

iCOS: Option-Implied COS Method*

Evgenii Vladimirov[†]

Amsterdam School of Economics, University of Amsterdam
and Tinbergen Institute

February 10, 2024

Abstract

This paper proposes the option-implied Fourier-cosine method, iCOS, for non-parametric estimation of risk-neutral densities, option prices, and option sensitivities. The iCOS method leverages the Fourier-based COS technique, proposed by Fang and Oosterlee (2008), by utilizing the option-implied cosine series coefficients. Notably, this procedure does not rely on any model assumptions about the underlying asset price dynamics, it is fully non-parametric, and it does not involve any numerical optimization. These features make it rather general and computationally appealing. Furthermore, we derive the asymptotic properties of the proposed non-parametric estimators and study their finite-sample behavior in Monte Carlo simulations. Our empirical analysis using S&P 500 index options and Amazon equity options illustrates the effectiveness of the iCOS method in extracting valuable information from option prices under different market conditions. Additionally, we apply our methodology to dissect and quantify observation and discretization errors in the VIX index.

Keywords: Derivatives; Option pricing; Risk-neutral distribution; Option Greeks; COS method.

JEL Classification: C13; C58; G13.

1 Introduction

Option prices play a crucial role in financial markets, providing investors and researchers with valuable insights into market expectations, risks, and dynamics associated with underlying instruments. This information is essential for effective risk management, portfolio optimization, and academic research aimed at understanding market dynamics. However, accurate extraction

*I am very grateful to Caio Almeida, Torben Andersen, Peter Boswijk, Zhenyu Cui, Fang Fang, Gustavo Freire, Roger Laeven, Olivier Scaillet, Viktor Todorov, and conference participants at the Financial Econometrics Conference at Lancaster University, the 2023 Annual SoFiE conference, and the 5th International Workshop in Financial Econometrics for helpful comments and suggestions. I am also thankful for the financial support received from the SoFiE travel grant and the Amsterdam University Fund.

[†]Correspondence to: Econometric Institute, Erasmus University Rotterdam, Burgemeester Oudlaan 50, 3062 PA, Rotterdam, The Netherlands. Email address: vladimirov@ese.eur.nl.

of this valuable information is a challenging task due to the complex nature of option prices and the presence of various sources of uncertainty.

Traditional option pricing models, commonly used to extract embedded information from options, rely on parametric assumptions about market dynamics and the distribution of asset returns. While these models have been widely used and have provided valuable insights, they often fail to capture the complexity and nuances of real-world market behavior and are often subject to model misspecification. To address these limitations, non-parametric methods have gained increasing attention in option pricing research. These methods aim to extract information about underlying dynamics directly from market data without making explicit assumptions about the underlying asset price dynamics.

In this paper, we propose a unified non-parametric estimation procedure for the risk-neutral density (RND), option prices, and option sensitivities, which are the key objects of interest in the option pricing literature. Our approach leverages the Fourier-based cosine technique, the COS method, proposed by Fang and Oosterlee (2008), in a model-free way by *implying* information from observed option contracts. Therefore, we refer to this method as the (option-)implied COS method, or, in short, iCOS. The proposed estimation method is fully non-parametric and does not require any optimization routines, offering a flexible and computationally appealing alternative to traditional techniques.

Fourier-based methods are widely used numerical techniques for option evaluation that exploit the relation between the probability density function (PDF) and the characteristic function (CF). Examples of these methods include those proposed by Carr and Madan (1999), Lewis (2001), and Fang and Oosterlee (2008). The popularity of these methods is due to the CFs – the Fourier transforms of the PDFs – that can be obtained in a (semi-)closed form for a large class of parametric models, such as the affine jump-diffusion class as defined in Duffie, Pan, and Singleton (2000). As such, the main requirement for these methods is a fully parametric specification of the CF, often induced from a parametrized asset price dynamics.

In contrast, the method proposed in this paper only requires the availability of observed option prices, which is a natural setup when working with option data. This allows us to bypass the unnecessarily restrictive requirement of a parametric CF and then utilize the flexibility of the COS machinery. In particular, we exploit the option spanning result of Carr and Madan (2001) to estimate the cosine series coefficients – the projection coefficients of the orthonormal Fourier-cosine basis. These implied cosine coefficients are the building blocks for constructing non-parametric estimators of the RND, option prices with strikes that are not observed in the market, and option sensitivities.

In particular, by utilizing the spanning result of Carr and Madan (2001) and the COS machinery, we provide portfolio representations for the RND, option prices and option deltas. These portfolios consist of option contracts with a continuum of strike prices and offer model-free representations of the objects of interest. Since, in practice, we observe only a finite number of option contracts subject to observation errors, we develop non-parametric estimators based on these spanning results. We derive the asymptotic properties of these non-parametric estimators in an asymptotic setting in which the mesh of the strike grid shrinks to zero, while maturity

and largest and smallest strike prices remain fixed. The resulting limiting distributions allow constructing confidence intervals for our estimates.

The existing semi- and non-parametric methods for the estimation or interpolation of option prices include parametric curve fitting (Shimko, 1993, Gatheral & Jacquier, 2014), local polynomial estimators (Aït-Sahalia & Duarte, 2003), (penalized) cubic splines (Bliss & Panigirtzoglou, 2002, Malz, 2014), and kernel-based methods (Aït-Sahalia & Lo, 1998, Grith, Härdle, & Schienle, 2012) among many others. These methods are also often used to obtain the RND via the famous Breeden and Litzenberger (1978) formula. In comparison, our iCOS approach does not rely on (local) parametric representations and can be considered as a ‘global’ non-parametric method. In other words, the developed estimators utilize all available option prices via the portfolio spanning results, while kernel-based approaches or local polynomial regression use only limited information. Furthermore, the proposed procedure provides a unified framework for estimating the RND, option prices and option sensitivities, without the need to take derivatives of estimated option functions.

There are also various methods that estimate the RND directly without invoking the Breeden and Litzenberger (1978) formula. Examples of such methods include the mixture of distributions (Melick & Thomas, 1997, Gemmill & Saflekos, 2000), the Gram–Charlier approximation (Jarrow & Rudd, 1982, Rompolis & Tzavalis, 2007), and the Hermite expansion (Xiu, 2014, Lu & Qu, 2021), among many others. For a comprehensive review of various methods, see Figlewski (2018). In this paper, we utilize the Fourier-cosine expansion of the RND and imply the cosine expansion coefficients directly from the observed option prices. This offers a flexible, computationally appealing, and model-free alternative to traditional techniques. A similar direction is taken in Cui and Yu (2021), Cui and Xu (2022), and Bossu, Carr, and Papanicolaou (2022), who also imply the expansion coefficients from the observed option prices. In fact, Cui and Yu (2021) use the Fourier cosine method in a model-free way to extract the RND. However, our paper goes beyond these studies by proposing a unified framework to estimate the RND together with the option prices and option sensitivities. Furthermore, we explicitly control the truncation of the RND to a finite interval, allow for observation errors in options, and derive asymptotic results for the proposed estimators.

Our paper also contributes to the literature on model-free Greeks estimation (Bates, 2005, Alexander & Nogueira, 2007). Similar to these studies, our approach enables non-parametric estimation of option sensitivities, such as option deltas. However, unlike the methods proposed by Bates (2005) and Alexander and Nogueira (2007), our approach does not require knowledge of the derivative of option prices with respect to strike and, hence, does not involve fitting the implied volatility curve. Instead, we estimate the option-implied deltas using an approach similar to the estimation of the RND and option prices, utilizing the portfolio spanning result based on the Fourier expansion. This eliminates the need for optimization and calibration to the market, thereby reducing the impact of model misspecification and calibration errors.

Our paper is also closely related to several studies that propose non-parametric approaches to estimate various risk measures from short-dated options, such as the Levy density in Qin and Todorov (2019), spot volatility in Todorov (2019), and jump variation in Todorov (2022).

Like these papers, our approach is fully non-parametric and accounts for errors in the observed option prices. Our analysis differs in terms of the specific information extracted from the option contracts, and our methodology is not restricted to short-dated options, making it more broadly applicable.

We conduct extensive simulation experiments to assess the finite-sample properties of the proposed non-parametric estimators. We consider the classical and well-understood Black and Scholes (1973) model and the more realistic ‘double-jump’ stochastic volatility model of Duffie et al. (2000) as data generating processes. We find good finite-sample performance of all three non-parametric estimators for different maturities and show superiority of our procedure to kernel-based smoothing methods.

Finally, in our empirical application, we demonstrate the effectiveness of the iCOS method across various market settings. For that, we analyze the performance of the method in the highly liquid and well-studied market of S&P 500 index (SPX) options. We also consider Amazon (AMZN) options during an Earnings Announcement Day (EAD) in a unique, high-volatility conditions with short maturities and a distinct bimodal density pattern due to the EAD effect. Additionally, we apply our methodology to dissect and quantify errors in the VIX index, one of most popular measure of market volatility. We find that observation errors in the VIX, that arise from the imperfect observation of option prices, are centered around zero and have a small magnitude, reaching up to 0.04 percentage points of the index value. In contrast, discretization errors lead to positive biases, reaching up to 0.7 percentage points during periods of high volatility.

The rest of the paper is organized as follows. In Section 2, we describe the option-implied COS method. The non-parametric estimators for the RND, option prices, and option deltas along with their asymptotic properties are discussed in Section 3. Section 4 provides the Monte Carlo simulation results. The empirical applications are in Section 5. Section 6 concludes the paper. The proofs of the propositions are collected in Appendix A and some additional results are in Appendix B.

2 Option-implied COS method

In this section, we start with the discussion of the COS method, proposed by Fang and Oosterlee (2008). Then, we show how the option-implied information can be incorporated into this machinery. The option-implied COS allows us to get the expansion expressions for the RND, option prices and option sensitivities that serve as the building blocks for the non-parametric estimators introduced in the next section.

2.1 The COS method

The COS method introduced by Fang and Oosterlee (2008) is based on the idea that the conditional density function $f(y)$ on an interval $[a, b] \subset \mathbb{R}$ can be represented via its Fourier

cosine series expansion:

$$\begin{aligned} f(y) &= \frac{1}{b-a} A_0 + \frac{2}{b-a} \sum_{m=1}^{\infty} A_m \cos\left(m\pi \frac{y-a}{b-a}\right) \\ &= \frac{2}{b-a} \sum_{m=0}^{\infty}{}' A_m \cos(u_m y - u_m a), \end{aligned} \quad (1)$$

where \sum' indicates the sum with the first term weighted by one-half, $u_m := \frac{m\pi}{b-a}$, and the cosine coefficients

$$A_m = \int_a^b \cos(u_m y - u_m a) f(y) dy, \quad m = 0, 1, \dots \quad (2)$$

Fang and Oosterlee (2008) showed that the cosine coefficients (2) can be calculated via the ('truncated') characteristic function (CF). In fact, let us denote the CF of the density function restricted to the interval $[a, b]$ by

$$\phi^{[a,b]}(u) := \int_a^b e^{iuy} f(y) dy.$$

Then, premultiplying $\phi^{[a,b]}(u_m)$ by $e^{-iu_m a}$ and taking the real part we find that

$$A_m = \Re \left\{ \int_a^b e^{iu_m(y-a)} f(y) dy \right\} = \Re \left\{ \phi^{[a,b]}(u_m) e^{-iu_m a} \right\}.$$

Working with the regular CF of the density $\phi(u_m)$ turns out to be more convenient than with its 'truncated' version $\phi^{[a,b]}(u_m)$ since the CF is often available in a semi-closed form for many parametric option pricing models. Fortunately, the truncated version is well approximated by its infinite counterpart $\phi(u_m)$ for sufficiently wide interval $[a, b]$, and so are the cosine coefficients:

$$A_m \approx \Re \left\{ \phi(u_m) e^{-iu_m a} \right\} =: \tilde{A}_m. \quad (3)$$

Therefore, the COS method allows for efficient pricing of options under any parametric model when a closed-form CF is available. In particular, consider a European-style option with a general payoff $v(y, T)$ as a function of the state variable y at a maturity time T . Define the cosine series coefficients of the payoff function $v(y, T)$ as

$$H_m := \frac{2}{b-a} \int_a^b v(y, T) \cos(u_m y - u_m a) dy. \quad (4)$$

Then, the COS formula to evaluate the price of this contract at time $t = 0$ is derived by plugging the Fourier cosine series expansion of the conditional density into the risk-neutral valuation. That is, under the risk-neutral measure \mathbb{Q} with the deterministic interest rate r , we have:

$$\begin{aligned} v_0 &= e^{-rT} \mathbb{E}^{\mathbb{Q}}[v(y, T)] \stackrel{(1)}{\approx} e^{-rT} \int_a^b v(y, T) f(y) dy \\ &= e^{-rT} \int_a^b v(y, T) \left(\frac{2}{b-a} \sum_{m=0}^{\infty}{}' A_m \cos(u_m y - u_m a) \right) dy \\ &= e^{-rT} \sum_{m=0}^{\infty}{}' A_m H_m \stackrel{(2)}{\approx} e^{-rT} \sum_{m=0}^{\infty}{}' \tilde{A}_m H_m \\ &\stackrel{(3)}{\approx} e^{-rT} \sum_{m=0}^{N-1}{}' \tilde{A}_m H_m, \end{aligned} \quad (5)$$

where by $\overset{(i)}{\approx}$ we denote the subsequent numerical approximation. Note that the product of the two functions $v(y, T)$ and $f(y)$ is represented by the product of their Fourier-cosine series coefficients A_m and H_m . The coefficients H_m can be calculated analytically for many types of options. In Appendix B, we provide the analytic formulas for call options.

The COS method allows for fast option evaluation using the Fourier cosine expansion. There are three numerical approximations involved: (1) truncation of the integration range in the risk-neutral expectation, (2) usage of the CF $\phi(u_m)$ (and hence, \tilde{A}_m) instead of the truncated counterpart $\phi^{[a,b]}(u_m)$, and (3) the cosine series truncation. Moreover, what is more important, this method requires a parametric model assumption, which is unknown a priori.

In contrast to the traditional COS method, the approach proposed in this paper does not rely on a parametric specification of the CF. Instead, we use a finite number of plain vanilla option prices observed in the market. Given these observable prices, we can extract the risk-neutral density, price European options with strike prices that are not observed in the market (i.e. perform interpolation), and compute the option sensitivities. Furthermore, as we show in the next subsection, our approach does not require a proper choice of the interval $[a, b]$, i.e., it entirely eliminates the first two numerical approximation errors.

2.2 Option-implied information

Let us denote by S_t the underlying price at time t for a stock or an index under consideration and by F_t the futures price at time t for this underlying asset with some fixed maturity. Let us further denote by $C_0(K)$ and $P_0(K)$ the call and put option prices at time $t = 0$ maturing at date $T > 0$ with a strike price K . Assuming the existence of an arbitrage-free financial market and denoting with \mathbb{Q} the risk-neutral measure, the prices of out-of-the-money (OTM) options at time $t = 0$ are given as the discounted risk-neutral conditional expectations of the corresponding payoff functions:

$$O_0(K) = \begin{cases} C_0(K) = e^{-rT} \mathbb{E}^{\mathbb{Q}}[\max(S_T - K, 0)], & \text{if } K > F_0, \\ P_0(K) = e^{-rT} \mathbb{E}^{\mathbb{Q}}[\max(K - S_T, 0)], & \text{if } K \leq F_0, \end{cases}$$

where r is a deterministic interest rate. Throughout the paper we consider option contracts with a fixed maturity, thus we eliminate the dependence on T in our notations.

The payoff spanning result of Carr and Madan (2001) allows expressing the value of any European-style contingent contract with a general payoff function $v(S_T)$ as a weighted portfolio of a risk-free bond and plain vanilla OTM options:

$$\begin{aligned} v_0 &= e^{-rT} \mathbb{E}^{\mathbb{Q}}[v(S_T)] = e^{-rT} \int_0^\infty v(S_T) f_S(S_T) dS_T \\ &= e^{-rT} v(F) + \int_0^\infty v''(K) O_0(K) dK, \end{aligned} \tag{6}$$

where $f_S(\cdot)$ is the risk-neutral density of the future underlying price, and $F := F_0$ is the futures price at time $t = 0$. This result has been extensively used in the literature to, e.g. construct the VIX index (CBOE, 2015), extract the risk-neutral expectations (Bakshi, Kapadia, & Madan,

2003), and imply the characteristic function in a model-free way (Todorov, 2019, Boswijk, Laeven, & Vladimirov, 2022).

For our purposes, it turns out to be convenient to modify this spanning formula to find values of contingent claims restricted to a finite interval $[\alpha, \beta] \subset \mathbb{R}$. For that, let us denote the risk-neutral valuation on a finite interval $[\alpha, \beta]$ as

$$v_0^{[\alpha, \beta]} = e^{-rT} \mathbb{E}^{\mathbb{Q}}[v(S_T) \mathbf{1}_{\{\alpha \leq S_T \leq \beta\}}] = e^{-rT} \int_{\alpha}^{\beta} v(S_T) f_S(S_T) dS_T.$$

Then we have the following spanning result for $v_0^{[\alpha, \beta]}$.

Proposition 1 *The risk-neutral expectation of a European-style option contract with a twice continuously-differentiable payoff function $v(S_T) \in C^2([\alpha, \beta])$ restricted to a finite interval $[\alpha, \beta]$ with $\alpha \leq F \leq \beta$ can be replicated as follows:*

$$\begin{aligned} v_0^{[\alpha, \beta]} &= e^{-rT} v(F) + \int_{\alpha}^{\beta} v''(K) O_0(K) dK \\ &\quad + v(\beta) C'_K(\beta) - v(\alpha) P'_K(\alpha) - v'(\beta) C_0(\beta) + v'(\alpha) P_0(\alpha), \end{aligned}$$

where $C'_K(x)$ and $P'_K(x)$ are derivatives of the call and put option prices with respect to strike evaluated at x .

The proof follows from the application of the general payoff spanning formula (6) to a function $v(S_T) \mathbf{1}_{\{\alpha \leq S_T \leq \beta\}}$ and using properties of the Dirac Delta function. Like the spanning result of Carr and Madan (2001), Proposition 1 allows to replicate $v_0^{[\alpha, \beta]}$ by constructing a portfolio of risk-free asset and plain vanilla OTM options with strike prices from the interval $[\alpha, \beta]$. Due to the restriction to the interval, the weights at the boundary option contracts, $C_0(\beta)$ and $P_0(\alpha)$, and risk-free asset position are adjusted. Note that with $\alpha \rightarrow 0$ and $\beta \rightarrow \infty$, the value $v_0^{[\alpha, \beta]}$ converges to the unrestricted contract value v_0 .

A natural choice for the truncated interval is the range of observable strike prices, i.e., we can set α to be the smallest observable strike price \underline{K} and β to be the largest observable strike \overline{K} . Alternatively, we can restrict the estimation to the interval with the most liquid options, i.e., $(\alpha, \beta) \subset (\underline{K}, \overline{K})$, which can be practically more appealing. This choice prevents the truncation errors, which are inevitable in the standard Carr-Madan spanning result (6).

Now we can find the replicating portfolio for the (discounted) cosine coefficients A_m . For that, we consider the transformed¹ variable $y = \log \frac{S_T}{x}$ (and, thus, $a = \log \frac{\alpha}{x}$ and $b = \log \frac{\beta}{x}$) with some $x > 0$ and notice that

$$\begin{aligned} A_m &= \int_a^b \cos(u_m y - u_m a) f(y) dy = \int_{\alpha}^{\beta} \cos\left(u_m \log \frac{S_T}{\alpha}\right) f_S(S_T) dS_T \\ &= \mathbb{E}^{\mathbb{Q}} \left[\cos\left(u_m \log \frac{S_T}{\alpha}\right) \mathbf{1}_{\{\alpha \leq S_T \leq \beta\}} \right]. \end{aligned} \quad (7)$$

Applying Proposition 1 to the function $v(S_T) = \cos\left(u_m \log \frac{S_T}{\alpha}\right)$ and denoting the second-order derivative of this function with respect to S_T as

$$\psi_m(s) := v''(s) = \frac{u_m}{s^2} \left(\sin\left(u_m \log \frac{s}{\alpha}\right) - u_m \cos\left(u_m \log \frac{s}{\alpha}\right) \right), \quad (8)$$

¹The motivation for this transformation comes from the COS method, where x is the option's strike price.

we get the (discounted) option-implied cosine coefficients as a portfolio of options:

$$e^{-rT}A_m = \underbrace{e^{-rT} \cos\left(u_m \log \frac{F}{\alpha}\right) + \int_{\alpha}^{\beta} \psi_m(K) O_0(K) dK}_{=:D_m} + \underbrace{\cos(m\pi) C'_K(\beta) - P'_K(\alpha)}_{=:b_m}. \quad (9)$$

Here, the term b_m adjusts D_m to account for the restriction to a finite interval, and with $\alpha \rightarrow 0$ and $\beta \rightarrow \infty$, the adjustment $b_m \rightarrow 0$. Therefore, we will also refer to D_m as the cosine coefficient. It is important to emphasize that the cosine coefficient A_m in equation (9) is completely model-free. This is in contrast to the standard COS method, where one needs to specify a parametric assumption on the dynamics of the underlying asset to get the cosine expansion coefficients via parametric CF. Furthermore, the option-implied coefficients A_m are exact, thus, approximation (2) in equation (5) is avoided.

After having implied the cosine coefficients, we can obtain the risk-neutral density (RND) using equation (1). For instance, setting $x = 1$ gives us the RND of the log future price $\log S_T$:

$$\begin{aligned} f(y) &= \frac{2}{b-a} \sum'_{m=0}^{\infty} A_m \cos(u_m y - u_m a) \\ &= \nu_f \sum'_{m=0}^{\infty} (D_m + b_m) \cos(u_m y - u_m \log \alpha), \end{aligned} \quad (10)$$

where $\nu_f := \frac{2e^{rT}}{\log(\beta/\alpha)}$. Equation (10) is an important representation of a portfolio of option prices with strike prices within the interval $[\alpha, \beta]$. Like the cosine coefficients A_m , the spanning result for the RND in (10) is exact and model-free. This serves as a basis for our non-parametric estimator of the RND, which we discuss in Section 3.

It is worth noting that equation (10) provides the values of the RND for any $y \in [a, b]$, although the density itself may have support on \mathbb{R} . This restriction to the finite interval is coherent with the availability of option data: if there are no options traded with strike prices $K < \alpha$, then it is difficult to infer information about the density accurately for $y < \log \alpha$ without making further (often parametric) assumptions.

2.3 Risk-neutral valuation for plain vanilla options

After having extracted the option-implied cosine coefficients A_m , the same COS machinery can be used to price options but in a model-free way. While there is generally no need to price options that are already observed in the market, the developed iCOS method can be used to further evaluate option prices with the strikes that are not listed in the market. In other words, this approach allows interpolating plain vanilla options within the interval $[\alpha, \beta]$ in a completely model-free way. Additionally, as we discuss in Section 3, option evaluation can be helpful in estimating unobserved quantities and constructing a feasible limiting distribution of the estimated RND and option sensitivities.

For the accurate pricing/interpolation, it is important to take into account the truncation levels, i.e., we shall separate the information available in the interval $[\alpha, \beta]$ from the information

outside of this range. For that, we can decompose the risk-neutral valuation of a contract with the general payoff function $v(S_T)$ as follows:

$$\begin{aligned}
v_0 &= e^{-rT} \mathbb{E}^{\mathbb{Q}}[v(S_T)] = e^{-rT} \int_0^{\infty} v(S_T) f_S(S_T) dS_T \\
&= e^{-rT} \left[\int_0^{\alpha} v(S_T) f_S(S_T) dS_T + \int_{\alpha}^{\beta} v(S_T) f_S(S_T) dS_T + \int_{\beta}^{\infty} v(S_T) f_S(S_T) dS_T \right] \\
&= v_0^{(0,\alpha)} + v_0^{[\alpha,\beta]} + v_0^{(\beta,\infty)}.
\end{aligned} \tag{11}$$

That is, we can express the price of the contract as a sum of values over the three non-overlapping intervals. The motivation for this is to approximate the infinite counterpart by a value on $[\alpha, \beta]$, while possibly taking into account the values outside this interval. In fact, the COS method of Fang and Oosterlee (2008) assumes that the value of a contract on $[\alpha, \beta]$, $v_0^{[\alpha,\beta]}$, represents the contract value v_0 well, i.e., it assumes that the values $v_0^{(0,\alpha)}$ and $v_0^{(\beta,\infty)}$ are negligible.

It turns out that for the plain vanilla options, the values outside this finite interval $[\alpha, \beta]$ can be well controlled, completely eliminating the integration range truncation errors, represented by approximation (1) in equation (5). In particular, for a call option with a strike price $x \in [\alpha, \beta]$, the value on the interval (β, ∞) is given by

$$\begin{aligned}
C_0^{(\beta,\infty)}(x) &= e^{-rT} \mathbb{E}^{\mathbb{Q}}[\max(S_T - x, 0) \mathbf{1}_{\{S_T > \beta\}}] = e^{-rT} \int_{\beta}^{\infty} \max(S_T - x, 0) f_S(S_T) dS_T \\
&= e^{-rT} (\beta - x) \int_{\beta}^{\infty} f_S(S_T) dS_T + e^{-rT} \int_0^{\infty} \max(S_T - \beta, 0) f_S(S_T) dS_T \\
&= -(\beta - x) C'_K(\beta) + C_0(\beta),
\end{aligned}$$

where $C_0(\beta)$ is the call price with the strike β and $C'_K(\beta)$ is its derivative with respect to the strike price evaluated at β . $C_0^{(\beta,\infty)}(x)$ is the price of a so-called *gap call option* with a strike price x and a trigger price β . A similar relation can be found with the *gap put option* for the value of a put option evaluated on the interval $(0, \alpha)$. See Appendix B for the details about the put options.

Therefore, we have the following relations for the call and put option prices with strike prices x such that $\alpha \leq x \leq \beta$:

$$C_0(x) = C_0^{[\alpha,\beta]}(x) + (x - \beta) C'_K(\beta) + C_0(\beta), \tag{12}$$

$$P_0(x) = P_0^{[\alpha,\beta]}(x) + (x - \alpha) P'_K(\alpha) + P_0(\alpha). \tag{13}$$

These decompositions allow us to take into account the truncation of the integral in the risk-neutral valuation. In particular, by setting α and β to the smallest and largest observable strike prices \underline{K} and \overline{K} respectively, we can account for the price of the gap call option using information from the observable range of strike prices. This means that our method is not affected by the choice of $[\alpha, \beta]$, and, hence, $[a, b]$, unlike the COS method.

Therefore, to interpolate call options, we compute $C_0^{[\alpha,\beta]}(x)$ using the COS machinery with the option-implied information from the corresponding interval and add the price of the gap call option. The value of the call contract truncated to the interval $[\alpha, \beta]$ is obtained using the

COS formula as

$$\begin{aligned}
C_0^{[\alpha, \beta]}(x) &= e^{-rT} \sum_{m=0}^{\infty} A_m H_m(x) = \sum_{m=0}^{\infty} D_m H_m(x) + \sum_{m=0}^{\infty} b_m H_m(x) \\
&= \sum_{m=0}^{\infty} D_m H_m(x) + \left(\sum_{m=0}^{\infty} (-1)^m H_m(x) \right) C'_K(\beta) - \left(\sum_{m=0}^{\infty} H_m(x) \right) P'_K(\alpha), \quad (14)
\end{aligned}$$

where $H_m(x)$ are the cosine series coefficients specific to the call payoff function with the strike price x . The closed-form expression for $H_m(x)$ is provided in Appendix B.

Therefore, the price of a call option with strike price $x \in [\alpha, \beta]$ can be represented as

$$\begin{aligned}
C_0(x) &= C_0^{[\alpha, \beta]}(x) + (x - \beta) C'_K(\beta) + C_0(\beta) \\
&= \underbrace{\sum_{m=0}^{\infty} D_m H_m(x)}_{=: \bar{C}_0(x)} + C_0(\beta) + \underbrace{\left(x - \beta + \sum_{m=0}^{\infty} (-1)^m H_m(x) \right)}_{=: Z_c(x)} C'_K(\beta) - \underbrace{\left(\sum_{m=0}^{\infty} H_m(x) \right)}_{=: Z_p(x)} P'_K(\alpha) \\
&= \bar{C}_0(x) + C_0(\beta) + Z_c(x) \theta_c + Z_p(x) \theta_p, \quad (15)
\end{aligned}$$

where we additionally denote $\theta_c := C'_K(\beta)$ and $\theta_p := P'_K(\alpha)$. Equation (15) represents the price of a call option as a portfolio of a continuum of OTM contracts with strike prices from the interval $[\alpha, \beta]$, with additional hedging terms due to the truncation on the finite interval. It is important to note that the decomposition (15) is exact, i.e., it does not involve any numerical approximations and integration range truncation errors thanks to the gap options. In practice, however, we only observe a finite number of OTM options and truncate the cosine series expansion with a finite number of terms N . We address these issues in the next section.

Although this decomposition is circular (to find the price of a single option we need to know the prices of a continuum of options), it is essential in practice, where we observe only a finite number of option prices but might be interested in pricing options with strikes that are not observed in the market, i.e., we use (15) to perform the interpolation.

Finally, the call and put price first-order derivatives with respect to the strike price, θ_c and θ_p , are not directly observable in the market. However, we can approximate them using, e.g., the finite-difference approach. Alternatively, as we show in the next section, we can estimate them from the observed cross-section of option contracts as they are linearly loaded on the call prices. After having estimated these derivatives, we can use them to non-parametrically estimate the RND and option sensitivities.

2.4 Implied delta

In a similar model-free way, we can replicate some option sensitivities such as the option delta, δ , the first derivative of the option price with respect to the underlying asset price S_0 . For that,

we first note that

$$\begin{aligned}
\frac{\partial A_m}{\partial S_0} &= \frac{\partial}{\partial S_0} \mathbb{E}^{\mathbb{Q}} \left[\cos \left(u_m \log \frac{S_T}{\alpha} \right) \mathbf{1}_{\{\alpha \leq S_T \leq \beta\}} \right] \\
&= \mathbb{E}^{\mathbb{Q}} \left[\frac{\partial}{\partial S_T} \left(\cos \left(u_m \log \frac{S_T}{\alpha} \right) \mathbf{1}_{\{\alpha \leq S_T \leq \beta\}} \right) \frac{\partial S_T}{\partial S_0} \right] \\
&= -\frac{u_m}{S_0} \mathbb{E}^{\mathbb{Q}} \left[\sin \left(u_m \log \frac{S_T}{\alpha} \right) \mathbf{1}_{\{\alpha \leq S_T \leq \beta\}} \right],
\end{aligned}$$

where we additionally assume that $\frac{\partial S_T}{\partial S_0} = \frac{S_T}{S_0}$, i.e. the solution to the stochastic process S_T is homogenous of degree one as a function of initial stock price S_0 .

Let us further denote $B_m := e^{-rT} \mathbb{E}^{\mathbb{Q}} \left[\sin \left(u_m \log \frac{S_T}{\alpha} \right) \mathbf{1}_{\{\alpha \leq S_T \leq \beta\}} \right]$. Then, B_m can be replicated similar to the cosine coefficients terms A_m using Proposition 1:

$$B_m = e^{-rT} \sin \left(u_m \log \frac{F}{\alpha} \right) + \int_{\alpha}^{\beta} \tilde{\psi}_m(K) O_0(K) dK - \frac{u_m}{\beta} (-1)^m C_0(\beta) + \frac{u_m}{\alpha} P_0(\alpha), \quad (16)$$

with

$$\tilde{\psi}_m(s) := -\frac{u_m}{s^2} \left(\cos \left(u_m \log \frac{s}{\alpha} \right) + u_m \sin \left(u_m \log \frac{s}{\alpha} \right) \right). \quad (17)$$

Hence, the delta of the call option restricted to the interval $[\alpha, \beta]$ is given by

$$\begin{aligned}
\delta^{[\alpha, \beta]}(x) &= \frac{\partial C_0^{[\alpha, \beta]}(x)}{\partial S_0} = e^{-rT} \sum_{m=0}^{\infty} \frac{\partial A_m}{\partial S_0} H_m(x) \\
&= -\frac{1}{S_0} \sum_{m=1}^{\infty} u_m B_m H_m(x).
\end{aligned}$$

Furthermore, the delta of the gap call option for $x < \beta$ can be expressed as

$$\begin{aligned}
\delta^{(\beta, \infty)} &= \frac{\partial C_0^{(\beta, \infty)}(x)}{\partial S_0} = e^{-rT} \mathbb{E}^{\mathbb{Q}} \left[\frac{\partial}{\partial S_0} \max(S_T - x, 0) \mathbf{1}_{\{S_T > \beta\}} \right] \\
&= \frac{1}{S_0} e^{-rT} \int_{\beta}^{\infty} \mathbf{1}_{\{S_T > x\}} S_T f_S(S_T) dS_T \\
&= \frac{1}{S_0} e^{-rT} \int_{\beta}^{\infty} (S_T - \beta) f_S(S_T) dS_T + \frac{1}{S_0} e^{-rT} \int_{\beta}^{\infty} \beta f_S(S_T) dS_T \\
&= \frac{1}{S_0} (C_0(\beta) - \beta C'_K(\beta)),
\end{aligned}$$

which does not depend on x . Therefore, combining two deltas, the implied delta for a European call is given as following:

$$\delta(x) = \delta^{[\alpha, \beta]}(x) + \delta^{(\beta, \infty)} = -\frac{1}{S_0} \sum_{m=1}^{\infty} u_m B_m H_m(x) + \frac{1}{S_0} (C_0(\beta) - \beta \theta_c). \quad (18)$$

As the RND expansion (10) and option evaluation (15), the replication result for the delta (18) is exact and model-free for all strike values $x \in [\alpha, \beta]$. A similar spanning result can be derived for option gamma, the second order sensitivity. However, one can also easily obtain gamma by properly scaling the option-implied RND (10) (see, e.g., Bates, 2005 and Alexander & Nogueira, 2007).

It is important to emphasize that all three spanning results (given by equations (10), (15), and (18)) are exact, and the restriction to the finite interval does not lead to the associated truncation errors due to the usage of the gap call option. Furthermore, the choice of the interval is data-driven. This is in contrast to the traditional COS method, where one has to select wide interval to minimize the integration range truncation error.

All three option-implied quantities depend on a continuum number of option contracts. In practice, however, we observe only a finite number of option contracts. Nevertheless, all three expressions are easy to approximate using a limited number of observable option prices. Additionally taking into account the observation errors, in the next section, we develop feasible non-parametric estimators for the RND, option prices, and deltas.

3 Estimation

In this section, we introduce the computationally feasible estimators based on the option-implied COS valuation method and derive their asymptotic properties.

3.1 The observation scheme

Unlike the COS method, our approach does not require a parametric specification of model dynamics under the risk-neutral measure. Instead, we use a cross-section of option prices observed in the market. However, these options are observed across a finite set of strike prices and prone to observation errors due to, e.g., bid-ask spread, tick sizes of quotes, and liquidity issues. Therefore, we first describe the observation option scheme.

Our data consists of n OTM option prices observed at time $t = 0$ and expiring at a fixed time $T > 0$ with a deterministic sequence of strike prices:

$$0 < \underline{K} := K_1 < K_2 < \cdots < K_n =: \overline{K} < \infty.$$

For the asymptotic analysis developed below, we assume that the smallest and the largest strike prices, \underline{K} and \overline{K} , are fixed, and the number of options n with strike prices between them goes to infinity by shrinking the strike mesh. We could also consider the joint asymptotic scheme as in Todorov (2019) and Boswijk et al. (2022), where $\underline{K} \rightarrow 0$ and $\overline{K} \rightarrow \infty$ at certain rates as $n \rightarrow \infty$. However, we fix $\alpha = \underline{K}$ to be the smallest strike price and $\beta = \overline{K}$ to be the largest strike price since the choice of the interval $[\alpha, \beta]$ in the spanning results does not introduce an integration range truncation error, as discussed in the previous section.

Assumption 1 *The smallest and largest strike prices are fixed at $\alpha = \underline{K}$ and $\beta = \overline{K}$, and the strike prices between them are equidistant, i.e.*

$$\Delta_n := \Delta K_i = K_i - K_{i-1} = \frac{\beta - \alpha}{n - 1}, \quad i = 2, \dots, n.$$

Assumption 1 can be relaxed to allow for a non-equidistant grid. However, using an equidistant grid simplifies our analysis and enables us to use various numerical approximation methods

under one umbrella. Table 1 summarizes several popular numerical integration methods for an equidistant grid. Later, we comment on our results with a non-equidistant grid of strike prices.

We emphasize again that we do not require an increasing range of strike prices since we account for the restriction to a finite interval. This is different from other methods that estimate RNDs using expansion series, such as in Lu and Qu (2021) and Cui and Yu (2021). Fixing the interval $[\alpha, \beta]$ comes as a great advantage in practice since it allows us to choose different intervals of strike prices without introducing additional approximation errors. For instance, we can consider an interval with only actively traded options.

Table 1: Numerical integration methods and their errors

Methods	Coefficients w_i	Error order
Left Riemann sum	$w_1 = \dots = w_{n-1} = 1, w_n = 0$	$\psi'(s)\mathcal{O}(n^{-1})$
Right Riemann sum	$w_1 = 0, w_2 = \dots = w_n = 1$	$\psi'(s)\mathcal{O}(n^{-1})$
Trapezoidal rule	$w_1 = w_n = \frac{1}{2}, w_2 = \dots = w_{n-1} = 1$	$\psi''(s)\mathcal{O}(n^{-2})$
Simpson's 1/3 rule	$w_1 = w_n = \frac{1}{3}, w_i = \frac{1}{3}(3 + (-1)^i), i = 2, \dots, n-1$	$\psi^{(4)}(s)\mathcal{O}(n^{-4})$

Note: This table provides the list of popular numerical approximations for definite interval in the form of $\int_{\alpha}^{\beta} \psi(x)dx \approx \sum_{i=1}^n w_i \psi(x_i) \Delta_n$. For each method, we report the coefficients w_i assuming equidistant grid Δ_n and corresponding orders of the error terms with s being some number between α and β .

Assumption 2 *Option prices are observed with an additive error term:*

$$O(K_i) = O_0(K_i) + \varepsilon_i, \quad i = 1, \dots, n,$$

where the observation errors ε_i are such that: (1) $\mathbb{E}[\varepsilon_i] = 0$, (2) $\mathbb{E}[\varepsilon_i^2] = \sigma_i^2$ are positive and finite-valued, (3) $\mathbb{E}[\varepsilon_i^4] < \infty$, and (4) ε_i and ε_j are conditionally independent whenever $i \neq j$.

As common in the option pricing literature, Assumption 2 imposes an additive error structure form with independent but possibly heteroskedastic error terms (see, e.g., Andersen, Fusari, & Todorov, 2015, Todorov, 2019, Boswijk et al., 2022). The independence assumption can be further relaxed by considering a spatial dependence as in Andersen, Fusari, Todorov, and Varneskov (2021) at the cost of more complex expressions for the limiting distributions. This would, however, play a secondary role in the developed estimation procedure.

Note that we drop the null index to denote the observed OTM prices. Furthermore, due to the put-call parity, the same observation errors translate into the counterpart in-the-money contracts, i.e., both $C(K_i) = C_0(K_i) + \varepsilon_i$ and $P(K_i) = P_0(K_i) + \varepsilon_i$ for the call and put contracts with the same strike price K_i and error term ε_i .

3.2 Option prices estimator

Using n observable option prices, we can estimate (part of the) cosine coefficients D_m defined in equation (9), by using a numerical approximation of the integral:

$$\hat{D}_m := e^{-rT} \cos\left(u_m \log \frac{F}{\alpha}\right) + \sum_{i=1}^n w_i \psi_m(K_i) O(K_i) \Delta_n, \quad (19)$$

where w_i are the coefficients of a chosen numerical integration method, as listed in Table 1.

The deviation of the estimated cosine expansion coefficient \widehat{D}_m from its true value D_m stems from the observation and discretization errors. These errors also arise in the VIX calculation (see, e.g., Jiang & Tian, 2005 and Jiang & Tian, 2007). Todorov (2019) and Boswijk et al. (2022) also analyze these errors in their estimation procedures along with the truncation errors that arise due to integration over a finite interval. However, in our setting, there are no truncation errors for the cosine coefficients D_m since we take this truncation further into account. See the discussion in Section 2.

Given the fixed smallest and largest strike prices α and β , we have the following asymptotic result for the cosine coefficients.

Proposition 2 *Under Assumptions 1–2, the computationally feasible estimator \widehat{D}_m with fixed $m > 0$ is such that*

$$\mathbb{E} \left[\widehat{D}_m - D_m \right] = \zeta_{m,n}^D,$$

where $\zeta_{m,n}^D = \mathcal{O} \left(\frac{m^{2+\iota}}{n^\iota} \right)$ is the discretization error with the order controlled by the chosen numerical integration scheme $\iota \geq 1$, and as $n \rightarrow \infty$

$$\frac{\widehat{D}_m - D_m}{\sigma_D(m)} \xrightarrow{d} \mathcal{N}(0, 1),$$

with $\sigma_D^2(m) = \sum_{i=1}^n w_i^2 \psi_m^2(K_i) \sigma_i^2 \Delta_n^2$.

The proof of Proposition 2 is provided in Appendix A. The proposition states that although the estimator \widehat{D}_m with fixed m based on a finite number of option prices is biased, it is asymptotically unbiased as the number of option prices n increases. This serves as a building block for the non-parametric estimators introduced below.

Next, we introduce the computationally feasible option-implied call price estimator $\widehat{\overline{C}}(x)$ of the error-free counterpart $\overline{C}_0(x)$, defined in equation (15), with a strike x and the payoff restricted to the interval $[\alpha, \beta]$. It can be expressed as a linear combination of asymptotically unbiased estimators \widehat{D}_m with $m = 1, \dots, N - 1$ as follows:

$$\widehat{\overline{C}}(x) := \sum_{m=0}^{N-1} \widehat{D}_m H_m(x), \quad (20)$$

where N is the number of expansion terms in the Fourier-cosine expansion and $\widehat{D}_0 = e^{-rT}$.

Unlike its error-free counterpart $\overline{C}_0(x)$, the option-implied call price estimator $\widehat{\overline{C}}(x)$ is based on a finite number of noisy option prices and is prone to three types of errors. These errors can be decomposed as follows:

$$\widehat{\overline{C}}(x) - \overline{C}_0(x) = \xi(x) + \zeta(x) + \overline{\eta}(x), \quad (21)$$

where $\xi(x)$, $\zeta(x)$, and $\overline{\eta}(x)$ are observation, discretization and series truncation errors, respectively, all formally defined in Appendix A. The series truncation error $\overline{\eta}(x)$ refers to the truncation of the cosine expansion to the finite number of terms N . To get an order of this truncation error, we additionally impose the assumption on the smoothness of the RND.

Assumption 3 The RND of the future prices $f_S(s) \in C^p([\alpha, \beta])$ with $p > 1$ and $[\alpha, \beta] \subset \mathcal{D}$, where $\mathcal{D} \subseteq \mathbb{R}^+$ is the support of the RND.

Proposition 3 Under Assumptions 1–3, the computationally feasible option-implied call price estimator $\widehat{C}(x)$ with a strike price $x \in [\alpha, \beta]$ is such that

$$\mathbb{E} \left[\widehat{C}(x) - \overline{C}_0(x) \right] = \zeta(x) + \overline{\eta}(x),$$

where the accumulated discretization error $\zeta(x) = \sum_{m=1}^{N-1} \zeta_{m,n}^D H_m(x) = \mathcal{O}\left(\frac{N^{1+\iota}}{n^\iota}\right)$ with $\iota \geq 1$ and the series truncation error $\overline{\eta}(x) = \mathcal{O}(N^{1-p})$. Furthermore, as $n \rightarrow \infty$ and $N \rightarrow \infty$ with $Nn^{-1/2} \rightarrow 0$, we have

$$\frac{\widehat{C}(x) - \overline{C}_0(x)}{\overline{\sigma}_c(x)} \xrightarrow{d} \mathcal{N}(0, 1),$$

where

$$\overline{\sigma}_c^2(x) = \sum_{i=1}^n w_i^2 \psi^2(x, K_i) \sigma_i^2 \Delta_n^2$$

with $\psi(x, K_i) := \sum_{m=1}^{N-1} \psi_m(K_i) H_m(x)$.

The proof of Proposition 3 is provided in Appendix A.

The evaluation of options introduces two types of biases: the discretization error bias $\zeta(x)$ and the series truncation bias $\overline{\eta}(x)$, which arises from the truncation of the cosine expansion as in the original COS method. The former vanishes with an increase in the number of option prices n , while the latter decreases with an increase in the number of expansion terms N as in the COS method. The joint asymptotic for n and N , with N increasing slower than \sqrt{n} , guarantees that the option-implied call price estimator $\widehat{C}(x)$ is asymptotically unbiased. However, in a finite setting with noisy option prices, an increase in expansion terms can lead to a higher variance of the estimators. We address this bias-variance trade-off in the next subsection by choosing an optimal number of expansion terms N^* .

We also note that, like in the COS method, for a sufficiently large interval $[\alpha, \beta]$, the estimator $\widehat{C}(x)$ gives a good approximation for the call price with a payoff unrestricted to this interval, $C_0(x)$. However, to accurately evaluate call options, we also need the first-order derivatives of call and put options θ_c and θ_p evaluated at the boundaries of this interval (see equation (15)). Although we could use finite differences to estimate the first-order derivatives non-parametrically, here we use a simple linear relation of observed option prices on these derivatives instead. In particular, for the observed call price with the strike price K_i , we can get the following decomposition:

$$\begin{aligned} C(K_i) &= C_0(K_i) + \varepsilon_i \\ &= \overline{C}_0(K_i) + C_0(\beta) + Z_c(K_i)\theta_c + Z_p(K_i)\theta_p + \varepsilon_i \\ &= \widehat{C}(K_i) - \xi(K_i) - \zeta(K_i) - \overline{\eta}(K_i) + C(\beta) - \varepsilon_n + Z_c(K_i)\theta_c + Z_p(K_i)\theta_p + \varepsilon_i \\ &= \widehat{C}(K_i) + C(\beta) - \zeta(K_i) - \overline{\eta}(K_i) - \varepsilon_n + Z_c(K_i)\theta_c + Z_p(K_i)\theta_p - \xi(K_i) + \varepsilon_i, \end{aligned}$$

where the second equality follows from equation (15) and the third one from the decomposition (21). Therefore, given n observed option prices, the first derivatives of call and put options θ_c and θ_p can be estimated using a simple linear regression of $C(K_i) - \widehat{C}(K_i) - C(\beta)$ on $Z_c^N(K_i)$ and $Z_p^N(K_i)$ with an intercept $\widehat{\theta}$, where $Z_c^N(K_i)$ and $Z_p^N(K_i)$ are the partial sum counterparts of $Z_c(K_i)$ and $Z_p(K_i)$, respectively, defined in (15). Adding an intercept into this regression reduces finite-sample biases due to the discretization and truncation errors.

Finally, the option-implied call price estimator for any strike price $x \in [\alpha, \beta]$ is given by

$$\widehat{C}(x) := \widehat{C}(x) + C(\beta) + Z_c^N(x)\widehat{\theta}_c + Z_p^N(x)\widehat{\theta}_p + \widehat{\theta}, \quad (22)$$

where $\widehat{\theta}$, $\widehat{\theta}_c$ and $\widehat{\theta}_p$ are the OLS estimates of the aforementioned regression. We emphasize here that while the OLS estimates are obtained using a finite number of observed option prices, the estimator (22) can be obtained for any strike price $x \in [\alpha, \beta]$. Therefore, the call price estimator $\widehat{C}(x)$ can be seen as an interpolation-approximation method, for which we have the following asymptotic result.

Proposition 4 *Under Assumptions 1–3, the computationally feasible option-implied call price estimator $\widehat{C}(x)$ with a strike price $x \in [\alpha, \beta]$ is such that as $n \rightarrow \infty$ and $N \rightarrow \infty$ with $Nn^{-1/2} \rightarrow 0$,*

$$\frac{\widehat{C}(x) - C_0(x)}{\sigma_c(x)} \xrightarrow{d} \mathcal{N}(0, 1),$$


with the variance $\sigma_c^2(x)$ given by equation (A.2) in Appendix A.

The proof of Proposition 4 can be found in Appendix A. Like the estimator $\widehat{C}(x)$, the call price estimator $\widehat{C}(x)$ is asymptotically unbiased when the number of expansion terms grows slower than \sqrt{n} .

The developed call price estimator is related to non-parametric kernel smoothing methods that are widely used in the literature (see, e.g., Ait-Sahalia & Lo, 1998, Grith et al., 2012, Dalderop, 2020), but it can be considered as a ‘global’ smoother. While kernel methods are typically local smoothers (bandwidth parameters control the locality of these estimators), our call price estimator uses all available option prices via the portfolio spanning result (15) discussed in Section 2. This difference allows our method to provide a more flexible approximation of option prices. In the simulation section, we compare these two approaches and demonstrate the superiority of our method.

To demonstrate the ‘global’ nature of our approach, in Figure 1 we display the weights of option portfolios for the at-the-money call option for different number of expansion terms N . The illustration is based on the Black-Scholes model and the simulation set-up is outlined in Section 4. As shown in the figure, our call price estimator for the strike price $K = F_0$ utilizes all available option contracts instead of restricting information to a local neighborhood. When the number of expansion terms increases, the weights concentrate more around the target strike price while still incorporating information from all available contracts.

Figure 1: iCOS weights for the at-the-money call price estimator



Note: This figure plots the weights of the option portfolios for the at-the-money call price estimator for different number of expansion terms N . The illustration is based on the simulated Black-Scholes model with strike prices between 85% and 110% of the spot price with equidistant increments and 201 option contracts. The further simulation details are discussed in Section 4

3.3 RND estimator

After having estimated the cosine coefficients \widehat{D}_m and the first order derivatives $\widehat{\theta}_c$ and $\widehat{\theta}_p$, we can get the non-parametric estimator for the RND of the log price:

$$\widehat{f}(y) = \nu_f \sum_{m=0}^{N-1} \left(\widehat{D}_m + (-1)^m \widehat{\theta}_c - \widehat{\theta}_p \right) \cos(u_m y - u_m \log \alpha), \quad (23)$$

where $\nu_f = \frac{2e^{rT}}{\log(\beta/\alpha)}$. The RND of the future price S_T is obtained by the appropriate transformation of the log price density (23). A similar asymptotic result carries over to the non-parametric RND estimator.

Proposition 5 *Under Assumptions 1–3, the computationally feasible option-implied RND estimator $\widehat{f}(y)$ is such that for any fixed $y \in [\log \alpha, \log \beta]$ as $n \rightarrow \infty$ and $N \rightarrow \infty$ with $Nn^{-1/6} \rightarrow 0$*

$$\frac{\widehat{f}(y) - f(y)}{\nu_f \sigma_f(y)} \xrightarrow{d} \mathcal{N}(0, 1),$$

where $\sigma_f^2(y)$ is the variance term formally defined in equation (A.3) in Appendix A.

The proof of Proposition 4 can be found in Appendix A as well. Unlike the option-implied call price estimator, here we require the number of expansion terms N to grow slower than $n^{-1/6}$ due to the different cosine function. Nevertheless, when this condition is met, the option-implied RND remains asymptotically unbiased.

We emphasize again that this estimator is not for the truncated density, but for the full RND evaluated at any y within the interval $[\log \alpha, \log \beta]$. If one wishes to estimate the RND outside of this interval, additional, often parametric assumptions have to be made about the behavior of the density of options in areas where no option prices are observed. For instance, one possible approach to estimating the RND outside of this interval is to extrapolate option prices beyond the observable range of strike prices using a parametric form based on no-arbitrage conditions.

This extrapolated data can then be used to estimate the RND based on the same estimator (23). In practice, for sufficiently liquid options, the observed range of strike prices covers almost an entire distribution.

3.4 Delta estimator

The non-parametric estimator for the option delta can also be derived in a similar way using the spanning result (18) given a finite number of option prices:

$$\widehat{\delta}(x) = -\frac{1}{S_0} \sum_{m=1}^{N-1} u_m \widehat{B}_m H_m(x) + \frac{1}{S_0} \left(C(\beta) - \beta \widehat{\theta}_c \right), \quad (24)$$

where

$$\widehat{B}_m := e^{-rT} \sin \left(u_m \log \frac{F}{\alpha} \right) + \sum_{i=1}^n w_i \widetilde{\psi}_m(K_i) O(K_i) \Delta_n - \frac{u_m}{\beta} (-1)^m C(\beta) + \frac{u_m}{\alpha} P(\alpha).$$

An analogous asymptotic result holds for the delta estimator (24), but with an additional assumption.

Assumption 4 *The solution to the stochastic process S_T is homogenous of degree one as a function of initial stock price S_0 , i.e., $\frac{\partial S_T}{\partial S_0} = \frac{S_T}{S_0}$.*

Assumption 4 is required to derive the expansion for the delta given in equation (18). The same assumption is imposed for the other non-parametric delta estimators (see, e.g., Bates (2005) and Alexander and Nogueira (2007)).

Proposition 6 *Under Assumptions 1-4, the computationally feasible option-implied delta estimator $\widehat{\delta}(x)$ is such that for any fixed $x \in [\alpha, \beta]$ as $n \rightarrow \infty$ and $N \rightarrow \infty$ with $Nn^{-1/2} \rightarrow 0$*

$$\frac{\widehat{\delta}(x) - \delta(x)}{\frac{1}{S_0} \sigma_\delta(y)} \xrightarrow{d} \mathcal{N}(0, 1),$$

where $\sigma_\delta^2(y)$ is the variance term formally defined in Appendix A.

It is worth noting, that in the current formulations, all limiting results are self-scaling. This implies that equidistant strike price Assumption 1 can be easily relaxed without affecting the limiting distributions.

3.5 Optimal number of expansion terms

As discussed earlier, the number of expansion terms N controls the bias-variance tradeoff in the developed non-parametric estimators. An increase in the number of terms reduces the bias resulting from the series truncation error, but increases the variance of the estimators. To find the optimal number of expansion terms N for the Fourier-cosine expansion in the iCOS method, we consider the expansion for the RND².

²Depending on the purposes, one could also determine the optimal N that minimizes the difference between the observed and estimated option prices. However, since option prices are observed with noise, this approach can potentially lead to severe arbitrage violations.

To assess the impact of truncation on the Fourier-cosine expansion, it is convenient to consider the fit of the density based on the Mean Integrated Squared Error (MISE), defined as

$$\text{MISE}_N := \mathbb{E} \left[\int_a^b (\hat{f}(y) - f(y))^2 dy \right], \quad (25)$$

where $\hat{f}(y)$ is the density estimate based on N expansion terms. Following Leitao, Oosterlee, Ortiz-Gracia, and Bohte (2018) and Kronmal and Tarter (1968), the MISE can be decomposed as follows:

$$\text{MISE}_N = \sum_{m=N}^{\infty} A_m^2 + \sum_{m=1}^{N-1} \mathbb{E} \left[(\hat{A}_m - A_m)^2 \right],$$

where A_m is the m -th Fourier-cosine coefficient, and \hat{A}_m is its estimate. Since the discretization errors contribute to the asymptotically vanishing bias, we consider the Asymptotic MISE (AMISE), where the second moment equals the variance of A_m . Hence, the optimal number of expansion terms N trades off the bias, given by the first part, and the variance of the estimator.

We can derive a recursive relationship in N as follows:

$$\begin{aligned} \text{AMISE}_{N+1} &= \sum_{m=N+1}^{\infty} A_m^2 + \sum_{m=1}^N \text{Var}(\hat{A}_m) \\ &= \sum_{m=N}^{\infty} A_m^2 - A_N^2 + \sum_{m=1}^{N-1} \text{Var}(\hat{A}_m) + \text{Var}(\hat{A}_N) \\ &= \text{AMISE}_N - \left(A_N^2 - \text{Var}(\hat{A}_N) \right), \end{aligned}$$

from which we can see that if $A_N^2 - \text{Var}(\hat{A}_N) > 0$, then $\text{AMISE}_{N+1} < \text{AMISE}_N$. We can use this inequality as a rule to determine the optimal number of expansion terms.

The variance of the cosine coefficient estimators depends on the estimates \hat{D}_N and $\hat{\theta}$, and is provided in a closed-form in Appendix B. A true value of A_N is, however, unknown a priori. Furthermore, due to the presence of discretization error bias in finite settings, this inequality may not accurately represent the MISE. Nevertheless, we can operationalize this inequality to obtain a rule-of-thumb for the optimal number of expansion terms N given the feasible estimates \hat{A}_N . We provide such a rule-of-thumb algorithm in Appendix B.

Finally, we note that the MISE given by equation (25) is for the Fourier-cosine expansion. Hence, the delta estimator, which essentially utilizes the Fourier-sine expansion, may require a different³ optimal choice of expansion terms \tilde{N} . In this case, a similar rule can be applied but with the sine coefficients \hat{B}_N instead.

4 Monte Carlo study

In this section, we investigate the finite-sample performance of the developed non-parametric estimators. We consider two models to generate data: the Black and Scholes (1973) model and

³The sine expansion is known to have a slower rate of convergence than the cosine series. In fact, this is the main reason for popularity of the Fourier cosine expansions rather than the Fourier or sine series.

the ‘double-jump’ stochastic volatility model of Duffie et al. (2000). The former has closed-form solutions for the true quantities of interests, while the latter offers a more realistic depiction of options data that features two stylized facts – stochastic volatility and jump components in returns and volatility.

For each model, we set the initial spot price $S_0 = 4000$, the interest rate $r = 0$, and the strike prices between 85% and 110% of the spot price with equidistant increments of 5, similar to the available S&P 500 index option data. This results in $n = 201$ option contracts for each maturity.

We distort the true option prices of each model by adding homoskedastic observation errors, i.e.,

$$O(K_i) = O_0(K_i) + 0.025 \cdot \epsilon, \quad i = 1, \dots, n,$$

where ϵ is an i.i.d. standard normal random variable. This error structure roughly matches the dispersion of errors in empirical applications, where option errors are typically within the tick size of \$0.05. Note that the smallest and the largest strike prices are fixed and the corresponding OTM options are strictly positive.

4.1 Black-Scholes model

First, we consider the Black-Scholes model with short (30 days) and long (1 year) maturities, and set the volatility parameter $\sigma = 0.3$. The true option prices are generated via the Black-Scholes formula and then distorted with the additive error terms as described above. Table 2 provides the simulation results of the estimated option call prices for a selection of strike prices, along with the estimates of $\hat{\theta}$. The latter includes the intercept $\bar{\theta}$ and the first-order derivatives θ_c and θ_p , which have closed-form solutions in the Black-Scholes model. The number of expansion terms is set to $N = 14$ for short maturity options and to $N = 7$ for long maturity options. This choice is motivated by the rule-of-thumb discussed in Section 3.5 and Appendix B. The numerical integration scheme is set to Simpson’s 1/3 rule throughout the simulations.

Although we estimate option prices from the set of OTM prices themselves, the simulation results indicate the convergence of the considered approach. Furthermore, the Monte Carlo results show no significant bias for all levels of strikes and maturities and a reduction in the variance of option prices. In fact, the Monte Carlo standard deviations of the estimated prices are smaller than the standard deviations used to simulate option errors, indicating the smoothing effect of the estimation procedure. The asymptotic standard deviations, defined as the square root of the average estimated asymptotic variance, roughly correspond to the Monte Carlo standard errors, which indicates the validity of the constructed standard errors.

The estimated parameters $\hat{\theta}$ also exhibit good finite-sample performance. The estimated intercept $\bar{\theta}$, which collects the average of finite-sample biases, indicates that these errors are of rather small order. The good finite-sample performance of the first-order derivatives θ_c and θ_p is crucial for the RND and delta estimators considered below.

The Monte Carlo simulation results for the RND estimates are reported in Table 3. Note that the RND estimator is for the log price $\log S_T$ and evaluated at a few strike levels after the

Table 2: Monte Carlo results for the call prices based on the Black-Scholes model

K/F_0	0.86	0.9	0.95	1.0	1.05	1.09	$\bar{\theta}$	θ_c	θ_p
$T = 30$ days									
$C_0(K)$	565.11	417.38	256.86	137.21	62.66	29.79	0.0	-0.125	0.032
MC bias	0.0002	0.0002	-0.00032	0.00031	0.0007	-0.00117	0.00347	0.00064	-0.00045
MC std	0.0087	0.0071	0.0066	0.0066	0.0072	0.0083	0.0205	0.0009	0.0011
As. std	0.0084	0.007	0.0068	0.0068	0.0068	0.008	0.0141	0.0008	0.001
$T = 1$ year									
$C_0(K)$	777.92	680.52	571.75	476.94	395.27	338.66	0.0	-0.32	0.348
MC bias	-0.00065	-0.00033	-0.00014	0.00052	0.00014	-0.00064	0.00157	0.00013	0.69513
MC std	0.0063	0.0055	0.0048	0.0049	0.0049	0.0058	0.0231	0.0003	0.0004
As. std	0.006	0.0054	0.0047	0.0049	0.005	0.0057	0.0113	0.0003	0.0004

Note: This table provides Monte Carlo simulation results for the call option price estimates, based on 1000 replications from the Black-Scholes model. Two settings with short (30 days) and long (1 year) maturities are considered. For different strike prices, each panel lists the true call value (C_0), the Monte Carlo bias (MC bias), the Monte Carlo standard deviation (MC std), and the asymptotic standard deviation (As. std), defined as the square root of the average estimated asymptotic variance. The number of expansion terms is $N = 14$ and $N = 7$ for $T = 30$ days and $T = 1$ year option contracts, respectively. The numerical integration scheme is Simpson's 1/3 rule.

Table 3: Monte Carlo results for the RND based on the Black-Scholes model

K/F_0	0.86	0.9	0.95	1.0	1.05	1.09
$T = 30$ days						
$f(\log K)$	1.07	2.31	3.98	4.63	3.85	2.69
MC bias	-0.002	-0.0029	0.0037	-0.0018	-0.0052	0.0083
MC std	0.0574	0.0239	0.02	0.0212	0.0241	0.0615
As. std	0.055	0.0235	0.0203	0.0213	0.0229	0.0571
$T = 1$ year						
$f(\log K)$	1.25	1.3	1.33	1.31	1.27	1.21
MC bias	0.0046	0.0009	0.0002	-0.001	0.0001	0.0074
MC std	0.0185	0.006	0.0033	0.0042	0.005	0.015
As. std	0.0176	0.0058	0.0033	0.0043	0.005	0.0149

Note: This table provides Monte Carlo simulation results for the RND estimates, based on 1000 replications from the Black-Scholes model. Two settings with short (30 days) and long (1 year) maturities are considered. For different strike levels, each panel lists the true RND value for log future price ($f(\log K)$), the Monte Carlo bias (MC bias), the Monte Carlo standard deviation (MC std), and the asymptotic standard deviation (As. std), defined as the square root of the average estimated asymptotic variance. The number of expansion terms is $N = 14$ and $N = 7$ for $T = 30$ days and $T = 1$ year option contracts, respectively. The numerical integration scheme is Simpson's 1/3 rule.

corresponding log transformation. The true RND for the Black-Scholes model is the normal distribution with mean $\log S_0 - \frac{1}{2}\sigma^2 T$ and variance $\sigma^2 T$. Similar to the option price estimates, the estimated RNDs exhibit good finite-sample performance. The asymptotic standard errors roughly match the Monte Carlo standard errors, and both tend to increase towards the bounds of the considered interval.

Finally, the simulation results for the delta estimates are reported in Table 4. The true values are obtained in the closed-form under the Black-Scholes assumptions. The number of

Table 4: Monte Carlo results for the call deltas based on the Black-Scholes model

K/F_0	0.86	0.9	0.95	1.0	1.05	1.09
$T = 30$ days						
$\delta(K)$	0.964	0.898	0.739	0.517	0.3	0.169
MC bias	-0.00622	-0.0064	-0.00622	-0.00634	-0.00669	-0.00762
MC std	0.00122	0.00118	0.0012	0.00116	0.00106	0.00125
As. std	0.0014	0.00138	0.00142	0.00137	0.00129	0.00141
$T = 1$ year						
$\delta(K)$	0.74	0.69	0.63	0.56	0.49	0.45
MC bias	-0.0027	-0.00313	-0.00308	-0.00312	-0.00324	-0.00387
MC std	0.00138	0.00134	0.00134	0.00135	0.00119	0.00131
As. std	0.00139	0.00136	0.00136	0.00136	0.00122	0.00133

Note: This table provides Monte Carlo simulation results for the delta estimates, based on 1000 replications from the Black-Scholes model. Two settings with short (30 days) and long (1 year) maturities are considered. Each panel lists, for different strike levels, the true delta value, the Monte Carlo bias and standard deviation. For different strike levels, each panel lists the true delta value ($\delta(K)$), the Monte Carlo bias (MC bias), the Monte Carlo standard deviation (MC std), and the asymptotic standard deviation (As. std), defined as the square root of the average estimated asymptotic variance. The number of expansion terms is $\tilde{N} = 25$ and the numerical integration scheme is Simpson’s 1/3 rule.

expansion terms for the delta estimator is set to $\tilde{N} = 25$ since the sine series exhibits a slower convergence rate (see discussion in Section 3.5). Unlike the RND and call price estimates, the estimation of the delta exhibits small bias terms, which are economically likely to be negligible.

4.2 SVCJ model

The ‘double-jump’ stochastic volatility model of Duffie et al. (2000), labeled as SVCJ, allows for stochastic volatility and jumps in returns and volatility, and under the risk-neutral measure \mathbb{Q} is given by the following system of stochastic differential equations:

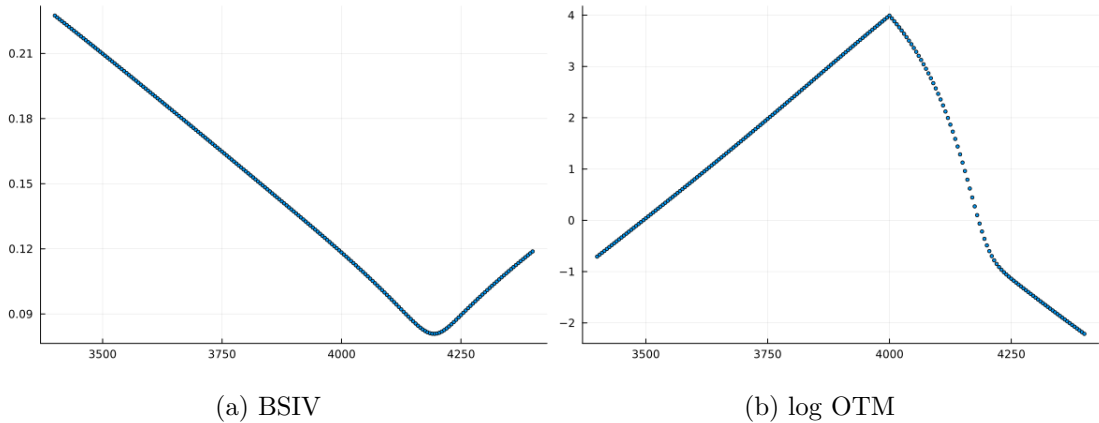
$$\begin{aligned} d \log S_t &= (r - \frac{1}{2}v_t - \mu\lambda)dt + \sqrt{v_t}dW_{1,t} + J_t dN_t, \\ dv_t &= \kappa(\bar{v} - v_t)dt + \sigma\sqrt{v_t}dW_{2,t} + J_t^v dN_t, \end{aligned}$$

where two Brownian motions W_1 and W_2 are correlated with coefficient ρ , and N_t is a Poisson jump process with intensity λ . The jump sizes in returns are Gaussian, $J \sim \mathcal{N}(\mu_j, \sigma_j^2)$ with the expected relative jump size in returns $\mu = \exp(\mu_j + \frac{1}{2}\sigma_j^2) - 1$, while the co-jump sizes in volatility are exponentially distributed, $J^v \sim \exp(1/\mu_v)$, and independent of jump sizes in returns. We choose the following parameter values for the simulation:

$$v_0 = 0.1^2, \kappa = 2.6, \bar{v} = 0.02, \rho = -0.95, \sigma = 0.3, \lambda = 1.0, \mu_j = -0.05, \sigma_j = 0.03, \mu_v = 0.05.$$

Since there is no closed-form solution for option prices under the SVCJ model, we simulate them using the COS method. We use the analytic solution for the CCF and a large number of expansion terms $N = 1024$ with $[a, b] = [-4\sqrt{T}, 4\sqrt{T}]$. We then add observation errors to these option prices as described previously.

Figure 2: Simulated option prices from the SVCJ model on BSIV and log-OTM spaces



Note: This figure illustrates the simulated option prices from the SVCJ model displayed on BSIV (Panel (a)) and log-OTM (Panel (b)) spaces. The simulation details are provided in the main text.

Figure 2 displays the simulated option prices from the SVCJ model on BSIV and log OTM spaces for $T = 30$ days. The chosen model parameters generate the so-called implied volatility ‘smile’, which is commonly observed in the market, particularly for short-dated options. Capturing such pronounced ‘smiles’ can be challenging for many parametric and non-parametric methods since they require the methods to be rather flexible. As a consequence, these methods often fail to accurately capture option prices, RND, and deltas.

For the SVCJ model, we compare the simulation results of the developed approach with the closest non-parametric and widely-used alternative, the kernel smoother. In fact, kernel smoothing methods are also model-free and do not require any optimization routines. In particular, we consider the Nadaraya–Watson kernel estimator with the Gaussian kernel applied to the BSIV space, as in, e.g., Aït-Sahalia and Lo (1998) and Grith et al. (2012). After smoothing BSIV observations, we convert them into price dimension to obtain call price estimates. We then calculate the second-order derivatives to obtain the RND estimates due to Breeden and Litzenberger (1978). Fitting option prices on implied volatility space is commonly used in practice (see, e.g., Aït-Sahalia and Lo (1998), Andersen, Fusari, and Todorov (2015) among many others).

Finding the bandwidth parameter h is crucial for the kernel smoothing methods as it controls the bias-variance tradeoff. Since we are interested in estimating both option prices and the RND, we consider the kernel bandwidths, as in Aït-Sahalia and Lo (1998), of the following form:

$$h = \frac{c}{\log n} n^{-\frac{1}{2p+1}}$$

with $p = 2$ and some constant $c > 0$. In practice, one can use the cross-validation to find the optimal bandwidth, but in simulations we vary the constant c to illustrate the bias-variance tradeoff.

Table 5 provides simulation results for the call price estimates under the simulated SVCJ model. We compare the results obtained using our proposed iCOS method with the widely-used kernel smoothing approach with different bandwidth parameters. First, we observe that

Table 5: Monte Carlo results for the call prices under the SVCJ model

	K/F_0	0.86	0.9	0.95	1.0	1.05	1.09
$C_0(K)$		560.66	402.23	210.81	54.13	0.61	0.14
iCOS	MC bias	-0.00075	0.00039	-0.00036	3.0e-5	-0.00207	-0.00055
	MC std	0.0102	0.00969	0.00941	0.00926	0.009	0.00933
	As. std	0.00981	0.00944	0.009	0.00909	0.00906	0.00893
KS, $c = 0.2$	MC bias	-0.07909	-0.00276	-0.00383	-0.10338	0.35918	-0.04309
	MC std	0.00442	0.00427	0.0044	0.00433	0.00637	0.00369
KS, $c = 0.1$	MC bias	-0.01325	2.0e-5	-0.00027	-0.02496	0.11357	-0.00912
	MC std	0.00621	0.00582	0.00603	0.00589	0.00661	0.00594
KS, $c = 0.05$	MC bias	-0.00078	0.0005	6.0e-5	-0.00627	0.03224	-0.00178
	MC std	0.00828	0.00811	0.00835	0.00831	0.00839	0.00834
KS, $c = 0.03$	MC bias	-0.00066	0.00068	0.00024	-0.00249	0.01229	-0.00133
	MC std	0.01043	0.01051	0.01063	0.01061	0.01068	0.01082

Note: This table provides Monte Carlo simulation results for the call option price estimates under the SVCJ model based on the iCOS approach and the kernel smoothing (KS) with different bandwidths values. The number of expansion terms is $N = 25$ and the numerical integration scheme is Simpson’s 1/3 rule.

the simulation results for the call price estimates under the SVCJ model using our approach are similar to the results based on the Black-Scholes model discussed in the previous subsection. Second, as expected, the biases for the kernel smoother decline as the bandwidth parameter decreases, but this comes at the cost of increased variance. Notably, the biases for the kernel smoother are especially pronounced at the moneyness level of 1.05, which roughly corresponds to the ‘turning’ point of the smile depicted in Figure 2. Finally, when comparing two methods, we notice that only the results with the parameter $c = 0.03$ for the kernel smoother are comparable to the iCOS approach in terms of biases. However, such a small bandwidth value results in a non-smooth RND as we discuss below.

Table 6 provides the Monte Carlo simulation results for the RND of log future prices. We again compare our approach with the kernel smoothing method using the same set of bandwidth parameters. The developed iCOS approach demonstrates a good finite-sample performance, similar to the Black-Scholes case. In contrast, the results for the kernel smoothing approach are often biased or exhibit large variances, which is an indication of non-smooth and, hence, arbitrage-violating RND estimates.

Analogously, Table 7 provides the Monte Carlo results for the call delta estimates under the SVCJ model. The non-parametric iCOS method yields insignificant biases but larger variances than the kernel smoothing method. The latter, however, again exhibits biases at the moneyness level of 1.05 except for the parameter $c = 0.03$, which corresponds to a non-smooth RND.

Overall, when comparing two approaches, we observe that the kernel smoothing method fails to fully capture the shape of the observed option data, which results in the biased estimates of the call prices, RND, and deltas. Decreasing the bandwidth parameter reduces the biases but at the cost of a non-smooth RND with large variance. In contrast, the iCOS method is

Table 6: Monte Carlo results for the RND under the SVCJ model

	K/F_0	0.86	0.9	0.95	1.0	1.05	1.09
$f(\log K)$		0.13	0.49	2.76	11.37	3.11	0.03
iCOS	MC bias	0.0211	0.0032	0.0058	-0.0066	0.0711	0.0216
	MC std	0.1503	0.1138	0.0994	0.0972	0.0954	0.0937
	As. std	0.1462	0.1099	0.095	0.0933	0.091	0.0854
KS, $c = 0.2$	MC bias	0.002	0.0511	0.0169	0.319	-0.2341	0.0161
	MC std	0.0033	0.005	0.006	0.0702	0.0075	0.0043
KS, $c = 0.1$	MC bias	-0.0999	-0.0373	-0.1284	-0.2761	-0.1874	-0.0584
	MC std	0.0211	0.0263	0.0286	0.0316	0.0327	0.0285
KS, $c = 0.05$	MC bias	-0.0209	-0.0203	-0.0609	-0.1186	-0.1087	-0.0186
	MC std	0.1364	0.1469	0.1541	0.1627	0.1743	0.1844
KS, $c = 0.03$	MC bias	0.0063	-0.0174	-0.0296	0.0006	-0.0703	-0.0091
	MC std	0.469	0.4929	0.5117	0.5418	0.5904	0.6125

Note: This table provides Monte Carlo simulation results for the RND estimates under the SVCJ model based on the iCOS approach and the kernel smoothing (KS) with different bandwidths values. The number of expansion terms is $N = 25$ and the numerical integration scheme is Simpson's 1/3 rule.

able to simultaneously capture the shape of the observed option data, RND and deltas with insignificant biases and small variances.

5 Empirical applications

In this section, we first demonstrate the proposed estimation procedures in two empirical illustrations. Then, we apply the developed method to analyse errors in the VIX index.

5.1 SPX options

In the first empirical illustration, we consider options on the S&P 500 stock market index (SPX) obtained from the Chicago Board Options Exchange (CBOE), which are commonly used in the literature. We consider a snapshot of options at 3:45 pm ET with a time-to-maturity $T = 29$ days traded on April 1, 2021. The forward price implied from the put-call parity is $F = \$4008.5$. For the developed estimation procedure, we utilize the mid-quote prices of OTM contracts. Additionally, we compare the pricing accuracy of our approach with the bid-ask spread of the corresponding contracts, which allows us to assess the performance of our method in the real-world options market context.

We focus our analysis on the interval $[\alpha, \beta] = [2950, 4400]$, i.e., we use options with strikes from this interval. Although there are a few option contracts with strike prices beyond this range, they tend to be less liquid and spaced further apart⁴. We do not filter out any options, except for those with zero bid prices. This results in a total number of 239 OTM options with

⁴The distance between strike prices outside the considered interval is \$50 or \$100, while the distance for close to ATM options is only \$5.


Table 7: Monte Carlo results for the delta under the SVCJ model

	K/F_0	0.86	0.9	0.95	1.0	1.05	1.09
$\delta(K)$		0.9959	0.9852	0.9195	0.5959	0.0156	0.0011
iCOS	MC bias	5.6e-5	9.8e-5	7.6e-5	0.000129	0.000177	3.8e-5
	MC std	0.0014	0.0012	0.0013	0.0012	0.0014	0.0011
	As. std	0.0019	0.0019	0.002	0.0019	0.002	0.0019
KS, $c = 0.2$	MC bias	-0.000997	-2.1e-5	8.5e-5	0.001071	0.009063	0.000448
	MC std	3.7e-5	5.2e-5	5.7e-5	7.9e-5	8.6e-5	4.5e-5
KS, $c = 0.1$	MC bias	-0.000641	-1.0e-5	-4.1e-5	0.000165	0.003354	0.000397
	MC std	0.000135	0.000143	0.000152	0.00016	0.000182	0.000168
KS, $c = 0.05$	MC bias	-3.1e-5	-0.0	-3.5e-5	1.9e-5	0.001163	1.7e-5
	MC std	0.000383	0.000382	0.00041	0.000432	0.000474	0.00047
KS, $c = 0.03$	MC bias	1.0e-6	7.0e-6	-4.2e-5	7.0e-6	0.000594	-1.1e-5
	MC std	0.000772	0.000774	0.000833	0.000891	0.000948	0.00096

Note: This table provides Monte Carlo simulation results for the call delta estimates under the SVCJ model based on the iCOS approach and the kernel smoothing (KS) with different bandwidths values. The number of expansion terms is $\tilde{N} = 30$ and the numerical integration scheme is Simpson's 1/3 rule.

non-equidistant strike prices.

Figure 3: Estimated cosine coefficients



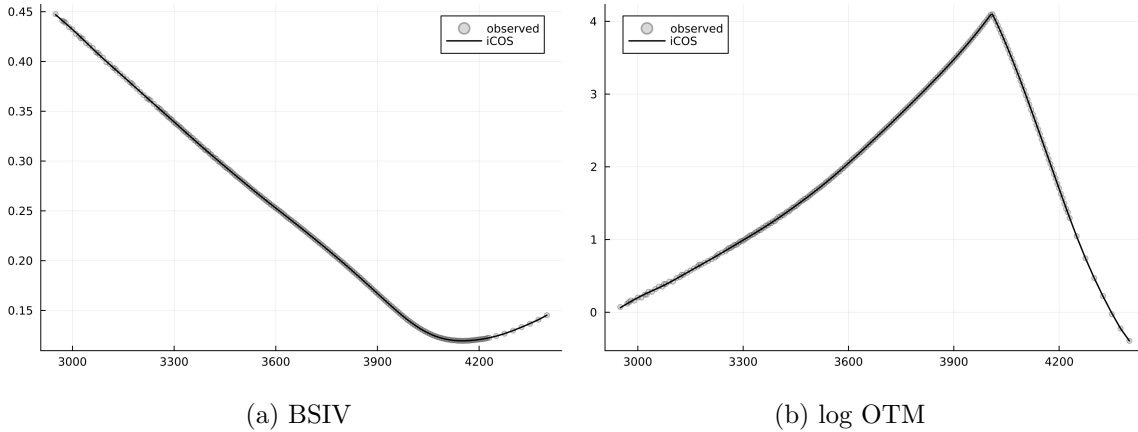
Note: This figure plots the option-implied cosine coefficient estimates \hat{D}_m and \hat{A}_m based on Simpson's 1/3 rule. The SPX options are with 29 days-to-maturity traded on April 1, 2021.

To approximate the integrals in the option-implied cosine coefficients (19), we use Simpson's 1/3 rule. Figure 3 displays the estimated coefficients \hat{D}_m and \hat{A}_m plotted against the term number m . The coefficients \hat{A}_m are displayed after applying a logarithmic transformation, along with their corresponding standard deviations σ_A . As expected, these coefficients converge towards zero up to a certain point but then exhibit a divergent pattern due to the discretization errors and increased variance. The rule-of-thumb algorithm, detailed in Appendix B, selects $N^* = 23$ as the optimal number of terms. This number roughly corresponds to the point where the cosine coefficients stabilize and their standard deviations surpass the magnitude of

the coefficients themselves. We use this number of terms for the subsequent analysis.

Figure 4 depicts the option-implied call price estimates $\widehat{C}(x)$ given by equation (22) for the considered data. The figure displays the market prices in terms of the BSIV and the logarithm of OTM prices, although the estimation is performed in terms of dollar-amount OTM prices. As observed, the estimates closely capture the shape of the implied volatility smile and the log OTM prices. Note that the solid line does not pass exactly through all market prices, but provides an approximation of them.

Figure 4: SPX option price estimates based on the iCOS method



Note: This figure illustrates the call price estimation of the proposed method based on the SPX options traded on April 1, 2021 with 29 days-to-maturity. Panel (a) displays the result of the estimation along with the observed option prices on BSIV space, while Panel (b) plots displays it on log OTM space. The estimation is conducted in terms of the dollar-amount OTM prices.


To investigate the accuracy of our estimation procedure, we plot in Figure 5(a) the pricing errors as the difference between the call option price estimates $\widehat{C}(x)$ and the market observed call prices $C(x)$ for $x \in \{K_1, \dots, K_n\}$. Most of these differences range between \$-0.05 and \$0.05, corresponding to two ticks of size \$0.05 in this dataset. Notably, these differences exhibit a nearly homoskedastic pattern with respect to strike prices. In other words, the error terms do not vary much with the price level of the options and strike prices, contrary to what is often assumed in the literature. It suggests that the primary source of observation errors for these highly liquid SPX options can be attributed to the minimum tick size.

Panel (b) in Figure 5 displays the same pricing differences divided by the half-spread, which is defined for each option contract with strike price x as

$$HS(x) = \frac{AskO(x) - BidO(x)}{2},$$

where $AskO(x)$ and $BidO(x)$ are the ask and bid prices for the OTM option with strike price x , respectively. The half-spread is calculated based on put option quotes if $x < F$ and on call options otherwise. Since we use mid-quote prices $O(x) = \frac{AskO(x) + BidO(x)}{2}$ for our estimation, these pricing errors indicate the percentage distance between the mid-quote price and the bid (if negative) or ask (if positive) prices. As shown in Figure 5(b), most pricing errors fall within half of the bid-ask spread, indicating a very good pricing performance of the method.

Figure 5: Pricing errors for call estimates



Note: This figure illustrates the pricing errors of the proposed method based on the SPX options traded on April 1, 2021 with 29 days-to-maturity. Panel (a) plots the pricing errors $\widehat{C} - C$. Panel (b) plots pricing errors relative to the half-spread $\frac{\widehat{C} - C}{HS}$.

Finally, Figure 6 plots the non-parametric estimates for the RND of future asset price and call deltas for the SPX options data. The RND estimate of a future price is obtained from the RND estimate of a log future price given by equation (23). Given the derived asymptotic theory (and the appropriate delta-rule), the confidence interval for the RND is displayed in a gray area. It is very narrow in the main part of the distribution and slightly widens towards the ends of the considered interval. We also note that we do not impose any arbitrage-free conditions on our RND and option price estimators. Thus, the RND estimates have negative values for some values of strikes. However, such minor arbitrage violations are unlikely to have any practical implications since all corresponding call price estimates fall within the minimum tick size and bid-ask spread.


The estimated call deltas $\widehat{\delta}$ are displayed alongside deltas based on the Black-Scholes model, δ_{BS} . As observed, the Black-Scholes deltas can substantially underestimate the in-the-money call deltas. This might potentially result in hedging errors as discussed in Bates (2005) and Alexander and Nogueira (2007).

5.2 AMZN options

In our second application, we examine equity options on Amazon with a very short time-to-maturity of $T = 1$ day. These options are traded on the Earning Announcement Day (EAD) of April 26, 2018, prior to the announcement itself. Compared to the SPX options, Amazon options are less liquid and are prone to larger observation errors due to their very short maturity. Moreover, the EAD introduces extra uncertainty about the stock price at the expiry.

Similar to the SPX options, we use mid-quote prices of OTM contracts and filter out only zero-bid contracts. We concentrate our analysis on the interval $[\alpha, \beta] = [1250, 1760]$, which corresponds to approximately 18% below and 16% above the underlying spot price of \$1518.96 on this EAD. The availability of such a wide interval for short-dated options is attributed to the information uncertainty surrounding the EAD. Based on the rule-of-thumb for the optimal

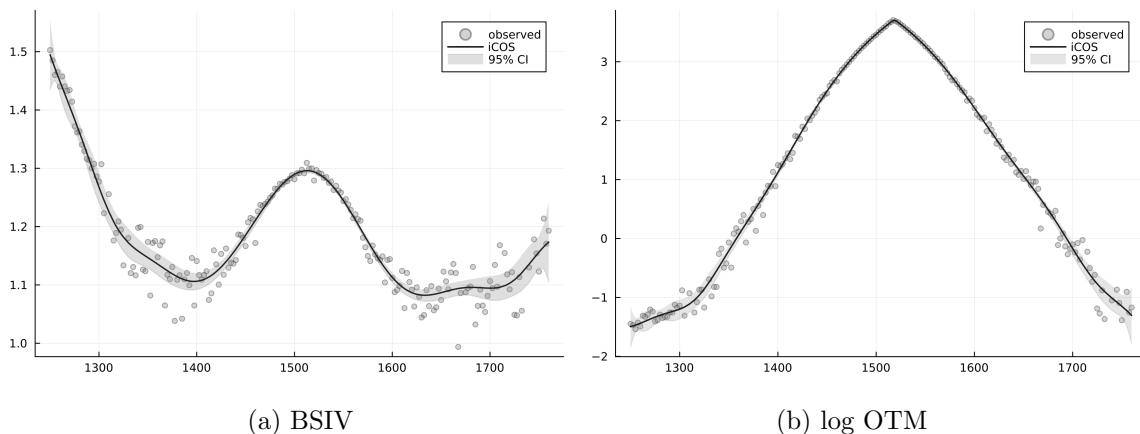
Figure 6: Option-implied RND and call delta for SPX options



Note: This figure plots the estimated RND (Panel (a)) with the 95% confidence interval and call deltas (Panel (b)) for the SPX options with 29 days-to-maturity traded on April 1, 2021. Panel (b) also depicts the Black-Scholes call deltas, δ_{BS} .

number of expansion terms, we set $N = 13$.

Figure 7: AMZN option price estimates based on the iCOS method



Note: This figure illustrates the call price estimation of the proposed method based on the AMZN options traded on April 26, 2018 with 1 day-to-maturity. Panel (a) displays the result of the estimation along with the observed option prices on BSIV space, while Panel (b) displays it on log OTM space.


Figure 7 presents the estimation result for the option-implied call prices displayed on BSIV and log OTM spaces. Notably, the BSIVs are exceptionally high, reaching approximately 130% for ATM options expiring in just one day. Furthermore, these BSIVs exhibit a distinctive W-shaped pattern, which is atypical for implied volatility curves⁵. The W-shape arises from the anticipation of a significant stock price jump following the earnings announcement release. Alexiou, Goyal, Kostakis, and Rompolis (2021) document frequent concave patterns in implied volatilities prior to the EAD for equity options.

Additionally, we observe a large dispersion of option prices. However, our estimation pro-

⁵Note that most parametric curves and models commonly used in the literature would fail to capture this pattern, leading to large estimated errors.

cedure effectively smoothes out the noisy data, resulting in accurate price estimates. For this example, in Figure 7 we also display the 95% confidence interval around the estimated call prices, obtained by applying the appropriate delta rules for the derived asymptotic results. We emphasize that this confidence interval reflects the uncertainty around the estimates and not the observation errors in option prices. Therefore, it does not and need not cover the observed prices.


Figure 8: Pricing errors for AMZN call estimates



Note: This figure illustrates the pricing errors of the proposed method based on the AMZN options traded on April 26, 2018 with 1 day-to-maturity. Panel (a) plots the pricing errors $\hat{C} - C$. Panel (b) plots pricing errors relative to the half-spread $\frac{\hat{C} - C}{HS}$.

Figure 8 displays the pricing errors for call price estimates. As with the SPX options, we plot the pricing errors and the errors relative to the half-spread. Consistent with Figure 7, the pricing errors are larger than those for the SPX options but are still centered around zero.

Figure 9: Option-implied RND and call delta for AMZN options



Note: This figure illustrates the estimated RND (Panel (a)) and call deltas (Panel (b)) with 95% confidence intervals for the 1 day-to-maturity AMZN options traded on April 26, 2018.

Finally, Figure 9 displays the estimated RND and deltas for Amazon options. The consequence of the W-shaped implied volatility curve is a bimodal RND, reflecting the market's

anticipation of two possible outcomes. The two modes of the estimated RND are at \$1442 and \$1590, which corresponds to around 5% down and 4.7% up from the spot underlying stock price, respectively. After the announcement, the next day's opening price for Amazon was \$1634 and it closed at \$1574 (3.62% up from the spot price). The bimodality of the RND is also reflected in the estimated deltas. As shown in Figure 9(b), the Black-Scholes deltas overestimate the deltas for the strikes around the first mode and underestimate them for strike prices close to the second mode.

5.3 Errors in VIX

The developed methodology allows us to analyse errors embedded in the VIX index. As noted by Jiang and Tian (2007), the construction of the VIX is prone to several types of approximation errors, including truncation and discretization errors. The former arises from truncating the real line to the range of observed strike prices, and the latter is due to the discreteness of strike prices. On top of that, option prices used in the VIX construction are subject to observation errors since the true prices are not observed perfectly, as argued in Section 3.1. This results in observation errors in the VIX index. Our methodology enables us to estimate and disentangle observation and discretization errors in the VIX.

In particular, the CBOE calculates the VIX index as⁶

$$\text{VIX} = 100 \cdot \sqrt{\frac{2}{T} e^{rT} \sum_{i=1}^n \frac{\Delta K_i}{K_i^2} O(K_i) - \frac{1}{T} \left(\frac{F}{K_0} - 1 \right)^2}, \quad (26)$$

where K_0 is the largest strike price below the forward level F , $\Delta K_i = \frac{1}{2}(K_{i+1} - K_{i-1})$ for $i = 2, \dots, n-1$, and $T = 30$ days. For more details, see the CBOE (2015) white paper. Since the OTM option prices are observed with noise, the VIX itself contains measurement error. Given the consistent estimator of option prices $\widehat{O}(K_i)$, we can (re-)construct the VIX using these estimates and obtain estimates of the observation errors in the VIX. That is, we calculate

$$\widehat{\text{VIX}} = 100 \cdot \sqrt{\frac{2}{T} e^{rT} \sum_{i=1}^n \frac{\Delta K_i}{K_i^2} \widehat{O}(K_i) - \frac{1}{T} \left(\frac{F}{K_0} - 1 \right)^2} \quad (27)$$

and define $\widehat{\xi}_{\text{VIX}} := \text{VIX} - \widehat{\text{VIX}}$ as the estimator of the observation error in the VIX index.

On the other side, the VIX is developed to approximate the model-free implied volatility. Since our method allows us to further consistently interpolate between observed strike prices, we can construct the measure of the model-free corridor implied volatility (CIV) as

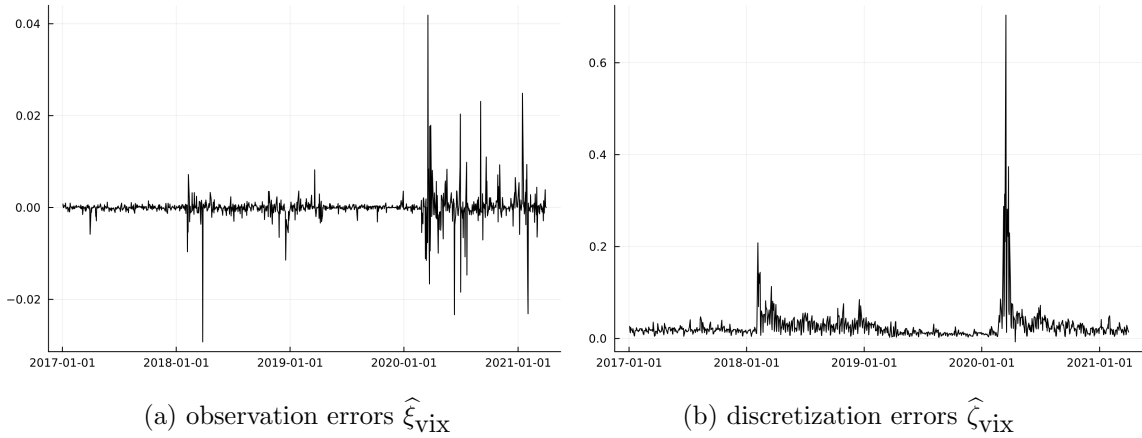
$$\widehat{\text{CIV}} = 100 \cdot \sqrt{\frac{2}{T} e^{rT} \int_{\alpha}^{\beta} \frac{1}{x^2} \widehat{O}(x) dx}. \quad (28)$$

The VIX can be seen as a measure of CIV with barriers fixed at the lowest and highest strike prices that the CBOE uses for calculating the index (Andersen, Bondarenko, & Gonzalez-Perez,

⁶For simplicity of notation, the exposition is based on a single maturity of 30 days. The CBOE averages (in total variances) the two VIX measures constructed using the near-term and the next-term options. In our empirical application, we follow the same procedure.

2015). Therefore, we define $\widehat{\zeta}_{\text{vix}} := \widehat{\text{VIX}} - \widehat{\text{CIV}}$ as the estimator of the discretization error in the VIX index.


Figure 10: Errors in the VIX



Note: This figure plots the time series of the estimated observation (Panel a) and discretization (Panel b) errors in the VIX index.

To estimate the observation and discretization errors in the VIX, we consider the SPX options obtained from the CBOE from January 3, 2017 until April 1, 2021. We follow the exact same procedure for the construction of the index as outlined in their white paper (CBOE, 2015). To reduce the finite-sample bias in the iCOS procedure due to the discreteness of the observed option strikes, for each tenor, we interpolate option prices using cubic splines applied to implied volatilities. This can be seen as a bias-reduction procedure as motivated in Boswijk et al. (2022) in the context of option-implied CCFs.

Figure 11: Histograms of percentage errors in the VIX



Note: This figure plots the normalized histograms of the percentage estimated errors in the VIX index. The percentage observation errors (Panel a) are defined as $100 \cdot \widehat{\xi}_{\text{vix}}/\widehat{\text{VIX}}$ and the percentage discretization errors (Panel b) are given by $100 \cdot \widehat{\zeta}_{\text{vix}}/\widehat{\text{VIX}}$.

Figure 10 displays the time series plots of the estimated observation and discretization errors in the VIX over the period of more than four years. Figure 11 complements it with histograms of the percentage errors over the same time period. We notice that the observation errors are

centered around zero, while the discretization errors are mainly positive. This is expected since observation errors in option prices do not introduce biases in the VIX, while the discreteness of the strikes leads to a finite-sample bias in the constructed index. In fact, the sample average of the percentage observation errors is nearly zero, and the average of the percentage discretization errors is estimated at around 0.135%. Furthermore, the magnitude of observation errors is rather low, reaching in absolute terms up to 0.04 percentage points. The discretization errors, on the other hand, can result in a substantial overestimation of the index, with the differences up to 0.7 percentage points during high volatility periods.

6 Conclusion

In this paper, we proposed a non-parametric estimation procedure for option prices, RND, and option sensitivities. This method is based on the combination of Fourier-based cosine technique and the option spanning result of Carr and Madan (2001). This combination allows for a flexible and accurate estimation of the density, option prices and option sensitivities without imposing parametric assumptions on the dynamics of underlying asset and on the shape of implied volatility surface. We have also established the asymptotic properties of the proposed estimators and demonstrated the finite sample properties through the Monte Carlo simulations.

The usage of the proposed method is illustrated in empirical applications using options data on the S&P 500 stock market index and Amazon equity options on the Earning Announcement Day. The empirical analysis demonstrates the effectiveness of the iCOS method in accurately estimating option prices and capturing important market features in different market conditions. Additionally, we demonstrated the usefulness of our methodology to dissect and quantify errors in the VIX index, one of most popular measure of market volatility. We found that observation errors in the VIX are centered around zero and have a small magnitude, while discretization errors can lead to positive and substantial biases in the VIX index.

Appendix A Proofs

Proof of Proposition 2:

Given the option observation scheme outlined in Assumptions 1 and 2, the total measurement error in the option-implied cosine coefficients \widehat{D}_m can be decomposed as follows:

$$\begin{aligned}\widehat{D}_m - D_m &= \sum_{i=1}^n w_i \psi_m(K_i) O(K_i) \Delta_n - \int_{\alpha}^{\beta} \psi_m(K) O_0(K) dK \\ &= \underbrace{\sum_{i=1}^n w_i \psi_m(K_i) \varepsilon_i \Delta_n}_{=: \xi_{m,n}^D} + \underbrace{\sum_{i=1}^n w_i \psi_m(K_i) O_0(K_i) \Delta_n - \int_{\alpha}^{\beta} \psi_m(K) O_0(K) dK}_{=: \zeta_{m,n}^D},\end{aligned}$$

where the error term $\xi_{m,n}^D$ represents the observation error due to the noisy observation of option prices, and $\zeta_{m,n}^D$ is the discretization error resulting from the numerical approximation of the integral.

Under Assumptions 1 and 2, the observation error has zero mean $\mathbb{E}[\xi_{m,n}^D] = 0$ and $\xi_{m,n}^D = \mathcal{O}_p(n^{-1/2})$ since

$$\begin{aligned}\mathbb{E}[|\xi_{m,n}^D|^2] &\leq \sum_{i=1}^n w_i^2 \psi_m^2(K_i) \mathbb{E}[\varepsilon_i^2] \Delta_n^2 \\ &= \sum_{i=1}^n w_i^2 \psi_m^2(K_i) \sigma_i^2 \Delta_n^2 \\ &\leq \mathcal{C} n^{-1},\end{aligned}$$

where \mathcal{C} is some constant. Furthermore, this implies convergence to zero in probability, i.e., $\xi_{m,n}^D = o_p(1)$.

To invoke the Lyapunov CLT for non-identical but independent random variables, we first verify the Lyapunov's condition. For that, we note that the ratio

$$\begin{aligned}\frac{\sum_{i=1}^n \mathbb{E}[|w_i \psi_m(K_i) \varepsilon_i|^{2+\omega}]}{(\sum_{i=1}^n w_i^2 \psi_m^2(K_i) \sigma_i^2)^{1+\frac{\omega}{2}}} &= \frac{\sum_{i=1}^n \mathbb{E}[|w_i \psi_m(K_i) \varepsilon_i|^{2+\omega}] \Delta_n}{(\sum_{i=1}^n w_i^2 \psi_m^2(K_i) \sigma_i^2 \Delta_n)^{1+\frac{\omega}{2}} (\Delta_n)^{-\frac{\omega}{2}}} \\ &\leq (\Delta_n)^{\frac{\omega}{2}} \frac{\sum_{i=1}^n |w_i \psi_m(K_i)|^{2+\omega} \mathbb{E}[|\varepsilon_i|^{2+\omega}] \Delta_n}{(\sum_{i=1}^n w_i^2 \psi_m^2(K_i) \sigma_i^2 \Delta_n)^{1+\frac{\omega}{2}}} \\ &\leq \mathcal{C} n^{-\frac{\omega}{2}},\end{aligned}$$

for some $\omega > 0$ and another constant \mathcal{C} . The last inequality follows since the summations in the numerator and denominator converge to definite integrals as $n \rightarrow \infty$, and the smallest and largest strike prices (which define the integration range) are fixed by Assumption 1. Hence, the Lyapunov's condition is satisfied and we can use the Lyapunov CLT, which yields as $n \rightarrow \infty$

$$\frac{\sum_{i=1}^n w_i \psi_m(K_i) \varepsilon_i}{\sqrt{\sum_{i=1}^n w_i^2 \psi_m^2(K_i) \sigma_i^2}} = \frac{\xi_{m,n}^D}{\sqrt{\sum_{i=1}^n w_i^2 \psi_m^2(K_i) \sigma_i^2 \Delta_n}} \xrightarrow{d} \mathcal{N}(0, 1).$$

Finally, the discretization error $\zeta_{m,n}^D$ does not depend on the stochastic observation errors in option prices but introduces bias in the estimation of \widehat{D}_m :

$$\mathbb{E} \left[\widehat{D}_m - D_m \right] = \zeta_{m,n}^D.$$

This bias, however, can be controlled by the choice of the numerical integration scheme and is of order $\mathcal{O} \left(\frac{m^{2+\iota}}{n^\iota} \right)$ for $\iota \geq 1$. The latter follows since $\psi_m^{(\iota)}(x) = \mathcal{O}(m^{2+\iota})$ and numerical integration error bounds. Therefore, for the fixed m and $n \rightarrow \infty$, we obtain the limiting distribution for the option-implied cosine coefficients in the form:

$$\frac{\widehat{D}_m - D_m}{\sigma_D(m)} \xrightarrow{d} \mathcal{N}(0, 1)$$

with $\sigma_D^2(m) := \sum_{i=1}^n w_i^2 \psi_m^2(K_i) \sigma_i^2 \Delta_n^2$.

□

Proof of Proposition 3:

We start by decomposing the difference between the option-implied call price estimator $\widehat{C}(x)$ and its error-free counterpart $\overline{C}_0(x)$ for a fixed strike price $x \in [\alpha, \beta]$:

$$\begin{aligned} \widehat{C}(x) - \overline{C}_0(x) &= \sum_{m=0}^{N-1} \widehat{D}_m H_m(x) - \sum_{m=0}^{\infty} D_m H_m(x) \\ &= \sum_{m=0}^{N-1} \sum_{i=1}^n w_i \psi_m(K_i) O(K_i) \Delta_n H_m(x) \\ &\quad - \sum_{m=0}^{N-1} \int_{\alpha}^{\beta} \psi_m(K) O_0(K) dK H_m(x) - \sum_{m=N}^{\infty} D_m H_m(x) \\ &= \sum_{i=1}^n w_i \underbrace{\sum_{m=1}^{N-1} \psi_m(K_i) H_m(x) O(K_i) \Delta_n}_{=:\psi(x, K_i)} \\ &\quad - \underbrace{\int_{\alpha}^{\beta} \sum_{m=1}^{N-1} \psi_m(K) H_m(x) O_0(K) dK}_{=:\psi(x, K)} - \underbrace{\sum_{m=N}^{\infty} D_m H_m(x)}_{=:\overline{\eta}(x)} \\ &= \sum_{i=1}^n w_i \psi(x, K_i) O(K_i) \Delta_n - \int_{\alpha}^{\beta} \psi(x, K) O_0(K) dK + \overline{\eta}(x), \end{aligned}$$

where the difference between the first two terms involves the discretization and observation errors, and the third term is due to the truncation of the cosine series. Similar to the proof of Proposition 2, we have the following decomposition for the difference between the first two terms:

$$\begin{aligned} &\sum_{i=1}^n w_i \psi(x, K_i) O(K_i) \Delta_n - \int_{\alpha}^{\beta} \psi(x, K) O_0(K) dK \\ &= \underbrace{\sum_{i=1}^n w_i \psi(x, K_i) \varepsilon_i \Delta_n}_{=:\xi(x)} + \underbrace{\sum_{i=1}^n w_i \psi(x, K_i) O_0(K_i) \Delta_n - \int_{\alpha}^{\beta} \psi(x, K) O_0(K) dK}_{=:\zeta(x)}, \end{aligned}$$

where $\xi(x)$ represents the observation and $\zeta(x)$ is the discretization error in the call price estimator $\widehat{C}(x)$.

To get the orders of these errors that depend both on the number of options n and the number of expansion terms N , we first notice that $\psi_m(x) = \mathcal{O}(m^2)$ and $H_m(x) = \mathcal{O}(m^{-2})$. Hence, $\psi(x, K_i) = \sum_{m=1}^{N-1} \psi_m(K)H_m(x) = \mathcal{O}(N)$. Therefore, similar to $\xi_{m,n}^D = \mathcal{O}_p(n^{-1/2})$, the order of the observation error $\xi(x) = \mathcal{O}_p(Nn^{-1/2})$.

As for the discretization error, we can further express it as

$$\begin{aligned} \zeta(x) &= \sum_{i=1}^n w_i \psi(x, K_i) O(K_i) \Delta_n - \int_{\alpha}^{\beta} \psi(x, K) O_0(K) dK \\ &= \sum_{i=1}^n w_i \sum_{m=1}^{N-1} \psi_m(K_i) H_m(x) O_0(K_i) \Delta_n - \int_{\alpha}^{\beta} \sum_{m=1}^{N-1} \psi_m(K) H_m(x) O_0(K) dK \\ &= \sum_{m=1}^{N-1} \left(\sum_{i=1}^n w_i \psi_m(K_i) O_0(K_i) \Delta_n - \int_{\alpha}^{\beta} \psi_m(K) O_0(K) dK \right) H_m(x) \\ &= \sum_{m=1}^{N-1} \zeta_{m,n}^D H_m(x). \end{aligned}$$

That is, the discretization error in the call price estimator arises from the discretization errors in the option-implied cosine coefficients \widehat{D}_m with $m = 1, \dots, N-1$. Therefore, the order of $\zeta(x) = \mathcal{O}\left(\frac{N^{1+\iota}}{n^\iota}\right)$ is also determined by the chosen numerical integration scheme with $\iota \geq 1$.

Finally, the order of the truncation error $\bar{\eta}(x)$ depends on the smoothness of the density function. From Fang and Oosterlee (2008), Lemma 4.2 and Lemma 4.3, we have that $\bar{\eta}(x) = \mathcal{O}(e^{-N\tilde{\nu}})$ with $\tilde{\nu} > 0$ in the case of infinitely differentiable density functions and $\bar{\eta}(x) = \mathcal{O}(N^{1-p})$ under Assumption 3.

Given the orders of these errors, we see that the truncation error decreases with N , while the observation and discretization errors can also increase with N . However, if N increases slower than \sqrt{n} , i.e., if $Nn^{-1/2} \rightarrow 0$ as $N \rightarrow \infty$ and $n \rightarrow \infty$, the observation and discretization errors (asymptotically) converge to zero. In other words, under this condition, the option-implied call price estimator $\widehat{C}(x)$ is asymptotically unbiased.

Finally, similar arguments as in the proof of Proposition 2 allows us to apply the Lyapunov CLT to the observation errors $\xi(x)$ and obtain the following asymptotic distribution for $\widehat{C}(x)$:

$$\frac{\widehat{C}(x) - \bar{C}_0(x)}{\bar{\sigma}_c(x)} \xrightarrow{d} \mathcal{N}(0, 1)$$

as $n \rightarrow \infty$ and $N \rightarrow \infty$ with $Nn^{-1/2} \rightarrow 0$, where

$$\bar{\sigma}_c^2(x) := \sum_{i=1}^n w_i^2 \psi^2(x, K_i) \sigma_i^2 \Delta_n^2 = \sum_{i=1}^n w_i^2 \left(\sum_{m=1}^{N-1} \psi_m(K_i) H_m(x) \right)^2 \sigma_i^2 \Delta_n^2.$$

□

Lemma 1 *The OLS estimator $\widehat{\theta}$ of the linear regression of $C(K_i) - \widehat{C}(K_i) - C(\beta)$ on $Z_c^N(K_i)$ and $Z_p^N(K_i)$ with an intercept is consistent, asymptotically unbiased, and normally distributed as $n \rightarrow \infty$ and $N \rightarrow \infty$ with $Nn^{-1/2} \rightarrow 0$.*

Proof: For ease of notation, let us define a vector of dependent variables as

$$\mathbf{c} := \left(C(K_1) - \widehat{C}(K_1) - C(\beta), \dots, C(K_n) - \widehat{C}(K_n) - C(\beta) \right)',$$

a $n \times 3$ matrix of regressors as

$$\mathbf{Z} := (\mathbf{z}_1, \dots, \mathbf{z}_n)' \text{ with } \mathbf{z}_i := (1, Z_c^N(K_i), Z_p^N(K_i))', \quad i = 1, \dots, n,$$

a parameter vector as $\boldsymbol{\theta} := (\bar{\theta}, \theta_c, \theta_p)'$, and vectors of observation, discretization and series truncation errors:

$$\boldsymbol{\epsilon} := (\varepsilon_1, \dots, \varepsilon_n)',$$

$$\boldsymbol{\xi} := (\xi_1, \dots, \xi_n)', \quad \xi_i := \xi(K_i) = \sum_{j=1}^n w_j \psi_{ij} \varepsilon_j \Delta_n, \quad \psi_{ij} := \psi(K_i, K_j),$$

$$\boldsymbol{\zeta} := (\zeta_1, \dots, \zeta_n)', \quad \zeta_i := \zeta(K_i) = \sum_{j=1}^n w_j \psi(K_i, K_j) O_0(K_j) \Delta_n - \int_{\alpha}^{\beta} \psi(K_i, K) O_0(K) dK,$$

$$\boldsymbol{\eta} := (\eta_1, \dots, \eta_n)', \quad \eta_i := \eta(K_i) = -\bar{\eta}(K_i) + \left(\sum_{m=N}^{\infty} (-1)^m H_m(K_i) \right) \theta_c - \left(\sum_{m=N}^{\infty} H_m(K_i) \right) \theta_p.$$

Note that the series truncation error

$$\eta(x) = -\bar{\eta}(x) + \left(\sum_{m=N}^{\infty} (-1)^m H_m(x) \right) \theta_c - \left(\sum_{m=N}^{\infty} H_m(x) \right) \theta_p = e^{-rT} \sum_{m=N}^{\infty} A_m H_m(x).$$

Let the intercept $\bar{\theta}$ capture the average of finite-sample biases due to the discretization and truncation errors, i.e.,

$$\bar{\theta} = \frac{1}{n} \sum_{i=1}^n \eta_i - \frac{1}{n} \sum_{i=1}^n \zeta_i.$$

We note that $\bar{\theta}$ converges to zero asymptotically, but we include it to reduce the finite-sample bias. Then the OLS regression of the form

$$\mathbf{c} = \mathbf{Z}\boldsymbol{\theta} + \mathbf{e}, \text{ with } \mathbf{e} := \boldsymbol{\epsilon} - \varepsilon_n - \boldsymbol{\xi} - \boldsymbol{\zeta} + \boldsymbol{\eta} - \bar{\theta}, \quad (\text{A.1})$$

gives consistent and asymptotically unbiased estimator:

$$\widehat{\boldsymbol{\theta}} = (\mathbf{Z}'\mathbf{Z})^{-1} \mathbf{Z}'\mathbf{c} = \boldsymbol{\theta} + (\mathbf{Z}'\mathbf{Z})^{-1} \mathbf{Z}'(\boldsymbol{\epsilon} - \varepsilon_n - \boldsymbol{\xi} - \boldsymbol{\zeta} + \boldsymbol{\eta} - \bar{\theta}).$$

Indeed, the consistency follows since each of the error term (asymptotically) converges to zero under the joint asymptotic scheme, i.e. as $n \rightarrow \infty$ and $N \rightarrow \infty$ with $Nn^{-1/2} \rightarrow 0$. In fact, for the observation errors $\boldsymbol{\xi}$ we have that $\frac{1}{n}\mathbf{Z}'\boldsymbol{\xi} = \mathcal{O}_p(Nn^{-1/2})$ due to the same order of each component ξ_i . Furthermore, we have that

$$\begin{aligned} \frac{1}{n}\mathbf{Z}'\boldsymbol{\xi} &= \frac{1}{n} \sum_{i=1}^n \mathbf{z}_i \xi_i = \frac{1}{n} \sum_{i=1}^n \mathbf{z}_i \left(\sum_{j=1}^n w_j \psi_{ij} \varepsilon_j \Delta_n \right) \\ &= \frac{1}{n} \sum_{j=1}^n \sum_{i=1}^n \mathbf{z}_i w_j \psi_{ij} \varepsilon_j \Delta_n = \frac{1}{n} \sum_{j=1}^n w_j \left(\sum_{i=1}^n \mathbf{z}_i \psi_{ij} \Delta_n \right) \varepsilon_j \xrightarrow{a.s.} 0. \end{aligned}$$

Also, since $\zeta_i = \mathcal{O}\left(\frac{N^{1+\iota}}{n^\iota}\right)$ with $\iota \geq 1$, which is controlled by the choice of the numerical integration (see Table 1), we have

$$\frac{1}{n} \mathbf{Z}' \boldsymbol{\zeta} - \frac{1}{n} \sum_{i=1}^n \zeta_i = \frac{1}{n} \sum_{i=1}^n (z_i - 1) \zeta_i \rightarrow 0$$

as $n \rightarrow \infty$ and $N \rightarrow \infty$ with $Nn^{-1/2} \rightarrow 0$. Finally, as the number of expansion terms increases, for the truncation errors we have

$$\frac{1}{n} \mathbf{Z}' \boldsymbol{\eta} - \frac{1}{n} \sum_{i=1}^n \eta_i = \frac{1}{n} \sum_{i=1}^n (z_i - 1) \eta_i \rightarrow 0.$$

Therefore, under Assumption 2 on the observation error scheme, the OLS estimates $\hat{\boldsymbol{\theta}}$ are consistent and asymptotically unbiased.

To derive the limit distribution, we first introduce a matrix $\tilde{\boldsymbol{\Psi}} := \{w_j \psi_{ij} \Delta_n\}_{i,j=1,\dots,n}$ with its j -th column denoted as $\tilde{\boldsymbol{\Psi}}_{\cdot j}$. Rearranging the terms, we have:

$$\mathbf{Z}' \boldsymbol{\xi} = \sum_{i=1}^n z_i \xi_i = \sum_{j=1}^n \sum_{i=1}^n z_i w_j \psi_{ij} \Delta_n \varepsilon_j = \sum_{j=1}^n \mathbf{Z}' \tilde{\boldsymbol{\Psi}}_{\cdot j} \varepsilon_j = \mathbf{Z}' \tilde{\boldsymbol{\Psi}} \boldsymbol{\epsilon}.$$

Note that each term $\mathbf{Z}' \tilde{\boldsymbol{\Psi}}_{\cdot j} = \sum_{i=1}^n z_i w_j \psi_{ij} \Delta_n = \mathcal{O}(1)$ since it converges to a finite interval $w_j \int_{\alpha}^{\beta} z(x) \psi(x, K_j) dx$. Finally, to incorporate the observation error of the last call option price ε_n , we modify the last column of the matrix $\tilde{\boldsymbol{\Psi}}$ by adding the unit vector $\mathbf{1}_n$ of length n , i.e.,

$$\boldsymbol{\Psi} := \left[\tilde{\boldsymbol{\Psi}}_{\cdot 1:n-1}, \tilde{\boldsymbol{\Psi}}_{\cdot n} + \mathbf{1}_n \right].$$

Therefore, the limit distribution of $\hat{\boldsymbol{\theta}}$ is given by the CLT as

$$n^{-1/2} \mathbf{Z}' (\boldsymbol{\epsilon} - \varepsilon_n - \boldsymbol{\xi}) = n^{-1/2} \mathbf{Z}' (\boldsymbol{\epsilon} - \boldsymbol{\Psi} \boldsymbol{\epsilon}) \xrightarrow{d} \mathcal{N}(0, \mathbf{M}),$$

with

$$\mathbf{M} = \text{plim } n^{-1} \mathbf{Z}' (\boldsymbol{\epsilon} - \boldsymbol{\Psi} \boldsymbol{\epsilon}) (\boldsymbol{\epsilon} - \boldsymbol{\Psi} \boldsymbol{\epsilon})' \mathbf{Z} = \lim n^{-1} \mathbf{Z}' (\boldsymbol{\Sigma}_{\boldsymbol{\epsilon}} - 2\boldsymbol{\Psi} \boldsymbol{\Sigma}_{\boldsymbol{\epsilon}} + \boldsymbol{\Psi} \boldsymbol{\Sigma}_{\boldsymbol{\epsilon}} \boldsymbol{\Psi}') \mathbf{Z},$$

where $\boldsymbol{\Sigma}_{\boldsymbol{\epsilon}} := \text{diag}(\sigma_1^2, \dots, \sigma_n^2)$. Denoting $\mathbf{M}_{\mathbf{z}\mathbf{z}} := \lim n^{-1} \mathbf{Z}' \mathbf{Z}$, we get the limit distribution

$$\sqrt{n}(\hat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \xrightarrow{d} \mathcal{N}(0, \mathbf{M}_{\mathbf{z}\mathbf{z}}^{-1} \mathbf{M} \mathbf{M}_{\mathbf{z}\mathbf{z}}^{-1}).$$

Hence, the variance of $\hat{\boldsymbol{\theta}}$ is given by

$$\text{Var}(\hat{\boldsymbol{\theta}}) = (\mathbf{Z}' \mathbf{Z})^{-1} \mathbf{Z}' (\boldsymbol{\Sigma}_{\boldsymbol{\epsilon}} - 2\boldsymbol{\Psi} \boldsymbol{\Sigma}_{\boldsymbol{\epsilon}} + \boldsymbol{\Psi} \boldsymbol{\Sigma}_{\boldsymbol{\epsilon}} \boldsymbol{\Psi}') \mathbf{Z} (\mathbf{Z}' \mathbf{Z})^{-1}.$$

□

Lemma 2 *Feasible estimates for the variance matrix of the observation errors in option prices can be obtained from the residuals of the OLS regression $\hat{\boldsymbol{\epsilon}} := \mathbf{c} - \hat{\mathbf{c}}$ from Lemmas 1 as $\hat{\boldsymbol{\Sigma}}_{\boldsymbol{\epsilon}} = \frac{n}{\nu} \text{diag}(\hat{\varepsilon}_1^2, \dots, \hat{\varepsilon}_n^2)$ with $\nu = \text{tr}(\mathbf{Q}) - 2\text{tr}(\mathbf{Q}\boldsymbol{\Psi}) + \text{tr}(\mathbf{Q}\boldsymbol{\Psi}\boldsymbol{\Psi}')$ with $\mathbf{Q} = \mathbf{I} - \mathbf{Z}(\mathbf{Z}' \mathbf{Z})^{-1} \mathbf{Z}'$.*

Proof: In the standard linear regression setting, the feasible version of the covariance matrix can be obtained using the OLS residuals defined as $\widehat{\boldsymbol{e}} := \boldsymbol{c} - \boldsymbol{Z}\widehat{\boldsymbol{\theta}}$. However, some finite sample correction might be needed to obtain unbiased estimates of the error variances. Using the projection matrix argument with $\boldsymbol{Q} = \boldsymbol{I} - \boldsymbol{Z}(\boldsymbol{Z}'\boldsymbol{Z})^{-1}\boldsymbol{Z}'$, we can express the OLS residuals as:

$$\widehat{\boldsymbol{e}} = \boldsymbol{c} - \widehat{\boldsymbol{c}} = \boldsymbol{Q}\boldsymbol{e} = \boldsymbol{Q}(\boldsymbol{I} - \boldsymbol{\Psi})\boldsymbol{\epsilon},$$

where we ignore asymptotically vanishing terms. If option errors were homoskedastic with an error variance term σ_ϵ^2 , the sum of squared residuals would estimate

$$\mathbb{E}[\widehat{\boldsymbol{e}}'\widehat{\boldsymbol{e}}] = \mathbb{E}[\text{tr}(\boldsymbol{\epsilon}\boldsymbol{\epsilon}'(\boldsymbol{I} - \boldsymbol{\Psi})'\boldsymbol{Q}(\boldsymbol{I} - \boldsymbol{\Psi})))] = \nu\sigma_\epsilon^2$$

with

$$\nu := \text{tr}((\boldsymbol{I} - \boldsymbol{\Psi})'\boldsymbol{Q}(\boldsymbol{I} - \boldsymbol{\Psi})) = \text{tr}(\boldsymbol{Q}) - 2\text{tr}(\boldsymbol{Q}\boldsymbol{\Psi}) + \text{tr}(\boldsymbol{Q}\boldsymbol{\Psi}\boldsymbol{\Psi}'),$$

where $\text{tr}(\boldsymbol{Q}) = n - 3$, and the traces of the last two terms can be easily computed in practice.

Hence, using the degrees of freedom correction term n/ν yields unbiased estimates for the variance of observation errors under the homoskedastic error assumption. Although we do not impose homoskedasticity, in practice, we adjust the squared residuals for degrees of freedom ν to obtain estimates of the error variances in finite samples, that is, we use $\widehat{\boldsymbol{\Sigma}}_\epsilon := \frac{n}{\nu}\text{diag}(\widehat{e}_1^2, \dots, \widehat{e}_n^2)$. This is similar to the one conventionally used to obtain heteroskedasticity-consistent standard errors of parameter estimates in the context of linear regression (MacKinnon & White, 1985). Importantly, the matrix $\boldsymbol{\Psi}$ and, hence, its trace are functions of the number of expansion terms N . Therefore, we correct for degrees of freedom with every cosine expansion term. \square

Proof of Proposition 4:

Using results and notations defined in Lemma 1, we have

$$\begin{aligned} \widehat{C}(x) - C_0(x) &= \left(\widehat{C}(x) + C(\beta) + \boldsymbol{z}(x)'\widehat{\boldsymbol{\theta}} \right) - \left(\overline{C}_0(x) + C_0(\beta) + Z_c(x)\theta_c + Z_p(x)\theta_p \right) \\ &= \widehat{C}(x) - \overline{C}_0(x) + \boldsymbol{z}(x)'\widehat{\boldsymbol{\theta}} + \bar{\theta} - \bar{\theta} + \varepsilon_n - Z_c(x)\theta_c - Z_p(x)\theta_p \\ &= \xi(x) + \varepsilon_n + \zeta(x) - \eta(x) + \bar{\theta} + \boldsymbol{z}(x)'(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}). \end{aligned}$$

The last equality follows since

$$\begin{aligned} -\eta(x) &= \bar{\eta}(x) - \left(\sum_{m=N}^{\infty} (-1)^m H_m(x) \right) \theta_c + \left(\sum_{m=N}^{\infty} H_m(x) \right) \theta_p \\ &= - \sum_{m=N}^{\infty} D_m H_m(x) - \sum_{m=N}^{\infty} b_m H_m(x) = -e^{-rT} \sum_{m=N}^{\infty} A_m H_m(x). \end{aligned}$$

Given Proposition 3 and Lemma 1, we obtain a consistent and asymptotically unbiased call price estimator $\widehat{C}(x)$. A finite-sample bias in this estimator arises from the deviation of the discretization and truncation errors for an option with a strike price x from the average bias across all contracts $\bar{\theta} = \frac{1}{n} \sum_{i=1}^n \eta_i - \frac{1}{n} \sum_{i=1}^n \zeta_i$. Therefore, the magnitude of the bias term is sufficiently smaller than that in the $\widehat{C}(x)$ estimator. The order, however, is the same since $\bar{\eta}(x)$

dominates the other term in $\eta(x)$ as the density is smoother than the payoff function of the call option. Therefore, we have that $\eta(x) - \frac{1}{n} \sum_{i=1}^n \eta_i = \mathcal{O}(N^{1-p})$ and $\zeta(x) - \frac{1}{n} \sum_{i=1}^n \zeta_i = \mathcal{O}\left(\frac{N^{1+\iota}}{n^\iota}\right)$ with $\iota \geq 1$. That is, the call price estimator is in fact asymptotically unbiased.

The limiting distribution of the call price estimator is determined by the first and the last terms since $\sqrt{n}\xi(x) = \mathcal{O}_p(1)$ and so is the OLS estimator. Hence, we need to combine the limiting distributions resulting from $\widehat{C}(x)$ and $\widehat{\theta}$. For that, let us first denote

$$\boldsymbol{\psi}_w(x) := (w_1\psi(x, K_1)\Delta_n, \dots, w_n\psi(x, K_n)\Delta_n + 1).$$

That is, if $x = K_i$ for some $i = 1, \dots, n$, then $\boldsymbol{\psi}_w(K_i) = \boldsymbol{\Psi}_i$ is the i -th row of the matrix $\boldsymbol{\Psi}$. Then, we can express the first and the last terms as

$$\begin{aligned} \sqrt{n} \left(\xi(x) + \varepsilon_n + \mathbf{z}(x)'(\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}) \right) &= \sqrt{n} \left(\boldsymbol{\psi}_w(x)\boldsymbol{\epsilon} + \mathbf{z}(x)'(\mathbf{Z}'\mathbf{Z})^{-1} \mathbf{Z}'(\mathbf{I} - \boldsymbol{\Psi})\boldsymbol{\epsilon} \right) \\ &= \sqrt{n} \left(\boldsymbol{\psi}_w(x) + \mathbf{z}(x)'(\mathbf{Z}'\mathbf{Z})^{-1} \mathbf{Z}'(\mathbf{I} - \boldsymbol{\Psi}) \right) \boldsymbol{\epsilon}, \end{aligned}$$

which converges by the CLT to the normal distribution with mean zero and the asymptotic variance matrix

$$\begin{aligned} \mathbf{M}_c &= \text{plim } n \left[\boldsymbol{\psi}_w(x) + \mathbf{z}(x)'(\mathbf{Z}'\mathbf{Z})^{-1} \mathbf{Z}'(\mathbf{I} - \boldsymbol{\Psi}) \right] \boldsymbol{\epsilon} \boldsymbol{\epsilon}' \left[\boldsymbol{\psi}_w(x)' + (\mathbf{I} - \boldsymbol{\Psi}')\mathbf{Z}(\mathbf{Z}'\mathbf{Z})^{-1} \mathbf{z}(x) \right] \\ &= \lim n \boldsymbol{\psi}_w(x) \boldsymbol{\Sigma}_\boldsymbol{\epsilon} \boldsymbol{\psi}_w(x)' + \mathbf{z}(x)' \mathbf{M}_{zz}^{-1} \mathbf{M} \mathbf{M}_{zz}^{-1} \mathbf{z}(x) \\ &\quad + 2 \lim n \mathbf{z}(x)' (\mathbf{Z}'\mathbf{Z})^{-1} \mathbf{Z}'(\mathbf{I} - \boldsymbol{\Psi}) \boldsymbol{\Sigma}_\boldsymbol{\epsilon} \boldsymbol{\psi}_w(x)'. \end{aligned}$$

Therefore, the variance matrix of the call price estimator is given by

$$\begin{aligned} \sigma_c^2(x) &:= \boldsymbol{\psi}_w(x) \boldsymbol{\Sigma}_\boldsymbol{\epsilon} \boldsymbol{\psi}_w(x)' + \mathbf{z}(x)' \text{Var}(\widehat{\boldsymbol{\theta}}) \mathbf{z}(x) + 2\mathbf{z}(x)' (\mathbf{Z}'\mathbf{Z})^{-1} \mathbf{Z}'(\mathbf{I} - \boldsymbol{\Psi}) \boldsymbol{\Sigma}_\boldsymbol{\epsilon} \boldsymbol{\psi}_w(x)' \\ &= \overline{\sigma}_c^2(x) + \mathbf{z}(x)' \text{Var}(\widehat{\boldsymbol{\theta}}) \mathbf{z}(x) + 2\mathbf{z}(x)' (\mathbf{Z}'\mathbf{Z})^{-1} \mathbf{Z}'(\mathbf{I} - \boldsymbol{\Psi}) \boldsymbol{\Sigma}_\boldsymbol{\epsilon} \boldsymbol{\psi}_w(x)'. \end{aligned} \quad (\text{A.2})$$

The feasible version of this variance can be obtain by using $\widehat{\boldsymbol{\Sigma}}_\boldsymbol{\epsilon}$ from Lemma 2. □

Proof of Proposition 5:

The proof closely follows the proofs of Propositions 3 and 4. Therefore, for ease of notation, we first define analogous objects:

$$\begin{aligned} H_m^f(y) &:= \cos(u_m y - u_m \log \alpha), \quad \psi^f(y, K) := \sum_{m=1}^{N-1} \psi_m(K) H_m^f(y), \\ \boldsymbol{\psi}_w^f(y) &:= \left(w_1 \psi^f(y, K_1) \Delta_n, \dots, w_n \psi^f(y, K_n) \Delta_n \right), \\ \eta^f(y) &:= e^{-rT} \sum_{m=N}^{\infty} A_m H_m^f(y), \quad \xi^f(y) := \sum_{i=1}^n w_i \psi^f(x, K_i) \varepsilon_i \Delta_n = \boldsymbol{\psi}_w^f(y) \boldsymbol{\epsilon}, \\ \zeta^f(y) &:= \sum_{i=1}^n w_i \psi^f(x, K_i) O_0(K_i) \Delta_n - \int_{\alpha}^{\beta} \psi^f(x, K) O_0(K) dK = \sum_{m=1}^{N-1} \zeta_{m,n}^D H_m^f(y), \\ Z_c^f(y) &:= \sum_{m=0}^{N-1} (-1)^m H_m^f(y), \quad Z_p^f(y) := \sum_{m=0}^{N-1} H_m^f(y), \quad \mathbf{z}^f(y) := \left(0, Z_c^f(y), Z_p^f(y) \right)'. \end{aligned}$$

Repeating the steps from the proof of Proposition 3, we can decompose the (scaled) difference between the non-parametric RND estimator, defined in equation (23) and the true RND given by the relation (10) for a given point y :

$$\begin{aligned}
\frac{\widehat{f}(y) - f(y)}{\nu_f} &= \sum_{m=0}^{N-1} \left(\widehat{D}_m + (-1)^m \widehat{\theta}_c - \widehat{\theta}_p \right) \cos(u_m y - u_m \log \alpha) \\
&\quad - e^{-rT} \sum_{m=0}^{\infty} A_m \cos(u_m y - u_m \log \alpha) \\
&= \sum_{m=0}^{N-1} \widehat{D}_m H_m^f(y) - \sum_{m=0}^{N-1} D_m H_m^f(y) + \eta^f(y) + Z_c^f(y) (\widehat{\theta}_c - \theta_c) + Z_p^f(y) (\widehat{\theta}_p - \theta_p) \\
&= \xi^f(y) + \zeta^f(y) - \eta^f(y) + \mathbf{z}^f(y)' (\widehat{\boldsymbol{\theta}} - \boldsymbol{\theta}).
\end{aligned}$$

We notice, however, differently from Proposition 3, the order of the cosine coefficient $H_m^f(y) = \mathcal{O}(1)$ and, hence, $\psi^f(y, K) = \mathcal{O}(N^3)$. This further results in different orders for the observation and discretization errors, namely $\xi^f(y) = \mathcal{O}_p(N^3 n^{-1/2})$ and $\zeta^f(y) = \mathcal{O}(N^{3+\iota} n^{-\iota})$, respectively. Therefore, to guarantee the asymptotic unbiasedness of the RND estimator, we require N to grow slower than $n^{-1/6}$, i.e. $N n^{-1/6} \rightarrow 0$.

Finally, using similar arguments as in the proof of Proposition 4, we can establish the limiting distribution of the non-parametric RND estimator:

$$\frac{\widehat{f}(y) - f(y)}{\nu_f \sigma_f(y)} \xrightarrow{d} \mathcal{N}(0, 1),$$

as $n \rightarrow \infty$ and $N \rightarrow \infty$ with $N n^{-1/6} \rightarrow 0$ and the covariance matrix

$$\sigma_f^2(y) := \boldsymbol{\psi}_w^f(y)' \boldsymbol{\Sigma}_\epsilon \boldsymbol{\psi}_w^f(y) + \mathbf{z}^f(y)' \text{Var}(\widehat{\boldsymbol{\theta}}) \mathbf{z}^f(y) + 2 \mathbf{z}^f(y)' (\mathbf{Z}' \mathbf{Z})^{-1} \mathbf{Z}' (\mathbf{I} - \boldsymbol{\Psi}) \boldsymbol{\Sigma}_\epsilon \boldsymbol{\psi}_w^f(y)'. \quad (\text{A.3})$$

□

Proof of Proposition 6:

Similarly, we can decompose the scaled difference between the non-parametric delta estimator and the true value given by the spanning result (18) under Assumption 4:

$$\begin{aligned}
S_0(\widehat{\delta}(x) - \delta(x)) &= - \sum_{m=1}^{N-1} u_m \widehat{B}_m H_m(x) + (C(\beta) - \beta \widehat{\theta}_c) + \sum_{m=1}^{\infty} u_m B_m H_m(x) - (C_0(\beta) - \beta \theta_c) \\
&= - \sum_{m=1}^{N-1} u_m \sum_{i=1}^n w_i \widetilde{\psi}_m(K_i) O(K_i) \Delta_n H_m(x) + \sum_{m=1}^{N-1} u_m \int_{\alpha}^{\beta} \widetilde{\psi}_m(K) O_0(K) dK H_m(x) \\
&\quad + \sum_{m=N}^{\infty} u_m B_m H_m(x) + C(\beta) - C_0(\beta) - \beta (\widehat{\theta}_c - \theta_c) \\
&\quad + \sum_{m=1}^{N-1} \frac{u_m^2}{\beta} (-1)^m H_m(x) (C(\beta) - C_0(\beta)) - \sum_{m=1}^{N-1} \frac{u_m^2}{\alpha} H_m(x) (P(\alpha) - P_0(\alpha)),
\end{aligned}$$

which can further be written as

$$\begin{aligned}
S_0 \left(\widehat{\delta}(x) - \delta(x) \right) &= - \sum_{i=1}^n w_i \underbrace{\sum_{m=1}^{N-1} u_m \widetilde{\psi}_m(K_i) H_m(x) O(K_i) \Delta_n}_{=:\widetilde{\psi}(x, K_i)} + \int_{\alpha}^{\beta} \underbrace{\sum_{m=1}^{N-1} u_m \widetilde{\psi}_m(K) H_m(x) O_0(K) dK}_{=:\widetilde{\psi}(x, K)} \\
&\quad + \underbrace{\sum_{m=N}^{\infty} u_m B_m H_m(x)}_{=:\widetilde{\eta}(x)} + \underbrace{\left(1 + \sum_{m=1}^{N-1} \frac{u_m^2}{\beta} (-1)^m H_m(x) \right)}_{=:\widetilde{Z}_c^N(x)} \varepsilon_n - \underbrace{\sum_{m=1}^{N-1} \frac{u_m^2}{\alpha} H_m(x)}_{=:\widetilde{Z}_p^N(x)} \varepsilon_1 \\
&\quad - \beta(\widehat{\theta}_c - \theta_c) \\
&= - \sum_{i=1}^n w_i \widetilde{\psi}(x, K_i) O(K_i) \Delta_n + \int_{\alpha}^{\beta} \widetilde{\psi}(x, K) O_0(K) dK + \widetilde{\eta}(x) \\
&\quad + \widetilde{Z}_c^N(x) \varepsilon_n + \widetilde{Z}_p^N(x) \varepsilon_1 - \beta(\widehat{\theta}_c - \theta_c).
\end{aligned}$$

The difference between the first two terms can be similarly decomposed into the observation and discretization errors. Therefore, we have

$$S_0 \left(\widehat{\delta}(x) - \delta(x) \right) = \widetilde{\xi}(x) + \widetilde{\zeta}(x) + \widetilde{\eta}(x) + \widetilde{Z}_c^N(x) \varepsilon_n + \widetilde{Z}_p^N(x) \varepsilon_1 - \beta(\widehat{\theta}_c - \theta_c).$$

Similar to the proofs of Propositions 4 and 5, we can establish the limiting distribution of the non-parametric delta estimator:

$$\frac{\widehat{\delta}(x) - \delta(x)}{\frac{1}{S_0} \sigma_{\delta}(y)} \xrightarrow{d} \mathcal{N}(0, 1)$$

as $n \rightarrow \infty$ and $N \rightarrow \infty$ with $Nn^{-1/2} \rightarrow 0$ since the order of $\widetilde{\psi}_m(x)$ is the same as of $\psi_m(x)$ and $\widetilde{\eta}(x)$ decreases with an increase of N . Here, the variance matrix estimator of $\widehat{\delta}(x)$ is given by

$$\sigma_{\delta}^2(x) := \widetilde{\psi}_w(x) \Sigma_{\varepsilon} \widetilde{\psi}_w(x)' + \beta^2 \text{Var}(\widehat{\theta}_c) + 2\mathbf{z}^{\delta}(x)' (\mathbf{Z}' \mathbf{Z})^{-1} \mathbf{Z}' (\mathbf{I} - \Psi) \Sigma_{\varepsilon} \widetilde{\psi}_w(x)'$$

with $\mathbf{z}^{\delta}(x)' = (0, -\beta, 0)$ and

$$\widetilde{\psi}_w(x) := \left(w_1 \widetilde{\psi}(x, K_1) \Delta_n + \widetilde{Z}_p^N(x), w_2 \widetilde{\psi}(x, K_2) \Delta_n, \dots, w_{n-1} \widetilde{\psi}(x, K_{n-1}) \Delta_n, w_n \widetilde{\psi}(x, K_n) \Delta_n + \widetilde{Z}_c^N(x) \right).$$

□

Appendix B Additional results

B.1 Put options

In this section we provide the results for put options. In particular, for the plain vanilla put option with a strike price $x > \alpha$, the value on the interval $(0, \alpha)$ is given by

$$\begin{aligned}
P_0^{(0, \alpha)}(x) &= e^{-rT} \mathbb{E}^{\mathbb{Q}}[\max(x - S_T, 0) \mathbf{1}_{\{S_T < \alpha\}}] = e^{-rT} \int_0^{\alpha} \max(x - S_T, 0) f_S(S_T) dS_T \\
&= e^{-rT} (x - \alpha) \int_0^{\alpha} f_S(S_T) dS_T + e^{-rT} \int_0^{\infty} \max(\alpha - S_T, 0) f_S(S_T) dS_T \\
&= (x - \alpha) P_K'(\alpha) + P_0(\alpha),
\end{aligned}$$

where $P_0(\alpha)$ is the put price with the strike α and $P'_K(\alpha)$ is its derivative with respect to the strike price evaluated at α . $P_0^{(0,\alpha)}(x)$ represents the price of a so-called *gap put option* with a strike price x and a trigger price α .

Therefore, for the put option price with strike price x such that $\alpha \leq x \leq \beta$, we have:

$$P_0(x) = P_0^{[\alpha,\beta]}(x) + (x - \alpha)P'_K(\alpha) + P_0(\alpha). \quad (\text{B.1})$$

The value of the put contract with a payoff restricted to the interval $[\alpha, \beta]$ is given by the COS formula, similar to the call counterpart, as

$$P_0^{[\alpha,\beta]}(x) = \sum_{m=0}^{\infty} D_m H_m^p(x) + \left(\sum_{m=0}^{\infty} (-1)^m H_m^p(x) \right) \theta_c - \left(\sum_{m=0}^{\infty} H_m^p(x) \right) \theta_p, \quad (\text{B.2})$$

where $H_m^p(x)$ is the cosine series coefficient for the put payoff function with the strike price x and $\theta_c = C'_K(\beta)$ and $\theta_p = P'_K(\alpha)$ are the first-order derivatives of the call and put options evaluated at the boundaries of the interval. Therefore, the put option can be represented as the following portfolio

$$\begin{aligned} P_0(x) &= P_0^{[\alpha,\beta]}(x) + (x - \alpha)\theta_p + P_0(\alpha) \\ &= \underbrace{\sum_{m=0}^{\infty} D_m H_m^p(x)}_{=: \bar{P}_0(x)} + P_0(\alpha) + \underbrace{\left(\sum_{m=0}^{\infty} (-1)^m H_m^p(x) \right)}_{=: Z_c^p(x)} \theta_c + \underbrace{\left(x - \alpha - \sum_{m=0}^{\infty} H_m^p(x) \right)}_{=: Z_p^p(x)} \theta_p \\ &= \bar{P}_0(x) + P_0(\alpha) + Z_c^p(x)\theta_c + Z_p^p(x)\theta_p. \end{aligned} \quad (\text{B.3})$$

This spanning result for the put option is analogous to the call price representation given by equation (15). This result can be used to further construct the computationally feasible non-parametric put price estimator similar to the call options.

B.2 Cosine coefficients

The cosine coefficients H_m that correspond to the cosine transformation of the call payoff with the strike price x and transformed variable $y = \log \frac{S_T}{x}$ are given by

$$\begin{aligned} H_m(x) &= \frac{2}{b-a} \int_a^b x \max\{e^y - 1, 0\} \cos(u_m y - u_m a) dy \\ &= \frac{2}{b-a} \int_0^b x(e^y - 1) \cos(u_m y - u_m a) dy. \end{aligned}$$

Using the result of Fang and Oosterlee (2008, eq. 24) and the interval bounds $a = \log \frac{\alpha}{x}$ and $b = \log \frac{\beta}{x}$ from Section 2, we have for $m > 0$

$$\begin{aligned} H_m(x) &= \frac{2x}{b-a} \left[\frac{1}{1+u_m^2} \left(\cos(m\pi)e^b - \cos(-u_m a) - u_m \sin(-u_m a) \right) + \frac{1}{u_m} \sin(-u_m a) \right] \\ &= \frac{2x}{\log \frac{\beta}{\alpha}} \left[\frac{1}{1+u_m^2} \left(\cos(m\pi) \frac{\beta}{x} - \cos \left(u_m \log \frac{\alpha}{x} \right) + u_m \sin \left(u_m \log \frac{\alpha}{x} \right) \right) - \frac{1}{u_m} \sin \left(u_m \log \frac{\alpha}{x} \right) \right] \\ &= \frac{2x}{u_m(1+u_m^2) \log \frac{\beta}{\alpha}} \left((-1)^m \frac{u_m \beta}{x} - u_m \cos \left(u_m \log \frac{\alpha}{x} \right) - \sin \left(u_m \log \frac{\alpha}{x} \right) \right). \end{aligned}$$

A similar result can be obtained for the put options. These cosine coefficients are deterministic functions of the strike price and the fixed interval bounds α and β .

B.3 Optimal number of expansion terms

In this section we provide a rule-of-thumb algorithm for the optimal number of expansion terms N . The algorithm operationalizes the result that if $A_N^2 - \text{Var}(\widehat{A}_N) > 0$, then $\text{AMISE}_{N+1} < \text{AMISE}_N$ following the discussion in Section 3.5.

The estimated cosine coefficients \widehat{A}_m are given by

$$\widehat{A}_m = e^{rT} \left(\widehat{D}_m + (-1)^m \widehat{\theta}_c - \widehat{\theta}_p \right), \quad (\text{B.4})$$

where \widehat{D}_m is given by equation (19), and $\widehat{\theta}_c$ and $\widehat{\theta}_p$ are the OLS estimates from regression (A.1). Following the proof steps and notations in Lemma 1 and Proposition 4, the variance of the cosine coefficients is given by

$$\begin{aligned} \sigma_A^2(m) &:= \text{Var} \left(\widehat{A}_m \right) = \boldsymbol{\psi}_m \boldsymbol{\Sigma}_\epsilon \boldsymbol{\psi}_m' + \mathbf{z}'_A \text{Var}(\widehat{\boldsymbol{\theta}}) \mathbf{z}_A + 2\mathbf{z}'_A (\mathbf{Z}'\mathbf{Z})^{-1} \mathbf{Z}'(\mathbf{I} - \boldsymbol{\Psi}) \boldsymbol{\Sigma}_\epsilon \boldsymbol{\psi}_m' \\ &= \sigma_D^2(m) + \mathbf{z}'_A \text{Var}(\widehat{\boldsymbol{\theta}}) \mathbf{z}_A + 2\mathbf{z}'_A (\mathbf{Z}'\mathbf{Z})^{-1} \mathbf{Z}'(\mathbf{I} - \boldsymbol{\Psi}) \boldsymbol{\Sigma}_\epsilon \boldsymbol{\psi}_m', \end{aligned} \quad (\text{B.5})$$

where $\mathbf{z}_A := (0, (-1)^m, -1)'$, $\boldsymbol{\psi}_m := (w_1 \psi_m(K_1) \Delta_n, \dots, w_n \psi(K_n) \Delta_n)$, and $\sigma_D^2(m)$ is the variance of the estimator \widehat{D}_m given in Proposition 2. The feasible variances are obtained by using the feasible covariance matrix $\widehat{\boldsymbol{\Sigma}}_\epsilon$ as discussed in the proof of Lemma 1.

For a given number of expansion terms N , we collect the estimated cosine coefficients and their feasible standard deviations into the vectors:

$$\widehat{\mathbf{A}} := \left(\widehat{A}_1, \dots, \widehat{A}_N \right)' \quad \text{and} \quad \widehat{\boldsymbol{\sigma}}_A := (\widehat{\sigma}_A(1), \dots, \widehat{\sigma}_A(N))'.$$

In the algorithm below, we denote the calculation of these vectors $\widehat{\mathbf{A}}$ and $\widehat{\boldsymbol{\sigma}}_A$ by the function $i\text{COSA}(\mathbf{K}, \mathbf{C}, N)$, which takes as inputs the data vectors \mathbf{K} and \mathbf{C} of strike prices and call prices, and the number of expansion terms N .

Since the true value of the cosine coefficient A_N is unknown, we use the estimated value \widehat{A}_N instead, averaged over the adjacent coefficients. In particular, we notice that the true cosine coefficients A_m exponentially decay to zero as m increases with a rate depending on the smoothness of the density. Therefore, the coefficients are averaged out after taking the logarithm of their absolute values. This average is then compared to the logarithm of the standard deviation of the coefficients. This rule-of-thumb is sketched in Algorithm 1.

Algorithm 1 Optimal number of expansion terms

Input: \mathbf{K}, \mathbf{C}

Output: N^*

$N = 5, N_{max} = 50$

$\bar{a} = 1, s_a = 0$

while $\bar{a} > s_a$ & $N < N_{max}$ **do**

$N = N + 1$

$\hat{\mathbf{A}}, \hat{\boldsymbol{\sigma}}_{\mathbf{A}} = iCOSA(\mathbf{K}, \mathbf{C}, N)$

$\bar{a} = \frac{1}{3}(\log |\hat{\mathbf{A}}_{[N-2]}| + \log |\hat{\mathbf{A}}_{[N-1]}| + \log |\hat{\mathbf{A}}_{[N]}|)$

$s_a = \log \hat{\boldsymbol{\sigma}}_{\mathbf{A}[N-1]}$

end while

$N^* = N - 1$

References

- Aït-Sahalia, Y., & Duarte, J. (2003). Nonparametric option pricing under shape restrictions. *Journal of Econometrics*, 116(1-2), 9–47.
- Aït-Sahalia, Y., & Lo, A. W. (1998). Nonparametric estimation of state-price densities implicit in financial asset prices. *The Journal of Finance*, 53(2), 499–547.
- Alexander, C., & Nogueira, L. M. (2007). Model-free hedge ratios and scale-invariant models. *Journal of Banking & Finance*, 31(6), 1839–1861.
- Alexiou, L., Goyal, A., Kostakis, A., & Rompolis, L. (2021). Pricing event risk: Evidence from concave implied volatility curves. *Swiss Finance Institute Research Paper*(21-48).
- Andersen, T. G., Bondarenko, O., & Gonzalez-Perez, M. T. (2015). Exploring return dynamics via corridor implied volatility. *The Review of Financial Studies*, 28(10), 2902–2945.
- Andersen, T. G., Fusari, N., & Todorov, V. (2015). Parametric inference and dynamic state recovery from option panels. *Econometrica*, 83(3), 1081–1145.
- Andersen, T. G., Fusari, N., Todorov, V., & Varneskov, R. T. (2021). Spatial dependence in option observation errors. *Econometric Theory*, 37(2), 205–247.
- Bakshi, G., Kapadia, N., & Madan, D. (2003). Stock return characteristics, skew laws, and the differential pricing of individual equity options. *The Review of Financial Studies*, 16(1), 101–143.
- Bates, D. S. (2005). Hedging the smirk. *Finance Research Letters*, 2(4), 195–200.
- Black, F., & Scholes, M. (1973). The pricing of options and corporate liabilities. *Journal of Political Economy*, 81(3), 637–654.
- Bliss, R. R., & Panigirtzoglou, N. (2002). Testing the stability of implied probability density functions. *Journal of Banking & Finance*, 26(2-3), 381–422.
- Bossu, S., Carr, P., & Papanicolaou, A. (2022). Static replication of European standard dispersion options. *Quantitative Finance*, 22(5), 799–811.
- Boswijk, H. P., Laeven, R. J., & Vladimirov, E. (2022). Estimating option pricing models using a characteristic function-based linear state space representation. *Tinbergen Institute Discussion Paper 2022-075/III*.
- Breeden, D. T., & Litzenberger, R. H. (1978). Prices of state-contingent claims implicit in option prices. *Journal of Business*, 621–651.
- Carr, P., & Madan, D. (1999). Option valuation using the fast Fourier transform. *Journal of Computational Finance*, 2(4), 61–73.
- Carr, P., & Madan, D. (2001). Optimal positioning in derivative securities. *Quantitative Finance*, 1(1), 19–37.
- CBOE. (2015). *VIX white paper*. <https://cdn.cboe.com/resources/vix/vixwhite.pdf>.
- Cui, Z., & Xu, Y. (2022). A new representation of the risk-neutral distribution and its applications. *Quantitative Finance*, 22(5), 817–834.
- Cui, Z., & Yu, Z. (2021). A model-free Fourier cosine method for estimating the risk-neutral density. *The Journal of Derivatives*, 29(2), 149–171.

- Dalderop, J. (2020). Nonparametric filtering of conditional state-price densities. *Journal of Econometrics*, *214*(2), 295–325.
- Duffie, D., Pan, J., & Singleton, K. (2000). Transform analysis and asset pricing for affine jump-diffusions. *Econometrica*, *68*(6), 1343–1376.
- Fang, F., & Oosterlee, C. W. (2008). A novel pricing method for European options based on Fourier-cosine series expansions. *SIAM Journal on Scientific Computing*, *31*(2), 826–848.
- Figlewski, S. (2018). Risk-neutral densities: A review. *Annual Review of Financial Economics*, *10*(1), 329–359.
- Gatheral, J., & Jacquier, A. (2014). Arbitrage-free SVI volatility surfaces. *Quantitative Finance*, *14*(1), 59–71.
- Gemmell, G., & Saffekos, A. (2000). How useful are implied distributions? *The Journal of Derivatives*, *7*(3), 83–91.
- Grith, M., Härdle, W. K., & Schienle, M. (2012). *Nonparametric estimation of risk-neutral densities*. Springer.
- Jarrow, R., & Rudd, A. (1982). Approximate option valuation for arbitrary stochastic processes. *Journal of Financial Economics*, *10*(3), 347–369.
- Jiang, G. J., & Tian, Y. S. (2005). The model-free implied volatility and its information content. *The Review of Financial Studies*, *18*(4), 1305–1342.
- Jiang, G. J., & Tian, Y. S. (2007). Extracting model-free volatility from option prices: An examination of the VIX index. *The Journal of Derivatives*, *14*(3), 35–60.
- Kronmal, R., & Tarter, M. (1968). The estimation of probability densities and cumulatives by Fourier series methods. *Journal of the American Statistical Association*, *63*(323), 925–952.
- Leitao, Á., Oosterlee, C. W., Ortiz-Gracia, L., & Bohte, S. M. (2018). On the data-driven COS method. *Applied Mathematics and Computation*, *317*, 68–84.
- Lewis, A. L. (2001). A simple option formula for general jump-diffusion and other exponential Lévy processes. Available at SSRN 282110.
- Lu, J., & Qu, Z. (2021). Sieve estimation of option-implied state price density. *Journal of Econometrics*, *224*(1), 88–112.
- MacKinnon, J. G., & White, H. (1985). Some heteroskedasticity-consistent covariance matrix estimators with improved finite sample properties. *Journal of econometrics*, *29*(3), 305–325.
- Malz, A. M. (2014). A simple and reliable way to compute option-based risk-neutral distributions. *FRB of New York Staff Report*(677).
- Melick, W. R., & Thomas, C. P. (1997). Recovering an asset’s implied PDF from option prices: an application to crude oil during the Gulf crisis. *Journal of Financial and Quantitative Analysis*, *32*(1), 91–115.
- Qin, L., & Todorov, V. (2019). Nonparametric implied Lévy densities. *The Annals of Statistics*, *47*(2), 1025–1060.
- Rompolis, L. S., & Tzavalis, E. (2007). Retrieving risk neutral densities based on risk neutral moments through a Gram–Charlier series expansion. *Mathematical and Computer*

- Modelling*, 46(1-2), 225–234.
- Shimko, D. (1993). Bounds of probability. *Risk*, 6(4), 33–37.
- Todorov, V. (2019). Nonparametric spot volatility from options. *The Annals of Applied Probability*, 29(6), 3590–3636.
- Todorov, V. (2022). Nonparametric jump variation measures from options. *Journal of Econometrics*, 230(2), 255–280.
- Xiu, D. (2014). Hermite polynomial based expansion of European option prices. *Journal of Econometrics*, 179(2), 158–177.