0DTE option pricing

This draft: March 15, 2024   First draft: July 7, 2023

Abstract

Trading in 0-Days-To-Expiry (0DTE) options has grown exponentially over the last few years. After describing this exploding market, we present novel closed-form pricing formulae that accurately capture the 0DTE implied-volatility surface. We use a local-in-time approach, relying on Edgeworth-like expansions of the log-return characteristic function, explicitly suited to price ultra-short-tenor instruments. The expansions provide skewness and kurtosis adjustments which depend on the underlying non-affine return characteristics in closed form. We show significant improvements in pricing and hedging as compared to state-of-the-art models. We conclude by providing suggestive results on nearly *instantaneous predictability* by estimating 0DTE-based risk premia.

Keywords: Zero days-to-expiry options (0DTEs), pricing, hedging, instantaneous return/variance risk premia.

JEL classification: C51, C52, G12, G13.

1 Introduction

The market for 0DTE options has seen a meteoric increase in traded volume over the last few years. We focus on SPX (i.e., S&P 500) options. For the years 2014 to 2023, Fig. 1 represents the percentage of daily volume (in terms of number of contracts traded) associated with SPX options expiring during the same trading day.

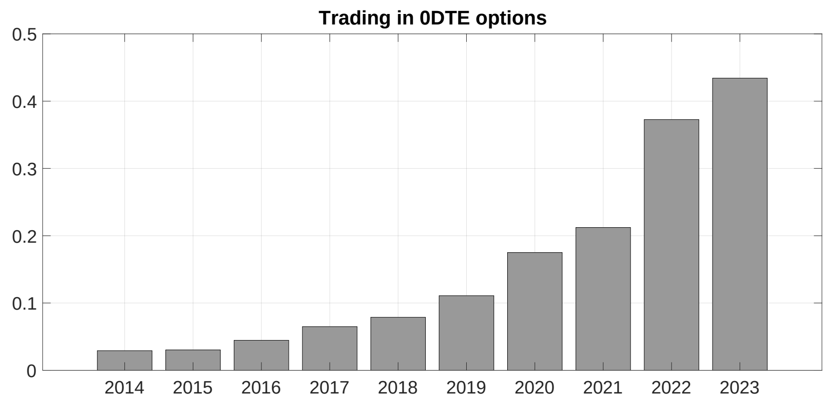

Figure 1: *The figure reports volume (i.e., number of contracts traded) in 0DTE SPX options as a fraction of total volume for options with maturity up to one year. The data covers the period from 2014 to 2023. Data source: OptionMetrics.*

Daily transaction volume in 0DTE SPX options is now over 43% of overall daily option volume, up more than 100% as compared to 2021 levels. According to J.P. Morgan Chase's estimates released in February 2023, this figure corresponds to a daily notional dollar volume around $1 trillion.

While academic work on 0DTEs is still very limited, investors have been paying attention. So, have exchanges. The CBOE is currently believed to derive more than 56% of its revenues from its option business.[1] This figure justifies the CBOE's incentive to progressively increase the listing frequency of weekly options (weeklies) to a daily frequency, a phenomenon that has lead to the current access to 0DTEs.[2] On each individual day, 0DTEs are - in essence - weeklies listed a week earlier.

[1]See, e.g., the April 2023 Ambrus Capital's report "Dispelling false narratives about 0DTE options."

[2]The CBOE Friday weeklies, as well as weeklies expiring on other days, are listed on the following web page: `https://www.cboe.com/available_weeklys/`.

0DTEs offer new opportunities to traders who can now capitalize on very short-term directional bets, e.g. around macroeconomic announcements,[3] among other more complex strategies. These new opportunities have appeared as especially appealing in an uncertain (for equities) environment of increasing interest rates. The sell side of the market, however, faces considerable "convexity". Because of their short maturity, the gamma of 0DTEs (the second derivative of the option price with respect to the price of the underlying) is large. It is particularly large for at-the-money or near at-the-money positions. This feature, among others, has led some to warn about unintended consequences.[4]

Our interest is in the pricing (and hedging) of short-tenor options, including those in the ultra short-tenor segment, i.e. 0DTEs. We achieve this objective by imposing mild assumptions on the dynamics of the price process (X_t). The price process features both continuous and discontinuous components (X_t^c and X_t^d, respectively). All characteristics driving its dynamics, including, e.g., volatility, leverage (the correlation between price changes and volatility changes), volatility-of-volatility, the intensity of the jumps and the moments of the jump sizes are unrestricted processes only required to satisfy technical, but innocuous, smoothness conditions.

[3] Descriptive evidence is provided in Section 5.

[4] On February 15 2023, J.P. Morgan Chase's chief global markets strategist, Marko Kolanovic, warned that the growing size of the 0DTE segment may lead to sharp market swings as large as $30 billion dollars, particularly in the current low liquidity environment. Referring to the spike of volatility which affected short volatility strategies in February 2018, Mr. Kolanovic famously suggested the possibility of a "Volmageddon 2.0" in the 0DTE market. These swings are viewed as a feedback effect caused by option sellers rushing to hedge their (high gamma) positions after a large intra-daily market move. His opinion is shared by others. On the same day, Saqib Iqbal Ahmed of Reuters quoted the SpotGamma founder Brent Kochuba as saying: "(This market) could draw large, sudden hedging requirements from options dealers. This could be particularly dangerous around an unexpected news event that catches people offsides." He added: "Overall, we feel that 0DTEs pose the potential to create a flash crash at the index level." On February 21 2023, Garrett DeSimone, head of quantitative research at OptionMetrics, was quoted by Joseph Adinolfi of Market Watch as saying "We haven't seen the systemic risks present themselves yet, but there's a concern that if you have a big daily swing, like what we saw during March 2020, that we really don't know how the market-making mechanism is going to react." In the same article, Joseph Adinolfi cited Charlie McElligott, managing director of cross-asset strategy and global equity derivatives at Nomura, as indicating that he would be shocked if regulators weren't already trying to gauge the systemic risks associated with these products. Not everyone agrees with these concerns. On February 23 2023, Lu Wang of Bloomberg News quoted Nitin Saksena, strategist at Bank of America, as writing in a note: "Some are raising the alarm that directional end-users are net short out-of-the-money 0DTEs, thus sowing the seeds for a 'tail wags the dog' event akin to the February 18 'Volmageddon' ... The evidence so far suggests that 0DTEs positioning is more balanced/complex than a market that is simply one-way short tails." This recent debate is indicative of the opportunities and risks that the exponentially-growing 0DTE market has been quickly generating.

Pricing short-tenor options requires the evaluation of the $\mathbb{Q}$-characteristic function of the logarithmic price process over a short horizon τ ($\mathbb{C}^{\log X}(u,\tau)$). This function depends, of course, on the characteristic function associated with the continuous portion of the process ($\mathbb{C}^{\log X^c}(u,\tau)$) and on the characteristic function of the discontinuous portion of the process ($\mathbb{C}^{\log X^d}(u,\tau)$).

Importantly, our assumed price process represents a significant departure from specifications, e.g. Lévy or affine, for which $\mathbb{C}^{\log X}(u,\tau)$ is known in closed form. In spite of the generality of the price process, we provide two *closed-form expressions* for the price characteristic function over small-τ intervals. Theoretically, the resulting expressions are ideally suited, because of their local nature, to price instruments with short tenors, like 0DTEs. Computationally, they retain the tractability of characteristic functions known in closed form.

In our first (baseline) specification, we model $\mathbb{C}^{\log X^d}(u,\tau)$ as a known characteristic function and express $\mathbb{C}^{\log X^c}(u,\tau)$ as a small-τ Edgeworth-like expansion around the Gaussian characteristic function. Local Gaussianity is, of course, induced by the process' driving Brownian motion. Once more, we make no assumptions, other than smoothness, on the nature of the price characteristics (volatility, leverage and volatility-of-volatility, *inter alia*) their own $\mathbb{Q}$-dynamics and risk premia. We document that the expansion depends on leverage and the volatility-of-volatility, among other characteristics. The former affects the at-the-money skew, the latter impacts the at-the-money convexity of the implied volatility surface. By tilting locally (in τ) the conditional Gaussian characteristic function in such a way as to introduce negative return skewness (through leverage) as well as thicker return tails (through, e.g., the volatility-of-volatility), the proposed expansion is shown to generate model-implied volatilities which adapt effectively to the at-the-money (and near at-the-money) 0DTE implied volatilities in the data. The result is accurate pricing *within* the bid/ask for 80% of our options and superior performance as compared to state-of-the-art specifications.

The importance of discontinuous variation in assisting out-of-the-money short-term pricing is well-known (e.g., Andersen, Fusari, and Todorov, 2017). We instead formalize - in closed form - the role played by continuous variation in yielding at-the-money (and near at-the-money) implied volatilities which adapt to those in the data and, as a result, superior fit of the entire implied volatility surface. We document that daily volume in 0DTE SPX options is distributed evenly around at-the-money, something which is consistent with reasonably sophisticated trading strategy and the evidence that a relatively small portion of these instruments are traded by retail investors.[5] Because our proposed expansions are designed to capture the implied-volatility skew and convexity, large traded volume at-the-money (or near at-the-money) translates into significant price improvements as compared to state-of-the-art competitors. Conversely, because of traded volume around at-the-money in the ultra short segment of this market - that occupied by 0DTEs - the center of the moneyness range is of particular economic interest. Finally, all else equal, superior pricing at-the-money and near at-the-money is expected to improve pricing, as well as hedging, along the full implied volatility surface, even in the tails. We provide evidence that this is, in fact, the case.

In our second specification, we work with a small-τ Edgeworth-like expansion of the characteristic function of the *full* process $\mathbb{C}^{\log X}(u,\tau)$, one in which $\mathbb{C}^{\log X^c}(u,\tau)$ is expressed as in the baseline case *and* the jump portion of the characteristic function ($\mathbb{C}^{\log X^d}(u,\tau)$) is also expanded. The full expansion allows for jumps in prices, as in the baseline case, but also for jumps in volatility. The price and the volatility jumps are permitted to be idiosyncratic and joint. In other words, the expansion introduces - again, in closed form - an additional source of skewness (through discontinuous leverage induced by the price/volatility co-jumps) and an additional source of kurtosis (through the volatility jumps). We show that expanding the characteristic function of the jump in prices (without jumps in volatility) does not add in any economically-meaningful way to the baseline specification in which a known jump characteristic function is used. We also show that, in spite of their potential contribution to both skewness and kurtosis, the jumps in volatility have a marginal impact on our reported short-term pricing. We, of course, do not exclude that a more constrained (possibly affine)

[5]On February 21, 2023, Joseph Adinolfi of Market Watch referred to J.P. Morgan Chase data as suggesting that retail investing only amounts to about 5.6% of daily 0DTE option trading. Supporting interest in institutional investors, in an April 2023 note, Ambrus Capital wrote: "Post-March 2020 many institutions began advocating for option overlay programs. During 2021 many of those mandates were approved and put into action. Along with generic hedging mandates, large RIAs also began implementing yield-generating programs. During 2022, rates began to move higher and equities sank. This left wealth managers looking for other sources of yield. With equity volatility remaining muted, these yield programs started to attract more investors towards the end of Q1 of 2022. The addition of short-dated options made these programs very attractive to advisors and other institutions. The shrinking equity risk premium due to the move in rates left investors flocking to the growing volatility risk premium. This is why there was a substantial increase in 0DTE volume around Q2 of 2022. Additionally, for volatility hedge funds and market makers, this new tenor allowed a cleaner way to hedge gamma and theta risk. Furthermore, speculative hedge funds that are not derivative-focused began using 0DTEs to hedge macroeconomic-driven event risk."

specification would attribute a bigger role to these discontinuities. We, instead, find that flexible diffusive dynamics - such as the ones we assume - contribute to the returns' higher moment dynamics in ways which attenuate the pricing relevance of the volatility jumps.

We proceed as follows. In Section 2, we offer more motivation for focusing on 0DTE options. Section 3 further positions our contribution. In Section 4 we detail the two proposed pricing models. Empirical work begins with the first (baseline) specification before turning to the second specification in Section 9. Section 5 is about details of implementation and the data. Particular attention is devoted to the features of the *intra-daily* cross-sectional option data we use. Pricing performance is documented in Section 6. Section 7 evaluates robustness. Among other exercises, we slice pricing performance along a variety of dimensions: moneyness, volatility states and tenor. Section 8 reports on profits and losses from Δ- and Γ-hedging strategies. Section 9 turns to the full expansion of the process' characteristic function in our second model specification. In Section 10 we exploit the identification potential of the proposed pricing model to estimate instantaneous notions of the return/volatility risk premia. In other words, we study (nearly) *instantaneous predictability*. Section 11 concludes.

2 More on motivation

The rise of trading in short-term option contracts is generally associated with the introduction of SPX weekly options (weeklies) on October 28, 2005. These original weeklies were listed on Friday and were expiring at 4pm on the Friday of the subsequent week.[6] More than ten years later - on February 23 2016 - the CBOE introduced weeklies expiring on Wednesday of each week. Shortly thereafter, on August 15 2016, weeklies with Monday expirations were introduced. The CBOE advertised them as a new effective instrument to hedge "over-the-weekend" risks. Options with expirations on Tuesday and Thursday were offered on April 18 2022 and on May 11 2022. The result has been the generation of a very liquid market of option contracts allowing investors to trade ultra short-tenor instruments during *each* trading day and at *every* point in time within the trading day.

Fig. 2 reports the percentage of volume associated with SPX options with different tenors (0-7 days, 7 days to 1 month, 1 month to 3 months, 3 months to 6 months, and 6 months to

[6]The exception were weeklies which would have been set to expire on the third Friday of the month (the day of expiration of standard SPX option contracts). These options were not listed.

12 months). Within each year, the columns sum up to 1. It is apparent that all segments of the SPX market have shrunk - in relative terms - over time with the exception of the longer maturities, which have remained rather stable, and the shortest 0 to 7 day maturities, which have seen considerable growth. This said, total trading activity has increased over time and all maturity segments have witnessed a corresponding increase in absolute terms.

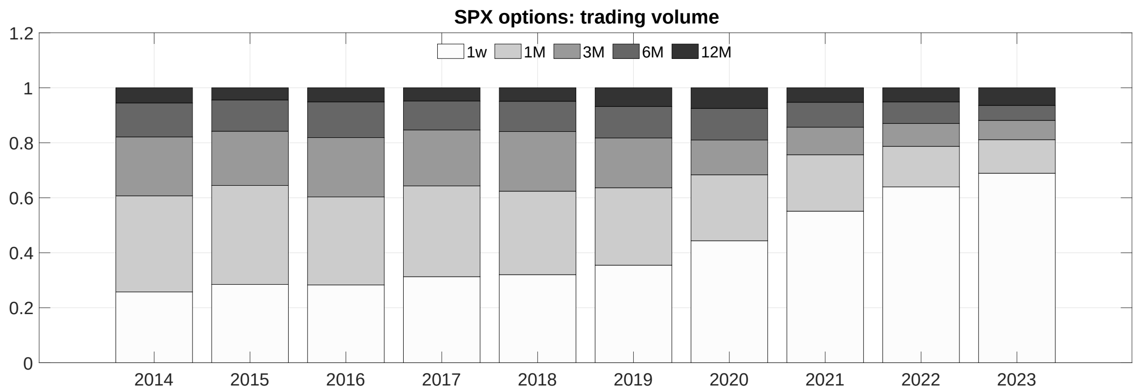

Figure 2: *The figure reports the percentage of volume (i.e., number of contracts traded) associated with SPX options with different tenors (0-7 days, 7 days to 1 month, 1 month to 3 months, 3 months to 6 months, and 6 months to 12 months). Within each year, the columns sum up to 1. The data covers the period from 2014 to 2023. Data source: OptionMetrics.*

Zooming into the 0 to 7 day maturities in Fig. 3, we notice that they have all remained reasonably stable over the time period 2014 to 2023 with the exception of the shortest maturities, namely those covered by 0DTE options, and the maturities immediately longer. The former, in particular, have witnessed an exponential growth.

Since the first draft of this article, the geographical diffusion of these instruments has also increased. So has their use in structured products. On August 28 2023, Deutsche Boerse began trading 0DTE options on the Euro Stoxx 50 index (0EXP). In its press release on August 22 2023, the exchange wrote: "Institutional demand for options with short-term expiries has strongly increased as investors seek to react quickly and precisely to specific market events".[7] On September 13 and 19 2023, Defiance ETFs launched QQQY and JEPY, two actively-managed funds designed to gain exposure to the Nasdaq 100 Index and the S&P

[7] While retail trading has seen an uptick in Europe over the last couple of years, its diffusion is far from that experienced in the US and China. The reference to "institutional demand" reflects this reality as well as the specificities of instruments which, even in the US, have drawn considerable institutional interest.

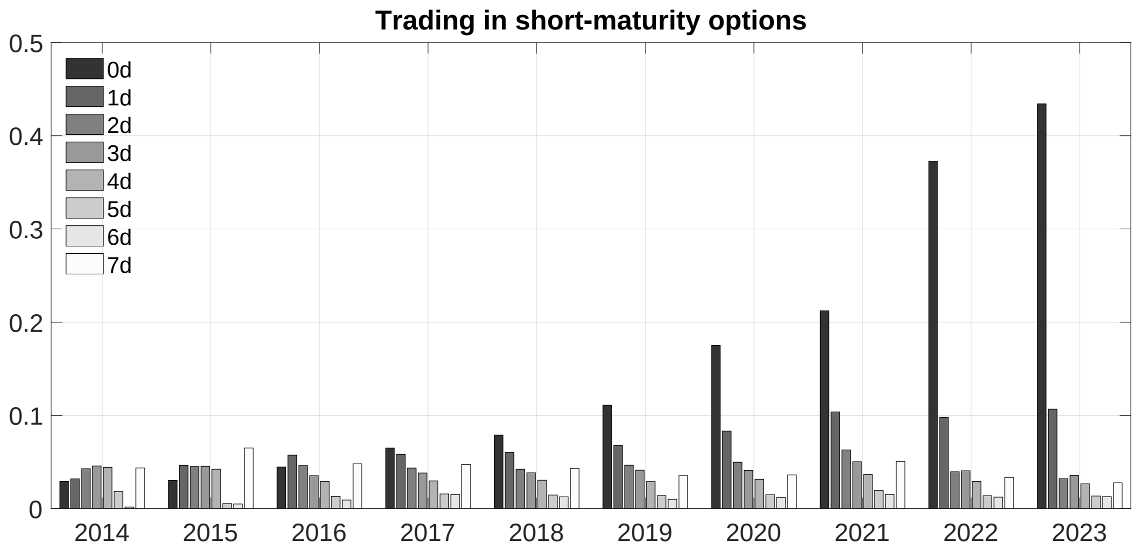

Figure 3: *The figure reports volume (i.e., number of contracts traded) associated with SPX options across various short tenors (ranging from 0 days to 7 days to maturity) as a fraction of total volume for options with maturity up to one year. The data covers the period from 2014 to 2023. Data source: OptionMetrics.*

500 index and generate income through 0DTE option writing. On October 13, after 30 days of trading, the combined AUM of the two funds was reported in a Defiance ETFs press release to be over $100 million.

In sum, the 0DTE market is growing by the day. Its growth, in turn, justifies academic attention to it. We address this need by focusing on issues of valuation.

3 Positioning

This article contributes to various strands of the literature.

Considerable work has been devoted to expansions (of various nature and of various quantities, from implied volatilities, to transition densities, to characteristic functions) for the purpose of valuing structured products, doing inference using the information in structured products, or both. We refer the reader, e.g., to Jarrow and Rudd (1982), Corrado and Su (1996), Sircar and Papanicolaou (1999), Lee (2001), Kunitomo and Takahashi (2001), Carr and Wu (2003), Durrleman (2004), Medvedev and Scaillet (2007), Bentata and Cont (2012), Gatheral, Hsu, Laurence, Ouyang, and Wang (2012), Forde and Jacquier (2011), Forde, Jacquier, and Lee (2012), Takahashi and Yamada (2012), Xiu (2014), Jacquier and Lorig (2015), Bandi and Renò (2017), Lorig, Pagliarani, and Pascucci (2017), Aït-Sahalia, Li, and Li (2021), Todorov (2021) and the references therein. We introduce a novel procedure to derive pricing expansions for short-tenor derivative instruments. The expansions are local in time and, therefore, particularly suited to value 0DTEs, our objective in this article.

A successful literature in financial econometrics has been devoted to the estimation of a variety of equity characteristics: from spot volatility, to spot leverage, to spot volatility-of-volatility, among other quantities. The bulk of this work has made use of the information contained in intra-daily price data for identification. For spot volatility, see, e.g., Fan and Wang (2008), Mykland and Zhang (2008), Kristensen (2010), Zu and Boswijk (2014), Mancini, Mattiussi, and Renò (2015), Bandi and Renò (2018), Bibinger, Hautsch, Malec, and Reiss (2019) and the references therein. Regarding spot leverage, see, e.g., Bandi and Renò (2012), Aït-Sahalia, Fan, and Li (2013), Wang and Mykland (2014), Wang, Mykland, and Zhang (2017), Kalnina and Xiu (2017), Aït-Sahalia, Fan, Laeven, Wang, and Yang (2017) and the references therein. For spot volatility-of-volatility, see, e.g., Vetter et al. (2015), Sanfelici, Curato, and Mancino (2015), Barndorff-Nielsen and Veraart (2012) and the references therein. There is also work which employs the information in short-term options for identification: Andersen, Fusari, and Todorov (2017) (spot volatility and tail characteristics), Todorov (2019) (spot volatility), Todorov (2021) (leverage) and Chong and Todorov (2023) (leverage and volatility-of-volatility). Identifying the equity characteristics is not our emphasis. We may, in fact, use existing estimation methods and employ the resulting estimates as inputs in the proposed pricing expansion. Having made this point, for reasons that will be discussed in Section 5, there is value in identifying the (needed, for valuation) equity characteristics *jointly*. We view joint identification of the characteristics as a contribution of this article of independent interest.

Turning to the 0DTE market in particular, academic work focusing specifically on this market is very limited. Brogaard, Han, and Won (2023) focus on spillovers from the 0DTE segment of the short-tenor option market to the market of the underlying. They relate 0DTE volume to the volatility of the underlying, showing how the former may influence (due, e.g., to informational effects and/or hedging) the latter. The destabilizing impact of the recent surge in 0DTE trading is disputed by Dim, Eraker, and Vilkov (2024) who instead highlight the strategic role of 0DTEs, e.g., as bets on resolution of uncertainty (or large price moves) around announcements. Even though the broader option market is generally viewed as a playground for sophisticated institutional investors, retail investors have shown particular interest in the 0DTE segment. Beckmeyer, Branger, and Gayda (2023) emphasize that while still limited, relative to institutional trading, the share of retail trading in the 0DTE segment (around 6%) is higher than that (2% to 4%) in the longer segment of the market. 0DTEs are, therefore, attractive to retail traders too. Beckmeyer, Branger, and Gayda (2023) report that roughly 75% of retail trades in S&P 500 options in the first quarter of 2023 were in 0DTE contracts. These trades were, however, generally unsuccessful, a result which mirrors existing results for longer, i.e. weekly, expiries (Bryzgalova, Pavlova, and Sikorskaya, 2022). Vilkov (2023) analyses a variety of trading strategies in the 0DTE space. He documents a tendency for the profit-and-losses of popular strategies to be negative and with considerable volatility. He concludes by cautioning traders lacking a "nuanced understanding" of the market.

We are interested in valuing 0DTEs. The topic is important for the buy side looking to understand fair values and place, e.g., directional bets. It is important for the sell side and market-makers attempting to quantify risks (including, but not limited to, gamma risk) and hedge them. It is, also, important for regulators wishing to assess the systemic implications of a new and extremely fast-growing market. Using rich affine models with self-exciting jumps, Bates (2019) provides an early study on the valuation of options with expiries equal to one day (and longer).[8] He emphasizes that his most successful model specification would still fail to capture the option skew and its time variation.[9] Bates (2019) advocates for non-affine specifications as a solution to this issue, a solution which is consistent with our proposed valuation method and its non-affine nature.

[8] As discussed in Section 5, our shortest expiry in this article is 5.5 hours but we have experimented with expiries as short as 1 hour with very similar findings.

[9] His model captures the skew over short maturities better than over long maturities. However, even over short maturities, there appears to be a significant, unexplained skew time variation.

4 The pricing model

We assume an adapted, real stochastic price process X_t defined on the filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ expressed as follows:

$$\begin{aligned} dX_t &= \underbrace{\mu_t X_t dt + \sigma_t X_t dW_t}_{dX_t^c} + \underbrace{(e^{x_t} - 1)X_t dN_t}_{dX_t^J}, \\ d\sigma_t &= \alpha_t dt + \beta_t dW_t + \beta_t' dW_t', \\ d\mu_t &= \gamma_t dt + \delta_t dW_t + \delta_t' dW_t^*, \\ d\beta_t &= \zeta_t dt + \eta_t dW_t + \eta_t' dW_t^{**}, \end{aligned} \tag{1}$$

where $\sigma_t > 0$, $\beta_t > 0$, $\beta_t' > 0$, $\delta_t > 0$, $\delta_t' > 0$, $\eta_t > 0$ and $\eta_t' > 0$, almost surely, $\forall t \geq 0$, and the quantities W_t, W_t', W_t^* and W_t^{**} are independent Brownian motions. N_t is a Poisson process and x_t is a random jump size. X_t^c and X_t^J are independent.

The characteristics σ_t and μ_t and, more generally, all other characteristics (including, e.g., β_t and β_t') are adapted processes. Because of the presence of a common Brownian motion W_t, they are allowed to be correlated among themselves and with X_t. The jump sizes x_t and the Poisson intensity λ_t are also adapted processes. Only technical smoothing conditions - to which we will return - are imposed on all processes.

The price process is, of course, neither Lévy nor - as said - affine. Given the generality of its assumed dynamics, obtaining a closed-form characteristic function (to be inverted for the purpose of pricing) is, therefore, infeasible, in general. Consistent with our objective to price short-tenor instruments, we address this issue by proposing two closed-form expansions of the characteristic function over short horizons, to which we now turn.

4.1 The first (baseline) specification

We are interested in the $\mathbb{Q}$-characteristic function of the logarithmic price process, namely $\mathbb{C}^{\log X}(u, \tau)$. For reasons of benchmarking with a model which has been successful in the pricing of short-term options (with tenors longer than one day), i.e., the one proposed by Andersen, Fusari, and Todorov (2017), we write

$$\mathbb{C}^{\log X}(u, \tau) = \mathbb{E}_t^{\mathbb{Q}}[e^{iu(\log X_{t+\tau} - \log X_t)}] = \mathbb{C}^{\log X^c}(u, \tau) \times \mathbb{C}^{\log X^d}(u, \tau). \tag{2}$$

In Andersen, Fusari, and Todorov (2017), parametric jumps (represented by $\mathbb{C}^{\log X^d}(u,\tau)$) are added to an independent diffusive component (represented by $\mathbb{C}^{\log X^c}(u,\tau)$) which is assumed to be conditionally (on spot volatility, σ_t) Gaussian. The intensity of the jumps, the moments of the jump sizes and spot volatility (the sole diffusive characteristic in Andersen, Fusari, and Todorov, 2017) are adapted processes.

In order to derive the cleanest possible comparison with this approach - and, in the process, in order to highlight the importance of diffusive dynamics - we work with the same model as in Eq. (2) but *solely* modify $\mathbb{C}^{\log X^c}(u,\tau)$. Rather than assuming a conditionally (on σ_t) Gaussian characteristic function (as in Andersen, Fusari, and Todorov, 2017), we employ a small-τ expansion of the characteristic function of $\log X^c$ around the Gaussian characteristic function. Among other effects, the expansion adds skewness to the Gaussian distribution (through a time-varying leverage process) and kurtosis (through, e.g., a time-varying volatility-of-volatility process). We will show that these tilts are central to better fit at-the-money, near at-the-money and beyond. Next, we provide details.

In the baseline model, consistent with the approach in Andersen, Fusari, and Todorov (2017), we parametrize $\mathbb{C}^{\log X^d}(u,\tau)$ and assume - at first - conditional Gaussian jump sizes x_t. Gaussianity is, of course, a classical assumption on the distribution of the jump sizes in logarithmic prices.[10] Any parametric assumption on the density of the jump sizes (leading to a known characteristic function) is, however, allowed. In Section 7, we therefore relax it and work with the tempered stable family.[11] The conditional mean and the standard deviation of the jump sizes are $\mu_{j,t}$ and $\sigma_{j,t}$, respectively. The infinitesimal intensity of the jumps is λ_t. Thus, all jump characteristics are allowed to be processes. In this sense, the assumed specification is more flexible than models in which these quantities are tightly parametrized (i.e., the many specifications in the tradition of Heston, 1993). It is also more flexible than models in which the quantities are assumed to be nonparametric functions of the state variables (as in Bandi and Renò, 2016, and Bandi and Renò, 2022). The intensity λ_t, for example, could be self-exciting.

[10] See, e.g., Bates (2000), Duffie, Pan, and Singleton (2000), Pan (2002), Eraker (2004) and Broadie, Chernov, and Johannes (2007).

[11] Andersen, Fusari, and Todorov (2017) estimate a conditional Gaussian jump specification before turning to the (conditional) tempered stable class (e.g., Carr, Geman, Madan, and Yor, 2003, and Carr and Wu, 2003) and its sub-cases, like the double-exponential model of Kou (2002) or the variance gamma model of Madan, Carr, and Chang (1998).

We write

$$\mathbb{C}^{\frac{\log X^d}{\sigma_t\sqrt{\tau}}}(u,\tau) = e^{\tau\lambda_t(e^{iu\frac{\mu_{j,t}}{\sigma_t\sqrt{\tau}} - u^2\frac{\sigma_{j,t}^2}{2\sigma_t^2\tau}} - 1 - iu\bar{\mu}_{j,t})},$$

where $\bar{\mu}_{j,t}$ is a $\mathbb{Q}$-compensator expressed as $e^{\frac{\mu_{j,t}}{\sigma_t\sqrt{\tau}} + \frac{1}{2}\frac{\sigma_{j,t}}{\sigma_t^2\tau}^2} - 1$. Given $\mathbb{C}^{\log X^c}(u,\tau)$, to which we now turn, the presence of the compensator $\bar{\mu}_{j,t}$ will guarantee that $\mathbb{C}^{\log X}(-i,\tau) = \mathbb{E}_t^{\mathbb{Q}}[e^{(\log X_{t+\tau} - \log X_t)}] = \mathbb{E}_t^{\mathbb{Q}}[X_{t+\tau}/X_t] = e^{r_t\tau} + O(\sqrt{\tau})$, as required by risk-neutral pricing.

We express $\mathbb{C}^{\log X^c}(u,\tau)$ as local (in τ) Edgeworth-like expansion of the $\mathbb{Q}$-characteristic function of $\log X_t^c$ based on the $\mathbb{Q}$-counterpart to the $\mathbb{P}$-process in Eq. (1) above. The $\mathbb{Q}$-process has the same dynamic structure as the $\mathbb{P}$-process.[12] As emphasized, within the adopted family of continuous-time semi-martingales for $\log X_t^c$, these dynamics are unrestricted and all characteristics are treated as general semi-martingales themselves. Hence, our proposed expression for $\mathbb{C}^{\log X^c}(u,\tau)$ is nonparametric in nature.

In order to separate the volatility-of-volatility from leverage, we adopt the classical Heston's specification (Heston, 1993) and suitably re-label quantities:

$$\beta_t = \widetilde{\beta}_t\rho_t,$$

$$\beta_t' = \widetilde{\beta}_t\sqrt{1-\rho_t^2}.$$

As a result, $\widetilde{\beta}_t$ now represents the full volatility-of-volatility and ρ_t is time-varying leverage. The separation between the two quantities, and their relative role, will be central to the economic interpretation of our findings.

Write $\widetilde{\mu}_t = r_t - \frac{\sigma_t^2}{2}$, where r_t is the risk-free rate. Assume $(\widetilde{\mu}_t, \sigma_t) \in D_W^{(4)}$.[13] By virtue of Theorem 1 in Bandi and Renò (2017) (see Appendix for details), the $\mathbb{Q}$-characteristic function of the diffusive component of the logarithmic price process, i.e., $\mathbb{C}^{\log X^c}(u,\tau)$ with $u \in \mathbb{C}$ and $\tau = T - t$, can be expressed as a small-τ Edgeworth-like expansion given by

[12] The relation between (some of) the $\mathbb{P}$-characteristics and the corresponding $\mathbb{Q}$-characteristics is made explicit in Section 10 by virtue of a compatible (non-monotonic) measure change.

[13] The notation signifies W-differentiability to the fifth order of the $\mathbb{Q}$-$\log X$ process, i.e. $\log X_t^{\mathbb{Q}}$. This is simply a smoothness condition on the characteristics of $\log X_t^{\mathbb{Q}}$ and, thus, on $\log X_t^{\mathbb{Q}}$ itself. It is asking $\log X_t^{\mathbb{Q}}$ to have a drift and a diffusion which depend on the main Brownian motion W (something that was specified explicitly in Eq. (1)). Their characteristics should also depend on the main Brownian motion W and so on, two more times. Bandi and Renò (2017) show formally that the smoothness properties of the process are related to the integrability of its characteristic function, something which is needed to derive the implied density and price.

$$\mathbb{C}^{\frac{\log X^c}{\sigma_t\sqrt{\tau}}}(u,\tau) = e^{iu\frac{\widetilde{\mu}_t\tau}{\sigma_t\sqrt{\tau}}-\frac{u^2}{2}}\left(1 \quad \underbrace{-iu^3\frac{\widetilde{\beta}_t\rho_t}{2\sigma_t}\sqrt{\tau}}_{\text{third moment adjustment}}\right.$$

$$\left.\underbrace{-u^2\left(\frac{(\alpha_t^{\mathbb{Q}}+\widetilde{\delta}_t)}{2\sigma_t}+\frac{\widetilde{\beta}_t^2}{4\sigma_t^2}\right)\tau+\frac{1}{24}\frac{\widetilde{\beta}_t^2}{\sigma_t^2}u^2\left(4u^2-\rho_t^2u^2\left(3u^2-8\right)\right)\tau+\frac{\eta_t}{6\sigma_t}u^4\tau}_{\text{second, fourth and sixth moment adjustment}}\right), \qquad (3)$$

where $\alpha_t^{\mathbb{Q}}$ is the $\mathbb{Q}$-drift of the volatility process and $\widetilde{\delta}_t$ is the W-volatility of the $\mathbb{Q}$-drift in logarithmic diffusive prices ($\widetilde{\mu}_t$).

Eq. (3) is an expansion around the conditional Gaussian density. The expansion provides both closed-form skewness (through ρ_t) and kurtosis (through, e.g., $\widetilde{\beta}_t$) adjustments. The first-order (in $\sqrt{\tau}$) term in the expansion is, in fact, a third moment adjustment capturing skewness. The second-order (in τ) term is, instead, an even (second, fourth and sixth) moment adjustment. This term captures kurtosis (as well as the high-order contribution of squared skewness to the sixth moment and the contribution of time-varying characteristics to the second moment). We will show that both adjustments are central to the reported price improvements.

In order to appreciate the flexibility of our proposed approach, we derive the implied parameters/processes entering the expansion in the case of two well-known specifications: Heston's affine model (Heston, 1993) and the logarithmic volatility model used, e.g., in Chernov, Gallant, Ghysels, and Tauchen (2003).

Heston's model reads as follows:

$$d\log X_t^c = (\mu - \sigma_t^2/2)dt + \sqrt{\sigma_t^2}dW_t,$$

$$d\sigma_t^2 = \kappa(\omega - \sigma_t^2)dt + \xi\sqrt{\sigma_t^2}\rho dW_t + \xi\sqrt{\sigma_t^2}\sqrt{1-\rho^2}dW_t',$$

where $\mu, \kappa, \omega, \xi, \rho$ are constant parameters. Using Ito's lemma repeatedly, we obtain

$$\alpha_t = \frac{1}{8\sigma_t}(4\kappa(\omega - \sigma_t^2) - \xi^2),$$

$$\widetilde{\beta}_t = \frac{1}{2}\xi,$$

$$\rho_t = \rho,$$

$$\widetilde{\delta}_t = -\frac{1}{2}\sigma_t \xi \rho,$$

and

$$\eta_t = 0.$$

Thus, Heston's model implies a constant leverage and a constant volatility-of-volatility (and, therefore, no volatility of the volatility-of-volatility). However, the volatility of the drift is time-varying (and depends, linearly, on the volatility itself).

The logarithmic volatility model reads, instead, as follows:

$$d \log X_t^c = (\mu - \sigma_t^2/2)dt + \sigma_t dW_t,$$

$$d \log \sigma_t = \kappa(\omega - \log \sigma_t)dt + \xi \rho dW_t + \xi \sqrt{1-\rho^2} dW_t',$$

where, again, $\mu, \kappa, \omega, \xi, \rho$ are constant parameters. Using, again, Ito's lemma, we obtain

$$\alpha_t = \frac{1}{2}\sigma_t(\xi^2 + 2\kappa(\omega - \log \sigma_t),$$

$$\widetilde{\beta}_t = \xi \sigma_t,$$

$$\rho_t = \rho,$$

$$\widetilde{\delta}_t = -\sigma_t^2 \xi \rho,$$

and

$$\eta_t = \sigma_t \xi^2 \rho^2.$$

In this model specification, leverage continues to be constant. The volatility-of-volatility and the volatility of the volatility-of-volatility are, instead, proportional to the volatility. We conclude that both models impose constraints on skewness (through a constant drift) and kurtosis (through, e.g., a constant or tightly parametrized volatility-of-volatility). In this article, we dispense with these (and other) constraints by working with unrestricted processes whose time-t realizations are estimated, period-by-period, using a cross-section of options (as discussed in Section 5).

In essence, capturing the volatility dynamics in a flexible way is central to our approach. If we were to freeze volatility (as in Andersen, Fusari, and Todorov, 2017), the model would amount to a conditionally (on time-varying spot volatility) Gaussian specification with jumps (a specification which we will later dub BSM). This is our $\tau = 0$ case. Because we do not freeze volatility, the remaining ($\sqrt{\tau}$ and τ) terms in the expansion (i.e., the skewness and the kurtosis correction) arise as a function of the volatility features (e.g., ρ_t, $\widetilde{\beta}_t$ and $\alpha_t^{\mathbb{Q}}$). The nonparametric nature of these features make it infeasible to employ known closed-form characteristic functions. We, instead, propose a closed-form expansion which is naturally suited to price short-tenor instruments. Below, we add discontinuities in volatility to the expansion.

4.2 The second specification

The second specification is a small-τ Edgeworth-like expansion of the $\mathbb{Q}$-characteristic function of the full logarithmic price process $\log X_t$ inclusive of finite-activity jumps in volatility:

$$\mathbb{C}^{\frac{\log X}{\sigma_t\sqrt{\tau}}}(u,\tau) = e^{iu\frac{\widetilde{\mu}_t^{\text{full}}\tau}{\sigma_t\sqrt{\tau}}-\frac{u^2}{2}}\left(1 - iu^3\frac{\widetilde{\beta}_t\rho_t}{2\sigma_t}\sqrt{\tau} - u^2\left(\frac{(\alpha_t^{\mathbb{Q}}+\widetilde{\delta}_t)}{2\sigma_t}+\frac{\widetilde{\beta}_t^2}{4\sigma_t^2}\right)\tau \right. \tag{4}$$

$$+\frac{1}{24}\frac{\widetilde{\beta}_t^2}{\sigma_t^2}u^2\left(4u^2-\rho_t^2u^2\left(3u^2-8\right)\right)\tau+\frac{\eta_t}{6\sigma_t}u^4\tau \tag{5}$$

$$+\underbrace{\tau\int_{\mathbb{R}}\left(e^{iu\frac{x}{\sigma_t\sqrt{\tau}}}-1\right)\lambda_t^X f_t^X(dx)}_{\text{idiosyncratic price jumps}} \tag{6}$$

$$+\underbrace{\tau\int_{\mathbb{R}^2}\int_0^1\left(e^{iu\frac{x}{\sigma_t\sqrt{\tau}}}e^{-\frac{u^2}{2}\left(\frac{s^2}{\sigma_t^2}+2\frac{s}{\sigma_t}\right)v}-1\right)\lambda_t^{X,\sigma}f_t^{X,\sigma}(dx,ds)dv}_{\text{joint price/volatility jumps}} \tag{7}$$

$$\left.+\underbrace{\tau\int_{\mathbb{R}}\int_0^1\left(e^{-\frac{u^2}{2}\left(\frac{s^2}{\sigma_t^2}+2\frac{s}{\sigma_t}\right)v}-1\right)\lambda_t^{\sigma}f_t^{\sigma}(ds)dv}_{\text{idiosyncratc volatility jumps}}\right), \tag{8}$$

with

$$\widetilde{\mu}_t^{\text{full}} = r_t - \frac{\sigma_t^2}{2} - \sigma_t\sqrt{\tau}\left(\int_{\mathbb{R}}\left(e^{\frac{x}{\sigma_t\sqrt{\tau}}}-1\right)\lambda_t^X f_t^X(dx) + \int_{\mathbb{R}}\left(e^{\frac{x}{\sigma_t\sqrt{\tau}}}-1\right)\lambda_t^{X,\sigma} f_t^{X,\sigma}(dx)\right), \quad (9)$$

where λ_t^X, $\lambda_t^{X,\sigma}$ and λ_t^σ, are the time-t intensities of the idiosyncratic price jumps, the joint price/volatility jumps and the idiosyncratic volatility jumps, respectively. As earlier, all intensities are processes. The quantities $f_t^X(.)$, $f_t^{X,\sigma}(.)$ and $f_t^\sigma(.)$ are, instead, the time-t densities of the price jump sizes, of the joint price/volatility jump sizes, and of the volatility jump sizes, respectively. Theorem 2 in Bandi and Renò (2017) justifies this expansion without applying it to pricing, the subject of this article. The Appendix provides details on pricing. Relative to the baseline specification, the price jump compensations are now folded into a re-defined term $\widetilde{\mu}^{\text{full}}$. As earlier, they insure that $\mathbb{E}_t^{\mathbb{Q}}[X_{t+\tau}/X_t] = e^{r_t\tau} + O(\sqrt{\tau})$, as required by risk-neutral pricing.

The volatility jumps have the potential to contribute to both overall skewness (through negative cross-moments between the contemporaneous price/volatility jump sizes in Eq. (7)) and to overall kurtosis (through a positive mean of the idiosyncratic volatility jump sizes in Eq. (8)). The conceptual importance of the first effect in contributing to total leverage (and, therefore, to overall skewness) is central to Bandi and Renò, 2016. The second effect is well-understood.

Our empirical work (in Section 5 through 8) begins with the first (baseline) specification. As emphasized, this choice is motivated by the need to relate our findings to existing benchmark(s) in the literature. It is also motivated by our interest in understanding whether a richer specification (specifically one which allows for jumps in volatility and, therefore, for added sources of skewness and kurtosis) leads to meaningful improvements over a successful, but more parsimonious, specification. To this extent, Section 4.2 documents that, in spite of its generality, our second specification does not contribute to the baseline meaningfully. Volatility jumps, in particular, will be shown to hardly affect the price of 0DTEs.

5 Implementation and data

Given values of the characteristics (σ_t, $\widetilde{\beta}_t$, and ρ_t, among others), Fourier inversion of the characteristic function in Eq. (2) would give us prices. We price options using the traditional quadrature method, as in Heston (1993). We first discuss how the inputs (i.e., the characteristics) are obtained for the purpose of implementing the proposed pricing model. We then turn to data. We will emphasize the richness of CBOE *intra-daily* option data as a source of granular information allowing pricing at any point in time *within* the trading day.

5.1 Estimating the equity characteristics

In order to price, we could feed time-series sample analogues of the "historical" characteristics into $\mathbb{C}^{\log \widetilde{X}}(u,\tau)$ and Fourier invert. This procedure is, however, delicate - for at least two reasons. First, while spot volatility (σ_t) can be measured accurately using time-series data, quantities based on spot volatility (like $\widetilde{\beta}_t$ and ρ_t) are known to be problematic. Their time-series sample analogues, in fact, entail *differences* of spot volatility estimates. While the measurement error in the point volatility estimates may be limited, differencing point estimates has the potential to exacerbate noise. They also entail squares/products of these differences, which are bound to yield biases in the presence of measurement error. Second, a large literature (c.f. the discussion in Section 3) has focused on the estimation of σ_t, $\widetilde{\beta}_t$ and ρ_t using high-frequency time-series data. We are, however, not aware of work on the remaining characteristics entering the characteristic function, i.e., $\alpha_t^{\mathbb{Q}}$, $\widetilde{\delta}_t$ and η_t. Even though sample analogues to these quantities can be constructed, given their nature we believe that they would be hard to identify reliably using time-series data.

In light of these considerations, we estimate all characteristics jointly by minimizing the average squared distance between the market Black and Scholes implied volatilities and the model-based implied volatilities. The criterion is, therefore, in the spirit of Bates (1996) and Andersen, Fusari, and Todorov (2015). The latter, in particular, estimate jump features and spot volatility (ρ_t and $\widetilde{\beta}_t$ not appearing in their framework). Estimating ρ_t and $\widetilde{\beta}_t$ requires an alternative model specification, which we introduce here. The resulting joint identification of the characteristics addresses both of the issues described in the previous paragraph. On the one hand, as far as $\widetilde{\beta}_t$ and ρ_t are concerned, we avoid the biases which would be induced by two-step estimation of volatility functionals using time-series sample analogues. On the other hand, joint identification also gives us estimates of the residual characteristics ($\widetilde{\alpha}_t = \alpha_t^{\mathbb{Q}} + \widetilde{\delta}_t$, as a sum, and η_t). Because our focus is on both pricing and hedging, below we evaluate the accuracy of the minimum distance estimates by using suitable pricing and hedging metrics.

Define $\Theta_t = (\rho_t, \widetilde{\beta}_t, \sigma_t, \widetilde{\alpha}_t, \eta_t, \lambda_t, \mu_{j,t}, \sigma_{j,t})$, the set of model characteristics at time t. Their estimates are obtained as the following argmin:

$$\widehat{\Theta}_t = \operatorname{argmin} \left\{ \frac{1}{N_t} \sum_{i \in \tau_t} \left[BSIV^{mkt}_{k_i,\tau_t} - BSIV_{k_i,\tau_t}(\Theta_t) \right]^2 \right\}, \tag{10}$$

where τ_t is the shortest available tenor on day t and N_t is the number of out-of-the money options used at time t with maturity τ_t. The optimization outcome is daily time series of estimates for all characteristics, i.e., $\widehat{\Theta}_t$.

5.2 Data

We employ CBOE *intra-daily* option data from January 2 2014 to May 11, 2023. Estimation is conducted at a single point in time during the trading day (10:30am) using the available cross section of options with the *shortest* maturity.

We note that, because the literature has not focused on the shortest segment of the option market, it is customary to use OptionMetrics data, estimate parameters/characteristics at the end of the trading day (3:59pm), and price options which expire on the next day or later (as in, e.g., Andersen, Fusari, and Todorov, 2017). OptionMetrics data is, however, updated less frequently than CBOE data. Also, since in OptionMetrics the option cross section is sampled at 3:59pm, 0DTEs would only be sampled with 1 minute to expiration. We, therefore, exploit the granularity of CBOE *intra-daily* option price data, volume information in Figs. 1, 2 and 3 being the only OptionMetrics data used in this study. In line with our objective (valuing and hedging 0DTEs), we estimate and price *during* the trading day.[14]

Because of our use of intra-daily option data, the selection of 10:30am as the estimation time is a choice. Such choice is dictated by the following observation. As shown in Fig. 4, the opening of the day (i.e., between 9:30am and 10:10am) is characterized by high mean trading volume (top left panel) but, also, by significant variation in volume (bottom left panel). The choice of 10:30am strikes a sensible compromise between high enough mean

[14]In light of well-known intra-day periodicity, the volatility of the underlying is expected to be higher at the beginning and at the end of the trading day. This heightened volatility has a positive impact on the option spreads set by market makers engaged in hedging activities perceived as riskier. Thus, an additional advantage of using cross sections of options in the middle of the trading day (as provided by CBOE data), rather than at the end of the trading day, is that the corresponding prices are likely less affected by execution costs.

volume and low enough variability, excessive variability being something which could affect the quality of the estimates across days.

We note that, even in this market, volume has a U-shaped pattern with higher levels both at the beginning and at the end of the trading day. On FOMC announcement days, the latter occur earlier (around 2pm) and are drastically steeper (top right panel). These effects confirm the importance - to which we refer in the Introduction - of directional trades around macroeconomic announcements in this market. While these trades are interesting in their own right, we do not pursue this line of inquiry in this article.[15]

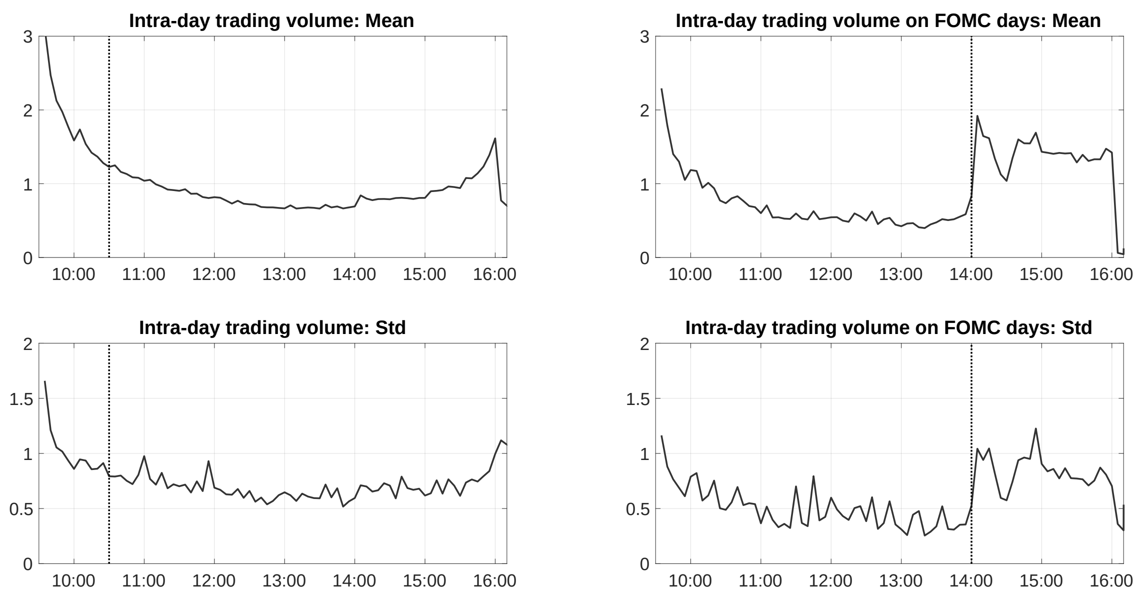

Figure 4: *The figure reports average intra-daily trading volume in ultra-short maturity options. The top panels depict the average standardized trading volume over the entire sample (top left) and during days corresponding to FOMC announcements (top right). Each day, we calculate trading volume within a specific time period (i.e., each 5-minute interval) and divide it by the average trading volume for that day (i.e., the line is centered around one by construction). We report averages across days. The bottom panels display the standard deviation associated with the estimates in the top panels. The time period is January 2, 2014, to May 11, 2023. Data source: CBOE.*

We apply minimal filters to the option data, following the approach outlined in Andersen,

[15] Our choice of 10:30am is also far enough from FOMC announcements and other macroeconomic announcements.

Fusari, and Todorov (2017). Specifically, we exclude options with zero bid prices, options for which the ratio between ask and bid prices exceeds 10 (considered highly illiquid) and options for which we cannot successfully compute the implied volatility. To address potential synchronicity issues between the option and the underlying, we derive the implied underlying forward price from put-call parity. Finally, as customary in the literature, we only retain out-of-the-money calls and puts since they are significantly more liquid than their in-the-money counterparts. This leaves us, each day, with an average of 45 option contracts (Fig. 5) across a broad moneyness range (Fig. 6). The result is a large average number of representative strikes giving us confidence in the reconstruction of the implied volatility surface, in our inference on the characteristics, and in the pricing/hedging evaluation.

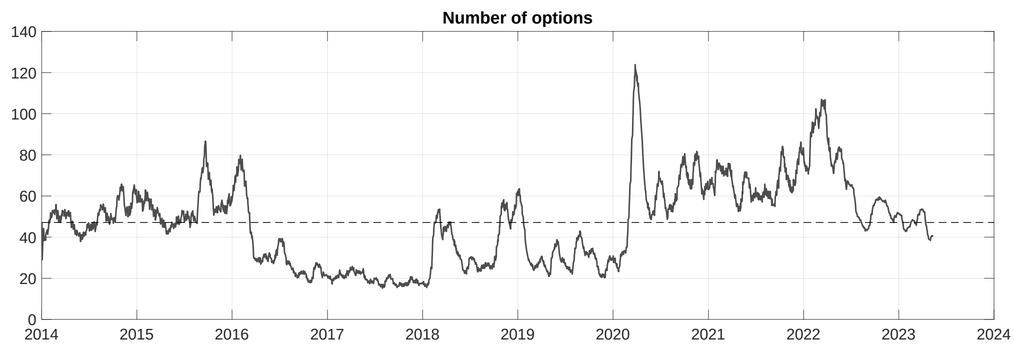

Figure 5: *The figure reports monthly averages of the daily number of options used for estimation and pricing. The time period is January 2, 2014, to May 11, 2023. Data source: CBOE.*

One observation is in order. During the first part of the sample, taking advantage of the shortest maturity may *still* not lead - for certain days - to the use of cross sections of 0DTEs. This is due to the absence (until May 11 2022) of continuous daily expiries. As indicated earlier, May 11 2022 is the day during which Thursday expiries were introduced, thereby completing the CBOE's menu of daily expiries. Because our data ends on May 11 2023, the last year in the data only uses 0DTEs. The next best alternative - to 0DTEs, when not available - is, instead, employed in previous years.

To illustrate this point, Fig. 7 reports the tenors used for estimation over the years, the shortest available tenor being invariably our choice. The vertical dashed line corresponds to the dates during which specific weeklies (i.e., weeklies expiring on specific days) were

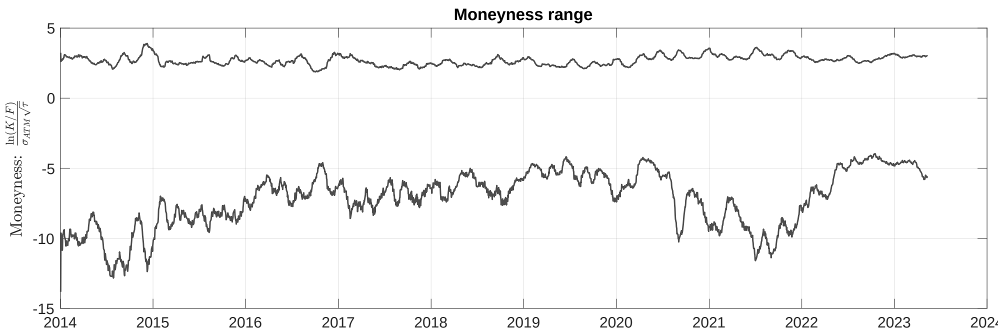

Figure 6: *The figure reports monthly averages of the standardized (by volatility) daily log-moneyness ranges. The time period is January 2, 2014, to May 11, 2023. Data source: CBOE.*

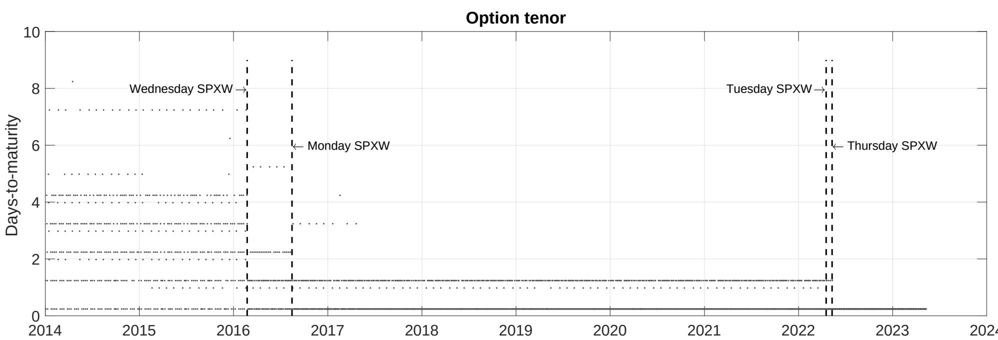

Figure 7: *The figure reports the options' tenors used for estimation and pricing. The time period is January 2, 2014, to May 11, 2023. Data source: CBOE.*

introduced. The horizontal dense points refer to options expiring at 4pm (weeklies). The horizontal sparse points refer, instead, to options expiring at 9am (these are regular monthly SPX options which settle on the third Friday of each month at 9am). Thus, at 10:30am of every trading day, i.e., the time of estimation, the shortest available tenor is the one of a weekly option expiring on the same day at 4pm, if that option exists. Alternatively, it is the one of a monthly option expiring at 9am on the subsequent trading day. In the absence of the latter, it is the one of a weekly option expiring at 4pm of the subsequent day, and so on. Needless to say, the introduction by the CBOE of a new weekly (each vertical line) is always accompanied by a deepening of the market and a reduction in the length of the available tenor. In the end, as made clear by the graph, we sometimes had to resort to maximum tenors of 7 days (at the beginning of the sample) but could, very frequently, use intra-daily tenors of 5.5 hours throughout the sample. These tenors were invariably used beginning on May 11, 2023. The increased liquidity associated with ultra-short-tenor options is the institutional feature which we exploit for estimation and pricing *within* the trading day.

5.3 The equity characteristics

Before turning to pricing, we report on the main characteristics, namely σ_t, $\widetilde{\beta}_t$, and ρ_t. Because large literatures have focused on their identification (generally one at a time), visualizing their dynamics is of separate interest. We note, also, that if our emphasis were on inference per se, not on pricing, our methods could be optimized. For instance, one could use the richness of CBOE intra-daily data, exploit multiple cross sections, and local average the estimates in order to enhance efficiency.

Fig. 8 reports the estimated (main) characteristics over time. Spot volatility behaves as expected, with a considerable spike in March 2020 due to the pandemic. Spot leverage is negative and hovers around -0.4. The correlation between spot volatility and spot leverage is negative and around -0.15, confirming previous findings (Bandi and Renò, 2012). Spot volatility-of-volatility is less understood in its dynamics but the reported increasing trend is consistent with the behavior of the VVIX during the time frame over which the VVIX is available. Given our results, the positive trend in the VVIX may, therefore, not be entirely attributable to increased risk premia.

6 Pricing performance

We now turn to pricing. We visualize it in Fig. 9. The figure reports the average (across days) smoothed empirical implied volatilities as compared to implied volatilities delivered by two models: the one we propose (described above) and a model which has been successful in pricing short-tenor instruments, namely the one suggested in Andersen, Fusari, and Todorov (2017). The horizontal axis is log-moneyness in units of implied volatility. Hence, a value of -4 and above, for instance, would define deep out-of-the-money puts. A value between -2 and -4 would represent out-of-the-money puts, and so on.

The model in Andersen, Fusari, and Todorov (2017) assumes parametric jumps and con-

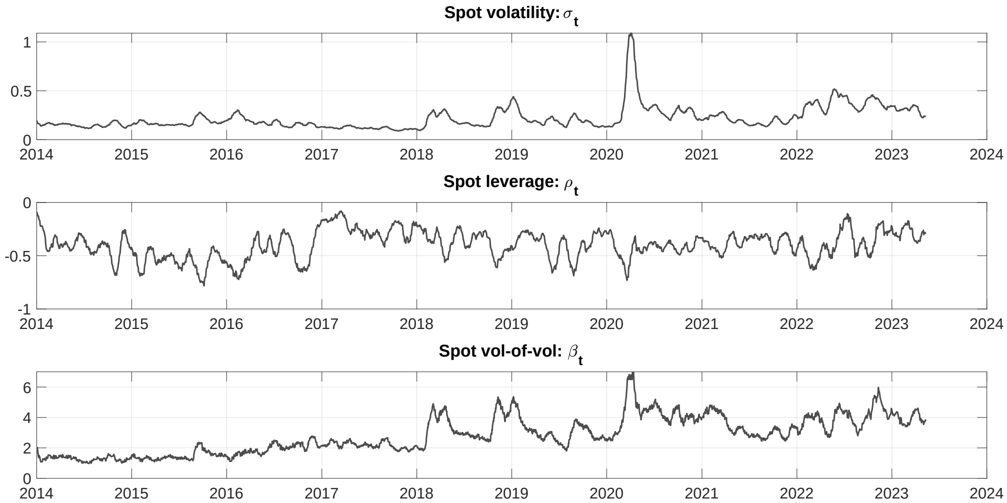

Figure 8: *The figure reports monthly averages of daily estimates of σ_t, ρ_t and $\widetilde{\beta}_t$. The time period is January 2, 2014, to May 11, 2023. Data source: CBOE.*

ditionally (on the time-varying volatility process) Gaussian diffusive dynamics. We dub it BSM, as in "Black-Scholes-Merton," but with the understanding that the spot volatility process is treated as an unrestricted process. Interestingly for the purposes of this comparison, the BSM model can be interpreted in two equivalent ways. It is a restricted version of the specification we propose with all of the diffusive parameters (in particular, ρ_t and $\widetilde{\beta}_t$) set to zero, with the exception of σ_t. It is also a model, otherwise equivalent to ours, in which one dispenses with the volatility dynamics, as captured - in particular - by ρ_t and $\widetilde{\beta}_t$. The latter observation will aid interpretation.

Fig. 9 shows that both models are effective in capturing deep out-of-the-money behavior. This is unsurprising, given that the focus in Andersen, Fusari, and Todorov (2017) is explicitly on tail dynamics and they are successful in capturing them. Because we assume the same jump specification (initially, conditional Gaussian jumps), we are expecting to perform similarly in deep out-of-the-money regions of the log moneyness range. While the jump specification could be modified across models (we do so in Section 7), allowing the intensity of the jumps and the moments of the jump size distribution to be time-varying (something that both Andersen, Fusari, and Todorov, 2017, and we do) renders conditional Gaussian jumps rather flexible.

Fig. 9 also shows that the model we propose is especially effective in capturing dynamics around at-the-money. Because the convexity (res. slope) of the implied volatility surface around at-the-money depends intimately on the volatility-of-volatility (resp. leverage) adding kurtosis to the conditional return distribution (through β_t and other quantities) and skewness (through ρ_t) justifies the superior fit. This is where our proposed (small τ) expansion is naturally suited to price ultra short-tenor options. The expansion skews the return distribution and fattens its tails. It does so in a granular fashion, i.e., by shifting probability mass relative to an otherwise locally Gaussian (because of the driving Brownian motions) distribution. Improvement at-the-money will translate into improvements over the full log-moneyness range, as we document below.

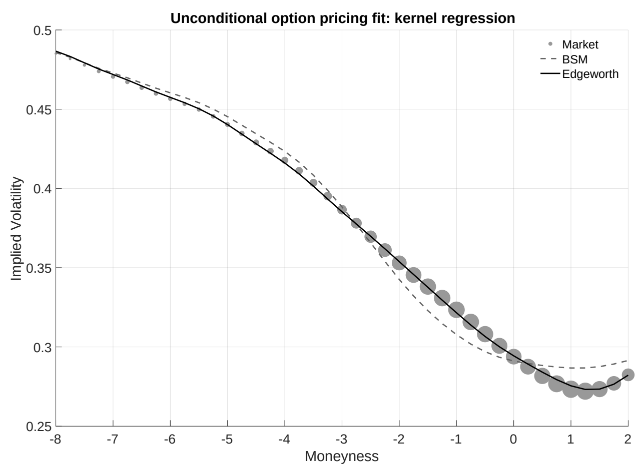

Figure 9: *The figure reports average (across days) smoothed empirical implied volatilities and implied volatilities based on two models (the BSM model in Andersen, Fusari, and Todorov, 2017, and our Edgeworth-like specification). The size of the circles (on the empirical implied volatilities) is proportional to traded volume for each log-moneyness level. The time period is January 2, 2014, to May 11, 2023. Data source: CBOE.*

Turning to a numerical assessment, we compute the RMSE of both models, namely the percentage square root of the average squared distance between the implied volatilities in the data and the implied volatilities delivered by each model. For notational simplicity, we define the model proposed in this article as *Edgeworth*. The metric is:

$$\text{RMSE}(model) = \frac{1}{T}\sum_{t=1}^{T} RMSE_t = \frac{1}{T}\sum_{t=1}^{T}\left[100\sqrt{\frac{1}{N_t}\sum_{i\in\tau_t}(BSIV^{mkt}_{k_i,\tau_t} - BSIV^{model}_{k_i,\tau_t})^2}\right],$$

where *model* is either *BSM* or *Edgeworth*.

We find that RMSE(BSM) $= 1.17$ and RMSE($Edgeworth$) $= 0.48$ In words, *BSM* yields an average error of about 1.2% in volatility terms. As an example, if the average implied volatility is 18%, *BSM* produces volatilities between (roughly) 17% and 19%. *Edgeworth*'s error is about 0.5%, with a gain of 60%. As emphasized previously, the improvement is due to the ability of *Edgeworth* to tilt the short-term conditional return distribution and give it the right degree of skewness and kurtosis.

The improvement, also, suggests that near at-the-money liquidity may be substantial. In turn, substantial near at-the-money liquidity would justify attention to that segment of the log-moneyness range. To this extent, Fig. 9 offers evidence on market liquidity around at-the-money (as well as in other segments of the log-moneyness range) by reporting circles on the empirical implied volatility surface. The size of the circles is proportional to the traded number of contracts. We show that volume is almost symmetrically distributed around at-the-money. Consistent with the 0DTE market being - like the broader option market - a playground for institutional investors, the reported symmetry around at-the-money is the result of speculative/hedging strategies which exploit the center of the log-moneyness range.

RMSE(model) is, of course, an average measure which does not take into account time variation. Next, we report the ratio $\frac{\text{RMSE}_t(Edgeworth)}{\text{RMSE}_t(BSM)}$ for every day in the sample (Fig. 10). The gain provided by *Edgeworth* relative to *BSM* fluctuates between about 30% and 65%. There is persistence in relative performance, with no obvious pattern (due to complex interactions between the diffusive characteristics).

The dynamic fit of the model can also be evaluated by looking at a metric which has proved to be challenging for otherwise successful models, such as the best model specifications in Bates (2019). Bates (2019) looks at time variation in a skew measure defined as the difference between the implied volatility of calls and puts for calls and puts one standard deviation out-of-the-money. He finds a "substantial and systematic gap between predicted and observed volatility smirks of less than one month's maturity" and argues that "exploring

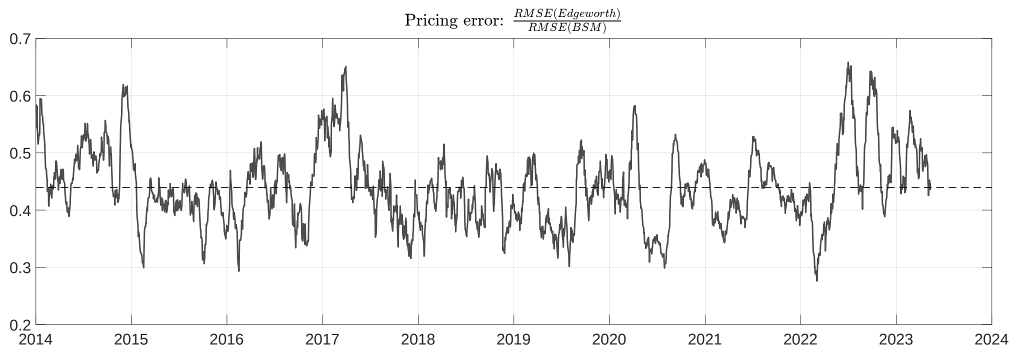

Figure 10: *The figure reports the ratio* $\frac{RMSE_t(Edgeworth)}{RMSE_t(BSM)}$ *for every day in the sample. The time period is January 2, 2014, to May 11, 2023. Data source: CBOE.*

alternative models of leverage and volatility feedback effects is the most promising area for future research". He concludes by advocating for more work on non-affine specifications.

We employ the same metric, which we visualize in the first panel of Fig. 11. The second panel of Fig. 11 reports, instead, on a second - equally natural - metric, one which focuses on convexity. The latter is constructed as the sum of the OTM put and call implied volatilities used to compute the skewness measure from which we subtract the implied volatility of the option closer to at-the-money (in essence, it is the second derivative of the implied volatility surface at-the-money.) We are able to replicate the time variation in the assumed implied-volatility skew and convexity measures very accurately. As discussed, the proposed expansions are explicitly designed to capture time-varying skewness and kurtosis in log-returns. They do so by freeing up - among other quantities - both leverage (ρ_t) and the volatility-of-volatility ($\widetilde{\beta}_t$) which are treated like unrestricted stochastic processes rather than, e.g., like affine specifications or, even, constant values. In this sense, we address - and support empirically - the call in Bates (2019) for "alternative models of leverage" and a deeper dive into non-affine specifications in both of its dimensions.

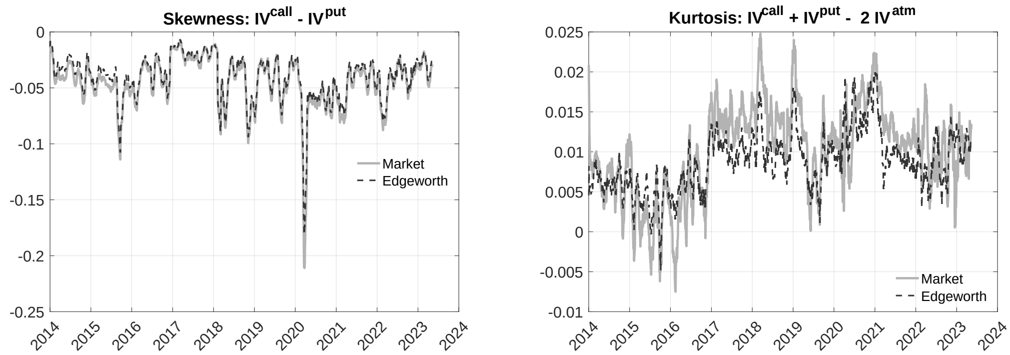

Figure 11: *The first panel reports the difference between the implied volatility of calls and puts for calls and puts one standard deviation out-of-the-money, an implied-volatility skew metric. The second panel looks at the at-the-money implied-volatility convexity. This is constructed as the sum of the OTM put and call implied volatilities used to compute the skew measure in the first panel from which we subtract the implied volatility of the option closer to at-the-money. The time period is January 2, 2014, to May 11, 2023. Data source: CBOE.*

7 Robustness

We now slice the previous results along several alternative dimensions, namely moneyness, volatility states and tenor. Next, we restrict the pricing model by setting to zero certain characteristics. We then price solely near at-the-money. Finally, we modify the jump distribution.

7.1 Pricing as a function of moneyness, volatility and tenor

Panel A of Table 1 reports on *moneyness*. The grey area, in particular, provides the ratio of the RMSEs of the two model specifications. Again, the lower the ratio, the better the performance of the *Edgeworth* model. As expected in light of our previous discussion, the largest improvement is at-the-money, where the *Edgeworth* RMSE is about a quarter of the *BSM* RMSE. Because continuous dynamics have a large impact at-the-money but, of course, affect the entire implied volatility surface (and because we employ the same jump specification as in Andersen, Fusari, and Todorov, 2017), we expect some improvements out-of-the-money too, although less pronounced. The numbers are consistent with this logic. Not only do they document strong performance for a model specification which skews the local return

Panel A: Moneyness					
	DOTMP	OTMP	ATM	OTMC	DOTMC
BSM	1.08	0.96	1.04	0.60	0.39
EDG	0.49	0.38	0.29	0.43	0.39
EDG/BSM	0.46	0.39	0.28	0.71	1.00
Panel B: Volatility					
	Q05	Q25	Q50	Q75	Q95
BSM	0.75	0.87	1.09	1.36	1.71
EDG	0.35	0.37	0.43	0.49	0.69
EDG/BSM	0.46	0.43	0.40	0.36	0.41
Panel C: Tenor					
	0D	1D	2D	3D	5D
BSM	1.04	1.28	1.27	1.33	1.40
EDG	0.45	0.48	0.53	0.57	0.79
EDG/BSM	0.43	0.37	0.42	0.43	0.56

Table 1: *Pricing performance.* Panel A reports the median RMSE for options with different moneyness level. Moneyness is defined as $m = \frac{\ln(K/F)}{IV_{\text{ATM}}\sqrt{\tau}}$. Specifically, DOTMP and OTMP represent deep out-of-the-money ($m < -4$) and out-of-the-money ($-4 < m < -2$) put options. ATM denotes at-the-money options ($-2 < m < 2$). OTMC and DOTMC represent out-of-the-money ($2 < m < 4$) and deep out-of-the-money ($m > 4$) call options. Panel B and C report the average RMSE over days with different levels of the ATM implied volatility (Q5, Q25, Q50, Q75 and Q95 correspond to the 5th, 25th, 50th, 75th and 95th quantile of the at-the-money implied volatility) and different option tenors (in days). The time period is January 2, 2014, to May 11, 2023. Data source: CBOE.

distribution and fattens its tails across moneyness levels, they also document *monotonically* increasing performance as one moves from the out-of-the-money ranges (whether for a call or a put) to the in-the-money range.

Panel B of Table 1 reports on performance across *volatility states.* In particular, we condition on implied volatility levels and compute the RMSEs during days in which implied volatility corresponds to the 5th, the 25th, the 50th, the 75th, and the 95th quantile of the at-the-money implied volatility distribution. A priori, we would want a successful model specification to do well across states. While both models depend on time-varying volatility and should adapt nicely to its levels, the distinguishing feature of the *Edgeworth* model is its ability to tilt the local return distribution through time-varying characteristics of the volatility process. This ability should depend in significant ways on the volatility level *only* to the extent that, e.g., the volatility-of-volatility and leverage are very highly correlated with volatility. So, while we expect some variability across volatility states, we also expect this variability not to be substantial. The numbers confirm this intuition. We report improvements hovering between 64% (when implied volatility is around its 75th quintile) and 54% (when implied volatility is around its 5th quintile).

Panel C of Table 1 is about *tenor*. The *Edgeworth* model leads to large improvements across maturities, the largest improvements being however associated with the shortest tenors (from 0DTEs to options with 3 days to maturity). This result is, again, in line with the local (i.e., short tenor) nature of the proposed expansion(s).

7.2 Restricting the expansion

We impose a natural restriction. We only consider the skewness adjustment and, therefore, dispense with all terms of order higher than $\sqrt{\tau}$. By doing so, we may evaluate the relative role played by skewness and kurtosis in the proposed expansion(s). We emphasize that the model is re-estimated.

The restriction leads to

$$\mathbb{C}^{\frac{\log X^c}{\sigma_t\sqrt{\tau}}}(u,\tau) \quad = \quad e^{iu\frac{\widetilde{\mu}_t\tau}{\sigma_t\sqrt{\tau}}-\frac{u^2}{2}}\left(1-\underbrace{iu^3\frac{\widetilde{\beta}_t\rho_t}{2\sigma_t}\sqrt{\tau}}_{\text{skewness adjustment}}\right).$$

We note that, in this case, $\widetilde{\beta}_t$ and ρ_t cannot be identified separately by minimizing the criterion in Eq. (10). Minimization of Eq. (10), however, leads to estimates of the product $\widetilde{\beta}_t\rho_t$. We find that this product correlates strongly (over 90% correlation) with the product of the same quantities obtained from the full minimization discussed earlier. The level of the estimates of the product is also very similar to the level of the product of the estimates.

More importantly for our purposes, the RMSE of this asymptotically first-order expansion is 0.75. Consistent with this figure, adjusting for skewness is of first-order importance and leads to a RMSE which is 63% that of the *BSM* model.

Having made this point, adding a kurtosis adjustment - as done in the second-order expansion in Eq. (3) - frees up the impact of $\widetilde{\beta}_t$ and ρ_t and improves fit further. The result is a RMSE ratio between *Edgeworth* and *BSM* of about 41% and, therefore, a substantial contribution of the kurtosis adjustment.

7.3 At-the-money pricing

We now focus on the at-the-money (or near at-the-money) range and consider log-moneyness levels between -2 and +2. Importantly, we do pricing by *only* including the continuous portion of the process and setting the jumps equal to zero. Generally speaking, this is intended to evaluate if at-the-money pricing can be conducted successfully by carefully tailoring the diffusive dynamics while dispensing with jump dynamics. More specifically, it is intended to study the relative role of ρ_t (and its tilt to skewness) and $\widetilde{\beta}_t$ (and its tilt to kurtosis) from another vantage point.

Consistent with previous results, we find that the expansion is successful in pricing at-the-money with an RMSE of 0.24. Renouncing the tilts provided by ρ_t and $\widetilde{\beta}_t$, i.e., using the conditionally Gaussian diffusive dynamics in BSM, yields an RMSE of 0.64, a value which is 2.7 times higher. We observe that BSM is implemented with price discontinuities. Hence, the comparison should be viewed as favoring BSM relative to the diffusive-only *Edgeworth* specification in this subsection. In spite of BSM's flexibility, *Edgeworth*'s relative performance continues to be very satisfactory.

Regarding inference on the characteristics, we note that the characteristics themselves may be identified in a fully nonparametric way. We may, in fact, work within log-moneyness levels between -2 and 2 without having to specify the jump size distribution. When doing so, we achieve high correlations between the nonparametric estimates in this subsection and the semi-parametric estimates in Section 4. Specifically, the correlations are 0.98, 0.8 and 0.95 for σ_t, ρ_t and $\widetilde{\beta}_t$, respectively. The levels are, also, very close.

7.4 The jump distribution

Assuming Gaussian jump sizes is a natural first step consistent with a rather large portion of the literature. In order to focus on diffusive dynamics, Gaussianity was - therefore - our initial choice both for BSM and for *Edgeworth*.

We now turn to the alternative (and preferred) tempered stable specification in Andersen,

	$\nu = -1$	$\nu = 0$	$\nu = 0.5$
BSM	1.08	0.93	0.90
Edgeworth	0.49	0.44	0.44
Edgeworth/BSM	0.45	0.47	0.49

Table 2: *Tempered stable jumps.* The table compares the RMSE of the conditional Gaussian model with tempered stable jumps in Andersen, Fusari, and Todorov (2017) with the *Edgeworth* model with tempered stable jumps. We set the parameter ν equal to three values associated with well-known specifications in the literature. The time period is January 2, 2014, to May 11, 2023. Data source: CBOE.

Fusari, and Todorov (2017), which we adopt for both models. Write the jump compensator as

$$\lambda_t \left\{ \frac{e^{-\psi^- |x|}}{|x|^{1+\nu}} 1_{\{x<0\}} + \frac{e^{-\psi^+ |x|}}{|x|^{1+\nu}} 1_{\{x>0\}} \right\} dtdx,$$

with $\nu < 2$. For $\nu < 0$, the jumps are of finite activity. If $\nu \in [0, 1)$, they are of infinity activity but finite variation. If $\nu \in [1, 2)$, they are of infinite variation. Popular models are sub-cases of this specification. When $\nu = -1$, one obtains the double-exponential model in Kou (2002). The case $\nu = 0$ corresponds to the variance gamma model of Madan, Carr, and Chang (1998).

Consistent with the findings in Andersen, Fusari, and Todorov (2017), this more flexible jump distribution (added to conditional Gaussian diffusive dynamics) is better able to capture out-of-the-money effects jointly with the at-the-money skew. The resulting RMSEs of *BSM* are, in fact, equal to 1.07 ($\nu = -1$), 0.93 ($\nu = 0$) and 0.90 ($\nu = 0.5$) and lower than in the case of Gaussian jumps (c.f. Table 2).

Because the fit of the *Edgeworth* model with Gaussian jumps is already satisfactory, while we anticipate some improvements due to the new jump specification, we do not expect the improvements to be quite as large as in a model with conditionally Gaussian diffusive dynamics, like *BSM*. Consistent with this intuition, the new *Edgeworth* (with tempered stable jumps) RMSEs are 0.49 ($\nu = -1$), 0.44 ($\nu = 0$) and 0.44 ($\nu = 0.5$), c.f. Table 2.

We conclude this analysis with the following observations. On the one hand, in Section 4 we have shown that Gaussian jumps are not excessively misspecified in a model in which the local return distribution is properly tilted by $\widetilde{\beta}_t$ and ρ_t. On the other hand, we have documented in this subsection that, even though a tempered stable specification may compensate somewhat for misspecified continuous dynamics, suitable tilting of the local return distribution continues to be beneficial. Consistent with these remarks, a *Edgeworth* model with tempered stable jumps performs extremely well.

8 Hedging

Options are routinely (delta-)hedged by the sell side. 0DTE options, however, face large "convexity" or gamma risk. As pointed out in the Introduction, this risk has been recently invoked as a potential cause of instability loops stemming from price changes in the underlying leading to large repositioning (due to hedging motives) which may, in turn, yield more volatility in the price of the underlying.

Our goal in this section is not to dispute or support the plausibility of these adverse loops, something which is admittedly of interest to regulators but would require a better institutional understanding of this new, fast-evolving, market. In agreement with our pricing comparisons, our objective is to evaluate the hedging abilities of our proposed model relative to alternative specifications. Effective hedging may, of course, lead to large price impacts for the underlying (or other options). However, the cash flow stability that effective hedging guarantees should not just be of interest to specialists. It should also be of interest to regulators, particularly when large option positions are taken by central institutions.

We consider both delta (Δ) and gamma (Γ) hedging. We report median profits and losses (P&Ls). Regarding Δ hedging:

$$\begin{aligned} P\&L_j^M(\Delta^M) = \quad & O_{j,t=10:30am} - \Delta^M_{j,t=10:30am} \times \widetilde{X}_{t=10:30am} \\ - \quad & O_{j,t=4:00pm} + \Delta^M_{j,t=10:30am} \times \widetilde{X}_{t=4:00pm}, \end{aligned}$$

where $O_{j,t}$ is the price of a specific option contract at time t, $\Delta^M_{j,t}$ is the option's Δ position implemented at time t based on the pricing model denoted by M and $\widetilde{X}_t$ is the price of the underlying at time t.

We work with three models: Black and Scholes (BS), BSM and *Edgeworth*. Both BSM and *Edgeworth* have Gaussian jumps. Consistent with the pricing evaluations in Sections 4 and 7, a hedging comparison with BSM is natural. In addition, we consider BS because it is, to this day, commonly employed in the industry. In industry practice, however, the constant BS volatility is replaced by the implied volatility for each specific level of log-moneyness, an adjustment that we also make. Barring this adjustment, BS may, of course, be viewed as another restricted version of *Edgeworth*, one in which $\rho_t = 0$ and $\widetilde{\beta}_t = 0$ (as in BSM) but, also, $\lambda_t = 0$.

Regarding Γ hedging, we compute a Γ-neutral portfolio before rendering it Δ-neutral. In order to Γ hedge OTM puts (resp. calls), we use the option with moneyness that is closest to $m = -0.5$ (resp. $m = 0.5$).[16] We denote by a subscript k the option that is used to perform Γ hedging (for a generic option j). In this case, the criterion is the median of the following P&Ls:

$$
\begin{aligned}
P\&L_j^M(\Gamma^M) = \quad & O_{j,t=10:30am} - \Gamma^M_{j,t=10:30am} \times O_{k,t=10:30am} \\
- \quad & (\Delta^M_{j,t=10:30am} - \Gamma^M_{j,t=10:30am} \times \Delta^M_{k,t=10:30am}) \times \widetilde{X}_{t=10:30am} \\
- \quad & O_{j,t=4:00pm} - \Gamma^M_{j,t=10:30am} \times O_{k,t=4:00pm} \\
+ \quad & (\Delta^M_{j,t=10:30am} - \Gamma^M_{j,t=10:30am} \times \Delta^M_{k,t=10:30am}) \times \widetilde{X}_{t=4:00pm}.
\end{aligned}
$$

Table 3 contains hedging results for several moneyness levels (first column). Columns 3 and 4 refer to Δ hedging. Columns 5 and 6 refer to Γ hedging. Needless to say, the best performing method should deliver the smallest average P&L across log-moneyness levels. In this sense, the reported findings are rather consistent. First, *Edgeworth* generally outperforms the two competing methods in both out-of-the-money directions. The near at-the-money range is, instead, one in which all methods perform similarly. Second, *Edgeworth*'s outperformance is shown to be economically and, generally, statistically significant across several moneyness levels. In situations in which *Edgeworth* does not outperform, the resulting figures are close across models, with a difference which is typically insignificant. Third, as expected, *BSM* performs better than *BS*, and more similarly to *Edgeworth*, particularly for deep out-of-the-money levels.

It is interesting to notice that *Edgeworth*'s hedging performance is complementary to

[16] For OTM puts (resp. calls) with $m = -0.5$ (resp. $m = 0.5$) we use $m = -1$ (resp. $m = 1$).

		Δ hedging		Γ hedging		Pricing
Moneyness	Volume	*Edg./BS*	*Edg./BSM*	*Edg./BS*	*Edg./BSM*	RMSE *Edg./BSM*
−5.0	0.005	0.681***	0.997**	0.933***	1.145***	0.480
−4.5	0.009	0.657***	0.952**	0.847***	1.053***	0.462
−4.0	0.013	0.631***	0.895	0.718***	0.914	0.480
−3.5	0.024	0.661***	0.859	0.647***	0.725***	0.514
−3.0	0.022	0.725***	0.882***	0.666***	0.659***	0.503
−2.5	0.041	0.801***	0.903***	0.760***	0.704***	0.428
−2.0	0.069	0.881***	0.934***	0.846***	0.734***	0.361
−1.5	0.094	0.912***	0.970	0.867***	0.779***	0.239
−1.0	0.096	0.965**	1.005	0.890***	0.810*	0.205
−0.5	0.100	0.984*	0.991	1.042***	0.850	0.293
0.0	0.090	1.019	1.028	1.052*	1.007	0.423
0.5	0.080	1.037	1.016	1.026	0.883	0.592
1.0	0.102	1.111	0.991	1.118	0.923	0.397
1.5	0.120	0.933	0.748	0.999*	0.838**	0.256
2.0	0.072	0.599	0.587*	0.686***	0.729***	0.323
2.5	0.036	0.580	0.752	0.509***	0.803***	0.535
3.0	0.027	0.551	0.861	0.492***	0.976***	0.564

Table 3: We consider both Δ-hedging and Γ-hedging (with Δ-hedged option positions). The table reports the ratios between the median P&Ls associated with three models (*Edgeworth*, *BSM* and *BS* with the constant volatility replaced by the corresponding implied volatility) for different levels of log-moneyness. For comparison with hedging performance (across log-moneyness levels), the last column reports pricing results, i.e., the ratios between the RMSEs of *Edgeworth* and *BSM*. *, ** and *** denote significance at the 10%, 5% and 1% level, respectively. The time period is January 2, 2014, to May 11, 2023. Data source: CBOE.

its pricing performance. In terms of pricing (res. hedging), *Edgeworth* has been shown to be effective across log-moneyness levels and, in particular, at-the-money (resp. out-of-the-money). Effective pricing, of course, hinges on accurate option prices. Hedging, on the other hand, relies on accurate *first* and *second derivatives* of option prices. In this sense, pricing is a static metric, hedging is dynamic. Under both a static and a dynamic metric, *Edgeworth* performs satisfactorily.

In order to better understand the role that greeks obtained from alternative models play across moneyness levels, Fig. 12 reports (in the upper left and right panels) the Δs and Γs of the three pricing models. In the lower left and right panels, instead, the figure reports the ratios between the *Edgeworth*'s greeks and the corresponding greeks obtained from the alternative specifications. Consistent with the findings in Table 3, the *Edgeworth*'s greeks diverge from the BS and the BSM greeks in the out-of-the-money range. While divergence from the latter is relatively more limited (albeit economically and statistically significant, c.f. Table 3), divergence from the former is large.

This is particularly true in the case of Γ. The upper right panel of Fig. 12 justifies this result. Because of $\widetilde{\beta}_t$ and ρ_t, i.e., the characteristics that are central to the price expansion, the *Edgeworth*'s Γ is peaked and left skewed. Fig. 13 provides an illustration. Increases in the characteristic $\widetilde{\beta}_t$ (treated here as a constant parameter, for illustration) yield increases in Γ. More negative values of the characteristic ρ_t (treated here, again, as a constant parameter) make, instead, Γ more left skewed. Combining the estimated average $\widetilde{\beta}_t$ and the estimated average ρ_t in the data (along with other model characteristics) leads to the graph in the upper right corner in Fig. 12 and the reported divergence between Γ values across models.

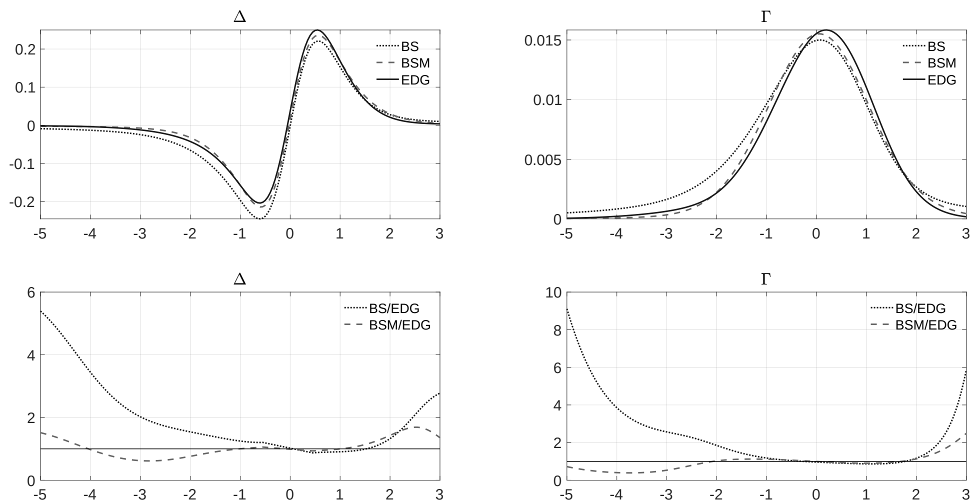

Figure 12: *The figure reports average (across days) smoothed* Δ*s and* Γ*s for three models: Edgeworth, BSM and BS. The top panels report absolute values, the bottom panels report ratios. The time period is January 2, 2014, to May 11, 2023. Data source: CBOE.*

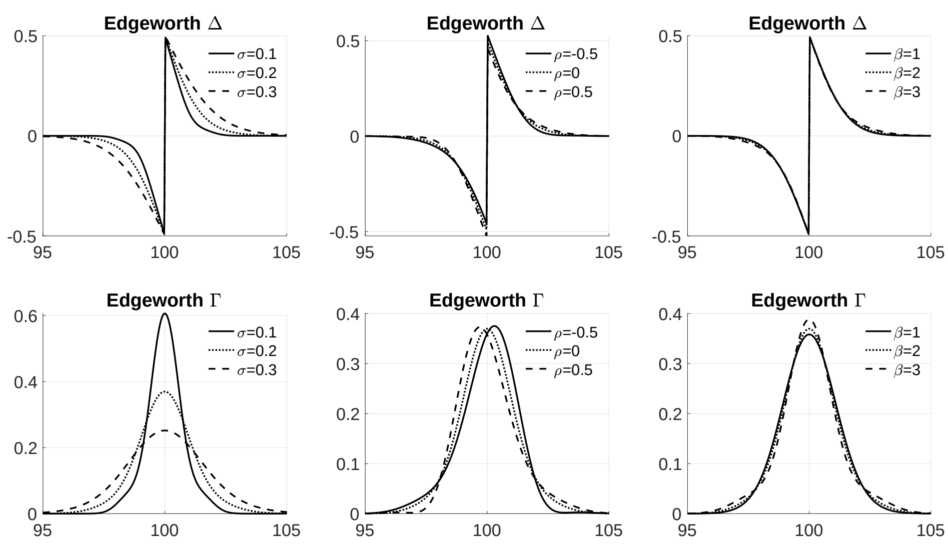

Figure 13: *The figure reports* Δ *and* Γ *for options with 1 day to maturity implied by the Edgeworth model with different values of the equity characteristics* σ_t, ρ_t, *and* $\widetilde{\beta}_t$. *The intensity* λ_t *is set to zero.*

9 Pricing using the second specification

We now turn to the full expansion in Subsection 4.2. Before discussing our findings, we emphasize that this specification is, understandably, computationally more intensive than the baseline specification. Differently from the baseline specification, in fact, the full expansion entails the computation of a bivariate and a trivariate integral, something which inevitably increases the computational complexity of the problem. For Fourier pricing, we therefore depend on a recently developed method (called SINC) introduced by Baschetti, Bormetti, Romagnoli, and Rossi (2022). The method offers three benefits: 1) it allows pricing of an entire cross-section of options with the same tenor and different strikes, 2) it is significantly faster than successful, competing methods (e.g., the COS procedure of Fang and Oosterlee, 2009) and 3) in our experiments, it has demonstrated superior (over competing methods) accuracy in pricing ultra short-tenor options, the subject of this article.

We begin with a restricted case in which we dispense with the volatility jumps (i.e., Eq. (7) and Eq. (8) are removed). This restricted case is, of course, immediately comparable to the baseline specification. Rather than writing the exact characteristic function of the jumps in prices, we expand the same in the tenor τ, for small τ. The expansion, therefore, simply amounts to ignoring terms of order higher than τ in the evaluation of $\mathbb{C}^{\log X_t^d}(u,\tau)$.

Two observations are in order. First, the expansion of the jump characteristic function inevitably requires the choice of a time-t jump measure which we assume to be conditionally (on the time-varying jump features $\mu_{j,t}$ and $\sigma^2_{j,t}$) Gaussian for comparability with the baseline, conditionally Gaussian, case. In this sense, while the expansion in $\mathbb{C}^{\log X^c}(u,\tau)$, i.e., the continuous portion of the process' characteristic function, is naturally viewed as been genuinely nonparametric, the expansion in $\mathbb{C}^{\log X^d}(u,\tau)$ is logically parametric because of the necessary distributional assumptions on the time-t jump size. Second, more importantly, given consistent distributional assumptions on the jump sizes, for very short tenors like the ones used in this study, the baseline specification and this second (restricted) specification are expected to yield very similar outcomes. We confirm that this is, in fact, the case. The RMSE (0.4918) is very close to that of the baseline specification (0.4754). In addition, all the equity characteristics are virtually identical to those reported in Subsection 5.3. As for the time-varying jump features, the median across days of the jump intensity, jump mean and jump standard deviation are 0.5369, -0.0603 and 0.0703 in the baseline specification.

They are 0.5474, -0.0557 and 0.0655 in the second (restricted) specification.

After establishing that expanding the characteristic function of the price jumps is immaterial for our short tenors, we turn to the full expansion. Our focus is now on the volatility jumps. As emphasized previously, we allow for joint and idiosyncratic contributions of these discontinuities. The result is additional effects on both skewness (through the term in Eq. (7)) and kurtosis (through the term in Eq. (8)).

As in the case of the expansion for the characteristic function of the price jumps, implementation now requires a choice of time-t volatility jump measure, both individually (in Eq. (7)) as well as in a bivariate specification with the price jumps (in Eq. (8)). Consistent with the work of Duffie, Pan, and Singleton (2000) and many others, we model the marginal density of the volatility jump sizes as being exponential. The density of the idiosyncratic price jump sizes is, as earlier, Gaussian with mean $\mu_{j,t}$ and variance $\sigma^2_{j,t}$. The conditional (on realizations of the volatility jump sizes) density of the joint price jump sizes is Gaussian with variance $\sigma^2_{j,t}$ and mean $\mu_{j,t} + \rho_J s_\sigma$, where s_σ is a specific volatility jump size realization and ρ_J is a free parameter. Thus, conditional on zero volatility jump sizes, the idiosyncratic price jumps and the joint price jumps have the same size density.

The variance jumps enter the process' characteristic function to the same order as the price jumps. Their impact on option payoffs is, however, further mediated by the driving Brownian motion of the price process. We are, therefore, expecting their contribution to short-term option valuation to be considerably more muted than that of the price jumps. The RMSE (0.4625) is now only marginally lower than the one delivered by the first (baseline) specification (0.4754). Once more, all of the equity characteristics are virtually identical to those reported in Subsection 5.3. We recall that the first (baseline) specification prices *within* the bid/ask for 80% of our options. It is, therefore, not surprising that adding volatility jumps is not leading to meaningful price improvements. It is, nonetheless, revealing to notice that the first specification was not unnecessarily loading on characteristics which would become less important once the variance jumps are introduced. In fact, the role played by these jumps is marginal in our data. On the one hand, we are not excluding that more constrained (possibly affine) diffusive specifications may attribute a more important role to them. On the other hand, these findings re-emphasize the ability of flexible diffusive dynamics - such as the ones we assume - to yield effective higher-moment adjustments to the return distribution and successful local pricing. We conclude that, relative to our assumed pricing and hedging metrics, a parsimonious and tractable specification - such as our baseline specification - fares extremely satisfactorily. We return to the first specification in the next section and evaluate it from a different angle.

10 0DTE risk premia

In this section we explore the informational content of 0DTEs for ultra short-term predictability. In particular, we focus on nearly instantaneous equity and variance risk premia.

In order to map the $\mathbb{P}$-dynamics of the price process in Eq. (1) into $\mathbb{Q}$-dynamics, we assume that the pricing kernel K_t is of the form:

$$K_t = K_0 \left(\frac{\widetilde{X}_t}{\widetilde{X}_0} \right)^{\phi} e^{\int_0^t \zeta_s ds + \psi(\sigma_t - \sigma_0)}, \tag{11}$$

where $\phi < 0$ and $\psi > 0$ are parameters controlling aversion to price risk and volatility risk, respectively. ζ_t is a stochastic process representing time preferences. A similar non-monotonic (in prices) specification for the pricing kernel has been studied by Christoffersen, Heston, and Jacobs (2013) and Bandi and Renò (2016), among others. In parametric option pricing models, Christoffersen, Heston, and Jacobs (2013) are emphatic about the role of non-monotonicity in leading to superior fit.

The instantaneous return and volatility risk premia are defined as the difference between the risk-neutral drift and the objective drift of the logarithmic price process and the volatility process, respectively. The following equations provide closed-form expressions for both premia given the price dynamics in Eq. (1) and the pricing kernel in Eq. (11). The equations specialize Proposition 8.1 in Bandi and Renò (2016) to our setting and are proved similarly.

In the case of the instantaneous return premium, we have

$$\underbrace{(\mu_t + \lambda_t \mathbb{E}_t \left[e^x - 1 \right])}_{\mathbb{E}_t\left[\frac{dX_t}{X_t dt}\right]} - \underbrace{(r_t - \frac{\sigma_t^2}{2})}_{\mathbb{E}_t^{\mathbb{Q}}\left[\frac{dX_t}{X_t dt}\right]} = -\phi \sigma_t^2 - \psi \rho_t \sigma_t \widetilde{\beta}_t - \lambda_t \mathbb{E}_t \left[(e^{\phi x} - 1)(e^x - 1) \right]. \tag{12}$$

As for the instantaneous volatility premium, we have

$$\alpha_t^{\mathbb{Q}} - \alpha_t = \phi \rho_t \sigma_t \widetilde{\beta}_t + \psi \widetilde{\beta}_t^2, \tag{13}$$

where $\alpha_t^{\mathbb{Q}}$ is the $\mathbb{Q}$-drift of the volatility process. Both are spot equations providing a mapping between time-t equity characteristics and time-t risk compensations. Our objects of interest are the parameters ϕ and ψ of the measure change K_t. We will identify them either from Eq. (12) or from Eq. (13). We will also do it jointly.

The intuition behind Eq. (12) and Eq. (13) is classical. The infinitesimal return risk premium depends on the covariance between returns and sources of variation in the measure change (returns and volatility). The terms σ_t^2 and $\lambda_t \mathbb{E}\left[(e^{\phi x} - 1)(e^x - 1)\right]$ capture the variability of returns (in its diffusive and jump portion) while the term $\rho_t \sigma_t \widetilde{\beta}_t$ captures the covariance between returns and volatility. The infinitesimal volatility risk premium depends, instead, on the covariance between volatility and, again, sources of variation in the measure change (returns and volatility). The term $\widetilde{\beta}_t^2$ captures the variability of volatility while the term $\rho_t \sigma_t \widetilde{\beta}_t$ captures the covariance between returns and volatility. This latter term is common between the two infinitesimal premia. Thus, given the assumed measure change, leverage affects both infinitesimal premia (c.f. Bandi and Renò, 2016, and Cheng, Renault, and Sangrey, 2023).

We adopt the following empirical strategy. In the case of Eq. (12), we approximate the jump portion of the return premium using a small-x Taylor expansion:

$$\lambda_t \mathbb{E}\left[(e^{\phi x} - 1)(e^x - 1)\right] \approx \phi \lambda_t (\mu_{j,t}^2 + \sigma_{j,t}^2).$$

The latter term is analogous to the jump variation of Andersen, Fusari, and Todorov (2017). Given the Taylor expansion, the return premium is now linear in ϕ and ψ. The parameters can, therefore, be estimated by running a linear regression of ultra short-term *excess* returns ($dX_t^{ex}/X_t dt$) on total (diffusive plus jump) variation, i.e., $\sigma_t^2 + \lambda_t(\mu_{j,t}^2 + \sigma_{j,t}^2)$, and the leverage term $\rho_t \sigma_t \widetilde{\beta}_t$. We compute (annualized) excess returns $dX_t^{ex}/X_t dt$ between 10:30 am (the time at which we sample options and estimate the equity characteristics) and 4:00 pm (the expiration time of the 0DTEs). Consistent with Eq. (12), we account for the convexity adjustment.

Turning to Eq. (13), we note that we cannot retrieve $\alpha_t^{\mathbb{Q}}$ from the option panel directly since, given the expansion, only the sum $\alpha_t^{\mathbb{Q}} + \widetilde{\delta}_t$ is identified. We instead use the VIX index. If we define by VIX_t^0 the VIX index computed using 0DTE options, we obtain:

$$
\begin{aligned}
\left(VIX_t^0\right)^2 &= \mathbb{E}_t^{\mathbb{Q}}\left[\frac{1}{T-t}\int_t^T \sigma_s^2 ds\right] \\
&= \sigma_t^2 + \mathbb{E}_t^{\mathbb{Q}}\left[\frac{1}{T-t}\int_t^T (\sigma_s^2 - \sigma_t^2) ds\right] \\
&= \sigma_t^2 + \mathbb{E}_t^{\mathbb{Q}}\left[\frac{1}{T-t}\int_t^T \left(\int_t^s 2\sigma_u\alpha_u + \widetilde{\beta}_u^2\right) duds\right] \\
&\approx \sigma_t^2 + \left[\alpha_t^{\mathbb{Q}}\sigma_t + \frac{1}{2}\widetilde{\beta}_t^2\right](T-t).
\end{aligned} \tag{14}
$$

The penultimate line in Eq. (14) follows from Itô's lemma since

$$
d\sigma_t^2 = (2\sigma_t\alpha_t + \widetilde{\beta}_t^2)dt + 2\sigma_t\left(\beta dW_t + \beta_t' dW_t'\right).
$$

The last line in Eq. (14) is a valid approximation if the equity characteristics are reasonably stable over the horizon $T-t$. The shorter the horizon, the better the approximation. In our case, $T-t$ is only 5.5 hours and, therefore, we expect the approximation to be empirically meaningful. Eq. (14) is used to recover $\alpha_t^{\mathbb{Q}}$ from the estimated equity characteristics and the VIX. As for its objective counterpart, i.e., α_t, we employ an estimate of $\sigma_{T'}$ at $T' > t$ and write

$$
\alpha_t = \frac{\sigma_{T'} - \sigma_t}{T' - t}.
$$

Because the equity characteristics are estimated each day at 10:30 am, we take $t =$ 10:30 am and $T' =$ 10:30 am on the next day.

We run the following regressions, individually and jointly:[17]

$$
\begin{aligned}
dX_t^{ex} + \frac{1}{2}\sigma_t^2 &= a + \phi\left(-\sigma_t^2 - \lambda_t(\mu_{j,t}^2 + \sigma_{j,t}^2)\right) + \psi\left(-\rho_t\sigma_t\widetilde{\beta}_t\right) + \varepsilon_t^r \\
\alpha_t^{\mathbb{Q}} - \alpha_t &= b + \phi\rho_t\sigma_t\widetilde{\beta}_t + \psi\widetilde{\beta}_t^2 + \varepsilon_t^\sigma.
\end{aligned}
$$

Estimation results, along with HAC-corrected t-statistics, are reported in Table 4. As ex-

[17] We use moving averages of estimates of $\alpha_t^{\mathbb{Q}} - \alpha_t$ and $\widetilde{\beta}_t$ with a window of 22 days to reduce the impact of estimation error.

	a	b	ϕ	ψ	$R^2(\%)$
Return premium	0.06	–	6.66	2.09^{***}	1.58
	(0.24)		(1.45)	(2.76)	
Volatility premium	–	7.65	-13.55^{***}	1.88^{***}	21.20
	-	(1.32)	(-6.75)	(3.72)	
Joint estimation	-1.55^{***}	7.85	-12.57^{***}	1.89^{**}	34.81
	(-4.68)	(0.79)	(-6.65)	(2.37)	

Table 4: *Instantaneous risk premia.* We estimate the model in Eq. (12) and (13) individually (by OLS) and jointly (by GMM). *, ** and *** denote significance at the 10%, 5% and 1% level, respectively. The time period is January 2, 2014, to May 11, 2023. Data source: CBOE.

pected, the return premium contains more limited pricing signal than the volatility premium. In the return premium, diffusive and jump variation do not lead to a significant estimate of ϕ, the parameter controlling the traditional risk-return trade-off. The covariance between shocks to prices and shocks to volatility (i.e., the leverage-related term $\rho_t \sigma_t \widetilde{\beta}_t$) is, instead, a significant predictor. The volatility risk premium is driven by both the volatility-of-volatility (which is consistent with the logic in Bollerslev, Tauchen, and Zhou, 2009) and the leverage term. The resulting R^2 is a remarkable 21.2% in a univariate specification (and 34.81% in a joint specification) and much higher than for the return premium (1.58%). Importantly, the parameter estimate of ψ in the two individual regressions are very close to each other in spite of the stark difference in identification strategy: through leverage, in the return equation, and through the volatility-of-volatility, in the volatility equation.

In principle, Eq. (12) requires a $\mathbb{P}$-jump variation on the right-hand side. Because of our identification strategy based on 0DTEs, we are forced to work with a $\mathbb{Q}$-jump variation. We, therefore, re-estimate the model by allowing the partial effect ϕ to change across diffusive and jump component, the parameter ϕ^c being now associated with diffusive variation (σ_t^2) and the parameter ϕ^d being associated with $\mathbb{Q}$-jump variation ($\lambda_t(\mu_{j,t}^2 + \sigma_{j,t}^2)$), c.f. Table 5. Econometrically, this separation readily accommodates a $\mathbb{Q}$-jump variation that is proportional to $\mathbb{P}$-jump variation (without using time-series data to estimate $\mathbb{P}$-jump variation). The associated slope estimate would be an estimate of ϕ^d divided by the proportionality

	a	ϕ^c	ϕ^d	ψ	$R^2(\%)$
Return premium	-0.76^{***}	12.02^{**}	-49.51^{***}	2.51^{***}	2.48
	(-2.67)	(2.42)	(-2.74)	(3.40)	
(without outliers)	-0.91^{***}	7.67	-41.25^{***}	2.64^{***}	1.58
	(-3.26)	(1.53)	(-3.15)	(2.89)	

Table 5: *Infinitesimal return risk premium.* We estimate the model in Eq. (12). *, ** and *** denote significance at the 10%, 5% and 1% level, respectively. The time period is January 2, 2014, to May 11, 2023. The second regression excludes instances in which the annualized volatility is larger than 100%. Data source: CBOE.

factor.

Importantly, the parameter ψ continues to be estimated (through the leverage-related term) in a robust way. Consistent with theory, the parameter ϕ^d is estimated to be negative (like ϕ in the volatility premium equation). The parameter ϕ^c is, instead, estimated to be positive but the removal of days in which annualized spot volatility is larger than 100% would make it statistically insignificant. We conclude that jump variation has the potential to drive price compensations in the return equation (which is consistent with Andersen, Fusari, and Todorov, 2017), the impact of diffusive variation being - once more - rather weak. Economically, the split between ϕ^c and ϕ^d would also justify a stochastic discount factor in which the effect of prices on risk depends on the continuous and discontinuous price components in a differential way. This line of inquiry - and the distinction between the two explanations - is better left for future work explicitly devoted to this subject. Our 0DTEs-based findings are, however, suggestive.

Two observations are in order. First, our results show that the leverage-related term is an effective predictor, irrespective of whether it is contained in the return premium (with weight ψ) or in the volatility premium (with weight ϕ). Using different conceptual frameworks, recent work has highlighted the identification potential of leverage for both of the parameters of the change of measure in Eq. (11), c.f. Bandi and Renò (2016) and Cheng, Renault, and Sangrey (2023). We validate their theory with data. Second, the dependence between predictors and return/volatility premia is, as implied by theory, effectively *instantaneous*. Table 4 is revealing of the fact that the observation of a panel of 0DTE options at 10:30 am facilitates prediction of the subsequent 5.5-hour return, as well as of the volatility premium.

For both of these predictions, a key role is played by the instantaneous covariance between variance and returns, a quantity whose components (σ_t, ρ_t and $\widetilde{\beta}_t$) are central to our proposed Edgeworth-like expansions.

11 Conclusions

Much of the growth in option trading over the last few years has been in the ultra short-tenor segment of the market, that occupied by 0-days-to-expiry or 0DTE options. 0DTEs are viewed by all as an opportunity and by some as a risk.

Our objective is to value 0DTEs, a process which is important for the buy side in order to understand fair values, for the sell side to hedge risks (possibly beyond the 0DTEs enhanced gamma risk) and for regulators to assess systemic implications.

We argue that tilting the conditional distribution of the return process over short horizons is central to the effective pricing (and hedging) of 0DTEs. We achieve tilting by invoking local Edgeworth-like expansions around the Gaussian distribution. The expansions add - locally - skewness (through continuous leverage) and kurtosis (through, e.g., the volatility-of-volatility), the contribution to skewness and kurtosis delivered by volatility jumps being empirically unimportant. The end result is accurate replication of both the implied volatility skew and convexity at-the-money, near at-the-money and beyond.

While focusing on 0DTEs is natural given the growth and large notional value of the 0DTE market, the proposed valuation method is of independent interest. Because it is based on local expansions of the characteristic function of the price process, it is - in fact - generally applicable to the pricing of any financial instrument with short-tenor payoffs.

Of independent interest is also the recovery of the time series of the underlying equity characteristics (σ_t, $\widetilde{\beta}_t$ and ρ_t, *in primis*). Large literatures have studied them (generally, one at a time). Using our approach to pricing, we identify them jointly, i.e., without the need for noisy two-step procedures requiring preliminary estimates of σ_t as an input for the identification of $\widetilde{\beta}_t$ and ρ_t. Importantly, we argued that their recovery does not have to hinge on the entire implied volatility surface. If conducted using near at-the-money options - along with our Edgeworth-like expansion for diffusive dynamics only - such recovery may dispense with a specification of the jump measures.

References

Aït-Sahalia, Y., J. Fan, R. J. Laeven, C. D. Wang, and X. Yang, 2017, "Estimation of the continuous and discontinuous leverage effects," *Journal of the American Statistical Association*, 112(520), 1744–1758.

Aït-Sahalia, Y., J. Fan, and Y. Li, 2013, "The leverage effect puzzle: Disentangling sources of bias at high frequency," *Journal of Financial Economics*, 109, 224–249.

Aït-Sahalia, Y., C. Li, and C. X. Li, 2021, "Implied stochastic volatility models," *The Review of Financial Studies*, 34(1), 394–450.

Andersen, T. G., N. Fusari, and V. Todorov, 2015, "Parametric Inference and Dynamic State Recovery from Option Panels," *Econometrica*, 83, 1081–1145.

——— , 2017, "Short-term market risks implied by weekly options," *The Journal of Finance*, 72(3), 1335–1386.

Bandi, F. M., and R. Renò, 2012, "Time-varying leverage effects," *Journal of Econometrics*, 169(1), 94–113.

Bandi, F. M., and R. Renò, 2016, "Price and volatility co-jumps," *Journal of Financial Economics*, 81(119), 107–146.

——— , 2017, "Local Edgeworth expansions," Working paper.

Bandi, F. M., and R. Renò, 2018, "Nonparametric stochastic volatility," *Econometric Theory*, 34(6), 1207–1255.

——— , 2022, "β in the tails," *Journal of Econometrics*, 227(1), 134–150.

Barndorff-Nielsen, O. E., and A. E. Veraart, 2012, "Stochastic volatility of volatility and variance risk premia," *Journal of Financial Econometrics*, 11(1), 1–46.

Baschetti, F., G. Bormetti, S. Romagnoli, and P. Rossi, 2022, "The SINC way: a fast and accurate approach to Fourier pricing," *Quantitative Finance*, 22(3), 427–446.

Bates, D. S., 1996, "Jumps and stochastic volatility: Exchange rate processes implicit in deutsche mark options," *The Review of Financial Studies*, 9(1), 69–107.

——— , 2000, "Post-'87 crash fears in the S&P 500 futures option market," *Journal of Econometrics*, 94(1-2), 181–238.

——— , 2019, "How crashes develop: intradaily volatility and crash evolution," *The Journal of Finance*, 74(1), 193–238.

Beckmeyer, H., N. Branger, and L. Gayda, 2023, "Retail Traders Love 0DTE Options... But Should They?," *Working paper, available on SSRN.*

Bentata, A., and R. Cont, 2012, "Short-time asymptotics for marginal distributions of semi-martingales," *arXiv preprint arXiv:1202.1302.*

Bibinger, M., N. Hautsch, P. Malec, and M. Reiss, 2019, "Estimating the spot covariation of asset prices: Statistical theory and empirical evidence," *Journal of Business & Economic Statistics*, 37(3), 419–435.

Bollerslev, T., G. Tauchen, and H. Zhou, 2009, "Expected stock returns and variance risk premia," *Review of Financial Studies*, 22(11), 4463–4492.

Broadie, M., M. Chernov, and M. Johannes, 2007, "Model specification and risk premia: Evidence from futures options," *Journal of Finance*, 62(3), 1453–1490.

Brogaard, J., J. Han, and P. Y. Won, 2023, "How does zero-day-to-expiry options Trading affect the volatility of underlying assets?," *Available at SSRN 4426358.*

Bryzgalova, S., A. Pavlova, and T. Sikorskaya, 2022, "Retail trading in options and the rise of the big three wholesalers," *Journal of Finance forthcoming.*

Carr, P., H. Geman, D. B. Madan, and M. Yor, 2003, "Stochastic volatility for Lévy processes," *Mathematical finance*, 13(3), 345–382.

Carr, P., and L. Wu, 2003, "What type of process underlies options? A simple robust test," *The Journal of Finance*, 58(6), 2581–2610.

Cheng, X., E. Renault, and P. Sangrey, 2023, “Identifying volatility risk price through leverage effect,” Working paper.

Chernov, M., A. R. Gallant, E. Ghysels, and G. Tauchen, 2003, “Alternative models for stock price dynamics,” *Journal of Econometrics*, 116(1-2), 225–257.

Chong, C. H., and V. Todorov, 2023, “Volatility of volatility and leverage effect from options,” Working paper.

Christoffersen, P., S. Heston, and K. Jacobs, 2013, “Capturing Option Anomalies with a Variance-Dependent Pricing Kernel,” *Review of Financial Studies*, 16(8), 1963–2006.

Corrado, C. J., and T. Su, 1996, “Skewness and kurtosis in S&P 500 index returns implied by option prices,” *Journal of Financial research*, 19(2), 175–192.

Dim, C., B. Eraker, and G. Vilkov, 2024, “0DTEs: Trading Gamma Risk and Volatility Propagation,” *Available at SSRN*.

Duffie, D., J. Pan, and K. Singleton, 2000, “Transform analysis and asset pricing for affine jump-diffusions,” *Econometrica*, 68, 1343–1376.

Durrleman, V., 2004, *From implied to spot volatilities*. Princeton University.

Eraker, B., 2004, “Do stock prices and volatility jump? Reconciling evidence from spot and option prices,” *Journal of Finance*, 59, 1367–1404.

Fan, J., and Y. Wang, 2008, “Spot volatility estimation for high-frequency data,” *Statistics and Its Interface*, 1, 279–288.

Fang, F., and C. W. Oosterlee, 2009, “A novel pricing method for European options based on Fourier-cosine series expansions,” *SIAM Journal on Scientific Computing*, 31(2), 826–848.

Forde, M., and A. Jacquier, 2011, “The large-maturity smile for the Heston model,” *Finance and Stochastics*, 15(4), 755–780.

Forde, M., A. Jacquier, and R. Lee, 2012, “The small-time smile and term structure of implied volatility under the Heston model,” *SIAM Journal on Financial Mathematics*, 3(1), 690–708.

Gatheral, J., E. P. Hsu, P. Laurence, C. Ouyang, and T.-H. Wang, 2012, “Asymptotics of implied volatility in local volatility models,” *Mathematical Finance*, 22(4), 591–620.

Heston, S. L., 1993, “A closed-form solution for options with stochastic volatility with applications to bond and currency options,” *Review of Financial Studies*, 6(2), 327–343.

Jacquier, A., and M. Lorig, 2015, “From characteristic functions to implied volatility expansions,” *Advances in Applied Probability*, 47(3), 837–857.

Jarrow, R., and A. Rudd, 1982, “Approximate option valuation for arbitrary stochastic processes,” *Journal of financial Economics*, 10(3), 347–369.

Kalnina, I., and D. Xiu, 2017, “Nonparametric estimation of the leverage effect: A trade-off between robustness and efficiency,” *Journal of the American Statistical Association*, 112(517), 384–396.

Kou, S. G., 2002, “A jump-diffusion model for option pricing,” *Management science*, 48(8), 1086–1101.

Kristensen, D., 2010, “Nonparametric filtering of the realised spot volatility: a kernel-based approach,” *Econometric Theory*, 26, 60–93.

Kunitomo, N., and A. Takahashi, 2001, “The asymptotic expansion approach to the valuation of interest rate contingent claims,” *Mathematical Finance*, 11(1), 117–151.

Lee, R. W., 2001, “Implied and local volatilities under stochastic volatility,” *International Journal of Theoretical and Applied Finance*, 4(01), 45–89.

Lorig, M., S. Pagliarani, and A. Pascucci, 2017, “Explicit implied volatilities for multifactor local-stochastic volatility models,” *Mathematical Finance*, 27(3), 926–960.

Madan, D. B., P. P. Carr, and E. C. Chang, 1998, “The variance gamma process and option pricing,” *Review of Finance*, 2(1), 79–105.

Mancini, C., V. Mattiussi, and R. Renò, 2015, “Spot volatility estimation using delta sequences,” *Finance and Stochastics*, 19(2), 261–293.

Medvedev, A., and O. Scaillet, 2007, "Approximation and calibration of short-term implied volatilities under jump-diffusion stochastic volatility," *The Review of Financial Studies*, 20(2), 427–459.

Mykland, P., and L. Zhang, 2008, "Inference for volatility-type objects and implications for hedging," *Statistics and Its Interface*, 1, 255–278.

Pan, J., 2002, "The jump-risk premia implicit in options: Evidence from an integrated time-series study," *Journal of financial economics*, 63(1), 3–50.

Sanfelici, S., I. V. Curato, and M. E. Mancino, 2015, "High-frequency volatility of volatility estimation free from spot volatility estimates," *Quantitative Finance*, 15(8), 1331–1345.

Sircar, K. R., and G. C. Papanicolaou, 1999, "Stochastic volatility, smile & asymptotics," *Applied Mathematical Finance*, 6, 107–145.

Takahashi, A., and T. Yamada, 2012, "An asymptotic expansion with push-down of Malliavin weights," *SIAM Journal on Financial Mathematics*, 3(1), 95–136.

Todorov, V., 2019, "Nonparametric spot volatility from options," *The Annals of Applied Probability*, 29(6), 3590–3636.

——— , 2021, "Higher-order small time asymptotic expansion of Itô semimartingale characteristic function with application to estimation of leverage from options," *Stochastic Processes and their Applications*, 142, 671–705.

Vetter, M., et al., 2015, "Estimation of integrated volatility of volatility with applications to goodness-of-fit testing," *Bernoulli*, 21(4), 2393–2418.

Vilkov, G., 2023, "0DTE Trading Rules," *Available at SSRN*.

Wang, C. D., and P. A. Mykland, 2014, "The estimation of leverage effect with high-frequency data," *Journal of the American Statistical Association*, 109(505), 197–215.

Wang, C. D., P. A. Mykland, and L. Zhang, 2017, "Estimating and forecasting volatility using leverage effect," *Available at SSRN 3083253*.

Xiu, D., 2014, "Hermite polynomial based expansion of European option prices," *Journal of Econometrics*, 179(2), 158–177.

Zu, Y., and H. P. Boswijk, 2014, "Estimating spot volatility with high-frequency financial data," *Journal of Econometrics*, 181(2), 117–135.

A Pricing

Consider a call option price, but a similar argument applies to puts:

$$\begin{aligned} C_{K,\tau} &= \mathbb{E}^{\mathbb{Q}}[\max(X_\tau - K, 0)] \exp^{-r\tau} \\ &= \left[\int_K^\infty (X_\tau - K) p^{\mathbb{Q}}(X_\tau) dX_\tau\right] \exp^{-r\tau}, \end{aligned}$$

where $p^{\mathbb{Q}}(X_T)$ is the risk-neutral density of the price process. Now, write the de-meaned standardized log price process as

$$Z_\tau = \frac{\ln(X_\tau) - \ln(X_t) - \left(r_t - \frac{1}{2}\sigma_t^2\right)\tau}{\sigma_t\sqrt{\tau}} = \frac{\ln(X_\tau) - a_\tau}{b_\tau}.$$

Note that $\ln(X_\tau) = a_\tau + b_\tau Z_\tau$ and, of course, $X_\tau = e^{a_\tau + b_\tau Z_\tau}$. Hence,

$$\begin{aligned} C_{K,\tau} &= \left[\int_K^\infty (X_\tau - K) p^{\mathbb{Q}}(X_\tau) dX_\tau\right] e^{-r_t\tau} \\ &= \left[\int_{\frac{\ln K - a_\tau}{b_\tau}}^\infty (e^{a_\tau + b_\tau Z_\tau} - K) p^{\mathbb{Q}}(Z_\tau) dZ_\tau\right] e^{-r_t\tau} \\ &= \underbrace{\left[\int_{\frac{\ln K - a_\tau}{b_\tau}}^\infty e^{a_\tau + b_\tau Z_\tau} p^{\mathbb{Q}}(Z_\tau) dZ_\tau\right] e^{-r\tau}}_{(1)} - \underbrace{\left[K \int_{\frac{\ln K - a_\tau}{b_\tau}}^\infty p^{\mathbb{Q}}(Z_\tau) dZ_\tau\right] e^{-r_t\tau}}_{(2)}, \end{aligned}$$

where $p^{\mathbb{Q}}(Z_\tau)$ is now the $\mathbb{Q}$-density of Z_τ. As in Black and Scholes, define

$$\frac{\ln K - a_\tau}{b_\tau} = -d_{2,\tau} = -\left(\frac{\ln\frac{X_t}{K} + (r_t - \frac{1}{2}\sigma_t^2)\tau}{\sigma_t\sqrt{\tau}}\right).$$

Now, by Fourier inversion, denoting by $\phi^{\mathbb{Q}}$ the moment generating function of Z_τ, we obtain

$$\begin{aligned} (2) =& K e^{-r_t\tau} \int_{-d_{2,\tau}}^\infty p^{\mathbb{Q}}(Z_\tau) dZ_\tau \\ =& K e^{-r_t\tau} \left(\frac{1}{2} + \frac{1}{\pi}\int_0^{+\infty} \mathbb{R}\left[\frac{e^{iud_{2,\tau}}\phi^{\mathbb{Q}}(iu)}{iu}\right] du\right). \end{aligned}$$

Similarly,

$$
\begin{aligned}
(1) =& e^{-r_t \tau} \int_{-d_{2,\tau}}^{\infty} e^{a_\tau + b_\tau Z_\tau} p^{\mathbb{Q}}(Z_\tau) dZ_\tau \\
=& S_t \left(\frac{1}{2} + \frac{1}{\pi} \int_0^{+\infty} \mathbb{R} \left[\frac{e^{iud_{2,\tau}} \phi^{\mathbb{Q}}(iu + b_\tau)}{iue^{r_t \tau}} \right] du \right).
\end{aligned}
$$

We use $\phi^{\mathbb{Q}}(iu) = \mathbb{E}_t[e^{iuZ_\tau}]$, as derived in Theorem 2 of Bandi and Renò (2017). We note that, when $\phi^{\mathbb{Q}}(iu)$ is Gaussian, the parameters are constant, and price jumps are absent, the pricing formula delivers the Black and Scholes model. When $\phi^{\mathbb{Q}}(iu)$ is (conditionally) Gaussian, volatility is time-varying and price jumps are present, the expression yields the model in Andersen, Fusari, and Todorov (2017). In our case, no distributional assumptions on Z_τ are made and all parameters are allowed to be time-varying. Both price and volatility jumps are permitted.