coefficients, see more details in Section B. Out-of-sample performance is evaluated by applying these trained coefficients to validation data: we compute the portfolio return $\mathcal{R}^{\pi}_{t+1}(\widehat{h}_{\lambda,\text{IS}})$ and evaluate performance through $\widehat{J}(\widehat{h}_{\lambda,\text{IS}})$ where $\widehat{J}$ is the objective function in (3.3).

This validation procedure differs fundamentally from standard linear out-of-sample forecasting. Rather than producing optimal portfolio weights directly, the training phase yields optimal coefficients of kernel basis functions. When evaluated at new market conditions, these coefficients generate state-contingent allocations through the kernel's similarity structure. Specifically, for a new observation at time $s+1$ with market state (m_{s+1}, v_{s+1}) comprising moneyness and conditioning variables, and training observations $\{(m_i, v_i)\}_{i=1}^{s}$, the out-of-sample function evaluation is

$$\widehat{h}_{\lambda,\text{OOS}}(m_{s+1}, v_{s+1}) = \sum_{i=1}^{s} \widehat{c}_{i,\text{IS}}\, k((m_{s+1}, v_{s+1}), (m_i, v_i)), \tag{4.1}$$

where $k(\cdot,\cdot)$ denotes the kernel defined in Appendix A, and $c_{i,\text{IS}}$ are the coefficients trained on the first s observations. This kernel-based evaluation mechanism enables natural generalization beyond the training sample by leveraging the smoothness properties encoded in the RKHS. The similarity structure $k((m_{s+1}, v_{s+1}), (m_i, v_i))$ provides an interpolation mechanism that weights training observations according to their proximity in the feature space, allowing the methodology to adapt to new market conditions while maintaining the learned structure of the allocation function.

Computationally, we implement the cross-validation procedure through efficient matrix decompositions that exploit the structure of the kernel matrix $\boldsymbol{K}$ and variance factor $\boldsymbol{K}$ (detailed in Appendix B). The hyperparameter optimization leverages the *Optuna* framework

(Akiba et al., 2019), which implements tree-structured Parzen estimator (TPE) sampling for efficient exploration of the non-convex λ space.

5 Empirical results

This section contains the empirical results obtained from our nonparametric estimator. Below, we first show that it correctly recovers parametric SDF in model-based economies. Subsequently, we apply it to real market data.

5.1 Simulation study for economies with monotone and U-shaped SDFs

To assess the robustness of our nonparametric approach, we implement our methodology to estimate the SDF in two canonical economies deriving from the models of Black-Scholes and Heston. Thus, we compare them with their known theoretical SDFs across different moneyness levels. This dual validation strategy allows us to test whether our nonparametric methodology can recover both monotonically decreasing and U-shaped SDFs. We first obtain the well-defined SDF curve from the Black-Scholes framework as specified in Appendix C, which exhibits the classic monotonically decreasing pattern in the stock price (Figure 5.1). For the variance-dependent economy, we implement the Christoffersen-Heston-Jacobs (CHJ) GARCH model (Christoffersen et al., 2013), where the SDF depends on both returns and variance. While monotonic in each variable individually, the projection onto returns yields a characteristic U-shape when the variance premium is negative (Figure 5.2), consistent with empirical findings.

Using each theoretical SDF as a benchmark, we run the machinery developed in Section 3 for 5000 periods and 31 OTM options, estimating through validation in the same manner as the baseline case, using only moneyness as an input feature. For the Black-Scholes economy, we employ reasonable values of μ and σ, while for the CHJ GARCH specification, we calibrate the variance dynamics parameters and the variance risk premium following Christoffersen et al. (2013). We repeat this process for 500 simulations in each economy and visualize the boxplot of the differences for the whole range of moneyness values.

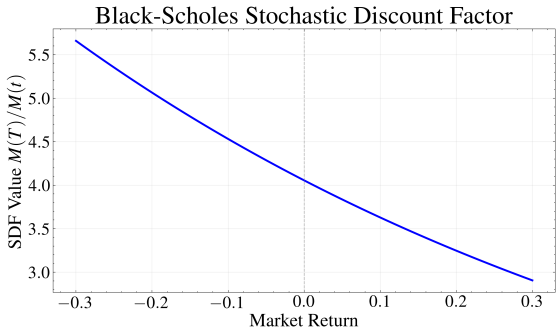

Figure 5.1: Black-Scholes SDF: monotonically decreasing in the market return (Black and Scholes (1973)).

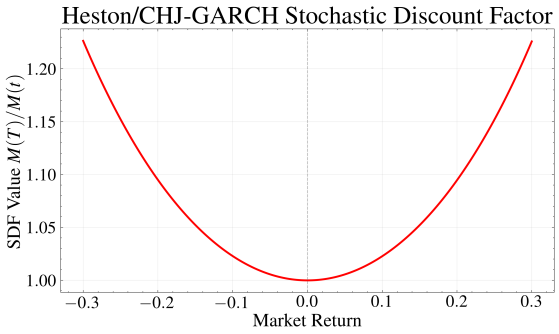

Figure 5.2: Heston SDF: U-shaped projection onto returns due to variance risk premium (Christoffersen et al. (2013)).

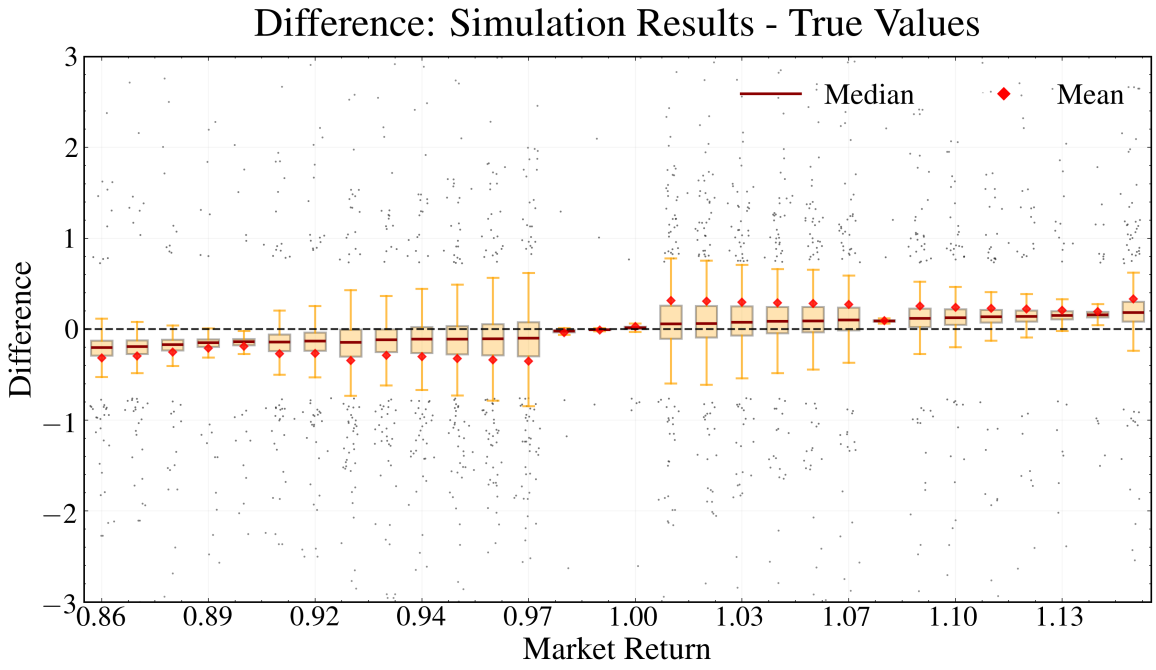

Figure 5.3: Distribution of differences between estimated and theoretical Black-Scholes stochastic discount factors across moneyness levels. The Black-Scholes economy assumes geometric Brownian motion dynamics with constant volatility under the physical measure, generating a monotonically decreasing SDF of the form in Equation (C.1). Boxplots show median, quartiles, and range of estimation errors across 500 simulations with 5000 periods and 31 OTM options each.

Figures 5.3 and 5.4 demonstrate that in both economies, the estimated SDF converges to the true one, confirming that the methodology successfully replicates the discount factors implied by fundamentally different SDF structures. The convergence holds despite the

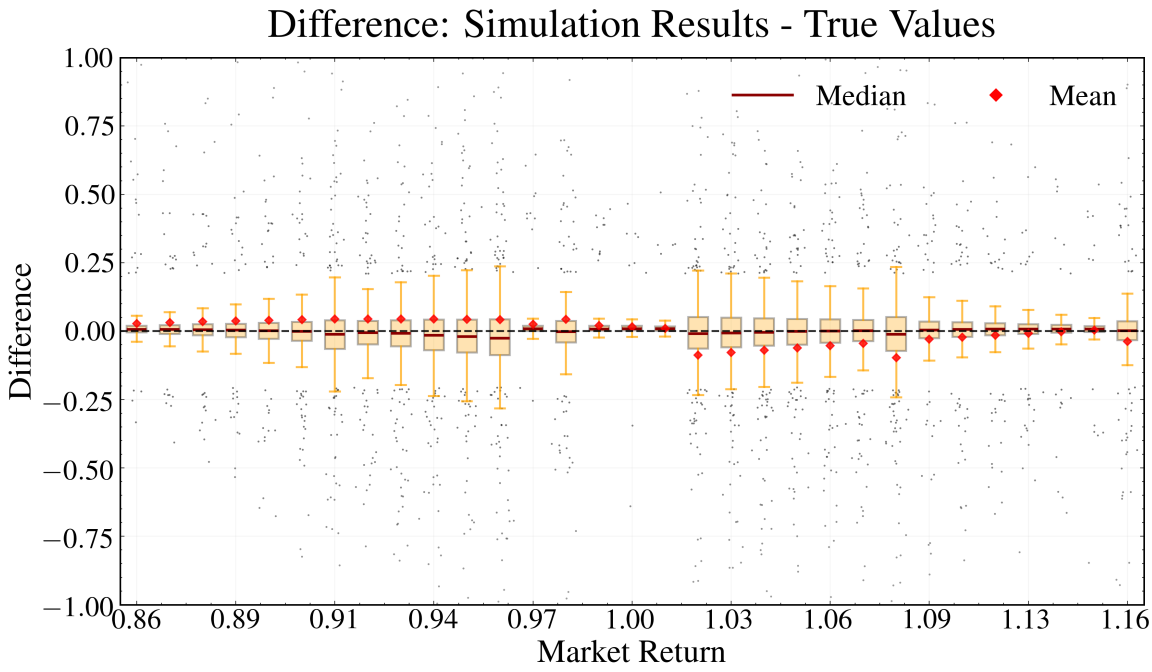

Figure 5.4: Distribution of differences between estimated and theoretical Heston stochastic discount factors across moneyness levels. The Heston economy implements the Christoffersen et al. (2013) GARCH specification with variance-dependent pricing kernel, where asset returns follow a GARCH process with time-varying conditional variance. Boxplots show median, quartiles, and range of estimation errors across 500 simulations with 5000 periods and 31 OTM options.

fixed number of options, indicating that the estimation error from the discretization of the integrals in the Carr-Madan formula remains controlled. Notably, our nonparametric approach learns the structure of the two economies without receiving input about them, except for the simulated observations, and recovers both the monotone decreasing and U-shaped patterns without imposing any structural assumptions about the SDF functional form, thereby validating its flexibility in capturing the complex risk preferences reflected in option markets.

5.2 Baseline results

Our empirical analysis reveals the non-linear behavior of the estimated SDF across different market states, providing insight into both the SDF itself and the implied optimal allocation strategies. These findings directly address the well-documented "SDF puzzle" - the empirical observation that estimated SDFs often violate the monotonically decreasing property predicted by standard economic theory with risk-averse agents.

‘ Figure 5.5 presents the estimated SDF as a function of market returns, conditioning on time-varying volatility as proxied by the VIX via the modalities specified in Section 3. The results demonstrate a near-linear pattern with a slight concavity, consistent with a heterogeneity implied by risk preferences. The SDF takes higher values during negative market returns (left side), with minimum values occurring for positive returns. This non-monotonic pattern contradicts the predictions of standard models with constant relative risk aversion, suggesting the presence of important non-linearities in investor preferences or market frictions that affect pricing. The first derivative of $\widehat{h}_\lambda$, shown in Figure 5.7, represents the optimal hedge component in our framework. The hedge component allocation is consistently positive and concave (i.e., the first derivative of the SDF is negative and U-shaped), consistent with proportionality to marginal utility and, from the perspective of a representative investor, in offsetting states of the economy in extreme states. The second derivative of the SDF illustrated in Figure 5.8 shows clearly the slight concavity/convexity effects present in the result. When flipped over the x-axis, the figure represents the optimal option allocation of the strategy. The pattern suggests a strong preference for OTM put options to hedge against crashes (positive values for negative returns) and an increasingly

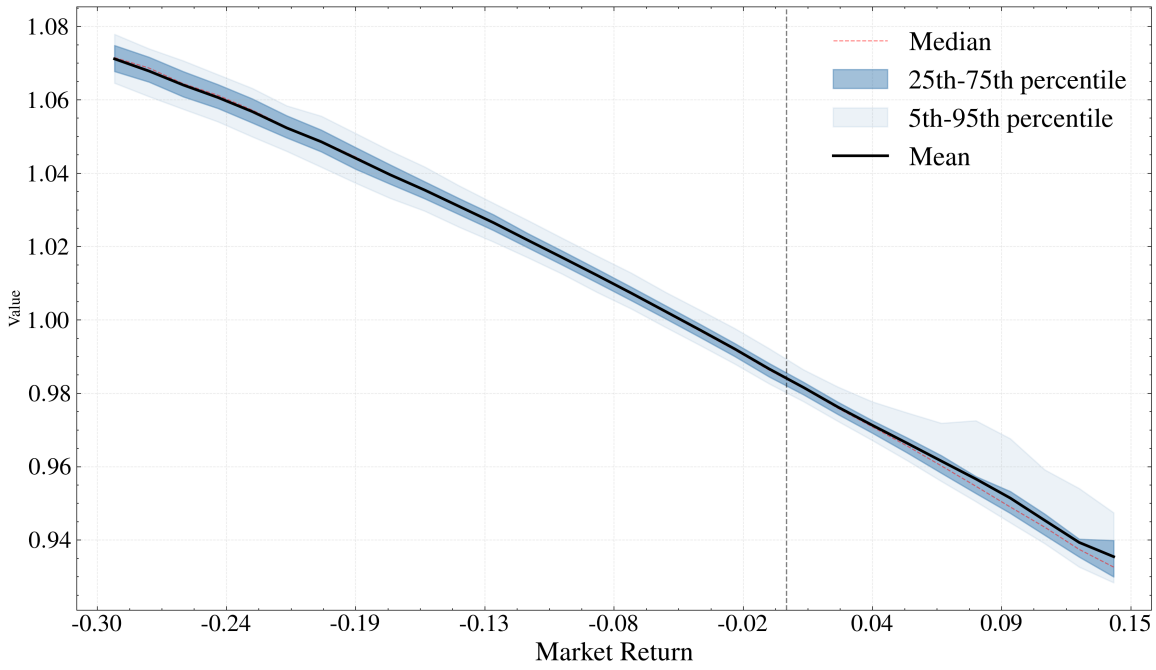

Figure 5.5: Distribution of the empirical stochastic discount factor as a function of S&P 500 market returns using monthly out-of-the-money European put and call options for the period 2000-2022. The figure displays the median (black solid line), mean (red dashed line), 25th–75th percentile ribbon (darker shaded region), and 5th-95th percentile ribbon (lighter shaded region) of SDF values across different return levels.

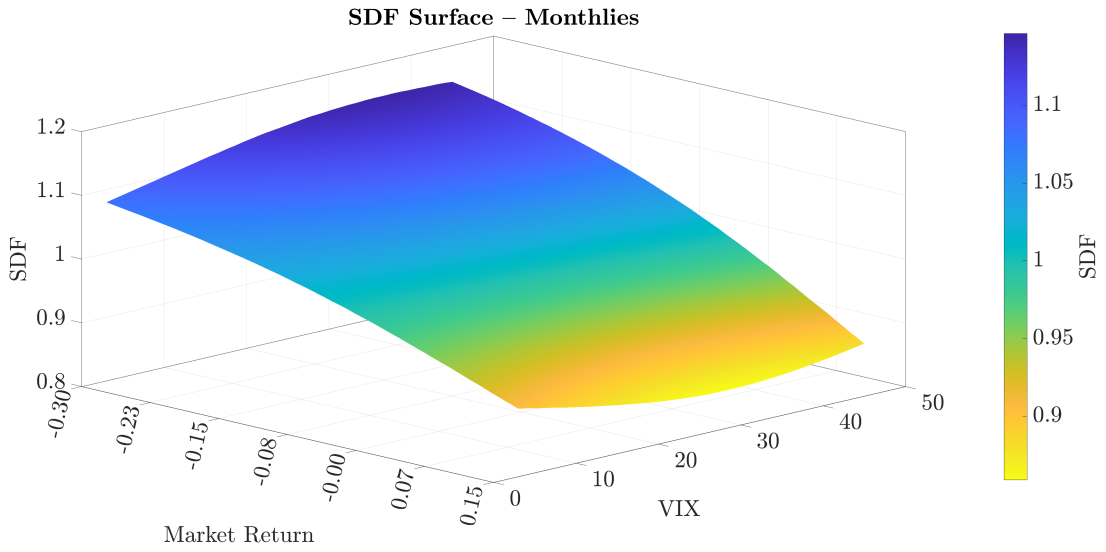

Figure 5.6: Surface of the stochastic discount factor as a function of S&P 500 market returns using monthly out-of-the-money European put and call options for the period 2000-2022, evaluated on a continuous grid of points.

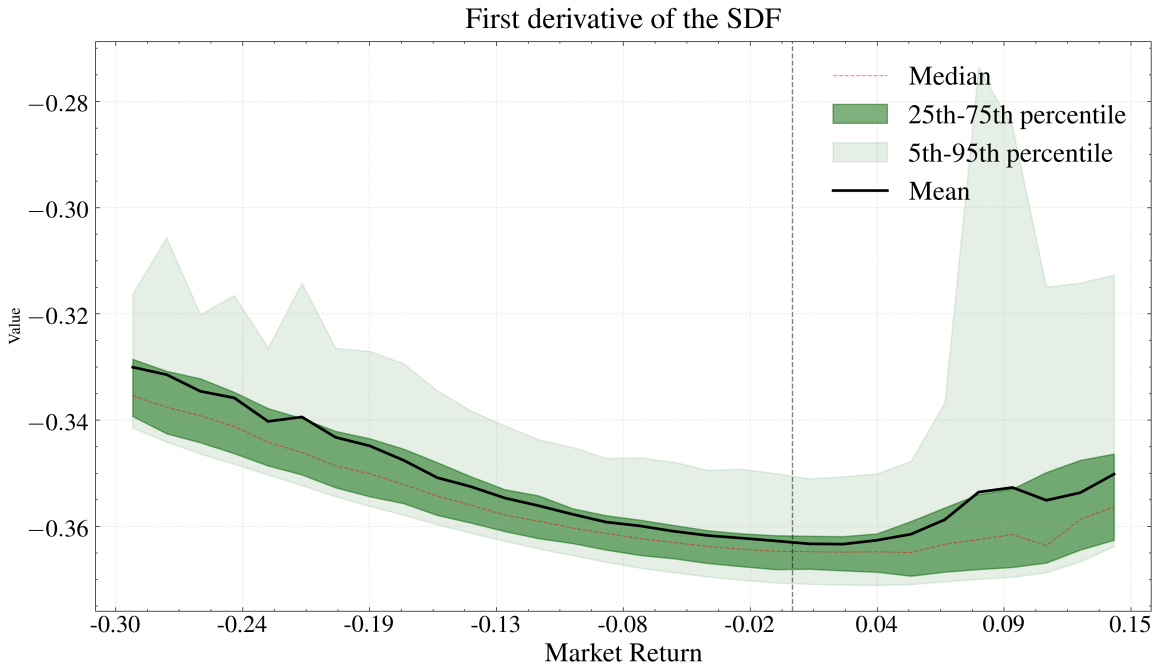

Figure 5.7: Distribution of the empirical first derivative of the SDF as a function of S&P 500 market returns using monthly out-of-the-money European put and call options for the period 2000-2022. The figure displays the median (black solid line), mean (red dashed line), 25th–75th percentile ribbon (darker shaded region), and 5th–95th percentile ribbon (lighter shaded region) of values across different return levels. Flipped over the x-axis, the figure represents the delta-hedge component of the optimal trading strategy.

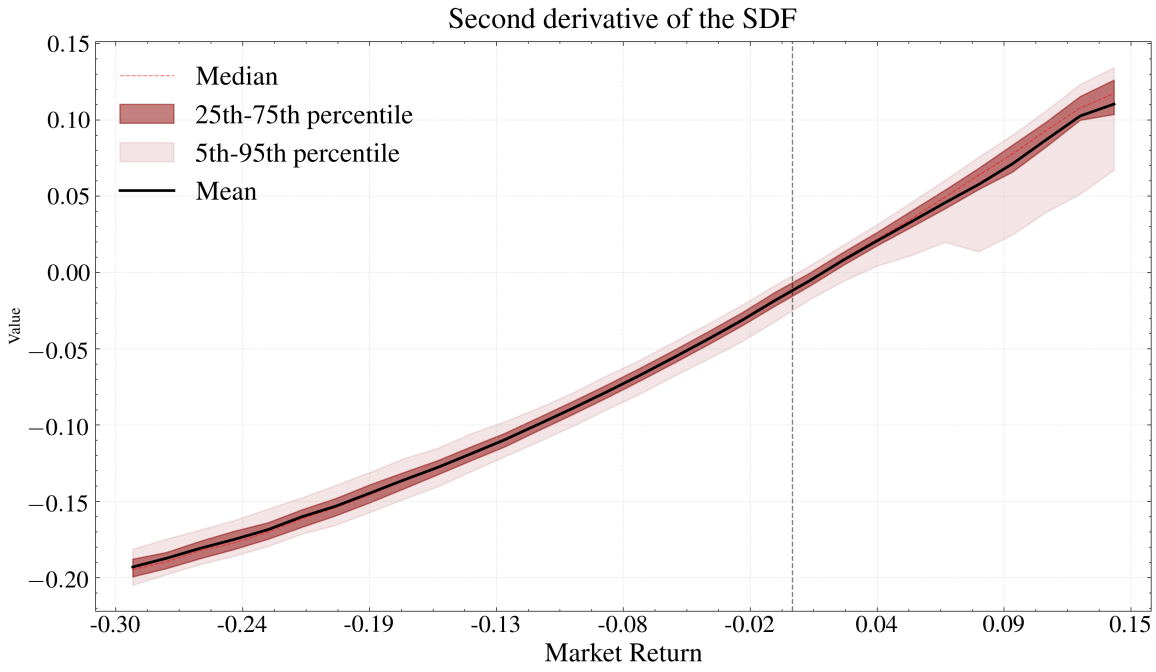

Figure 5.8: Distribution of the empirical second derivative of the SDF as a function of S&P 500 market returns using monthly out-of-the-money European put and call options for the period 2000-2022. The figure displays the median (black solid line), mean (red dashed line), 25th–75th percentile ribbon (darker shaded region), and 5th–95th percentile ribbon (lighter shaded region) of values across different return levels. Flipped over the x-axis, the figure represents the optimal trading strategy's option portfolio allocation.

negative allocation to OTM calls in good states, reflecting pessimistic behavior. If we analyze 5.8 from the perspective of the SDF (flipping over the x-axis), the nearly linear decline across market states reveals that the SDF exhibits concavity in bad states (negative returns) and convexity in good states (positive returns). This state-dependent curvature suggests heterogeneous investor preferences or behavioral biases that we see then implied in the main shape.

5.3 Conditioning the strategy on credit, volume, and CAY

The baseline specification learns the SDF through a trading strategy that conditions on forward moneyness as a state variable proxy and the VIX as a measure of time-varying volatility. While these factors capture substantial pricing variation documented in the empirical literature, our nonparametric framework also naturally accommodates higher-dimensional conditioning sets. We extend our analysis by incorporating additional factors known to affect risk premia. Following Nagel and Singleton (2011), we augment our specification with:

- **Credit risk**: Measured by the spread between AAA and BAA investment-grade bond yields, capturing flight-to-quality dynamics and systematic credit risk exposure.
- **Trading activity**: Proxied by the percentage of put options relative to total volume, reflecting market-wide demand for downside protection.
- **Macroeconomic conditions**: Incorporated through the CAY residual from Lettau and Ludvigson (2001), which captures deviations from the long-run consumption-wealth relationship.

To assess the marginal contribution of each factor, we conduct an ablation study using the sample period 2000-2019 and the previously validated regularization parameter $\hat{\lambda}$, i.e., ceteris paribus. We systematically remove each factor and estimate the resulting SDF, with Figure 5.9 presenting the comparative results.

The results reveal several key insights. First, the SDF pattern persists across all specifications, confirming the robustness of our main finding to the choice of conditioning variables. Second, while the extreme moneyness regions exhibit higher estimation uncertainty

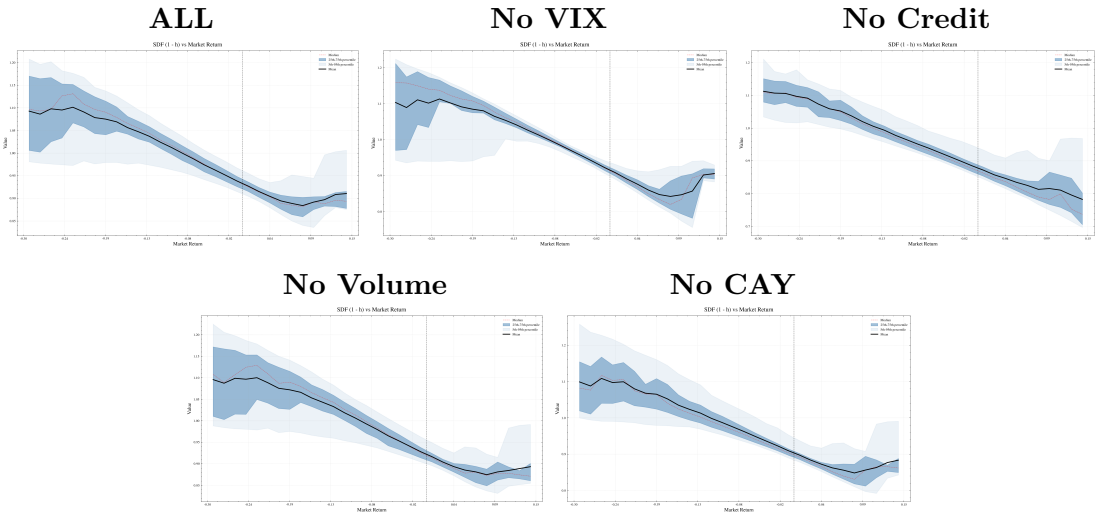

Figure 5.9: Feature importance analysis through systematic factor exclusion. Each panel displays the estimated SDF when removing the specified factor from the conditioning set, as a function of S&P 500 market returns using monthly out-of-the-money European put and call options for the period 2000-2019. Each figure displays the median (black solid line), mean (red dashed line), 25th–75th percentile ribbon (darker shaded region), and 5th–95th percentile ribbon (lighter shaded region) of values across different return levels. The conditioning set is composed of the VIX, the AAA-BAA credit spread, the option market volume, and the CAY factor from Lettau and Ludvigson (2001).

due to data sparsity, the core shape remains stable. Most notably, excluding the VIX—our primary control for time-varying volatility—introduces a subtle U-shaped deviation in the right tail. This pattern slightly hints at the behavior observed in the existing literature: when stochastic volatility is omitted from the conditioning set, the projection of the true SDF onto the return dimension can exhibit a local increase for extreme good states.

5.4 Weekly and ultra-short-maturity (0DTEs) options

To examine the economic validity of our nonparametric SDF estimation across different option maturities, we extend our analysis to weekly options (held from Friday to Friday, 2011-2022) and zero-days-to-expiration (0DTE) options (held from open to close, 2014-2023). The estimation follows our baseline methodology, using moneyness and VIX as features (with VIX open price for 0DTEs), and applying filters consistent with the monthly case. Figure 5.10 presents the estimated SDF from approximately 20,000 weekly option

contracts, revealing a pattern strikingly similar to our baseline monthly results. This consistency across weekly and monthly horizons reinforces our finding that the SDF reflects fundamental economic phenomena rather than horizon-specific artifacts.

In contrast, Figure 5.11 displays the SDF estimated from approximately 8,000 0DTE contracts, exhibiting markedly different behavior with pronounced volatility and potential non-monotonicities. This divergence from longer-maturity options is economically intuitive: at intraday horizons, microstructure effects dominate fundamental risk pricing. The 0DTE market attracts a distinct investor composition—predominantly market makers, high-frequency traders, and short-term speculators-whose trading motives differ substantially from those of the heterogeneous investors who price monthly or weekly risk. Moreover, the extreme time compression amplifies noise relative to signal, while our exogenous VIX factor (measured at open) cannot capture the rich intraday volatility dynamics that drive 0DTE pricing. Crucially, these intraday deviations do not contradict our main findings; rather, they delineate the temporal boundary where our framework transitions from capturing aggregate risk preferences to reflecting market microstructure. The persistence of the SDF shape at weekly and monthly horizons, where fundamental risk pricing dominates transient effects, validates our core contribution that properly estimated SDFs reflect genuine economic heterogeneity in preferences and beliefs.

5.5 Discussion

The SDF patterns in our results provide crucial insights into the SDF puzzle debate. Traditional explanations outside of the methodological joint hypothesis problem have centered on two main categories: behavioral mechanisms, including probability weighting and state-dependent preferences (Epstein and Schneider (2008); Grith et al. (2016)), and economic fundamentals such as heterogeneous investor preferences and state-dependent risk aversion (Bakshi et al. (2010)). Indeed, our findings suggest the puzzle reflects genuine economic phenomena rather than estimation artifacts. By circumventing the problematic $d\mathbb{Q}/d\mathbb{P}$ density ratio approach entirely through our RKHS framework, we eliminate the inverse-density noise that potentially generates spurious U-shaped patterns. This persistence aligns with recent theoretical advances, for instance, the fractional degree stochastic dominance frame-

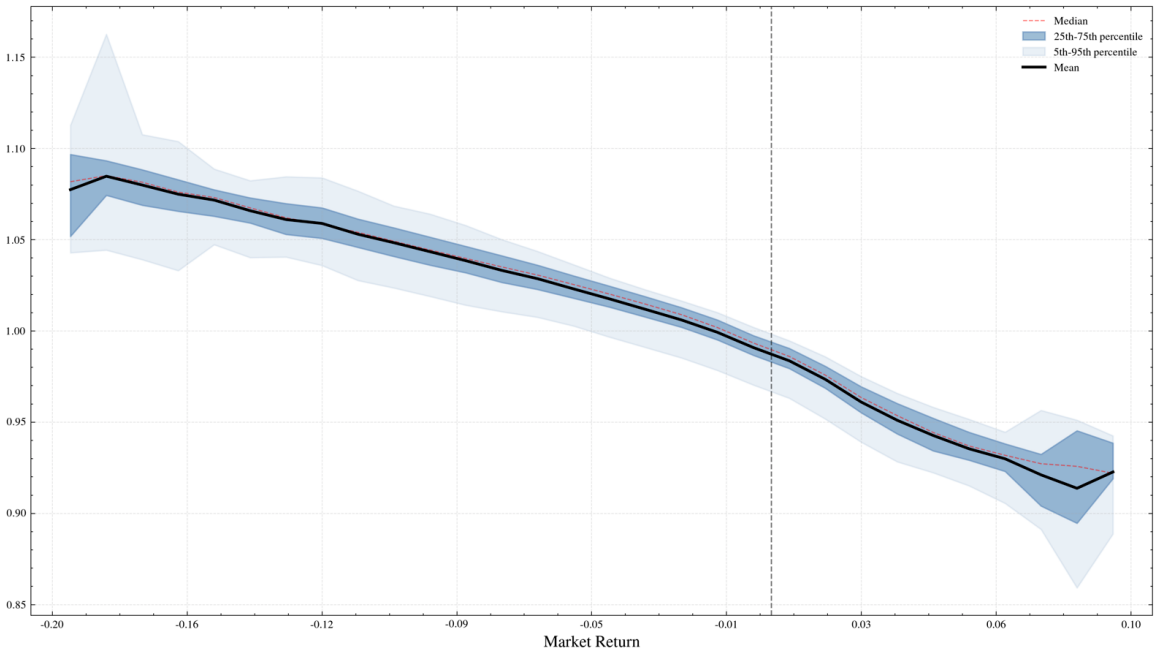

Figure 5.10: Distribution of the estimated stochastic discount factor as a function of S&P 500 market returns using weekly out-of-the-money European put and call options for the period 2011-2022. The figure displays the median (black solid line), mean (red dashed line), 25th–75th percentile ribbon (darker shaded region), and 5th–95th percentile ribbon (lighter shaded region) of SDF values across different return levels.

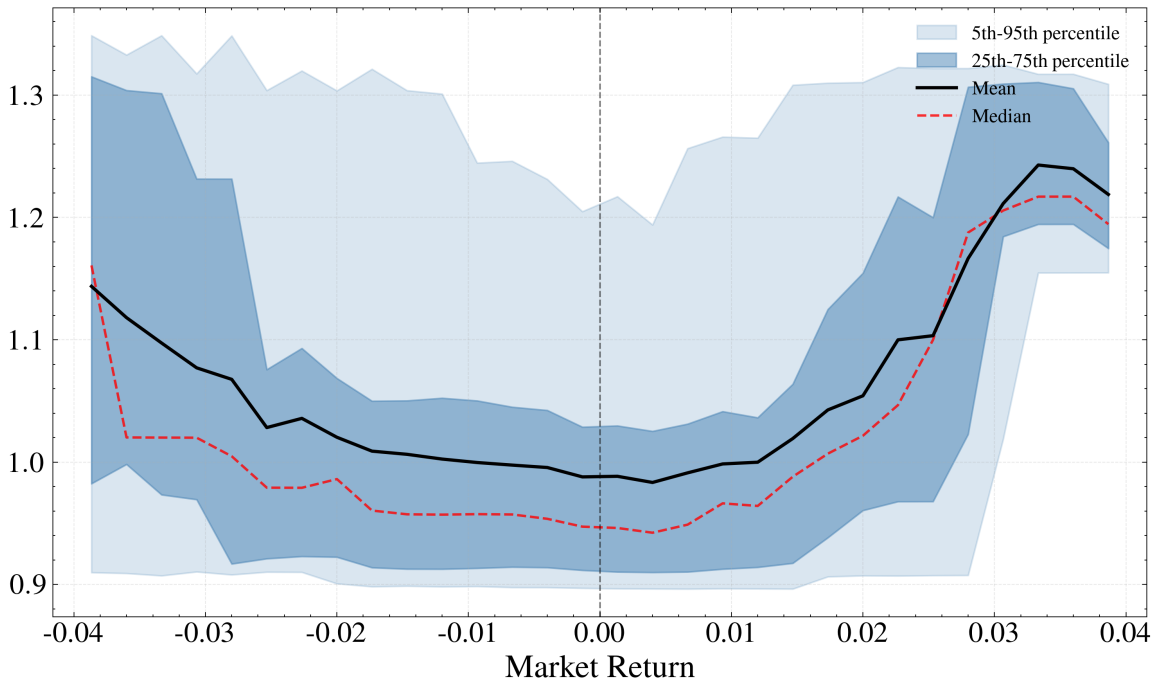

Figure 5.11: Distribution of the estimated stochastic discount factor as a function of S&P 500 market returns using ultra-short-maturity (0DTEs) out-of-the-money European put and call options for the period 2014-2023. The figure displays the median (black solid line), mean (red dashed line), 25th–75th percentile ribbon (darker shaded region), and 5th–95th percentile ribbon (lighter shaded region) of SDF values across different return levels.

work (Huang et al. (2020)), which provides rigorous foundations for concave-convex utility functions generating non-monotonic SDFs. Our nonparametric approach reveals that these patterns emerge directly from option market prices without imposing structural assumptions, suggesting that non-convex SDFs represent market participants' complex risk attitudes. The robustness of these patterns to our novel estimation methodology strengthens the case that the empirical SDF reflects fundamental aspects of how markets price risk under heterogeneous beliefs and preferences, rather than methodological shortcomings.

5.6 Nonparametric stochastic discount factor shape test

The properties of the SDF, especially its shape, carry profound economic implications. Monotonicity captures the fundamental principle of risk aversion: investors value consumption more in adverse states than in favorable ones, leading the SDF to decrease with market returns. Convexity reflects how rapidly marginal utility declines as the quantity of goods or services varies. Violations of these properties signal economic phenomena such as heterogeneous beliefs, incomplete markets, or state-dependent preferences that standard representative agent models fail to capture. Expected utility theory predicts that the SDF is monotonically decreasing and globally convex in the market returns. For differentiable functions, we can characterize monotonically decreasing behavior through a negative first derivative. Concurrently, for twice differentiable functions, we can characterize convexity as the positivity of the second derivative. To formally assess these shape properties of the estimated SDF, we perform the statistical tests as developed in Sen (2025).

5.6.1 Test description

We study two economically motivated shape restrictions on the stochastic discount factor (SDF):

1. Slope of the SDF: monotonically decreasing.

2. Curvature of the SDF: convex shape.

To separate preference-driven shape patterns from volatility regimes, we condition on market volatility and run the tests state-by-state. From the empirical data of both the monthly

and daily options (0DTEs), we first retrieve the maximum and minimum levels of the state-variable proxies, i.e., the forward moneyness and the VIX. We fix the volatility state at four levels based on the empirical VIX quartiles (low, low-mid, mid-high, high). Within each fixed VIX state, we treat v as exogenous and evaluate the SDF only along moneyness m; all derivatives are taken only with respect to m and not v. For each VIX state, we draw 100 independent grids $\mathcal{Z}$ of $n = 100$ moneyness points, generated uniformly over the OTM range used in our estimation for the respective datasets, i.e., both for monthly option data as well as 0DTEs, with ultra-short maturities. Randomizing the grids reduces sensitivity to any particular placement and provides distributional summaries. On each grid, we form a vector of local shape evaluations of the SDF: the first derivative for the monotonicity test and the second derivative for the convexity test. We then perform a joint test asking whether the entire vector can satisfy the relevant sign restriction. For each dataset and VIX state, we report: the empirical CDF (ECDF) of p-values across grids, see Figures 5.12 and 5.13, and table summaries, see Tables 5.1 and 5.2. As an additional robustness check, we also test for concavity later, and report our findings in Figure 5.14 and Table 5.3. As an additional robustness check, we also test for concavity later, and report our findings in Figure 5.14 and Table 5.3.

5.6.2 Test statistic

For each v representing a fixed VIX quartile and moneyness $(m_j)_{j=1}^n$ in grid $\mathcal{Z}$ where $n = 100$, we consider the following:

1. Monotonicity test: $-\frac{\partial}{\partial m} M(m_j, v) \geq 0$ for all $m_j \in \mathcal{Z}$;
2. Convexity test: $\frac{\partial^2}{\partial m} M(m_j, v) \geq 0$ for all $m_j \in \mathcal{Z}$.

For each pair of volatility level and moneyness grid $\mathcal{Z}$, we compute the test statistic W_T from Sen (2025, Theorem 4.6) and its critical value at the level $\alpha = 0.05$. To actually

compute the test statistic, we stack the derivative evaluations in a vector $\widehat{\boldsymbol{\theta}} \in \mathbb{R}^n$ as follows:

$$\text{Monotonicity test:} \quad \widehat{\boldsymbol{\theta}} = \left[-\frac{\partial}{\partial m} M(m_j, v) \right]_{j=1}^{n} \in \mathbb{R}^n;$$

$$\text{Convexity test:} \quad \widehat{\boldsymbol{\theta}} = \left[\frac{\partial^2}{\partial m} M(m_j, v) \right]_{j=1}^{n} \in \mathbb{R}^n;$$

Then, the test statistic W_T is a Wald-type "distance-to-feasibility" measure given as:

$$W_T := \min_{\boldsymbol{c} \geq \mathbf{0}} T(\widehat{\boldsymbol{\theta}} - \boldsymbol{c})^\top \widehat{\boldsymbol{\Omega}}_\lambda^{-1} (\widehat{\boldsymbol{\theta}} - \boldsymbol{c}), \tag{5.1}$$

where $\widehat{\boldsymbol{\Omega}}_\lambda$ is the $n \times n$ covariance matrix of deriavtive evaluations on the grid $\mathcal{Z}$. The W_T statistic measures the scaled projection error (scaled with $\widehat{\boldsymbol{\Omega}}_\lambda$) in the Mahalanobis distance to the feasibility region implied by the null hypothesis. It holds that under the *least favorable null*, the test statistic has the asymptotic distribution

$$W_T \xrightarrow{d} W \sim \chi^2_n - \bar{\chi}^2(\widehat{\boldsymbol{\Omega}}_\lambda, \mathbb{R}^n_+), \tag{5.2}$$

where $\bar{\chi}^2(\widehat{\boldsymbol{\Omega}}_\lambda, \mathbb{R}^n_+)$ is a mixture of χ^2-distributions, see Silvapulle and Sen (2001). In practice, the tail-probability of the asymptotic distribution of the test statistic is estimated via Monte Carlo replications, (Silvapulle and Sen (2001, Section 3.5)), where W_T is solved via a non-negative least squares problem, see Appendix B.2.

Remark 5.1. *In contrast to the conventional Wald test, this scenario involves a one-sided test. The test statistic, denoted as W_T, quantifies the* projection error *using the Mahalanobis distance derived from whitening by $\widehat{\boldsymbol{\Omega}}_\lambda^{-1/2}$. The asymptotic distribution of W_T under the null hypothesis H_0 is contingent upon which inequalities are active for $\boldsymbol{\theta}$. As $\boldsymbol{\theta}$ transitions further into the interior of $\mathbb{R}^n_+$, making the entries more positive, the test statistic W_T tends to decrease stochastically. The maximum value of the statistic, representing the least favorable scenario, happens when all constraints are binding, meaning all components of $\boldsymbol{\theta}$ are zero. Specifically, for any $\boldsymbol{\theta} \in \mathbb{R}^n_+$, the probability $\mathbb{P}_{\boldsymbol{\theta}}(W_T \geq c)$ is at most equal to $\mathbb{P}_{\boldsymbol{\theta}=\mathbf{0}}(W_T \geq c)$. Consequently, we set the critical values (or p-values) based on the least favorable null hypothesis, $\boldsymbol{\theta} = \mathbf{0}$.*

For completeness, we formally put the steps below:

1. **Fix the volatility state.** set v to one of the four VIX levels (low, low-mid, mid-high, high) and hold it fixed.

2. **Draw the evaluation grid.** Sample a grid $\mathcal{Z}$ of $n = 100$ moneyness points uniformly over the OTM range; repeat with 100 independent grids for robustness.

3. **Build the derivative vector.** Compute $\widehat{\boldsymbol{\theta}} \in \mathbb{R}^n$ by (i) evaluating the first derivative evaluations (for monotonicity test), or (ii) the second derivative evaluations (for convexity test), at all $m \in \mathcal{Z}$.

4. **Estimate dependence.** Compute the covariance matrix $\widehat{\boldsymbol{\Omega}}_\lambda$ of these derivative evaluations, see Appendix B.2 for details.

5. **Compute the test statistic.** Compute

$$W_T = \min_{\boldsymbol{c} \geq \mathbf{0}} \; T\,(\widehat{\boldsymbol{\theta}} - \boldsymbol{c})^\top \widehat{\boldsymbol{\Omega}}_\lambda^{-1} (\widehat{\boldsymbol{\theta}} - \boldsymbol{c}),$$

   i.e., the non-negative least-squares (cone) projection distance after "whitening" by $\widehat{\boldsymbol{\Omega}}_\lambda^{-1/2}$.

6. **Get inference and summarize.** Use the asymptotic distribution of the test statistic W_T to obtain p-values and critical values at $\alpha = 0.05$. Report, per VIX state: (i) acceptance rate (share of $p > \alpha$ across grids), (ii) median$(\log_{10} W_T)$, (iii) median$(\log_{10} \text{crit}_\alpha)$, and (iv) the gap median$(\log_{10} W_T - \log_{10} \text{crit}_\alpha)$, as well as the ECDFs of the p-values.

5.6.3 Test results

We put below the results of the monotonicity and convexity test of the empirical SDF. Later, as an additional sanity test, we test for concavity as well.

Monotonicity test. Figure 5.12 shows the empirical CDFs (ECDFs) of the p-values for the monotonicity tests for both datasets. For the baseline monthly options data (left panel),

the curves lie on $p = 1$ for all four volatility states; this indicates strong evidence of no rejection of a monotonically decreasing SDF across volatility states. In contrast, for 0DTEs (right panel), the ECDFs at low and low-mid volatility levels concentrate near zero, leading to rejections of the monotonically decreasing behavior of the SDF; while at mid-high level, it coincides with the critical level, and it shifts right and aligns with no rejection for the highest level. This points to evidence of nonmonotonic behavior and the presence of a possible U-shaped pattern. This behavior is also seen from Table 5.1 that summarizes the same test results. An additional insight is seen from the logarithm of the test statistic. Notice that it does not vary considerably across the volatility states. However, we notice that while that for the monthlies is comparatively much smaller than 0DTEs, indicating that (i) the SDF for the monthlies satisfies the null hypothesis (in this case: monotonically decreasing) comfortably, while that for the 0DTEs is either rejected/close to being rejected. In fact, in Figures 5.15 and 5.16, we have the actual plots of the empirical SDF standardized by volatility states, for both monthly options and 0DTEs.

Monotonicity				
Dataset	VIX	Accept. rate (%)	Median $\log_{10}(W_T)$	Median $\log_{10}(\text{crit})$
Monthlies	Low	100.000	-15.654	-15.654
	Low-Mid	100.000	-15.654	-15.654
	Mid-High	100.000	-15.654	-15.654
	High	100.000	-15.654	-15.654
0DTEs	Low	0.000	1.847	0.645
	Low-Mid	0.000	1.278	0.691
	Mid-High	98.000	0.718	0.730
	High	100.000	0.195	0.685

Table 5.1: Monotonicity test results across volatility states (Uniform grids, $n = 100$, $\alpha = 0.05$). For each dataset and VIX level (low, low-mid, mid-high, high) we report the acceptance rate at level $\alpha = 0.05$ (in percent), the median of $\log_{10}(W_T)$ across grids, and the median of $\log_{10}(\text{crit}_\alpha)$.

Convexity test. Figure 5.13 shows the empirical CDFs (ECDFs) of the p-values for the convexity tests for both datasets. For the baseline monthly options (left panel), we fail to

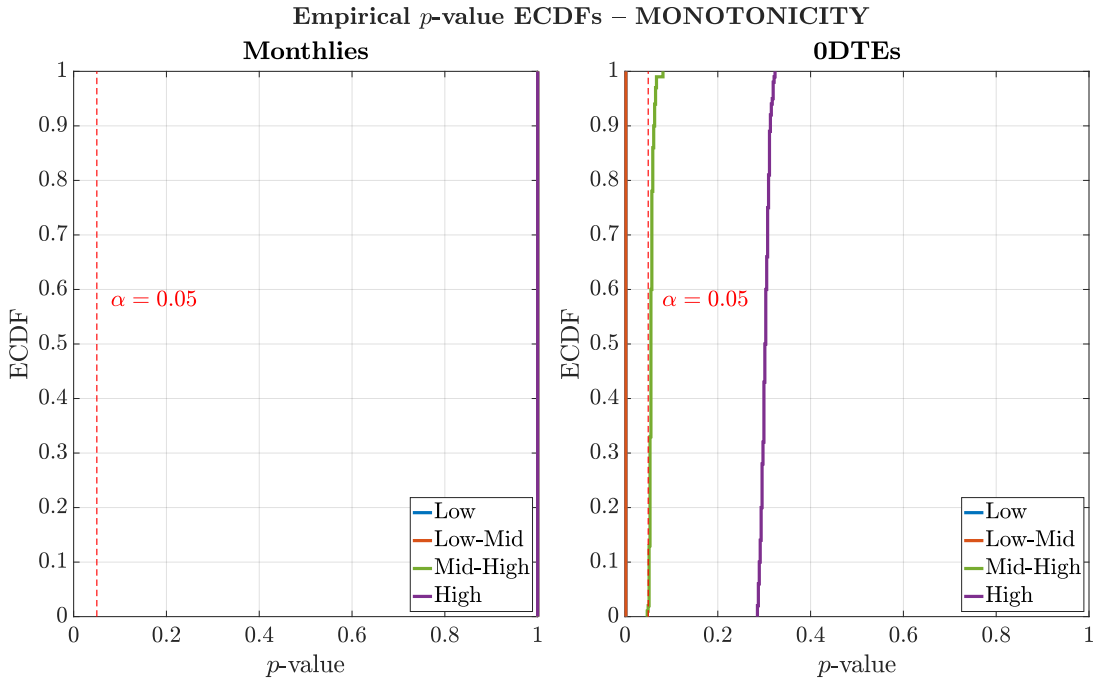

Figure 5.12: Empirical CDFs (ECDFs) of p-values for monotonicity of the SDF computed by the Wald-type test statistic W_T. The statistic is computed on 100 random grids of 100 points, drawn uniformly over the moneyness range of OTM monthly option data (2000–2022) and 0DTE option data (2014-2022), stratified by VIX quartiles. Monthly options (left) and 0DTE options (right), at level $\alpha = 0.05$, marked by the vertical red dashed line.

reject the null hypothesis for all the volatility states, except for the low case. We observe that the curves shift towards the right for increasing VIX. In contrast, for the 0DTEs (right panel), we find strong evidence of a convex shape as the ECDFs concentrate around $p = 1$, across all volatility states. This behavior is also seen from the results of Table 5.2. Finally, we see from Figure 5.15, that across the different VIX levels, the empirical SDF recovered using monthly options data does not have a convex shape, but rather an almost-linear, even an ever slightly concave shape; while for the 0DTEs, it is completely the converse case, since we see a pronounced U-shaped pattern, except for the highest volatility state, where it is almost linear.

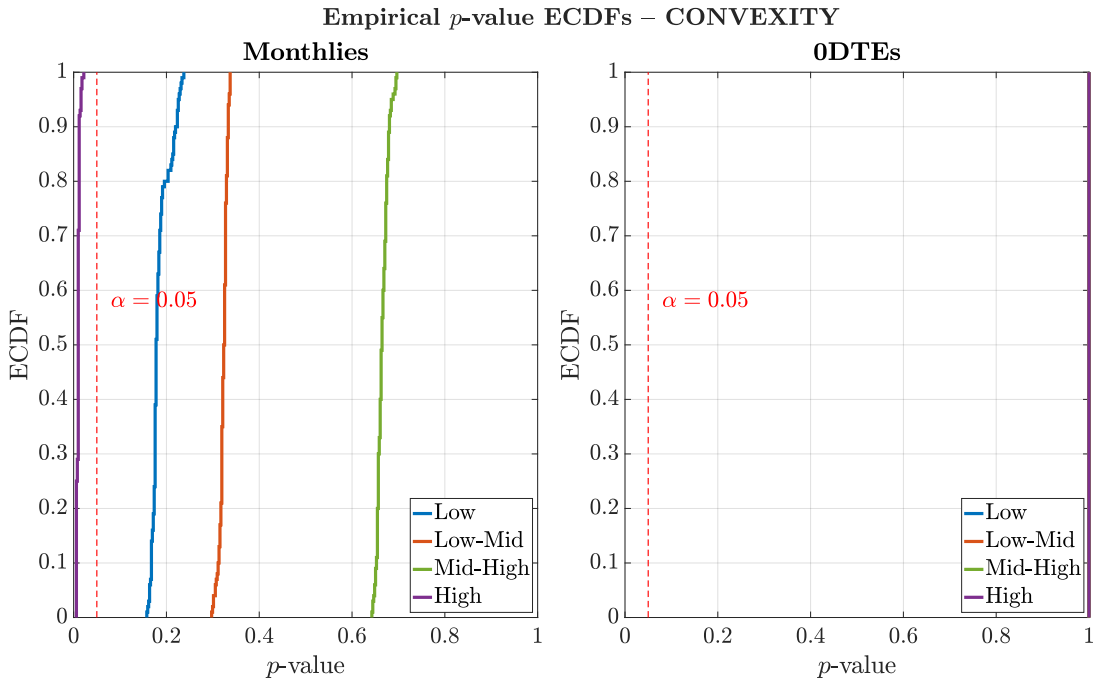

Figure 5.13: Empirical CDFs (ECDFs) of p-values for convexity of the SDF computed by the Wald-type test statistic W_T. The statistic is computed on 100 random grids of 100 points, drawn uniformly over the moneyness range of OTM monthly option data (2000–2022) and 0DTE option data (2014-2022), stratified by VIX quartiles. Monthly options (left) and 0DTE options (right), at level $\alpha = 0.05$, marked by the vertical red dashed line.

Concavity test. To further investigate the shape of the SDF, we now set the null hypothesis to concavity, to be that it is *concave*, that is, we set:

1. Concavity test: $-\frac{\partial^2}{\partial m} M(\cdot, v) \geq 0$ for all $m \in \mathcal{Z}$.

Convexity				
Dataset	VIX	Accept. rate (%)	Median $\log_{10}(W_T)$	Median $\log_{10}(\text{crit})$
Monthlies	Low	100.000	0.396	0.671
	Low-Mid	100.000	0.172	0.699
	Mid-High	100.000	-0.457	0.741
	High	0.000	0.947	0.717
0DTEs	Low	100.000	-15.654	-15.654
	Low-Mid	100.000	-15.654	-15.654
	Mid-High	100.000	-15.654	-15.654
	High	100.000	-15.654	-15.654

Table 5.2: Convexity test results across volatility states (Uniform grids, $n = 100$, $\alpha = 0.05$). For each dataset and VIX level (low, low-mid, mid-high, high) we report the acceptance rate at level $\alpha = 0.05$, the median of $\log_{10}(W_T)$ across grids, and the median of $\log_{10}(\text{crit}_\alpha)$.

With this null, we repeat the same experiment 100 times as in the previous case for both the monthly options, as well as the 0DTEs, and show our findings in Figure 5.14 and Table 5.3. In Figure 5.14, for the baseline monthly options data (left panel), the curves lie to the right of $\alpha = 0.05$, indicating no rejections of a concave SDF across volatility states. In contrast, for 0DTEs (rightpanel), the ECDFs concentrate near zero, except for the highest volatility level, leading to frequent rejections of the concave behavior of the SDF. Similar findings are also obtained from Table 5.3. Moreover, as discussed in the previous paragraph, the findings of the concavity test, when coupled with the plots in Figures 5.15 and 5.16 corroborate the empirical results that the SDF of the monthly options data is near-linear across all VIX levels, whereas for the 0DTEs, the SDF has a U-shape.

Thus, in conclusion, taken together, the evidence points to a monotonically decreasing and near-linear SDF, with negligible curvature at monthly horizons. Importantly, we do not find evidence of any U-shaped pattern; the situation changes for the ultra-short maturity options data, where we find the presence of U-shaped and non-monotonic behavior largely across most volatility states.

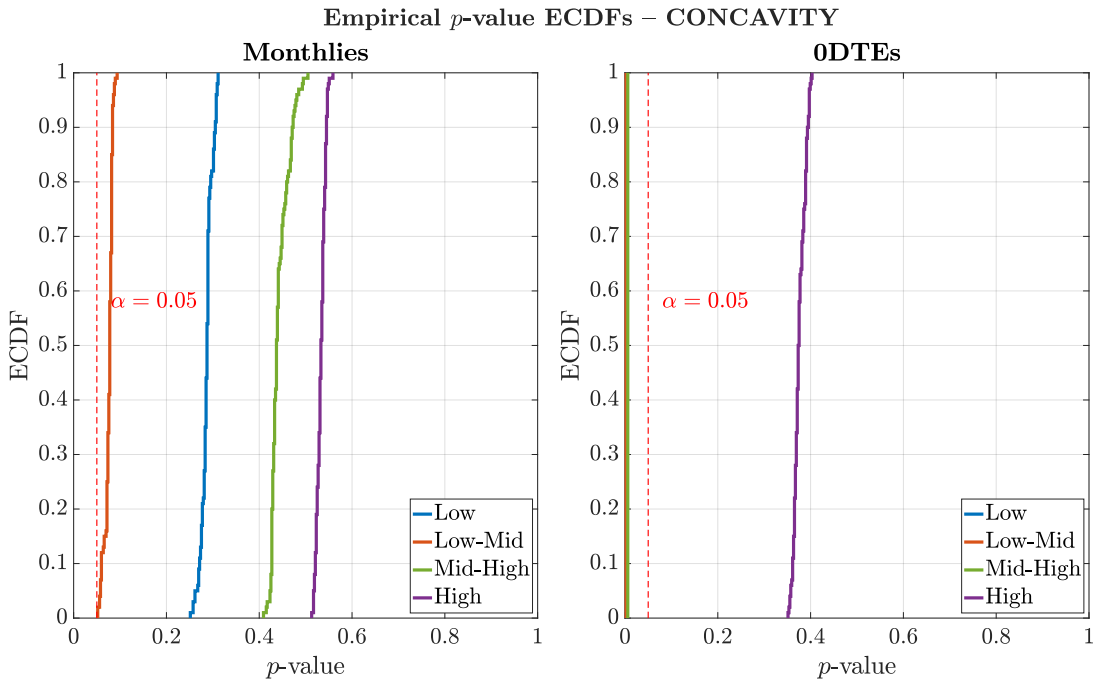

Figure 5.14: Empirical CDFs (ECDFs) of p-values for monotonicity of the SDF computed by the Wald-type test statistic W_T. The statistic is computed on 100 random grids of 100 points, drawn uniformly over the moneyness range of OTM monthly option data (2000–2022) and 0DTE option data (2014-2022), stratified by VIX quartiles. Monthly options (left) and 0DTE options (right), at level $\alpha = 0.05$, marked by the vertical red dashed line.

Concavity				
Dataset	VIX	Accept. rate (%)	Median $\log_{10}(W_T)$	Median $\log_{10}$(crit)
Monthlies	Low	100.000	0.247	0.671
	Low-Mid	100.000	0.626	0.699
	Mid-High	100.000	-0.027	0.741
	High	100.000	-0.141	0.717
0DTEs	Low	0.000	1.965	0.427
	Low-Mid	0.000	1.442	0.420
	Mid-High	0.000	1.003	0.475
	High	100.000	-0.591	0.414

Table 5.3: Concavity test results across volatility states (Uniform grids, $n = 100$, $\alpha = 0.05$). For each dataset and VIX level (low, low-mid, mid-high, high) we report the acceptance rate at level $\alpha = 0.05$, the median of $\log_{10}(W_T)$ across grids, and the median of $\log_{10}(\text{crit}_\alpha)$.

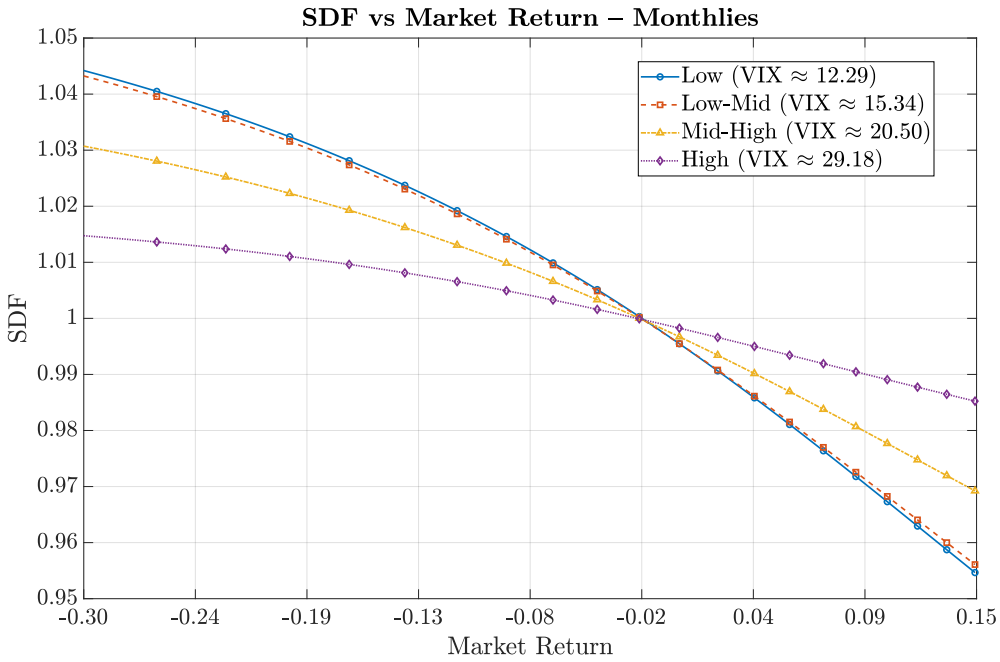

Figure 5.15: SDF across VIX quartiles as a function of monthly out-of-the-money European put and call options for the period 2000–2022, evaluated on a continuous grid of market returns.

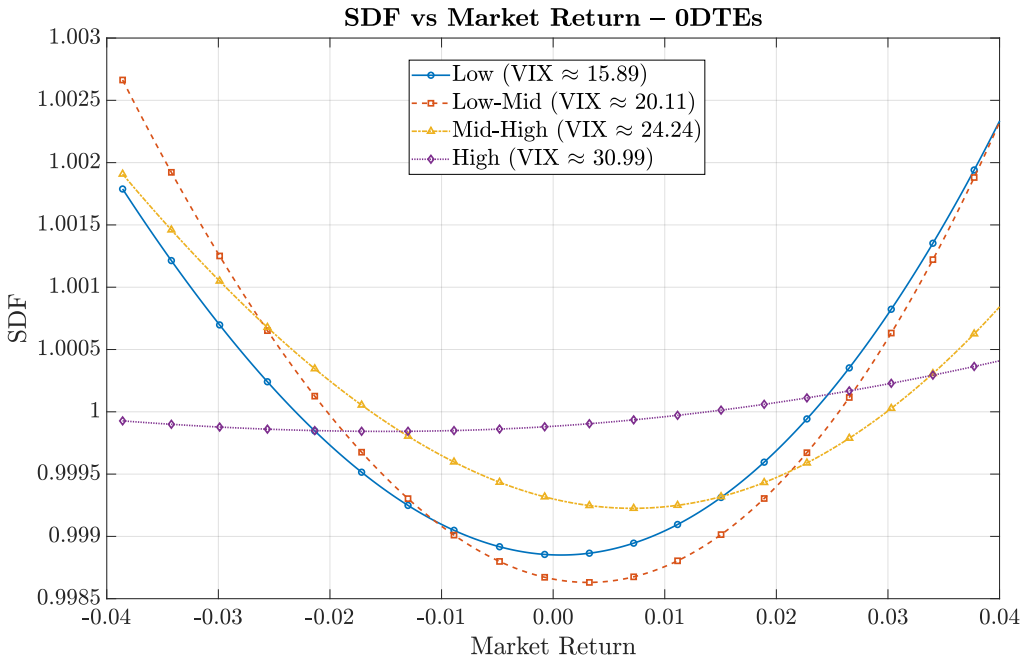

Figure 5.16: Ultra-short-term SDF across VIX quartiles as a function of monthly out-of-the-money European put and call options for the period 2014-2022, evaluated on a continuous grid of market returns.

5.7 Validating the fit for conditional moments and variance risk premium

A critical test of our recovered SDF is whether the implied physical probability measure $\mathbb{P}$ generates conditional expectations that match realized moments in-sample. This validation directly assesses whether our methodology correctly transforms the option-implied risk-neutral measure $\mathbb{Q}$ into an economically meaningful physical measure through $M(R_{t+1}, v_t) = 1 - h(R_{t+1}, v_t)$.

5.7.1 Extracting conditional expectations

Given the estimated SDF $M_t(R_{t+1}, v_t)$ at each date t, we recover conditional expectations under both the risk-neutral measure $\mathbb{Q}$ and the physical measure $\mathbb{P}$ via the fundamental pricing relationship. The SDF $M_t(R_{t+1}, v_t)$ represents the Radon–Nikodým derivative $d\mathbb{Q}/d\mathbb{P}$ conditioned on observables, satisfying

$$\mathbb{E}_t^{\mathbb{P}}[g(R_{t+1})] = \mathbb{E}_t^{\mathbb{Q}}\left[\frac{g(R_{t+1})}{M_t(R_{t+1}, v_t)}\right]. \tag{5.3}$$

This formula arises directly from the change of measure: for any measurable function g, the physical expectation equals the risk-neutral expectation weighted by the inverse SDF. We implement this transformation using a discrete representation of the risk-neutral distribution. Following Schneider (2019), we construct $k+1$ scenarios $\{x_{i,t}\}_{i=0}^{k}$ with risk-neutral probabilities $\{q_{i,t}\}_{i=0}^{k}$ that exactly match the first $2k+1$ risk-neutral moments via Gaussian quadrature. In our application we set $k = 2$, yielding three scenarios per period. The risk-neutral moments

$$\gamma_{j,t} = \mathbb{E}_t^{\mathbb{Q}}[R_{t+1}^j]$$

are extracted directly from the cross-section of option prices via the Carr–Madan replication formula (see also Schneider and Trojani (2019)). For gross returns $R_{t+1} = S_{t+1}/S_t$, the j-th moment is given by

$$\gamma_{j,t} = 1 + j(j-1)\left[\int_1^\infty m^{j-2}\frac{C_t(K)}{F}\,dk \;+\; \int_0^1 m^{j-2}\frac{P_t(K)}{F}\,dk\right], \tag{5.4}$$

where $C_t(K)$ and $P_t(K)$ denote call and put prices at moneyness m. We discretize these integrals using the trapezoidal rule over the available grid of option strikes. The martingale property

$$\gamma_{1,t} = \mathbb{E}_t^{\mathbb{Q}}[R_{t+1}] = 1$$

holds by no-arbitrage under the forward measure. Given the moments $\{\gamma_{j,t}\}_{j=0}^{2k+1}$, the quadrature algorithm recovers scenarios and probabilities by solving a system of polynomial equations (see Schneider (2019) for details). Physical expectations are then computed directly as

$$\mathbb{E}_t^{\mathbb{P}}[g(R_{t+1})] = \sum_{i=0}^{k} \frac{g(x_{i,t})}{M_t(x_{i,t}, v_t)} q_{i,t}, \tag{5.5}$$

where $M_t(x_{i,t}, v_t)$ denotes the SDF evaluated at scenario $x_{i,t}$. This formulation avoids the need to explicitly normalize physical probabilities, instead computing expectations directly via (5.3) in its discrete form. The approach naturally handles arbitrary test functions g without requiring separate density estimation under $\mathbb{P}$.

5.7.2 In-sample fit

We evaluate the in-sample fit of conditional expectations via the regression,

$$g(R_{t+1}) = \alpha + \beta \cdot \mathbb{E}_t^{\mathbb{P}}[g(R_{t+1})] + \varepsilon_{t+1}. \tag{5.6}$$

Table 5.5 reports results for four moment functions. Predictability of returns is in line with what is expected in the literature. Higher-order moments show substantially stronger predictability and suggest our SDF captures tail risk both in the monthly and weekly cases.

Table 5.4: In-Sample fit of Conditional Physical Expectations: Weekly Options. Regressions $g(R_{t+1}) = \alpha + \beta \cdot \mathbb{E}_t^{\mathbb{P}}[g(R_{t+1})] + \varepsilon_{t+1}$ for weekly S&P 500 options. Physical expectations constructed via equation (5.3). Sample: January 2011 to December 2022 ($N = 543$).

Moment	**R^2**	**N**
Excess Return	0.0084	503
Variance	0.3639	503
Third Moment	0.2002	503
Fourth Moment	0.2404	503

Moment	**R^2**	**N**
Excess Return	0.0155	236
Variance	0.1997	236
Third Moment	0.0170	236
Fourth Moment	0.1186	236

Table 5.5: In-Sample fit for Conditional Physical Expectations: Monthly Options. This table reports regressions $g(R_{t+1}) = \alpha + \beta \cdot \mathbb{E}_t^{\mathbb{P}}[g(R_{t+1})] + \varepsilon_{t+1}$ for realized moments $g(R_{t+1})$, where physical expectations are constructed via equation (5.3) using the estimated SDF. The sample comprises monthly S&P 500 options from 2000 to 2022 ($N = 236$). Variance is defined as $(R_{t+1} - 1)^2$; third and fourth moments are central moments.

Figures 5.17 and 5.18 present fitting error diagnostics. Residuals exhibit near-zero mean with volatility clustering during crisis periods (2008–2009, 2020). Variance residuals show stronger concentration around zero, reflecting superior fit, though with heteroskedasticity in high-volatility regimes.

5.7.3 Variance risk premium

The variance risk premium (VRP), defined here as the difference between real-world and risk-neutral variance expectations, measure the cost of hedging against variance,

$$\text{VRP}_t = \mathbb{E}_t^{\mathbb{P}}[(R_{t+1} - 1)^2] - \mathbb{E}_t^{\mathbb{Q}}[(R_{t+1} - 1)^2]. \tag{5.7}$$

Figure 5.19 shows that monthly VRP is predominantly negative throughout our sample as the literature would expect, averaging -0.5 basis points per month. Negative VRP in-

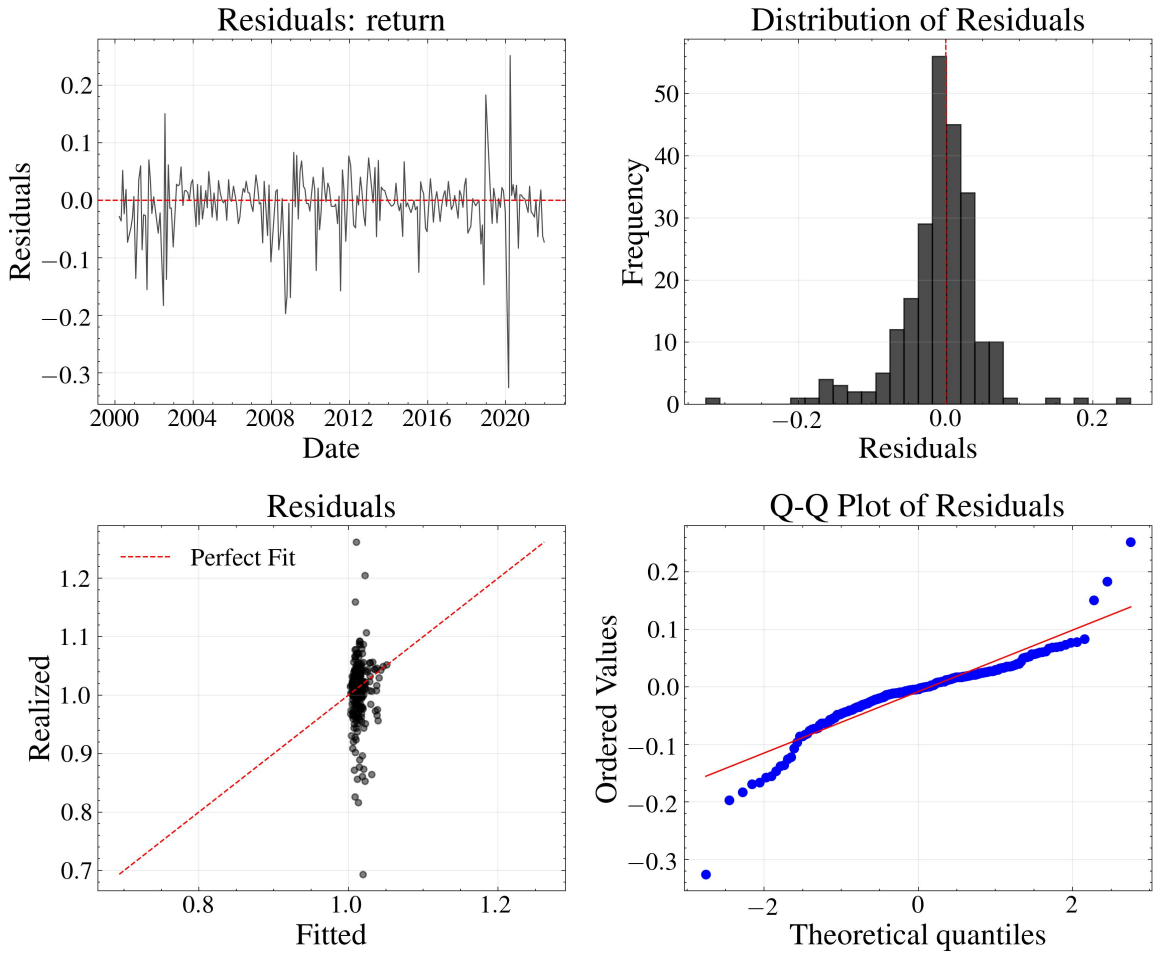

Figure 5.17: Return Residuals Diagnostics: Monthly Options. Residuals $\varepsilon_{t+1} = R_{t+1} - \mathbb{E}_t^{\mathbb{P}}[R_{t+1}]$ from monthly S&P 500 options. Top left: time series of fitting errors. Top right: error distribution (histogram). Bottom left: predicted versus realized returns with 45-degree reference line. Bottom right: quantile-quantile plot against normal distribution. Sample: 2000 to 2022.

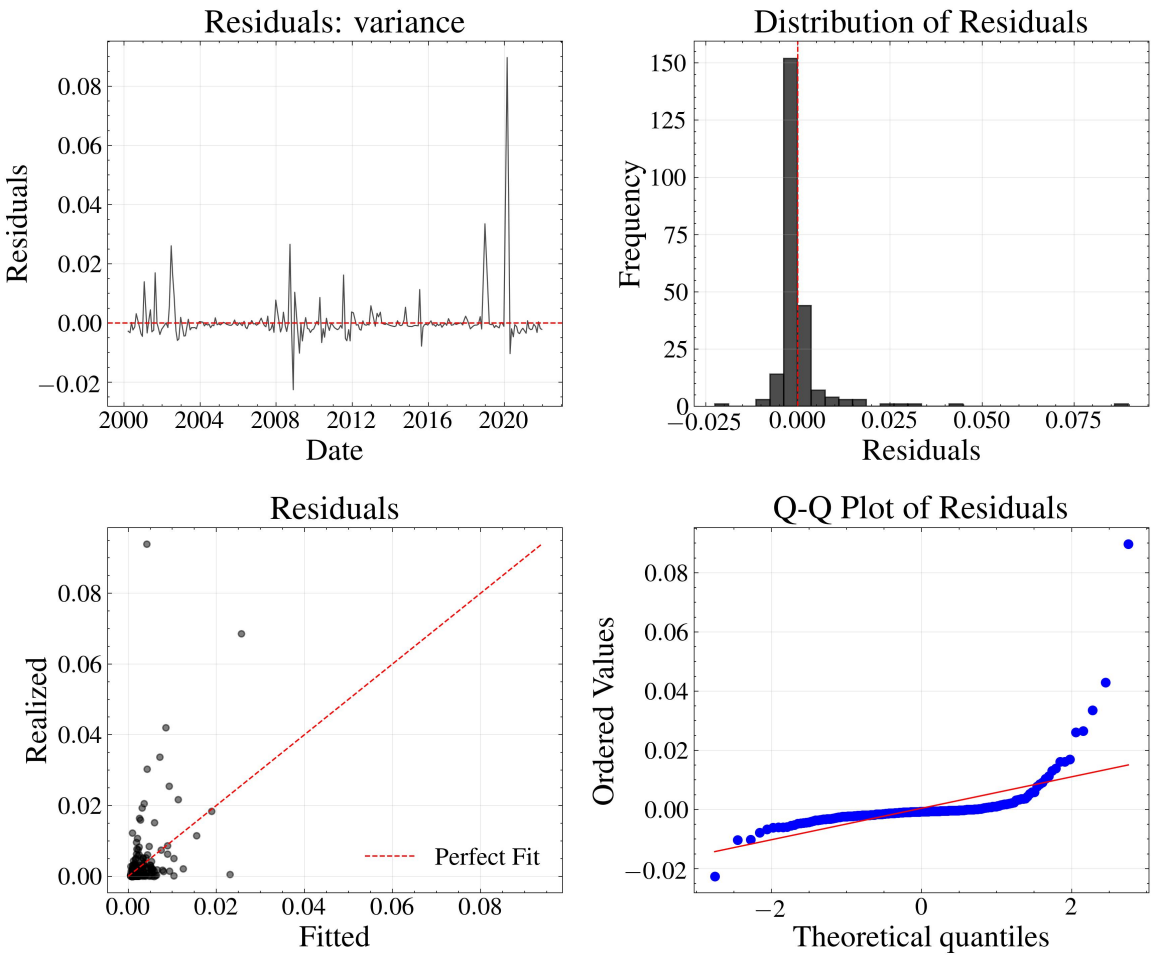

Figure 5.18: Variance Residuals Diagnostics: Monthly Options. Residuals $\varepsilon_{t+1} = (R_{t+1} - 1)^2 - \mathbb{E}_t^{\mathbb{P}}[(R_{t+1} - 1)^2]$ from monthly S&P 500 options. Top left: time series of fitting errors. Top right: error distribution. Bottom left: predicted versus realized variance with 45-degree reference line. Bottom right: quantile-quantile plot. Sample: 2000 to 2022.

dicates investors accept paying to hedge volatility, i.e. volatility being a priced risk factor. The magnitude and sign of the VRP reveal distinct dynamics across crisis types. During the 2008–2009 financial crisis, VRP becomes substantially negative (reaching −6 basis points), consistent with an endogenous crisis where markets anticipated the buildup of systemic risk and investors increasingly demanded volatility hedges. In contrast, the 2020 COVID-19 pandemic exhibits positive VRP episodes, reflecting the exogenous nature of the unanticipated shock. The positive VRP observed in the early sample period (2000–2005) likely reflects data limitations, as monthly option markets were less liquid and had fewer traded strikes during this period, leading to less precise estimation of risk-neutral moments. Given that our SDF estimation relies on cross-sectional option data at each date rather than time-series properties, this early-sample data quality does not affect our main results.

5.7.4 Scenario probability dynamics

Figure 5.20 displays the three-state discrete scenario representation ($k = 2$) underlying our conditional expectations. The model captures downside, middle, and upside scenarios with time-varying returns and probabilities. Scenario returns widen during market downturns reflecting increased return dispersion. During normal times, substantial probability mass concentrates in the middle scenario.

These validation results provide strong support for the economic consistency of our nonparametric SDF methodology. The conditional physical expectations demonstrates successful extraction of economically meaningful information from option prices. The predominantly negative variance risk premium confirms correct capture of investor risk aversion toward variance risk. Collectively, these results confirm that the physical probability measure implied by our estimated SDF generates conditional expectations consistent with realized moments, variance risk premia, reinforcing the reliability of our main SDF estimates. Further results for weekly options, which corroborate these findings, are presented in Appendix E.

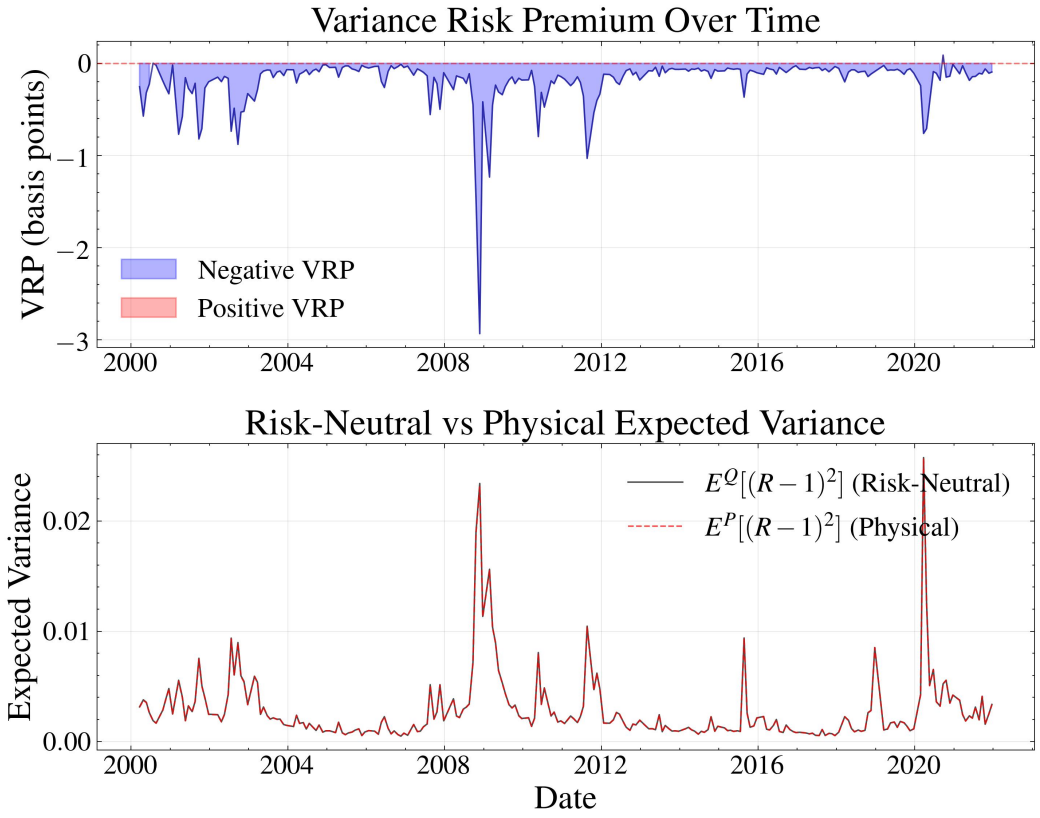

Figure 5.19: Variance Risk Premium Time Series: Monthly Options. Top panel: VRP defined as $\mathbb{E}_t^{\mathbb{P}}[(R_{t+1}-1)^2] - \mathbb{E}_t^{\mathbb{Q}}[(R_{t+1}-1)^2]$. Negative values (blue) indicate hedging demand for variance protection; positive values (red) indicate periods where physical variance expectations exceed risk-neutral expectations. Bottom panel: risk-neutral variance $\mathbb{E}_t^{\mathbb{Q}}[(R_{t+1}-1)^2]$ (solid) and physical variance $\mathbb{E}_t^{\mathbb{P}}[(R_{t+1}-1)^2]$ (dashed). Sample: monthly S&P 500 options, 2000 to 2022.

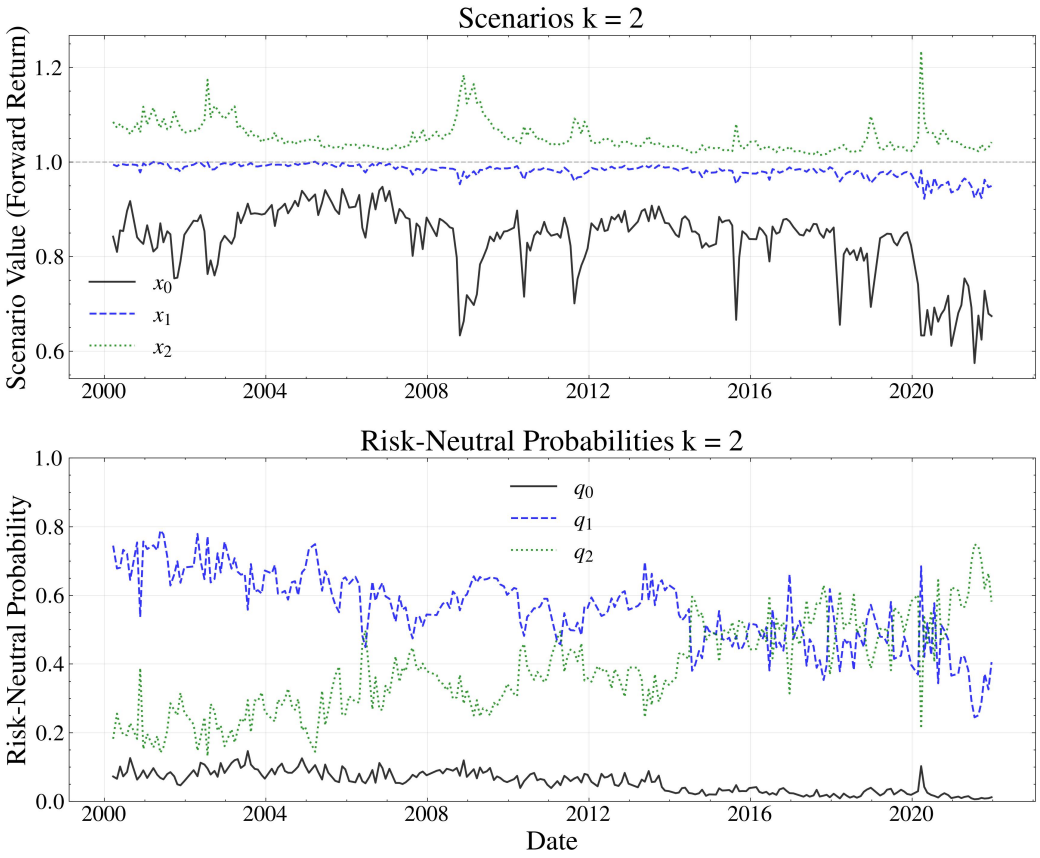

Figure 5.20: Three-State Scenario Representation: Monthly Options. Discrete scenario structure for the $k = 2$ model. Top panel: scenario returns x_0 (solid, downside), x_1 (dashed, middle), and x_2 (dotted, upside). Bottom panel: risk-neutral probabilities q_0, q_1, and q_2. By construction, $\sum_{i=0}^{2} q_i = 1$ and $\sum_{i=0}^{2} q_i x_i = 1$ (martingale condition). Sample: monthly S&P 500 options, 2000 to 2022.

6 Conclusion

This paper introduces a novel approach for the nonparametric estimation of the stochastic discount factor (SDF), significantly relaxing assumptions commonly made in the extant literature. As the key conceptual innovation, we cast the SDF estimation problem into a portfolio problem, by means of finding the allocation function that jointly determines the option allocation and hedge ratio that maximizes the empirical Sharpe ratio over all twice differentiable derivative claims in the S&P 500 market. Our methodology exploits the well-known duality between the Hansen-Jagannathan minimum-variance discount factor and the mean-variance efficient frontier, and by working directly in the payoff space avoids separate estimation of risk-neutral and physical densities.

Empirically, we recover SDF that are highly heterogeneous over different states of the world and different maturities. On average, they exhibit near-linear monotonically decreasing shapes for monthly horizons, and rapid non-monotonic swings between extreme concavity (convexity) when volatility is low (high) for ultra-short 0DTE horizons. We find only very weak evidence of the U-shaped patterns commonly reported in the literature for 0DTE horizons, and only when volatility is high. Monthly horizons do not accommodate a U-shape according to our evidence. Robustness analyses strengthen these core findings. Feature importance across multiple conditioning factors reveals that while the VIX remains the most critical control for time-varying volatility, the fundamental shape of the SDF persists even when incorporating credit risk, trading activity, and macroeconomic conditions.

Crucially, our methodology also allows for formal testing of the shapes of the estimated SDF. For monthly options, we can not reject the null hypothesis of monotonicity that goes along with the economic notion of risk aversion. From the near-linear shape, we can, however, neither reject convexity, nor concavity. In economic terms, we thus can not confirm risk prudence or downside risk aversion to be present in this market. For ultra-short maturity 0DTE options, the test rejects the null hypothesis of monotonicity, while it can not reject the null of convexity, and thus a possible U-shape, in particular for high volatility states of the world. The test reinforces the findings for 0DTEs, further validating evidence of an implicit distinction in risk preference heterogeneity when moving to higher

frequencies.

Our nonparametric approach allows important patterns of risk preferences to emerge directly from trading strategies without imposing restrictive functional forms. Future applications could comprise multivariate settings, for instance, in the foreign exchange market, and studies of the effect of transaction costs on market-implied risk preferences.

References

Yacine Aït-Sahalia and Jefferson Duarte. Nonparametric option pricing under shape restrictions. *Journal of Econometrics*, 116(1):9–47, 2003.

Yacine Aït-Sahalia and Andrew W. Lo. Nonparametric risk management and implied risk aversion. *Journal of Econometrics*, 94(1):9–51, 2000.

Takuya Akiba, Shotaro Sano, Toshihiko Yanase, Takeru Ohta, and Masanori Koyama. Optuna: A next-generation hyperparameter optimization framework, 2019.

Caio Almeida and René Garcia. Economic implications of nonlinear pricing kernels. *Management Science*, 63(10):3361–3380, 2017.

Caio Almeida, Gustavo Freire, and Rodrigo Hizmeri. 0DTE asset pricing. *SSRN Electronic Journal*, 2024.

David Backus, Mikhail Chernov, and Ian Martin. Disasters implied by equity index options. *The Journal of Finance*, 66(6):1969–2012, 2011.

Gurdip Bakshi, Dilip Madan, and George Panayotov. Returns of claims on the upside and the viability of u-shaped pricing kernels. *Journal of Financial Economics*, 97(1):130–154, 2010.

Federico M. Bandi, Nicola Fusari, and Roberto Renò. 0DTE option pricing. *SSRN Electronic Journal*, 2023.

Giovanni Barone-Adesi, Nicola Fusari, Antonietta Mira, and Carlo Sala. Option market trading activity and the estimation of the pricing kernel: A bayesian approach. *Journal of Econometrics*, 216(2):430–449, 2020.

Alain Berlinet and Christine Thomas-Agnan. *Reproducing Kernel Hilbert Spaces in Probability and Statistics*. Springer US, 2004.

Fischer Black and Myron Scholes. The pricing of options and corporate liabilities. *Journal*

of Political Economy, 81(3):637–654, 1973.

Oleg Bondarenko. Why are put options so expensive? *The Quarterly Journal of Finance*, 04(03):1450015, 2014.

Lotfi Boudabsa and Damir Filipović. Machine learning with kernels for portfolio valuation and risk management. *Finance Stoch.*, 26(2):131–172, 2022. ISSN 0949-2984. doi: 10.1007/s00780-021-00465-4. URL https://doi.org/10.1007/s00780-021-00465-4. [Vol. 26 (2021) on first page].

Michael W Brandt, Pedro Santa-Clara, and Rossen Valkanov. Parametric portfolio policies: Exploiting characteristics in the cross-section of equity returns. *The Review of Financial Studies*, 22(9):3411–3447, 2009.

Douglas T. Breeden and Robert H. Litzenberger. Prices of state-contingent claims implicit in option prices. *The Journal of Business*, 51(4):621–651, 1978.

P. Carr and D. Madan. Optimal positioning in derivative securities. *Quantitative Finance*, page 19–37, 2001.

Peter Carr and Liuren Wu. Variance risk premiums. *The Review of Financial Studies*, 22 (3):1311–1341, 2008.

Fousseni Chabi-Yo. Conditioning information and variance bounds on pricing kernels with higher- order moments: Theory and evidence. *The Review of Financial Studies*, 21(1): 181–231, 2007.

Carsten H. Chong and Viktor Todorov. Do equity and options markets agree about volatility? *SSRN Electronic Journal*, 2024.

Peter Christoffersen, Steven Heston, and Kris Jacobs. Capturing option anomalies with a variance-dependent pricing kernel. *The Review of Financial Studies*, 26(8):1963–2006, 2013.

Alessandro Crescini, Fabio Trojani, and Andrea Vedolin. Demand-based expected returns.

SSRN Electronic Journal, 2025.

Louis Eeckhoudt and Harris Schlesinger. Putting risk in its proper place. *The American Economic Review*, 96(1):280–289, 2006.

Larry G. Epstein and Martin Schneider. Ambiguity, information quality, and asset pricing. *The Journal of Finance*, 63(1):197–228, 2008.

Lawrence C. Evans. *Partial Differential Equations*. Graduate Studies in Mathematics, V. 19 GSM/19. American Mathematical Society, 1998.

Damir Filipović and Paul Georg Schneider. Kernel density machines. *SSRN Electronic Journal*, 2025.

Damir Filipović, Markus Pelger, and Ye Ye. Stripping the discount curve-a robust machine learning approach. *Swiss Finance Institute Research Paper*, 2022.

Maria Grith, Wolfgang K. Härdle, and Volker Krätschmer. Reference-dependent preferences and the empirical pricing kernel puzzle. *Review of Finance*, 21(1):269–298, 2016.

Lars Peter Hansen and Ravi Jagannathan. Implications of security market data for models of dynamic economies. *Journal of political economy*, 99(2):225–262, 1991.

Lars Peter Hansen and Scott F. Richard. The role of conditioning information in deducing testable restrictions implied by dynamic asset pricing models. *Econometrica*, 55(3):587–613, 1987.

Trevor Hastie, Robert Tibshirani, and Jerome Friedman. *The Elements of Statistical Learning*. Springer Series in Statistics. Springer New York Inc., New York, NY, USA, 2001.

Guanglian Hu and Yuguo Liu. The pricing of volatility and jump risks in the cross-section of index option returns. *Journal of Financial and Quantitative Analysis*, 57(6):2385–2411, 2022. doi: 10.1017/S0022109022000333.

Rachel J. Huang, Larry Y. Tzeng, and Lin Zhao. Fractional degree stochastic dominance. *Management Science*, 66(10):4630–4647, 2020.

Jens Carsten Jackwerth. Recovering risk aversion from option prices and realized returns. *The Review of Financial Studies*, 13(2):433–451, 2015.

Jens Carsten Jackwerth and Marco Menner. Does the Ross recovery theorem work empirically? *Journal of Financial Economics*, 137(3):723–739, 2020.

Ralph S. J. Koijen and Motohiro Yogo. A demand system approach to asset pricing. *Journal of Political Economy*, 127:1475–1515, 2019.

Kozak. Kernel trick for the cross-section. *SSRN Electronic Journal*, 2020.

Martin Lettau and Sydney Ludvigson. Consumption, aggregate wealth, and expected stock returns. *The Journal of Finance*, 56(3):815–849, 2001.

Matthew Linn, Sophie Shive, and Tyler Shumway. Pricing kernel monotonicity and conditional information. *The Review of Financial Studies*, 31(2):493–531, 2017.

Stefan Nagel and Kenneth J. Singleton. Estimation and evaluation of conditional asset pricing models. *The Journal of Finance*, 66(3):873–909, 2011.

Zhongjun Qu and Guang Zhang. Estimating state price densities implied by american options. *Journal of Business & Economic Statistics*, 0(0):1–14, 2025.

Steve Ross. The recovery theorem. *The Journal of Finance*, 70(2):615–648, 2015.

Robert Schaback and Holger Wendland. Kernel techniques: From machine learning to meshless methods. *Acta Numerica*, 15, 2006.

Schneider and Trojani. (Almost) model-free recovery. *The Journal of Finance*, 74(1):323–370, 2019.

Paul Schneider. Generalized risk premia. *Journal of Financial Economics*, 116(3):487–504, 2015.

Paul Georg Schneider. Ross recovery with time series information and economic constraints. *SSRN Electronic Journal*, 2019.

David Schreindorfer and Tobias Sichert. Conditional risk and the pricing kernel. *Journal of Financial Economics*, 171:104106, 2025.

Rohan Sen. Kernel-based nonparametric tests for shape constraints. *arXiv preprint arXiv:2510.16745*, 2025.

Mervyn J. Silvapulle and Pranab K. Sen. *Constrained Statistical Inference: Inequality, Order, and Shape Restrictions*. Wiley, 2001.

Karl N. Snow. Diagnosing asset pricing models using the distribution of asset returns. *The Journal of Finance*, 46(3):955–983, 1991.

Zhaogang Song and Dacheng Xiu. A tale of two option markets: Pricing kernels and volatility risk. *Journal of Econometrics*, 190(1):176–196, 2016.

Holger Wendland. *Scattered data approximation*, volume 17. Cambridge University Press, Cambridge, 2005.

Ding-Xuan Zhou. Derivative reproducing properties for kernel methods in learning theory. *Journal of Computational and Applied Mathematics*, 220(1):456–463, 2008.

A Reproducing kernel Hilbert spaces

We recap a fundamental notion in statistical machine learning, namely that of a reproducing kernel Hilbert space. For more background and applications, we refer the reader to Hastie et al. (2001); Berlinet and Thomas-Agnan (2004); Wendland (2005); Schaback and Wendland (2006). Let $\mathcal{X}$ be any arbitrary set and $\mathcal{H}$ be a Hilbert space of functions $h\colon \mathcal{X} \to \mathbb{R}$. $\mathcal{H}$ is called a *reproducing kernel Hilbert space (RKHS)* if, for any $x \in \mathcal{X}$, there exists a *canonical feature function* $\phi\colon \mathcal{X} \to \mathcal{H}$ such that the scalar product $\langle h, \phi(x)\rangle_{\mathcal{H}} = h(x)$ acts as evaluation at x for all $h \in \mathcal{H}$. The function $\phi(x)$ acts as the representer of the data point $x \in \mathcal{X}$ in $\mathcal{H}$ and is often denoted as $k(x, \cdot)$. The function $k\colon \mathcal{X} \times \mathcal{X} \to \mathbb{R}$ induced by

$$k(x, x') = \langle \phi(x), \phi(x')\rangle_{\mathcal{H}} = \langle k(x, \cdot), k(x', \cdot)\rangle_{\mathcal{H}},$$

is called the *reproducing kernel* of $\mathcal{H}$. It has the property that for any finite selection of points $x_1, \ldots, x_n \in \mathcal{X}$, the $n \times n$ matrix with elements $k(x_i, x_j)$ is symmetric and positive semi-definite. *Representer theorems* for regularized kernel-based regression problems facilitate finite-dimensional formulations. Our result, Theorem 3.1, is a variant thereof, adapted to the mean-variance loss.

Products of functions and kernels play an important role in multidimensional settings. Especially, in our setting, when we want to condition on temporal state variables, we seek a tensor product RKHS. It is known, see Berlinet and Thomas-Agnan (2004, Theorem 13, Section 4.6) that if $\mathcal{H}_1$ and $\mathcal{H}_2$ are two RKHS with respective reproducing kernels k_1 and k_2, then the tensor product $\mathcal{H}_1 \otimes \mathcal{H}_2$ is an RKHS with reproducing kernel $k := k_1 \otimes k_2$.

A.1 Hypothesis space specification

For our purpose, we seek a hypothesis space containing functions that are at least twice continuously differentiable in the moneyness. Since we want to compute the conditional SDF, we envisage conditioning the SDF on some exogenous state variables. Hence, we propose an RKHS $\mathcal{H}$ that is defined as:

$$k((m; v), (m'; v')) := k_{sob}(m, m') \cdot k_{sob}(v, v'), \tag{A.1}$$

where $k_{sob}(\cdot, \cdot)$ is the Sobolev kernel, see Wendland (2005, Chapter 10),

$$k_{sob}(x, y) := \frac{1}{8}\sqrt{\frac{\pi}{2}} e^{-|x-y|} \left((x-y)^2 + 3|x-y| + 3\right). \tag{A.2}$$

In this framework, we treat the m component (*moneyness*) as endogenous and the v component as exogenous. Therefore, in what follows, whenever we perform differentiation, it is only with the endogenous moneyness component and never with respect to the exogenous component. For economy of notation, we shall therefore adopt the following convention,

$$\begin{aligned} \phi(m; v) &:= k_{sob}(m, \cdot) \cdot k_{sob}(v, \cdot) \\ \phi'(m; v) &:= k'_{sob}(m, \cdot) \cdot k_{sob}(v, \cdot) \\ \phi''(m; v) &:= k''_{sob}(m, \cdot) \cdot k_{sob}(v, \cdot). \end{aligned} \tag{A.3}$$

The RKHS $\mathcal{H}$ defined using the kernel function in (A.1) satisfies $\mathcal{H} \subset \mathcal{C}^2(\mathcal{X})$ by the *Sobolev Imbedding Lemma*, see Evans (1998); this particular specification of the RKHS $\mathcal{H}$ ensures that it contains at least twice continuously differentiable functions in moneyness. This augurs well for our proposed methodology, given that the Carr-Madan spanning formula requires up to the second continuous derivative. This allows us to have reproducing properties for the derivatives of the kernel function, see Zhou (2008). Precisely, we have the following result:

$$h'(m;v) = \langle h, \phi'(m;v)\rangle_{\mathcal{H}}, \qquad h''(m;v) = \langle h, \phi''(m;v)\rangle_{\mathcal{H}} \quad \text{for any } h \in \mathcal{H}, \tag{A.4}$$

by which we can represent the optimal portfolio rules for both the delta-hedging component and the OTM options portfolio component in (2.7) as functions in $\mathcal{H}$.

A.2 Representer theorem

We now demonstrate how to arrive at the representer theorem for obtaining the optimal empirical solution to Problem 3.3. Towards this end, we follow the methodology developed in Sen (2025) and work in the RKHS $\mathcal{H}$ defined via the reproducing kernel in (A.1). For each period $t = 0, \ldots, T-1$, consider the observations given by $(\boldsymbol{X}_t, \boldsymbol{Y}_t)$, such that

$$\boldsymbol{X}_t := (X_{t,0}, X_{t,1}, \ldots, X_{t,N_K(t)}), \quad \boldsymbol{Y}_t := (Y_{t,0}, Y_{t,1}, \ldots, Y_{t,N_K(t)}). \tag{A.5}$$

The $\boldsymbol{X}_t$ are the features for the kernel function, while $\boldsymbol{Y}_t$ are the corresponding coefficients defined as:

$$\begin{aligned} X_{t,0} &:= (1, v_t), \qquad X_{t,i} := (m_{it}, v_t), \quad i = 1, \ldots, N_K(t) \\ Y_{t,0} &:= R_{t+1} - 1, \qquad Y_{t,i} := R^O_{i,t+1} w_{i,t}, \quad i = 1, \ldots, N_K(t). \end{aligned} \tag{A.6}$$

The reproducing derivatives property (A.4), coupled with the above definition, allows us to *represent* the payoff $\mathcal{R}^{\pi}_{t+1}(\cdot)$ from (2.5) as the following function in the RKHS $\mathcal{H}$:

$$\psi_t := \phi'(1; v_t)\,(R_{t+1} - 1) + \sum_{i=1}^{N_K(t)} \phi''(m_{it}; v_t) R^O_{i,t+1} w_{it}, \tag{A.7}$$

and it holds,

$$\mathcal{R}^{\pi}_{t+1}(h) = \langle h, \psi_t \rangle_{\mathcal{H}}, \quad \text{for } t = 0, \ldots, T-1. \tag{A.8}$$

This allows us to use Sen (2025, Theorem 2.4) to arrive at the claim in Theorem 3.1. In particular, the optimal solution lies in the finite-dimensional subspace

$$\mathcal{H}_0 := \operatorname{span}\left\{\{\phi'(1; v_t)\}, \{\phi''(m_{it}; v_t)\}\right\}, \quad i = 1, \ldots, N_K(t),\ t = 0, \ldots, T-1, \tag{A.9}$$

and for any pair (m, v), it holds,

$$\begin{aligned} h(m; v) &= \langle h, \phi(m, v) \rangle_{\mathcal{H}} \\ &= \sum_{t=0}^{T-1} c_{1,t} \langle \phi'(1; v_t), \phi(m, v) \rangle_{\mathcal{H}} + \sum_{t=0}^{T-1} \sum_{i=1}^{N_K(t)} c_{2,it} \langle \phi''(m_{it}; v_t), \phi(m, v) \rangle_{\mathcal{H}} \\ &= \sum_{t=0}^{T-1} c_{1,t} \frac{\partial}{\partial x} k_{sob}(1, m) \cdot k_{sob}(v_t, v) + \sum_{t=0}^{T-1} \sum_{i=1}^{N_K(t)} c_{2,it} \frac{\partial^2}{\partial x^2} k_{sob}(m_{it}, m) \cdot k_{sob}(v_t, v). \end{aligned} \tag{A.10}$$

B Implementation

This section outlines the matrix approach and implementation for the optimization problem in (3.5), along with the construction of the test statistic, and other supplemental details.

B.1 Matrix formulation

We fix the following notations for easier reading:

$$\begin{aligned} \mathcal{J}_t &:= \{1, \ldots, N_K(t)\}, \qquad \mathcal{T} := \{0, \ldots, T-1\} \\ N_K &:= \sum_{t=0}^{T-1} N_K(t), \qquad N := T + N_K. \end{aligned} \tag{B.1}$$

Define the *row vectors* of basis functions

$$\mathbf{\Phi}_1 := \left[\phi'(1; v_t)\right]_{t \in \mathcal{T}}, \qquad \mathbf{\Phi}_2 := \left[\phi''(m_{it}; v_t)\right]_{i \in \mathcal{J}_t,\, t \in \mathcal{T}}. \tag{B.2}$$

We can now define the basis functions of the working subspace $\mathcal{H}_0$ as the row vector

$$\boldsymbol{\Phi} := \Big[\boldsymbol{\Phi}_1,\, \boldsymbol{\Phi}_2\Big], \tag{B.3}$$

and the corresponding *kernel matrix* by taking the pairwise inner product:

$$\boldsymbol{K} := \langle \boldsymbol{\Phi}^\top, \boldsymbol{\Phi} \rangle_{\mathcal{H}} \in \mathbb{R}^{N \times N}. \tag{B.4}$$

Since we have the functional form of $\widehat{h}_\lambda$ as in (3.4), we can now write

$$\widehat{h}_\lambda = \boldsymbol{\Phi}\widehat{\boldsymbol{c}}, \qquad \widehat{\boldsymbol{c}} := \Big[[\widehat{c}_{1,t}]_{t\in\mathcal{T}}, [\widehat{c}_{2,it}]_{i\in\mathcal{J}_t,\, t\in\mathcal{T}}\Big]^\top \in \mathbb{R}^N. \tag{B.5}$$

We need to formulate the system of equations that solves for the optimal coefficients $\widehat{\boldsymbol{c}}$, which is done by writing each term in Problem 3.3 in terms of a matrix formulation. Towards that end, we construct the following vector of coefficients:

$$\boldsymbol{a}_t := \left[\begin{array}{c} \boldsymbol{a}_{1,t} \\ \hline \boldsymbol{a}_{2,t} \end{array}\right] \in \mathbb{R}^N, \tag{B.6}$$

where

$$\boldsymbol{a}_{1,t} := [0, \ldots, 0, R_{t+1} - 1, 0, \ldots, 0]^\top \in \mathbb{R}^T,$$

$$\boldsymbol{a}_{2,t} := \Big[0, \ldots, 0, \Big[R^O_{i,t+1} w_{it}\Big]_{i=1}^{N_K(t)}, 0, \ldots, 0\Big]^\top \in \mathbb{R}^{N_K}.$$

The vector $\boldsymbol{a}_t$ contains zero everywhere except at the entries corresponding to the time index t. Consider the mean vector

$$\bar{\boldsymbol{a}} := \widehat{\mathbb{E}}[\boldsymbol{a}_t] = \frac{1}{T}\sum_{t=0}^{T-1} \boldsymbol{a}_t \in \mathbb{R}^N,$$

and the centered vectors

$$\widetilde{\boldsymbol{a}}_t := \boldsymbol{a}_t - \bar{\boldsymbol{a}} \quad \text{for} \quad 0 \le t \le T-1.$$

Hence, for $t = 0, \dots, T-1$, we can write ψ_t from (A.7) as $\psi_t = \boldsymbol{\Phi} \boldsymbol{a}_t$, with $\widehat{\mu} = \widehat{\mathbb{E}}[\psi_t] = \widehat{\mathbb{E}}[\boldsymbol{\Phi} \boldsymbol{a}_t] = \boldsymbol{\Phi} \bar{\boldsymbol{a}}$ such that the mean term is

$$\widehat{\mathbb{E}}[\mathcal{R}^{\pi}_{t+1}(h)] = \langle h, \widehat{\mu} \rangle_{\mathcal{H}} = \langle \boldsymbol{c}^{\top} \boldsymbol{\Phi}^{\top}, \boldsymbol{\Phi} \bar{\boldsymbol{a}} \rangle_{\mathcal{H}} = \boldsymbol{c}^{\top} \boldsymbol{K} \bar{\boldsymbol{a}}.$$

Now, it holds,

$$\langle h, \psi_t - \widehat{\mu} \rangle_{\mathcal{H}} = \langle \boldsymbol{c}^{\top} \boldsymbol{\Phi}^{\top}, \boldsymbol{\Phi} \widetilde{\boldsymbol{a}}_t \rangle_{\mathcal{H}} = \boldsymbol{c}^{\top} \boldsymbol{K} \widetilde{\boldsymbol{a}}_t.$$

Therefore, the variance term reads

$$\widehat{\mathbb{V}}[\mathcal{R}^{\pi}_{t+1}(h)] = \widehat{\mathbb{E}}[\langle h, \psi_t - \widehat{\mu} \rangle^2_{\mathcal{H}}] = \widehat{\mathbb{E}} \left[(\boldsymbol{c}^{\top} \boldsymbol{K} \widetilde{\boldsymbol{a}}_t)^2 \right] = \boldsymbol{c}^{\top} \boldsymbol{K} \boldsymbol{\Sigma} \boldsymbol{K} \boldsymbol{c},$$

where $\boldsymbol{\Sigma} := \widehat{\mathbb{E}} \left[\widetilde{\boldsymbol{a}}_t \widetilde{\boldsymbol{a}}_t^{\top} \right]$. Finally, the regularization term can be written as

$$\langle h, h \rangle_{\mathcal{H}} = \langle \boldsymbol{c}^{\top} \boldsymbol{\Phi}^{\top}, \boldsymbol{\Phi} \boldsymbol{c} \rangle_{\mathcal{H}} = \boldsymbol{c}^{\top} \boldsymbol{K} \boldsymbol{c}.$$

Having computed the above terms, the matrix formulation of Problem 3.3 is given exactly by Problem 3.5

$$\widehat{\boldsymbol{c}} = \underset{\boldsymbol{c} \in \mathbb{R}^N}{\operatorname{argmin}} \; - \boldsymbol{c}^{\top} \boldsymbol{K} \bar{\boldsymbol{a}} + \frac{1}{2} \boldsymbol{c}^{\top} \boldsymbol{K} \boldsymbol{\Sigma} \boldsymbol{K} \boldsymbol{c} + \frac{\lambda}{2} \boldsymbol{c}^{\top} \boldsymbol{K} \boldsymbol{c}.$$

Problem 3.5 is convex and has a unique minimum, which, from the first-order conditions, can be obtained as

$$(\boldsymbol{K} \boldsymbol{\Sigma} \boldsymbol{K} + \lambda \boldsymbol{K}) \, \widehat{\boldsymbol{c}} = \boldsymbol{K} \, \bar{\boldsymbol{a}}. \tag{B.7}$$

It remains to discuss how to compute the kernel matrix $\boldsymbol{K}$ to solve for $\widehat{\boldsymbol{c}}$ using the above system of equations. We remark that $\boldsymbol{K}$ will have a block matrix structure. From (A.1) and (A.2), we can explicitly calculate the entries of each of the block matrices of $\boldsymbol{K}$.

For any $(x, y) \in \mathbb{R}^2$, we set $r := |x - y|$. We have the following set of formulae for the

partial derivatives of the Sobolev kernel:

$$\begin{aligned}
\frac{\partial^2}{\partial x \partial y} k_{sob}(x,y) &= -\frac{1}{8}\sqrt{\frac{\pi}{2}} e^{-r}(r^2 - r - 1), \\
\frac{\partial^3}{\partial x \partial y^2} k_{sob}(x,y) &= \frac{1}{8}\sqrt{\frac{\pi}{2}} e^{-r}(3 - r)(x - y), \\
\frac{\partial^4}{\partial x^2 \partial y^2} k_{sob}(x,y) &= \frac{1}{8}\sqrt{\frac{\pi}{2}} e^{-r}(r^2 - 5r + 3).
\end{aligned} \tag{B.8}$$

Using the above formulae, we can now explicitly compute the block matrices of the kernel matrix $\boldsymbol{K}$ as follows. We define

$$\begin{aligned}
\boldsymbol{K}_{1,1} &:= \langle \boldsymbol{\Phi}_1^\top, \boldsymbol{\Phi}_1 \rangle_{\mathcal{H}} = \left[\tfrac{\partial^2}{\partial x \partial y} k(1,1) \cdot k(v_t, v_s) \right]_{t,\, s \in \mathcal{T}} \in \mathbb{R}^{T \times T}, \\
\boldsymbol{K}_{1,2} &= \langle \boldsymbol{\Phi}_1^\top, \boldsymbol{\Phi}_2 \rangle_{\mathcal{H}} = \left[\tfrac{\partial^3}{\partial x \partial y^2} k(1, m_{js}) \cdot k(v_t, v_s) \right]_{t, s \in \mathcal{T},\, j \in \mathcal{J}_s} \in \mathbb{R}^{T \times N_K}, \\
\boldsymbol{K}_{2,2} &:= \langle \boldsymbol{\Phi}_2^\top, \boldsymbol{\Phi}_2 \rangle_{\mathcal{H}} = \left[\tfrac{\partial^4}{\partial x^2 \partial y^2} k(m_{it}, m_{js}) \cdot k(v_t, v_s) \right]_{t, s \in \mathcal{T},\, i \in \mathcal{J}_t,\, j \in \mathcal{J}_s} \in \mathbb{R}^{N_K \times N_K}.
\end{aligned} \tag{B.9}$$

We can write

$$\boldsymbol{K} = \left[\begin{array}{c|c} \boldsymbol{K}_{1,1} & \boldsymbol{K}_{1,2} \\ \hline \boldsymbol{K}_{1,2}^\top & \boldsymbol{K}_{2,2} \end{array}\right] \in \mathbb{R}^{N \times N}.$$

B.2 Construction of test statistic

In this section, we outline the steps (without proof) to construct the test statistic W_T from Section 5.6 to test the shape constraints of the SDF jointly on a finite grid, leveraging the sign of the derivative evaluations. For more details regarding the construction, we refer an interested reader to Sen (2025). For a fixed VIX level v and a finite testing grid, $\mathcal{Z} = \{m_j : 1 \le j \le n\}$, we first define the *row vector* of functions corresponding to the test:

$$\begin{aligned}
\text{Monotonicity test:} \quad & \boldsymbol{\Phi}_{\mathcal{Z}} := [-\phi'(m_1, v), \ldots, -\phi'(m_n, v)]; \\
\text{Convexity test:} \quad & \boldsymbol{\Phi}_{\mathcal{Z}} := [\phi''(m_1, v), \ldots, \phi''(m_n, v)].
\end{aligned} \tag{B.10}$$

and the corresponding kernel matrix for the test grid

$$\boldsymbol{K}_{\mathcal{Z}} := \langle \boldsymbol{\Phi}^\top, \boldsymbol{\Phi}_{\mathcal{Z}} \rangle_{\mathcal{H}} \in \mathbb{R}^{N \times n}, \tag{B.11}$$

where $\boldsymbol{\Phi}$ is defined in (B.3). Define the matrix of coefficients $\boldsymbol{A} := [\boldsymbol{a}_0, \ldots, \boldsymbol{a}_{T-1}] \in \mathbb{R}^{N \times T}$, where $\boldsymbol{a}_t$, $t = 0, \ldots, T-1$ are defined (B.6). We define the following row vector:

$$\boldsymbol{\Psi} := \boldsymbol{\Phi}\boldsymbol{A}, \tag{B.12}$$

The corresponding Gram matrix (in the $\boldsymbol{\Psi}$ basis) is:

$$\boldsymbol{G} := \langle \boldsymbol{\Psi}^\top, \boldsymbol{\Psi} \rangle_{\mathcal{H}} = \boldsymbol{A}^\top \langle \boldsymbol{\Phi}^\top, \boldsymbol{\Phi} \rangle_{\mathcal{H}} \boldsymbol{A} = \boldsymbol{A}^\top \boldsymbol{K} \boldsymbol{A} \in \mathbb{R}^{T \times T}. \tag{B.13}$$

Define the *centering matrix* $\boldsymbol{H} := \boldsymbol{I}_T - \frac{1}{T}\mathbf{1}\mathbf{1}^\top \in \mathbb{R}^{T \times T}$, where we define the *column vector* of ones $\mathbf{1} := [1, \ldots, 1]^\top \in \mathbb{R}^T$. Note that the matrix $\boldsymbol{H}$ is symmetric and *idempotent*, i.e., $\boldsymbol{H}^\top = \boldsymbol{H}$. We can define the following *row vector* of centered functions

$$\widetilde{\boldsymbol{\Psi}} := \boldsymbol{\Psi}\boldsymbol{H}, \tag{B.14}$$

and the matrix

$$\widetilde{\boldsymbol{G}} := \langle \widetilde{\boldsymbol{\Psi}}^\top, \boldsymbol{\Psi} \rangle_{\mathcal{H}} = \boldsymbol{H}\boldsymbol{G} \in \mathbb{R}^{T \times T}. \tag{B.15}$$

Now, we define the Gram matrix with respect to the centered basis functions as

$$\widetilde{\boldsymbol{G}}_{\mathcal{Z}} := \langle \widetilde{\boldsymbol{\Psi}}^\top, \boldsymbol{\Phi}_{\mathcal{Z}} \rangle_{\mathcal{H}} = \boldsymbol{H} \langle \boldsymbol{\Psi}^\top, \boldsymbol{\Phi}_{\mathcal{Z}} \rangle_{\mathcal{H}} = \boldsymbol{H}\boldsymbol{A}^\top \langle \boldsymbol{\Phi}^\top, \boldsymbol{\Phi}_{\mathcal{Z}} \rangle_{\mathcal{H}} = \boldsymbol{H}\boldsymbol{A}^\top \boldsymbol{K}_{\mathcal{Z}} \in \mathbb{R}^{T \times n}, \tag{B.16}$$

and the following vector:

$$\widetilde{\boldsymbol{h}} := \langle \widetilde{\boldsymbol{\Psi}}^\top, \boldsymbol{\Phi} \rangle_{\mathcal{H}} \widehat{\boldsymbol{c}} = \boldsymbol{H}\boldsymbol{A}^\top \langle \boldsymbol{\Phi}^\top, \boldsymbol{\Phi} \rangle_{\mathcal{H}} \widehat{\boldsymbol{c}} = \boldsymbol{H}\boldsymbol{A}^\top \boldsymbol{K} \widehat{\boldsymbol{c}} \in \mathbb{R}^T. \tag{B.17}$$

Then, we can compute the covariance matrix $\widehat{\boldsymbol{\Omega}}_\lambda$ as follows:

$$\widehat{\boldsymbol{\Omega}}_\lambda := \frac{1}{T}\boldsymbol{S}^\top \boldsymbol{S}, \qquad \boldsymbol{S} := \frac{1}{\lambda}(\boldsymbol{B} - \boldsymbol{V}^\top \boldsymbol{\Lambda}),$$

where

$$\begin{aligned}
\boldsymbol{B} &:= \left(\boldsymbol{I} - \operatorname{diag}(\widetilde{\boldsymbol{h}})\right)\widetilde{\boldsymbol{G}}_{\mathcal{Z}} + \frac{1}{T}\mathbf{1}\left(\widetilde{\boldsymbol{h}}^{\top}\widetilde{\boldsymbol{G}}_{\mathcal{Z}}\right) \in \mathbb{R}^{T\times n}, \\
\boldsymbol{V} &:= \left(\boldsymbol{I} - \operatorname{diag}(\widetilde{\boldsymbol{h}})\right)\widetilde{\boldsymbol{G}} + \frac{1}{T}\mathbf{1}\left(\widetilde{\boldsymbol{h}}^{\top}\widetilde{\boldsymbol{G}}\right) \in \mathbb{R}^{T\times T}, \\
\boldsymbol{\Lambda} &:= \frac{1}{T}\boldsymbol{H}\left(\lambda\boldsymbol{I}_T + \frac{1}{T}\boldsymbol{H}\boldsymbol{G}\boldsymbol{H}\right)^{-1}\widetilde{\boldsymbol{G}}_{\mathcal{Z}} \in \mathbb{R}^{T\times n}.
\end{aligned} \tag{B.18}$$

Having computed $\widehat{\boldsymbol{\Omega}}_\lambda$, we proceed to compute the test statistic W_T from (5.1) as follows. Consider the vector stacked evaluations of the derivative functional at the grid points $\widehat{\boldsymbol{\theta}}$. Set $\boldsymbol{b} := \widehat{\boldsymbol{\Omega}}_\lambda^{-1/2}\widehat{\boldsymbol{\theta}}$, where $\widehat{\boldsymbol{\Omega}}_\lambda^{-1/2}$ is a matrix root of $\widehat{\boldsymbol{\Omega}}_\lambda^{-1}$. Then, we can write:

$$W_T := T \min_{\boldsymbol{c}\geq\mathbf{0}} (\widehat{\boldsymbol{\theta}} - \boldsymbol{c})^{\top}\widehat{\boldsymbol{\Omega}}_\lambda^{-1}(\widehat{\boldsymbol{\theta}} - \boldsymbol{c}) = T \min_{\boldsymbol{c}\geq\mathbf{0}} \|\widehat{\boldsymbol{\Omega}}_\lambda^{-1/2}\boldsymbol{c} - \boldsymbol{b}\|_2^2.$$

The optimization problem has a unique minimizer

$$\boldsymbol{c}^\star := \min_{\boldsymbol{c}\geq\mathbf{0}} \|\widehat{\boldsymbol{\Omega}}_\lambda^{-1/2}\boldsymbol{c} - \boldsymbol{b}\|_2^2,$$

that can be solved as a non-negative least-squares program. Define the residuals $\boldsymbol{r} := \widehat{\boldsymbol{\Omega}}_\lambda^{-1/2}\boldsymbol{c}^\star - \boldsymbol{b}$. Then, we can compute the test statistic as $W_T = T\|\boldsymbol{r}\|_2^2$.

B.3 Further details

We put additional details regarding the implementation, including cross-validation, and the inclusion of additional exogenous variables. Depending on the selected exogenous features, the dimension of $\boldsymbol{K}$ and related objects will vary. We cover here three main cases:

- **Exogenous features are fixed** within date ($N := N_K + T$): this is our baseline specification with the VIX, and will follow from equations A.10.

Algorithm B.1 Time series cross-validation

INPUT: `train_data, validate_data, lambda_range`

OUTPUT: `best_lambda`

initialization: set `best_lambda` $\leftarrow$ `none`, `best_performance` $\leftarrow -\infty$
for `lambda` **in** `lambda_range` **do**
    ▷ *train model on training data with Sobolev kernel*
        `solver` $\leftarrow$ `ObjectiveSolver(sobolev_kernel)`
        `solver.prepare_panel_data(train_data)`
        `solver.compute_L_B()` ▷ pivoted Cholesky decomposition
        `solver.compute_Q()` ▷ variance factor
    ▷ *solve optimization with current lambda*
    if `use_constraint` **then**
            `c_opt` $\leftarrow$ `solver.solve_constrained(gamma, lambda, w)`
    else
            `c_opt` $\leftarrow$ `solver.solve_unconstrained(gamma, lambda)`
    ▷ *evaluate performance on validation data*
        `performance` $\leftarrow$ `calculate_OOS_performance(c_opt, validate_data, train_data)`
    if `performance` $>$ `best_performance` **then**
        `best_lambda` $\leftarrow$ `lambda`
        `best_performance` $\leftarrow$ `performance`
return `best_lambda`

- **No exogenous features** ($N := N_K + 1$):

$$h\left(m\right) = c_1 \frac{\partial}{\partial x} k_{sob}(1, m) + \sum_{t=0}^{T-1} \sum_{i=1}^{N_K(t)} c_{2,it} \frac{\partial^2}{\partial x^2} k_{sob}(m_{it}, m)$$

$$\mathcal{R}^{\pi}_{t+1}(h) = h'(1)(R_{t+1} - 1) + \sum_{i=1}^{N_K(t)} h''(m_{it}) R^O_{i,t+1} w_{it}$$

$$\boldsymbol{a}_t := \left[(R_{t+1} - 1), [R^O_{i,t+1} w_{it}]_{i=1}^{N_K(t)}, 0, \ldots, 0\right]^\top \in \mathbb{R}^{N_K+1}.$$

- **Exogenous features that vary** with each option ($N := N_K + T$): the hedge component embeds the mean $\bar{\sigma}$ as reference for the varying feature σ (e.g. implied volatility), since it cannot vary within date by construction.

$$h\left(m, v, \sigma\right) = \sum_{t=0}^{T-1} c_{1,t} \langle \phi'(1, v_t, \bar{\sigma}_t), \phi(m, v, \sigma) \rangle_{\mathcal{H}} + \sum_{t=0}^{T-1} \sum_{i=1}^{N_K(t)} c_{2,it} \langle \phi''(m_{it}, v_t, \sigma_{it}), \phi(m, v, \sigma) \rangle_{\mathcal{H}}.$$

$$R^{\pi}_{t+1}(h) = h'(1, v_t, \bar{\sigma}_t)(R_{t+1} - 1) + \sum_{i=1}^{N_K(t)} h''(m_{it}, v_t, \sigma_{it}) R^O_{i,t+1} w_{it}$$

$$\boldsymbol{a}_t := \left[0, \ldots, 0, (R_{t+1} - 1), [R^O_{i,t+1} w_{it}]_{i=1}^{N_K(t)}, 0, \ldots, 0\right]^\top \in \mathbb{R}^{M+T}.$$

C Trading the SDF in a Black-Scholes economy

Let $(\Omega, \mathcal{F}, \mathbb{P})$ and $(\Omega, \mathcal{F}, \mathbb{Q})$ be probability spaces under the $\mathbb{P}$ and the equivalent $\mathbb{Q}$ measure, respectively. In a Black-Scholes economy, the change of measure implies the discount factor, i.e., the Radon-Nikodym derivative:

$$\frac{d\mathbb{Q}}{d\mathbb{P}} = \mathcal{E}\left(\frac{r - \mu}{\sigma}\right) = \mathcal{E}\left(\text{-SR}\right) = e^{-rt} e^{-\text{SR} W_t - \frac{1}{2}\text{SR}^2 t},$$

where $\mathcal{E}\left(\int_0^t \phi_s dW_s\right) := \exp\left(\int_0^t \phi_s dW_s - \frac{1}{2}\int_0^t \phi_s^2 ds\right)$ is an exponential martingale and SR is the Sharpe ratio. To replicate the discount factor, we may proceed similarly to Schneider

(2015), i.e., using Itô's lemma to project M into the payoff space in terms of S_t via $W_t = \frac{\log\left(\frac{S_t}{S_0}\right)-(\mu-\frac{\sigma^2}{2})t}{\sigma}$. We also have

$$M(S_0, S_t) = \exp\left(\frac{\mu(\mu-\sigma^2)t - 2\log\left(\frac{S_t}{S_0}\right)\mu}{2\sigma^2}\right) = \left(\frac{S_t}{S_0}\right)^{-\frac{\mu}{\sigma^2}} \exp\left(\frac{\mu(\mu-\sigma^2)t}{2\sigma^2}\right), \tag{C.1}$$

which is a tradable claim. We can price it as:

$$\mathbb{E}^{\mathbb{Q}}[M(S_0, S_t)] = \mathbb{E}^{\mathbb{P}}[M(S_0, S_t)^2] = e^{-SR^2 t}\mathbb{E}^{\mathbb{P}}[e^{2SRW_t}] = e^{\frac{\mu^2}{\sigma^2}t}.$$

The maximum Sharpe ratio for the excess returns is:

$$\frac{\sigma^{\mathbb{P}}[e^{-\mathrm{SR}W_t - \frac{1}{2}\mathrm{SR}^2 t}]}{\mathbb{E}^{\mathbb{P}}[e^{--\mathrm{SR}W_t - \frac{1}{2}\mathrm{SR}^2 t}]} = \sqrt{e^{\frac{\mu^2}{\sigma^2}t} - 1}.$$

Now, let $s := \frac{S_T}{S_t}$, $\beta := -\frac{\mu}{\sigma^2}$ and $a(t,T) := \exp\left(\frac{\mu(\mu-\sigma^2)(T-t)}{2\sigma^2}\right)$ with $a(t,t) = 1$. Then, $M(1,s) = a(t,T)s^\beta$ and thus may be seen as a particular non-linear payoff, such that applying the Carr-Madan formula leads to $g''(s) = M''(1,s) = a(t,T)\beta(\beta-1)s^{\beta-2}$.

Corollary C.1. *In an economy with Brownian returns, we can replicate the stochastic discount factor via Carr-Madan allocations as:*

$$\begin{aligned}
&g(s) - g'(1)(s-1) - g(1) \\
&= a(t,T)(s^\beta - \beta(s-1)) - a(t,t) \\
&= a(t,T)\beta(\beta-1)\left(\int_0^1 K^{\beta-2}(K-s)^+ dK + \int_1^\infty K^{\beta-2}(s-K)^+ dK\right).
\end{aligned}$$

This example in the particular case of geometric Brownian motion (GBM) prices shows us that we may generate the discount factor payoff by entering $a(t,T)\beta(\beta-1)K^{\beta-2}$ OTM vanilla put and call options, delta hedging with hedge ratio $a(t,T)\beta$ and a bond investment of $a(t,t) = 1$. Now, let $O(x,K|x_0) := P(x,K)\mathbb{1}_{K\leq x_0} + C(x,K)\mathbb{1}_{K>x_0}$ be a homogeneous function representing the payoff of an OTM call or put option. Carr-Madan replication may

be rewritten as:

$$M\left(1,\frac{S_T}{S_t}\right)=1+M'\left(1,\frac{S_T}{S_t}\right)\left(\frac{S_T}{S_t}-1\right)+\int_0^\infty M''(1,K)O\left(\frac{S_T}{S_t},K|1\right)dK.$$

Hence, following from Corollary C.1:

$$\underbrace{\frac{S_T}{S_t}^{-\frac{\mu}{\sigma^2}}}_{\text{SDF}}=\underbrace{a(t,T)^{-1}}_{\text{Bond}}-\underbrace{\frac{\mu}{\sigma^2}\left(\frac{S_T}{S_t}-1\right)}_{\text{Hedge}}+\underbrace{\frac{\mu^2+\mu\sigma^2}{\sigma^4}\int_0^\infty K^{-\frac{\mu}{\sigma^2}-2}O\left(\frac{S_T}{S_t},K|1\right)dK}_{\text{OTM options portfolio}} \tag{C.2}$$

Remark C.2. *The OTM portfolio allocation includes the non-linear part of the replication. It can be seen as a risk-premia-adjusted version of a volatility swap allocation.*

$$\frac{\mu^2+\mu\sigma^2}{\sigma^4}K^{-\frac{\mu}{\sigma^2}-2}=\frac{\mu^2+\mu\sigma^2}{\sigma^4}K^{-\frac{\mu}{\sigma^2}}\cdot\underbrace{(K^{-2})}_{\textit{volatility swap allocation}}.$$

Simulating option portfolios within the particular framework of a Black-Scholes economy, we may compare the maximum discount factor Sharpe ratio to the one achieved by an investor replicating the strategy using Carr-Madan allocations, which will require a discretization of the integrals to build the portfolio gross return R^π. At a certain time t it follows from (C.2) that

$$R^\pi=\underbrace{\left(\frac{S_T}{S_t}\right)^{-\frac{\mu}{\sigma^2}}}_{\text{SDF}} \tag{C.3}$$

$$=\underbrace{a(t,T)^{-1}}_{\text{Bond}}-\underbrace{\frac{\mu}{\sigma^2}\left(\frac{S_T}{S_t}-1\right)}_{\text{Hedge}}+\underbrace{\frac{\mu^2+\mu\sigma^2}{\sigma^4}\sum_{i=1}^{N_K}K_i^{-\frac{\mu}{\sigma^2}-2}O\left(\frac{S_T}{S_t},K_i|1\right)w_i}_{\text{OTM options portfolio}}. \tag{C.4}$$

where N_K is the total number of strikes and w_i defines the discretized weights for the simulated strikes, which can correspond to any numerical integration scheme.

D Pricing in a Heston economy

Following Christoffersen et al. (2013), we adapt the continuous-time Heston model to a discrete-time GARCH framework with a variance-dependent pricing kernel. The CHJ spec-

ification introduces a crucial modification: while the pricing kernel remains monotonic in both stock returns and variance individually, its projection onto returns becomes non-monotonic when the variance premium is negative, generating the characteristic U-shaped pattern observed empirically. The discrete-time SDF ratio follows:

$$M(S_{t-1}, S_t) = \left(\frac{S_t}{S_{t-1}}\right)^{\phi} \exp\left(\delta + \eta h_t + \xi(h_{t+1} - h_t)\right), \tag{D.1}$$

where h_t denotes the conditional variance (not volatility), and the parameters ϕ, δ, η, and ξ are calibrated to ensure the Euler equation $\mathbb{E}_t[M_t/M_{t-1} \cdot R_t] = 1$ holds. The variance follows GARCH dynamics:

$$h_t = \omega + \beta h_{t-1} + \alpha(z_{t-1} - \gamma\sqrt{h_{t-1}})^2, \tag{D.2}$$

with z_t standard normal and γ capturing leverage effects. The key insight is that negative variance risk premiums ($\xi > 0$) create U-shaped SDFs when projected onto returns, reflecting the need to hedge against both bad and good times due to volatility dynamics.

E Conditional moments: Weekly Options

This appendix extends the conditional moment validation to weekly S&P 500 options (January 2011 to December 2022, $N = 503$). The methodology follows Section 5.7. Results are consistent with monthly options.

The weekly options analysis corroborates main findings from monthly options: our non-parametric SDF generates conditional physical expectations that successfully predict realized higher-order moments, produces economically sensible variance risk premia, and satisfies no-arbitrage restrictions. Consistency across monthly and weekly frequencies provides robust validation across maturities and sample periods.

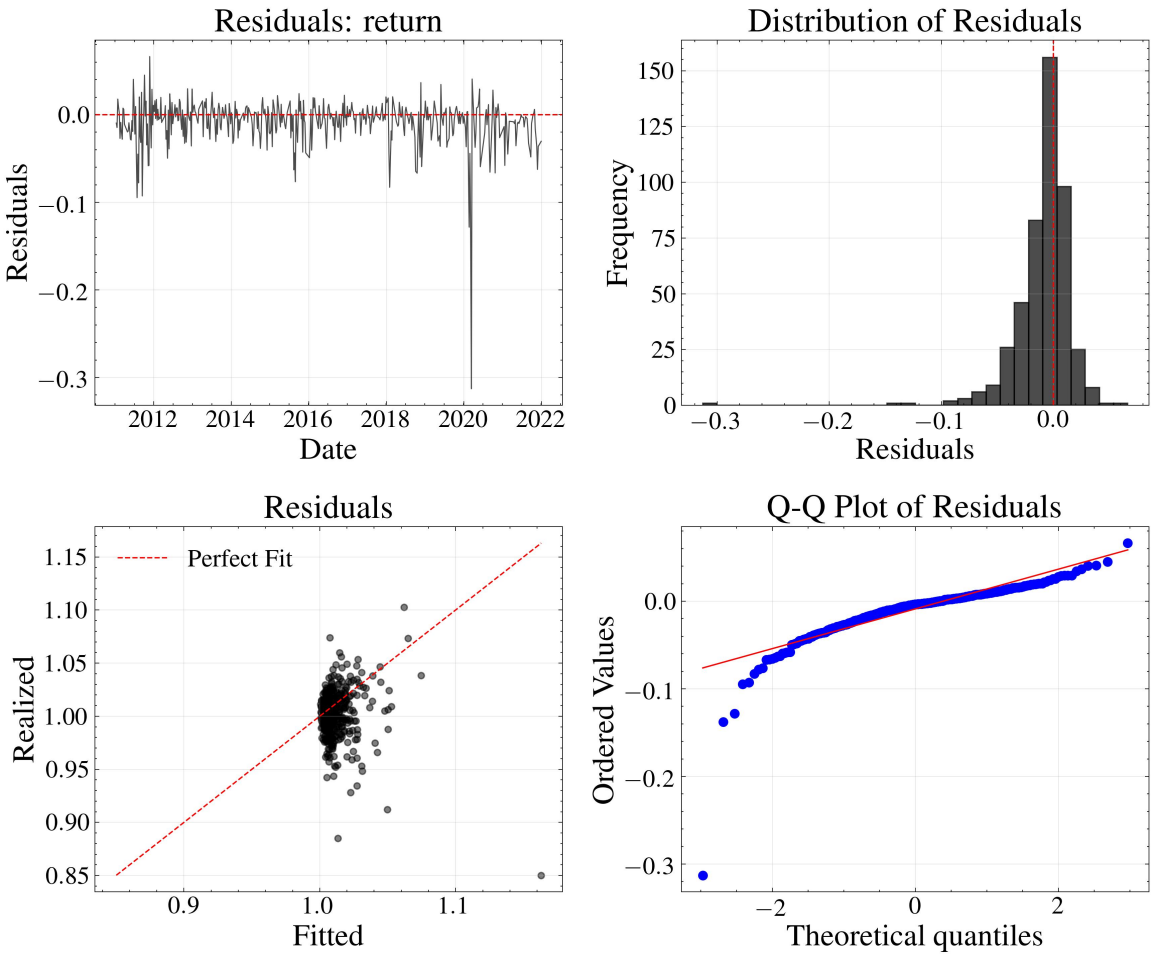

Figure E.1: Return Residuals Diagnostics: Weekly Options. Residuals $\varepsilon_{t+1} = R_{t+1} - \mathbb{E}_t^{\mathbb{P}}[R_{t+1}]$ from weekly S&P 500 options. Panel structure follows Figure 5.17. Sample: January 2011 to December 2022.

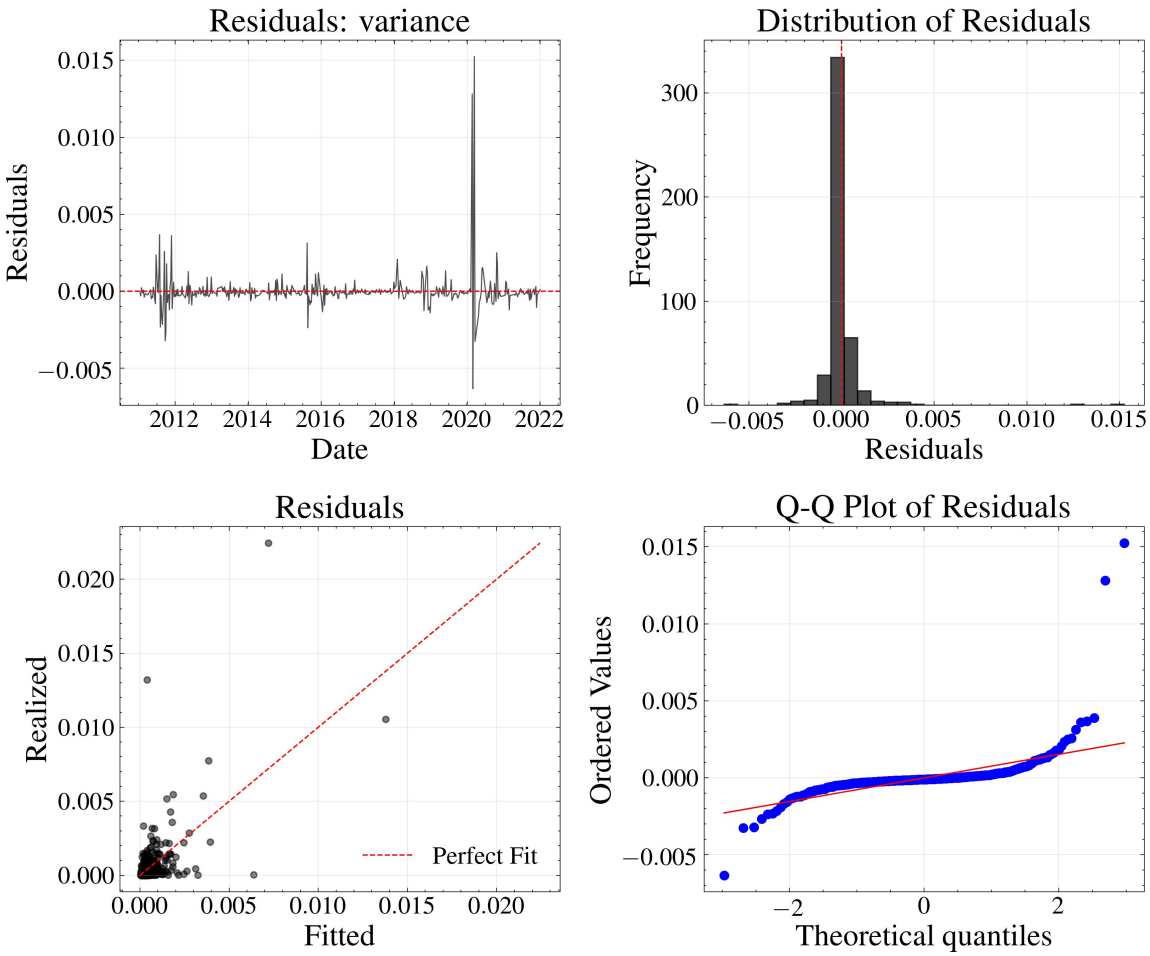

Figure E.2: Variance Residuals Diagnostics: Weekly Options. Residuals $\varepsilon_{t+1} = (R_{t+1} - 1)^2 - \mathbb{E}_t^{\mathbb{P}}[(R_{t+1} - 1)^2]$ from weekly S&P 500 options. Panel structure follows Figure 5.18. Sample: January 2011 to December 2022.

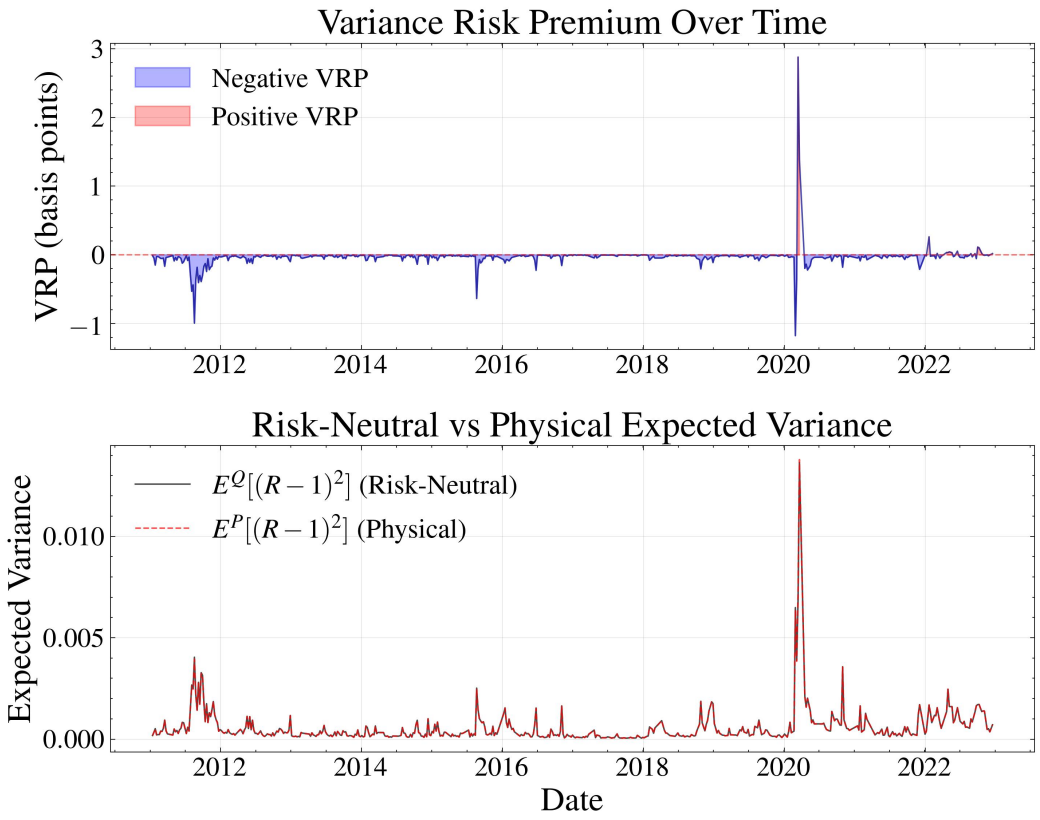

Figure E.3: Variance Risk Premium Time Series: Weekly Options. VRP and variance expectations from weekly S&P 500 options. Format follows Figure 5.19. Weekly VRP exhibits greater volatility than monthly, capturing intra-month variation in volatility compensation. Sample: January 2011 to December 2022.

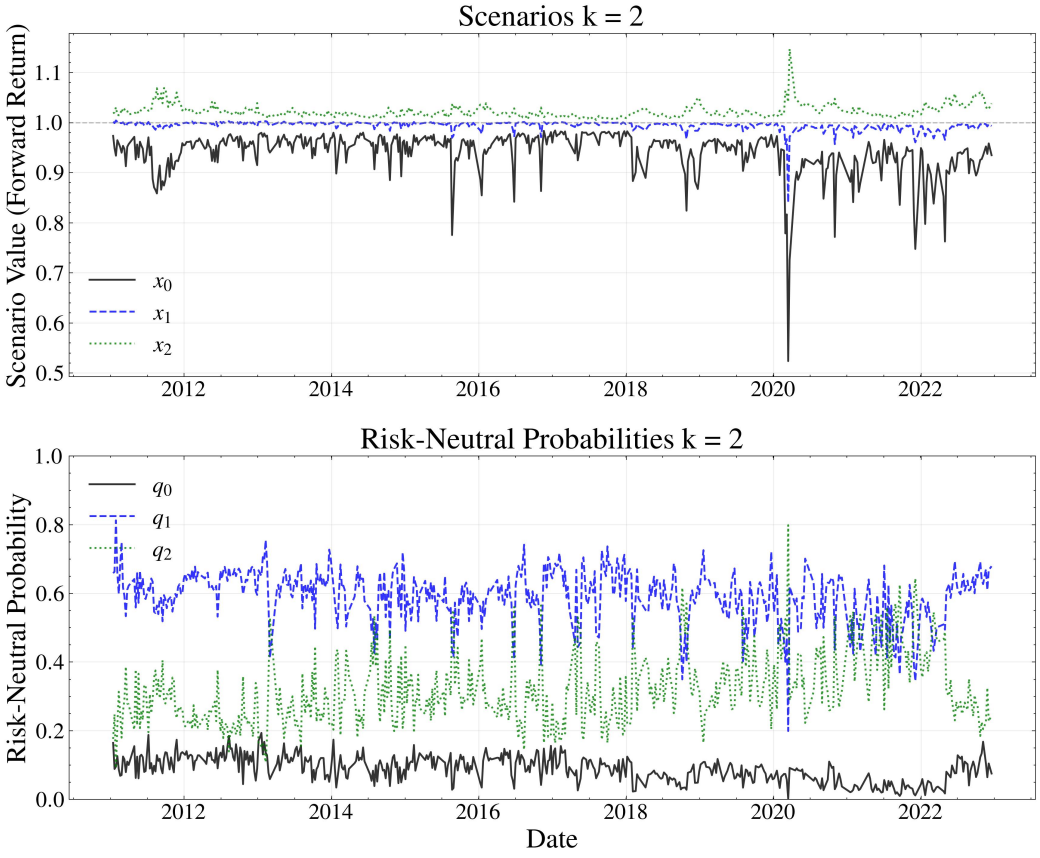

Figure E.4: Three-State Scenario Representation: Weekly Options. Discrete scenario structure for $k = 2$ model from weekly S&P 500 options. Format follows Figure 5.20. Sample: January 2011 to December 2022.