

Short-Term Market Risks Implied by Weekly Options*

Torben G. Andersen[†] Nicola Fusari[‡] Viktor Todorov[§]

June 14, 2016

Abstract

We study short-term market risks implied by weekly S&P 500 index options. The introduction of weekly options has dramatically shifted the maturity profile of traded options over the last five years, with a substantial proportion now having expiry within one week. Such short-dated options provide a direct way to study volatility and jump risks. Unlike longer-dated options, they are largely insensitive to the risk of intertemporal shifts in the economic environment. Adopting a novel semi-nonparametric approach, we uncover variation in the negative jump tail risk which is not spanned by market volatility and helps predict future equity returns. Incidents of tail shape shifts coincide with mispricing of standard parametric models for longer-dated options. As such, our approach allows for easy identification of periods of heightened concerns about negative tail events that are not always “signaled” by the level of market volatility and elude standard asset pricing models.

Keywords: Options, Jumps, Stochastic Volatility, Extreme Events, Time-Varying Tail Risk, Return Predictability.

JEL classification: C51, C52, G12.

*Andersen gratefully acknowledges support from CREATES, Center for Research in Econometric Analysis of Time Series (DNRF78), funded by the Danish National Research Foundation. The work is partially supported by NSF Grant SES-1530748. We also acknowledge financial support from the CME Group and access to data from the CME DataMine system. We are thankful to Ken Singleton (the editor) and anonymous referees, as well as seminar participants at the Booth School of Business, University of Chicago, the Kellogg School of Management, Northwestern University, Stanford University, Johns Hopkins University, the Second International Workshop in Financial Econometrics, Salvador, Brazil, October 2015, the European Econometric Society Meeting in Milan, December 2015, the Empirical Finance Workshop at ESSEC Business School, Office of Financial Research, Washington, D.C., Boston University and the Market Microstructure and High Frequency Data Conference at the Stevanovich Center, University of Chicago, May 2016. We also thank Ian Dew-Becker, Anna Cieslak, Ravi Jagannathan, Roger Lee, Jeff Russell, Fabio Trojani, and Dacheng Xiu for helpful comments. Yupeng Wang provided skillful research assistance. Finally, we are very grateful to John Angelos and Mike Warsh from the CBOE for detailed explanations regarding the institutional organization of trading in weekly options, and to Luca Benzoni and Ivana Ruffini from the Federal Reserve Bank of Chicago for facilitating this information exchange.

[†]Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208; e-mail: t-andersen@northwestern.edu.

[‡]The Johns Hopkins University Carey Business School, Baltimore, MD 21202; e-mail: nicola.fusari@jhu.edu.

[§]Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208; e-mail: v-todorov@northwestern.edu.

1 Introduction

Recent years have witnessed a rapid increase in the trading of short-dated options. For instance, S&P 500 option contracts with one week or less to maturity have seen their share of trading at the Chicago Board of Options Exchange (CBOE) rise steadily from about 10% in early 2010 to almost 30% in mid 2015. Furthermore, the volume in shorter-dated options is disproportionately skewed towards out-of-the-money (OTM) options relative to the pattern for longer-dated options. Similar developments are observed in many other index option markets and for options on individual names. This process has been facilitated by the introduction of a new option category, featuring sequential issuance of contracts expiring one week apart, the so-called weekly options, or “weeklies.”

The emergence of active trading in the weeklies represents a step towards market completion. This topic has a long history, with Ross (1976) emphasizing the enhanced spanning and the potential for efficiency gains from options trading, and Breeden and Litzenberger (1978) stressing the ability to replicate a wide range of payoffs through a static option portfolio.¹

Nevertheless, we face the question of why these particular contracts have been so successful. The primary distinguishing feature of these securities, relative to longer-dated options, is the intimate link between the pricing of options close to expiry and the state of the underlying asset return process.² When tenor is short, the expected volatility and jump intensity do not vary much over the remaining life of the option. This implies, in particular, that the relative prices of deep OTM options are largely independent of the level of diffusive volatility, and instead reflect the characteristics of the risk-neutral jump process. Likewise, the pricing of short-dated at-the-money (ATM) options depends primarily on current spot volatility. Hence, the weeklies improve market participants ability to acquire or lay off exposure to diffusive and jump price risks. Such arguments fail in the case of longer-dated instruments for which the expected variation in the future volatility and jump intensity cannot be ignored in valuation. In fact, realistic models for their joint dynamics invariably involve complex interactions, rendering semi-closed form option pricing feasible only under strong parametric assumptions, e.g., the specifications reside within the affine jump-diffusion model class of Duffie et al. (2000).

Reversing the above reasoning, we infer that prices for actively traded short-dated options simplify the task of identifying the concurrent spot volatility and pertinent features of the risk-neutral jump process, subject to minimal assumptions on the return generating process. Specifically,

¹The latter statement is more formally explored in Green and Jarrow (1987) and Nachman (1988). For further developments, see, e.g., Bakshi and Madan (2000) and Pan and Liu (2003), among many others.

²A similar type of connection is present for short-maturity bond prices. Collin-Dufresne et al. (2008) explore this relation to identify the state vector driving the short-rate dynamics in a model-free way.

short-maturity ATM options help pin down spot volatility, while the relative prices of deep OTM options assist in determining the intensity and distribution of jumps. The emergence of the weekly options moves this observation from the realm of theory to the domain of practical empirical work. The requisite quotes for short-dated options are now available on a daily basis. The goal of the current paper is to capitalize on the new opportunities afforded by the trading of weeklies to explore the characteristics of the risk-neutral distribution of equity-index returns as implied directly by the option data, largely avoiding reliance on parametric restrictions.

The current paper is, as far as we know, the first to explore the information content of weekly options for the underlying risk-neutral return dynamics in a systematic way. In part, this reflects the very recent emergence of near-continuous trading in short-dated options. Therefore, we first carefully review the basic features of our weekly option sample and provide detailed descriptions of our filtering procedures, imposed to control for excessive noise or errors.

In order to exploit information from all available short-dated options, we develop a new asymptotic option pricing approximation that is operative across all strikes, and not just for ATM and deep OTM options. It exploits the fact that, over short intervals of time, volatility and jump intensity do not change much in expectation. Using our new approximation, we proceed semi-nonparametrically, i.e., we impose only weak parametric restrictions on the jump distribution, while remaining silent about the dynamics of the volatility and jump intensity. The approach allows us to infer the spot characteristics for the risk-neutral return distribution exclusively from short-dated options. In particular, we generate separate estimates of the current volatility and jump intensity as well as the jump size distribution at the end of each trading day.

Our approach bears superficial resemblance to calibration procedures commonly applied in approximating option-implied volatility surfaces. The differences are critical and fundamental, however. We explicitly impose no-arbitrage constraints in estimation and synthesize the option price information into consistent estimates for spot volatility and key jump characteristics, amenable to formal econometric analysis. From the perspective of summarizing the state of the local risk-neutral distribution in a continuous-time setting, the spot volatility and jump characteristics constitute sufficient statistics. They fully characterize the local behavior of the underlying semimartingale representing the risk-neutral asset price process. Furthermore, our ability to extract consistent point estimates and generate suitable confidence regions for the state vector over time sets the stage for analysis of the dynamic properties of the system. In contrast, standard calibration delivers smoothed risk-neutral densities for specific horizons, but does not enforce any form of internal dynamic consistency and provides no guidance for the extraction of spot volatility, jump intensities

or jump size distributions. We label the most general version of our methodology “structural calibration” to emphasize the fact that it generates valid asymptotic inference for the key components of the state vector governing the evolution of the risk-neutral distribution at the point in time when we observe the option prices.³

Our semi-nonparametric procedure enables far more general modeling of the time variation in jump risk than in the prior literature. In particular, standard option pricing models allow no time variation in the jump distribution and only limited variation in the jump intensity. Relative to recent nonparametric approaches focused on tail estimation, we offer a far more comprehensive analysis. For example, Bollerslev and Todorov (2014) and Bollerslev et al. (2015) are strictly concerned with the tails and rely exclusively on very deep OTM contracts, obtained from regular option samples. In contrast, we estimate the entire jump distribution through a very different methodology and exploit short-dated options, taking full advantage of the introduction of weeklies.

1.1 Shifts in the Left Risk-Neutral Tail

To illustrate the potential importance of time variation in the risk-neutral jump distribution and the informational advantages provided by the short-dated options, Figure 1 depicts log prices for options with tenor around one and four weeks on three consecutive days in September 2012. The left panel concerns the shorter maturities. The three curves largely coincide for moneyness m in the range $(-2, 3)$, suggesting no major change in spot volatility or implied upside jump risk across these days.⁴ However, on September 20 and 21, the option prices are systematically higher for the deep OTM put options than on the preceding day. The gap emerges around $m = -2$ and increases steadily for further OTM puts. The finding is clearly robust, with close to twenty consecutive put prices being elevated relative to September 19. This change is suggestive of an isolated shift in the negative jump tail distribution, because a change in the spot volatility or jump intensity would boost the prices of options with moneyness closer to zero as well.

In the right panel of Figure 1, covering longer-dated options, one also observe a tendency towards an elevation in the prices of deep OTM puts, but the differences across the days are now smaller and there is considerable noise in the far OTM quotes, which generates non-monotonicity that further blurs the inference. This reflects the fact that longer-dated OTM puts are sensitive to not only the negative jump size distribution, but also to the pricing of the time-varying volatility and jump intensity as well as associated leverage-type effects. The risks due to time variation in

³See Jarrow and Kwok (2015) for an alternative method of disentangling model misspecification from estimation error through the imposition of minimal restrictions on the calibration of the option surface.

⁴Moneyness is formally defined in Section 2 as well as in the caption of Figure 1.

volatility and jump intensity, while having minimal impact over short intervals such as a week, have a nontrivial effect over longer periods, like a month. Thus, Figure 1 illustrates the informational advantages of short-dated options for studying the risk-neutral left jump tail.

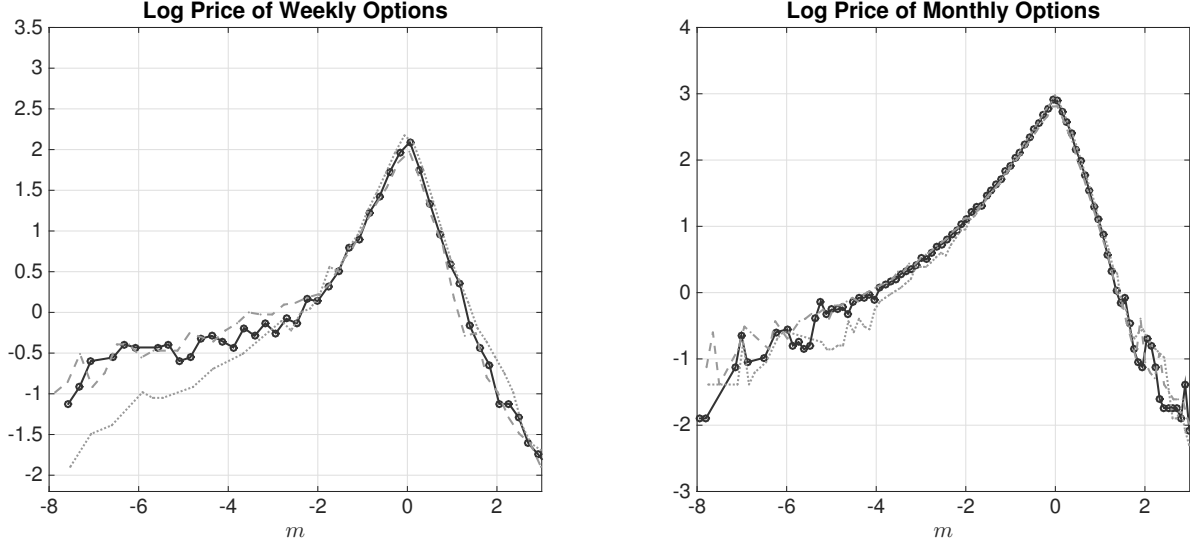


Figure 1: **Log Option Prices on September 19-21, 2012.** The log prices on September 20, 2012, are indicated by the full drawn line. On this date, the tenors for the short-dated and regular options were 7 and 28 calendar days, respectively. The dotted and dashed lines represent log prices for the corresponding option contracts on September 19 and 21, 2012. Moneyness m is defined as $\log(K/F_\tau)/(\sqrt{\tau}IV_{atm,\tau})$, where K is the strike price, τ is the tenor (in years), F_τ is the futures price for the underlying asset at option expiry, and $IV_{atm,\tau}$ is the at-the-money implied volatility for tenor τ .

Since traditional option pricing models preclude shifts in the jump distribution, it is revealing that Figure 1 suggests such events do occur and, in fact, may be identified empirically from short-dated options. Of course, an isolated shift in the jump distribution may be unusual. In most scenarios, we expect the jump distribution to change in conjunction with the volatility and jump intensity. Such simultaneous shifts complicate the identification of changes in the jump size distribution. This necessitates the development of formal tools for disentangling the effects in more general circumstances, and our structural calibration approach is designed to accomplish this task. At the same time, the type of event observed in Figure 1 is far from unique. We identify several similar instances of isolated shifts in the left jump tail across our four-year sample. We present a few illustrative cases in Appendix A.

The fact that our approach relies heavily on short-maturity options raises a potential concern. How can we guard against excessive noise or idiosyncratic pricing in short-dated options? This

question revolves around the strength of the statistical evidence and the economic plausibility of the implied tail shifts. These issues are recurrent themes throughout our analysis.

For now, we supplement our illustrative evidence with an account of the events surrounding September 20, 2012, which bolsters our presumption of a genuine shift in the risk-neutral tail. The backdrop is optimism due to an announcement, the preceding week, of additional monetary stimulus by the Federal Reserve. On September 20, even though the economic news were negative, the S&P 500 recovered most initial losses and experienced only a marginal decline. Likewise, on September 21, the economic news were at best mixed, and the S&P 500 ended flat.⁵ Hence, the broader outlook for equities remained fairly positive, but an element of uncertainty, or fear, may have entered the minds of some investors in the wake of the negative news updates which may have increased the pricing of contingencies associated with a sudden market downturn.

1.2 FOMC Announcements and Tail Movements

Figure 1 documents an isolated left tail shift in short-dated equity-index option prices which cannot be easily discerned from regular maturity options. Here, we illustrate how the tail characteristics of the risk-neutral distribution respond to economic news. To associate any shift with a specific event, it is useful to narrow the observation window to a short period surrounding the episode of interest. For that purpose, we turn to intraday quotes for short-maturity e-mini S&P 500 index options at the Chicago Mercantile Exchange Group (CME). We have data covering three consecutive announcements following regular FOMC meetings towards the end of 2013. This is a period when the question of whether the Federal Reserve would start to “taper” its bond purchases, implemented as the center piece of its quantitative easing program, was a major issue for the financial markets.

Figure 2 depicts the market and option price reactions to the FOMC announcements on September 18, October 30 and December 18, 2013. On the first date, the 2pm ET release of the FOMC statement surprised investors by declaring that the highly accommodative monetary policy would be continued, even as many observers had expected an initiation of tapering. As seen in the upper left panel, the S&P 500 jumped significantly on the news while the CBOE volatility index, VIX, dropped. This latter manifests itself in a small decline for all option prices in the moneyness range

⁵For September 20, the Zacks Stock Market News reports that the initial claims for unemployment benefits rose to a two-month high, U.S. leading economic indicators dropped, and manufacturing activity declined in the U.S., Europe and China, but stocks recovered as investors remained somewhat optimistic. For September 21, Zacks mentions that unemployment had risen in 26 states and the World Trade Organization lowered its forecast for international trade. At the same time, there were positive signs that the European Union was getting closer to a bail-out agreement with Spain.

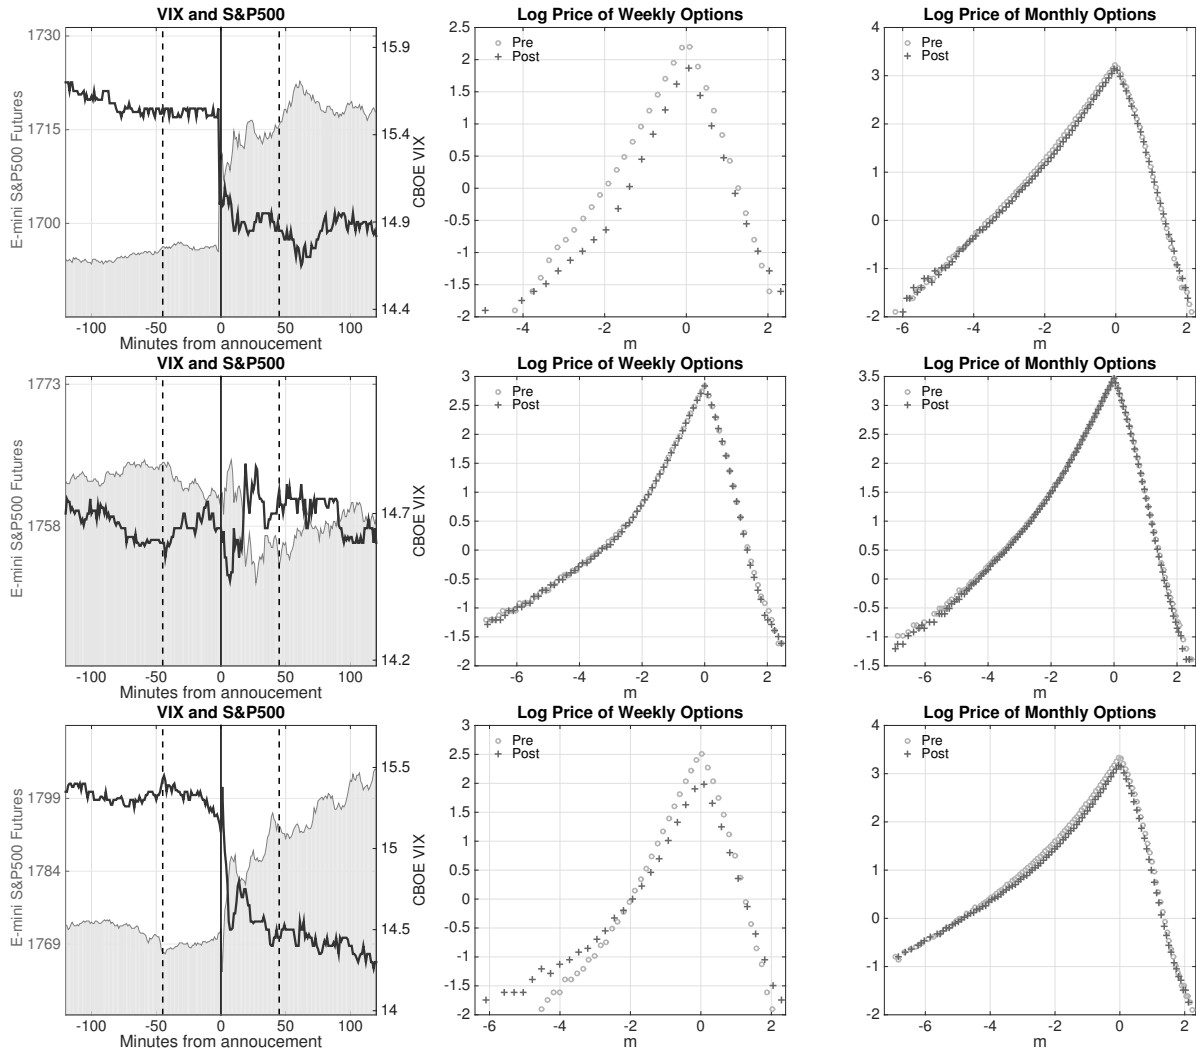


Figure 2: Option Prices around FOMC announcements. The rows correspond to FOMC announcements on September 18th 2013, October 30th 2013, and December 18th 2013, respectively. These scheduled announcements occur at 2 pm ET. The one-minute best bid and offer quotes are from the CME. The panels in the first column display the minute-by-minute evolution of the S&P 500 E-mini Futures price (shaded area) and the official CBOE VIX (dark line) from 120 minutes before to 120 minutes after the announcement (the one-minute CBOE VIX is from TickData: www.tickdata.com). The solid vertical line indicates the announcement time. The second and third panels in each row show the log option prices forty-five minutes before and after the announcement (corresponding to the dashed vertical lines in the first panel) for the shortest and second-shortest available tenors. On October 30th 2013, these maturities are equal to 16 and 51 days. For the two other announcements, the tenors are 2 days and 30 days.

$(-5, 1)$ of the upper right panel. However, these longer-dated options provide no indication of a change in the general pricing pattern. In contrast, the top middle panel provides a nuanced picture of the market reaction. The drop in the prices for ATM options is now readily discernible due to the relatively larger impact of volatility shifts for short-dated instruments. Moreover, around $m = -2$, there is a clear shift in the slope of the left tail, which is relatively thicker 45 minutes following the announcement compared to before. In fact, for $m = -4$, the two curves cross, implying that deep OTM put option prices are largely unaffected by the FOMC release. This suggests potential lingering concern, reflecting unease among some market participants about the prolonged quantitative easing or a sense that policy makers view the economic environment as weaker than investors anticipated.

The second row for October 30, 2013, captures an announcement with little market impact. There were relatively weak economic news earlier in the day and no change in policy was foreseen. The market drifted down slightly, and the VIX up, over our event window, but the change was minimal. This is also reflected by the absence of any noticeable shift in the option prices in the second and third column. In this case, the stability of the option prices corroborate the interpretation that market expectations were unaffected by the FOMC release.

Finally, the third row captures the FOMC release on December 18, 2013, that announces the initiation of tapering. The news were surprising, but not entirely unanticipated, and the market quickly soared following the news release. The associated drop in the VIX is again reflected in the drop in longer-dated option prices, evident in the bottom right panel. Yet, this is far from the full story. The middle panel captures a pronounced shift in the pattern of the option prices across moneyness, with the deeper OTM option prices (for $m < -2$) increasing significantly. One natural interpretation is that some investors view the change in the stance of monetary policy as having increased the possibility of market turbulence in the near future.

One potential concern with the evidence above is the possible effect of endogenous shifts in the market liquidity around scheduled announcements. Market makers post wide spreads prior to announcements to guard against the impact of jumps induced by the release. Likewise, the volatility is often temporarily elevated following the announcement which leads to wider spreads. Hence, we sample the option prices a full 45 minutes before and after the events, alleviating the impact from such potential distortions. In Appendix B, Figure 17 depicts the bid-ask spreads and quote update frequency for the options in the period surrounding the news releases. It is evident that the quotes are updated in real time and the mild elevation of the spreads post announcement is not systematically related to the pattern of tail shifts documented in Figure 2.

This evidence illustrates the point that information shocks, such as FOMC announcements, can have very different effects on perceived and priced market tail risks versus volatility. The short-maturity options provide an easy and robust way to identify such differences.

1.3 Interpretation and Overview

The above illustrations suggest that the risk-neutral jump tail may shift significantly in one direction, even if spot volatility remains stable or moves in the opposite direction. These cases exemplify an important general feature that we document in the body of the paper – the expected risk-neutral jump dynamics is not tightly linked to the concurrent return volatility.

The finding of jump variation, departing significantly from the dynamics of market volatility, runs counter to the common approach of modeling jump risk in no-arbitrage asset pricing models. In the latter, the jump intensity is proportional to volatility or its factors, see, e.g., Bates (2000) and Pan (2002). It is also contrary to the implications of standard consumption-based equilibrium models, such as Drechsler and Yaron (2011) and Wachter (2013), which, although flexibly specified, impose tight (affine) connections between the variation of market jump risk and volatility. Furthermore, to the extent the variation in the risk-neutral negative jump tail is unmatched by actual variation in the expected future negative market jumps, our evidence points to substantial nonlinearities in the part of the economy-wide pricing kernel that concerns “large” negative jumps.

On the constructive side, our semi-nonparametric approach provides an easy way to extract a measure for short-term downside tail risk. It can be used both for modeling and analyzing the importance of the latter in a robust way, and may serve as a useful input to policy making.⁶ In particular, we demonstrate that periods of heightened tail risk, as measured by short-dated options, are closely connected with periods in which traditional asset pricing models severely underestimate the expensiveness of longer-dated deep OTM put options. The latter have been used extensively in earlier empirical work and they constitute an important segment of the options included in the calculation of the VIX volatility index. In addition, we document that the pricing of the short-term negative tail variation is an important predictor for the equity risk premium, outperforming a number of metrics popular in the literature. Thus, downside tail pricing is an economically significant quantity with independent information content relative to the typical state variables employed in standard asset pricing models.

The rest of the paper is organized as follows. Section 2 describes the option data on the S&P

⁶“... policymakers can achieve better outcomes by basing their outlooks on risk-neutral probabilities derived from the prices of financial derivatives” (Narayana Kocherlakota, President of the Federal Reserve Bank of Minneapolis, from speech at the University of Michigan, June 8, 2012).

500 index and, in particular, its composition in terms of tenor. Our general setup and notation is presented in Section 3. In Section 4, we describe our semi-nonparametric approach to approximate the short-dated option prices. Our estimation and inference techniques are detailed in Section 5. Section 6 reviews the empirical findings for semi-nonparametric models with time-invariant and time-varying jump distributions. Section 7 displays the time series evolution of the spot variance and various jump risk measures implied by our preferred model. In Section 8, we investigate the integration in pricing between the weekly and longer-dated options and the link between tail risk measures and future equity returns. Section 9 concludes. A variety of additional details and supporting evidence are gathered in the Appendix.

2 S&P 500 Equity-Index Options

2.1 The SPX Options

The trading of equity-index options has grown rapidly in recent years, partially in response to the introduction of new contracts offering a more comprehensive set of expiration dates, especially at the short end of the maturity spectrum. The traditional S&P 500 equity-index (SPX) options have traded on the CBOE since 1987. They have one monthly expiration, at week’s end around the middle of any given month. These contracts have been complemented with so-called quarterly and weekly options over the last decade. Quarterly options (SPXQ) expire at the end of each quarter, providing four additional expiration dates per year. Weekly options (SPXW) expire at the end of the trading week, unless an expiration already exists close to that date. Since January 2014, the CBOE maintains six consecutive SPXW expiration dates, ensuring that a string of short maturity S&P 500 options, expiring one week apart, exists at all times. In particular, in our sample, an actively traded front-maturity option with expiry within nine calendar days is always present.

Figure 3 depicts the average number of SPXW options (“weeklies”) traded per day and their percentage relative to the overall trading of S&P 500 options on the CBOE. The average daily SPXW volume for the years 2011-2015 was around 71,000, 107,000, 198,000, 330,000 and 455,000 contracts, respectively. This raised their proportion from below 8% in early 2011 to above 50% by the middle of 2015. Since weeklies have a relatively short maturity upon issuance, this development has contributed to a fundamental shift in the maturity profile. Very short-dated options were thinly traded up till five years ago, and routinely excluded from empirical studies, but now they represent the most heavily traded segment of the maturity spectrum.

One may wonder what economic forces drive the increasing demand for short-dated options. A natural conjecture is that weeklies are used to hedge short-term exposures and, in particular, tail

risk, e.g., hedge funds protect the possibility of a large one-day loss. In fact, such type of products have long traded over-the-counter and are known as gap options. Pure jump type risk (big one-day losses) is much easier and cheaper to hedge with weeklies than with one-month options, as the latter load also on volatility risk. This may readily be illustrated by noting the large difference in the sensitivity of the prices for weekly versus 30-day options to jump events and shifts in the value of key jump parameters within standard parametric models. Moreover, from a practical perspective, the weeklies are attractive because they have very low time-value and hence managing them (delta-hedging) is inexpensive. This being said, there are no official data revealing who trade these instruments or which category of institutions are the ultimate end users.

We may, however, directly explore one alternative hypothesis. It is possible that the trading of weeklies cluster around scheduled news releases. As representatives of major economic releases, we studied periods surrounding the statement following FOMC meetings and the monthly employment report. We find no significant change in the trading of short-maturity options around these prescheduled announcements.⁷ As such, our findings suggest that the secular trend in trading activity reflects a general desire to manage and hedge short-term market-wide exposures.




Figure 3: **Weekly (SPXW) Options.** The bar plot (left y-axes) represents the average daily volume in weekly options. The light gray line (right y-axes) depicts the weekly options' volume as a percentage of the total option volume (the sum for SPX, SPXQ and SPXW).

The primary objective of our paper is to extract information about the current state of the risk-neutral dynamics (volatility and conditional tail probabilities) from short-maturity options. To allow meaningful inference, we require good coverage across the strike range at the front maturity.

⁷These results are available upon request.

This effectively determines the starting date for our sample, as the coverage for short-dated options is poor prior to 2011. Our empirical analysis is based on end-of-day S&P 500 equity-index option quotes from the CBOE, obtained via OptionMetrics for January 3, 2011 through August 28, 2015. We seek to contrast our findings from short-maturity options to what can be learned from an option panel composed of longer-dated options, resembling those routinely used in prior studies. To ensure a homogenous basis for comparison, we construct a common set of trading days for these separate option samples. Towards this end, we first extract quotes for all OTM SPX, SPXQ and SPXW options for 2011-2015.⁸ For each option with τ years to maturity and strike price K , we define moneyness as,

$$m = \frac{\ln(K/F_\tau)}{\sqrt{\tau} IV_{atm,\tau}},$$

where F_τ denotes the forward price for transactions τ years into the future, while $IV_{ATM,\tau}$ denotes the (annualized) implied volatility of the option with strike price closest to F_τ .

2.2 Characteristics of the Two Option Samples

Prior to analysis, we filter the option data to prevent noise, stemming from large bid-ask spreads or stale and illiquid contracts, from having an undue impact on the inference. The filtering procedure is detailed in Appendix C. Table 1 summarizes the composition of our final sample. The top panel reports the average number of option contracts in each moneyness-tenor category and the bottom panel provides the relative trading volume. The shaded column refers to our short-dated sample, and the next three to our regular sample.

The top panel reveals that short-dated options, consisting exclusively of front-maturity contracts, account for about 3.6% of all contracts in the sample. For options with tenor below 60 days, OTM put options constitute a slightly larger fraction than ATM options ($-2 \leq m \leq 2$), but this is reversed at longer maturities. Finally, there are relatively few quotes for OTM calls. The bottom panel shows that trading in the short-maturity sample, comprising almost 20% of the total volume, is more intense (per contract) than for regular options. This is also evident from Figure 4 as, at the end of our sample, almost forty percent of the volume is concentrated in contracts with tenor less than or equal to nine days. Furthermore, the volume is higher for ATM than OTM options, with the OTM call volume being extraordinarily light in the regular sample. Finally, the volume for tenors above 180 days is low, especially for far OTM options.

⁸The in-the-money options are invariably less liquid and have higher spreads. Thus, we obtain more accurate prices for these options by imputing them via put-call parity using the corresponding OTM options.

| | $\tau \leq 9$ | $9 < \tau \leq 60$ | $60 < \tau \leq 180$ | $\tau > 180$ | Total |
|-------------------------|---------------|--------------------|----------------------|--------------|-------|
| Number of contracts (%) | | | | | |
| $m < -2$ | 1.67 | 23.78 | 15.10 | 5.95 | 46.49 |
| $ m \leq 2$ | 1.82 | 20.56 | 19.01 | 10.34 | 51.73 |
| $m > 2$ | 0.15 | 1.45 | 0.13 | 0.05 | 1.78 |
| Total | 3.64 | 45.78 | 34.24 | 16.34 | 100 |
| Volume (%) | | | | | |
| $m < -2$ | 6.41 | 14.96 | 3.50 | 0.66 | 25.53 |
| $ m \leq 2$ | 13.09 | 40.45 | 15.75 | 3.97 | 73.26 |
| $m > 2$ | 0.43 | 0.74 | 0.03 | 0.00 | 1.21 |
| Total | 19.92 | 56.16 | 19.28 | 4.63 | 100 |

Table 1: **Option characteristics.** **Top Panel:** Percentage of option contracts over different combinations of tenor (τ , in days) and moneyness (m) between January 3, 2011, and August 28, 2015. **Bottom Panel:** Percentage of contract volume over different combinations of tenor (τ , in days) and moneyness (m) between January 3, 2011, and August 28, 2015.

Overall, our option samples provide good coverage across moneyness and maturity for short and medium tenors, except for the deep OTM calls. At longer maturities, the lack of liquidity for the latter instruments is a concern. Even if the information content of the options is excellent, it is inherently difficult to draw inference regarding the behavior of the upper right tail of the distribution across all maturities, and the problems compound for both tails at longer horizons.

Turning to the variation in coverage over time, Figure 5 depicts the maximum strike range captured across all tenors within our two categories. The left panel shows that, since 2012, the short-dated options almost always cover strikes below $m = -6$ on the downside, while the upside coverage invariably reaches $m = 2$, but only sporadically goes much beyond this level.

For brevity, the right panel of Figure 5 combines the options across all tenors in our regular sample. Hence, the extreme positive and negative moneyness is typically attained at different maturities. Moreover, the moneyness is not truncated, as is the case for the short-maturity sample. With this caveat in mind, we note the broadly similar coverage on the upside while, up through 2013, there are at least some tenors for which longer-dated options provide a broader strike range on the downside than afforded by our short-maturity sample. Nonetheless, the OTM put coverage is generally impressive for the short-dated options, reaching the threshold of -8 for more than one




Figure 4: **Short Maturity Options.** The bar plot (left y-axes) represents the average daily volume in short maturity options (i.e. with days to maturity less than or equal to nine). The light gray line (right y-axes) depicts the short maturity options' volume as a percentage of the total volume in SPX (sum of SPX, SPXQ and SPXW) options.




Figure 5: **Moneyness Range.** **Left Panel:** Moneyness range for short-maturity options. **Right Panel:** moneyness range for long-maturity options. The sample period is January 3, 2011 – August 28, 2015. Moneyness is computed as in the main text, $\log(K/F_\tau)/(\sqrt{\tau}IV_{atm,\tau})$.

third of the dates in the sample. In summary, our short-maturity sample covers an impressively broad range of strikes on the downside and largely matches the upside coverage provided by standard option samples.

Finally, Figure 6 depicts the average bid-ask spread for our two option categories. In the $[-4, 1]$ moneyness range, the relative spread is lower for the longer-dated options, but outside this range the liquidity of the short-dated options – partially reflecting our sample selection criteria – manifests itself in a relatively narrow spread. Hence, the reliability of the pricing for this critical segment of far OTM short-dated options is on par with, if not better than, that of the typical OTM options used in prior studies.




Figure 6: **Bid-Ask Spread.** Kernel regression of the relative (left axes) and absolute (right axes) bid-ask spread as a function of moneyness, m . The relative bid-ask spread is reported in implied volatility terms and computed as $(IV_{ask} - IV_{bid})/IV_{mid}$, where IV_{ask} , IV_{bid} , and IV_{mid} are implied volatilities computed using the ask, bid, and mid prices, respectively; the absolute bid-ask spread is reported in dollar terms and it is simply computed as the difference between the ask and the bid option prices. The sample period is January 3, 2011 – August 28, 2015. Moneyness is computed as in the main text, $\log(K/F_\tau)/(\sqrt{\tau}IV_{atm,\tau})$.

3 Setting and Notation

Throughout, we assume financial markets are arbitrage-free which, subject to mild regularity, implies the existence of a risk-neutral measure. The underlying price process X , i.e., the S&P 500 index, is governed by the following general risk-neutral dynamics,

$$\frac{dX_t}{X_{t-}} = (r_t - \delta_t)dt + \sqrt{V_t}dW_t + \int_{\mathbb{R}^2} (e^x - 1)\tilde{\mu}(dt, dx), \quad (1)$$

where r_t and δ_t are the instantaneous risk-free rate and dividend yield, which we will assume are deterministic;⁹ W is a Brownian motion; V is the diffusive stochastic variance process; μ is a counting jump measure with compensator $dt \otimes \nu_t(dx)$, and the difference $\tilde{\mu}(dt, dx) = \mu(dt, dx) - dt \otimes \nu_t(dx)$ is the associated martingale jump measure.¹⁰

The specification for X in equation (1) implies that the discounted cum-dividend gain process associated with X follows a local martingale. Subject to mild auxiliary restrictions, this is the minimal condition for the dynamics of X to preclude arbitrage.¹¹

Our estimation is based on a panel of options written on the asset X along with high-frequency data for X that are used to construct nonparametric volatility estimates. We denote the prices of European-style OTM options on X at time t by $O_{t,k,\tau}$. Assuming frictionless trading, option prices

⁹The unexpected variation in these series over the relevant maturities is small relative to that of other factors impacting the option prices. As such, this simplifying assumption has little practical effect on the results.

¹⁰We note that $\nu_t(\mathbb{R})$ might be infinite which implies jumps of infinite activity. Nevertheless, with slight abuse of notation, we will refer to $\nu_t(dx)$ as the jump distribution at time t .

¹¹One condition implicitly imposed in equation (1) is that X is an Itô semimartingale, i.e., its characteristics are absolutely continuous in time. This restriction is valid for nearly all prior parametric specifications in the literature.

equal the expected discounted payoffs under the risk-neutral measure,

$$O_{t,k,\tau} = \begin{cases} \mathbb{E}_t^{\mathbb{Q}} \left[e^{-\int_t^{t+\tau} r_s ds} (X_{t+\tau} - K)^+ \right], & \text{if } K > F_{t,t+\tau}, \\ \mathbb{E}_t^{\mathbb{Q}} \left[e^{-\int_t^{t+\tau} r_s ds} (K - X_{t+\tau})^+ \right], & \text{if } K \leq F_{t,t+\tau}, \end{cases} \quad (2)$$

where τ is the tenor, K the strike price, $F_{t,t+\tau}$ the futures price of the underlying asset X at time t for date $t + \tau$, and $k = \log(K/F_{t,t+\tau})$. As is common, we quote option prices in terms of their Black-Scholes implied volatility (BSIV). The option BSIV is denoted by $\kappa_{t,k,\tau}$.

4 Semi-nonparametric Models for Short-Maturity Option Pricing

A parametric approach to option pricing allows us to utilize the rich information in the entire option surface, covering all available maturities and strikes. On the other hand, option prices are highly nonlinear functions of the underlying state variables and parameters. Hence, the impact of even minor misspecification can be both significant and unpredictable. While this is largely unavoidable when pricing options across the full maturity range, the literature has developed partial remedies for short-dated options. One may adopt a semi-nonparametric approach in which parts of the pricing kernel are left unspecified, avoiding misspecification along particular dimensions. However, existing techniques apply only for a restrictive set of strikes. Below, we describe a novel semi-nonparametric procedure, which exploits a more robust local approximation to the prices of short-dated options, irrespective of the strike price.

Henceforth, we retain the general representation (1) for the risk-neutral dynamics of X , and impose parametric restrictions on the distribution of price jumps only. In other words, we stipulate that the jump intensity is of the form $dt \otimes \nu(dx; \mathbf{J}_t, \theta)$ for some parametric jump distribution $\nu(dx; \mathbf{J}_t, \theta)$, characterized by a time-invariant parameter vector θ and time-varying state vector \mathbf{J}_t .

Various asymptotic expansions have been applied for short-dated option prices in prior work. For example, Carr and Wu (2003) and Bollerslev and Todorov (2014) (see also Bentata and Cont (2012)) exploit the following OTM approximation,

$$\frac{O_{t,\tau,k}}{\tau X_{t-}} \longrightarrow \begin{cases} \int_{\mathbb{R}} (e^x - e^k)^+ \nu(dx; \mathbf{J}_t, \theta), & \text{if } k > 0 \\ \int_{\mathbb{R}} (e^k - e^x)^+ \nu(dx; \mathbf{J}_t, \theta), & \text{if } k < 0 \end{cases}, \quad \text{as } \tau \downarrow 0. \quad (3)$$

Hence, the valuation of short-dated OTM options is determined by the pricing of “big” jumps.¹² In practice, however, this approximation works well only for very deep OTM options, i.e., for k

¹²By terms such as “the pricing of future big jumps” or “the priced large jumps,” we refer to the conditional risk-neutral expectation of a given metric defined on such jumps (like their count or their quadratic variation). The (conditional) risk premium is the difference between the corresponding conditional expectations under \mathbb{P} and \mathbb{Q} . This convention applies throughout the paper for all other stochastic quantities of interest.

large in absolute value. Furthermore, once we move beyond very short maturities, the presence of a diffusion term and time-variation in volatility and jump intensity all have a non-trivial impact.

Another short-maturity option approximation concerns ATM options, see, e.g., Durrleman (2008).¹³ In particular, we have,

$$\kappa_{t,0,\tau} \longrightarrow \sqrt{V_t}, \quad \text{as } \tau \downarrow 0. \quad (4)$$

The approximations (3) and (4) work only for either deep OTM or ATM options, but not for the broad range of strikes between these extremes. Below, we implement our new approximation, which applies to *all* short-maturity options.

Our approach is based on computing short-dated option prices from an approximate representation of X with constant diffusive variance and jump intensity (as well as constant dividend yield and risk-free rate), denoted \tilde{X} . It is defined as,

$$\frac{d\tilde{X}_s}{\tilde{X}_{s-}} = (r_t - \delta_t) ds + \sqrt{V_t} dW_s + \int_{\mathbb{R}} (e^x - 1)(\mu(ds, dx) - ds \otimes \nu(dx; \mathbf{J}_t, \theta)), \quad s \geq t, \quad \tilde{X}_t = X_t. \quad (5)$$

\tilde{X} approximates X over a short horizon starting at time t . It is used to compute approximate prices at time t for short-maturity options written on X with any given strike price. The approximation freezes the characteristics of X , i.e., its drift, volatility, and jump intensity, at their values at time t . Therefore, conditional on \mathcal{F}_t , \tilde{X} is a Lévy process.¹⁴ The reason \tilde{X} approximates X well over short intervals is that, over such horizons, the volatility and the jump intensity of X are not expected to vary much. Moreover, risk-neutral parameter estimates typically imply more persistent volatility and jump intensity processes under \mathbb{Q} than \mathbb{P} , see, e.g., Broadie et al. (2007).

Pricing options written on \tilde{X} at time t is significantly easier than pricing options on X . Indeed, semi-closed option pricing formulas exist for large classes of Lévy processes, see, e.g., Cont and Tankov (2004).¹⁵ The theoretical price of an option, in terms of BSIV, at time t written on \tilde{X} with tenor τ and log-moneyness k depends on the parameter vector θ as well as the state vector, $\mathbf{S}_t = (V_t, \mathbf{J}_t)$. We label it $\tilde{\kappa}_{k,\tau}(\mathbf{S}_t, \theta)$. The formal analysis of the approximation error, $\tilde{\kappa}_{k,\tau}(\mathbf{S}_t, \theta) - \kappa_{t,k,\tau}$, relies on an asymptotic analysis of the price increment $\tilde{X}_{t+\tau} - X_{t+\tau}$ for $\tau \downarrow 0$. While the approximation \tilde{X} has been used previously by, e.g., Jacod and Protter (2012) to study the properties of nonparametric high-frequency volatility estimators, it is novel in the context of option pricing.

¹³Related, Medvedev and Scaillet (2007) consider short-maturity option price approximations around moneyness of zero for certain classes of jump-diffusion models.

¹⁴A Lévy process is a process with càdlàg paths and independent and stationary increments.

¹⁵Appendix D provides the requisite expressions for the conditional characteristic functions for our general semi-nonparametric models with Gaussian or generalized tempered stable jump specifications, as introduced subsequently in the paper. In turn, these expressions are used to price options via inverse Fourier techniques.

We will use our approximation, exploiting the simplified return dynamics in equation (5), to robustly identify the characteristics of the risk-neutral jump process implied by the pricing of short-dated options. Parametric models specify the entire risk-neutral law for X , including an explicit dynamic representation for the volatility and jump intensity processes, their relation, the interaction between volatility and return innovations, etc. By contrast, focusing solely on short-dated options and using the approximation $\tilde{\kappa}_{k,\tau}(\mathbf{S}_t, \theta)$, we need to specify only the jump distribution. Thus, the inference based on $\tilde{\kappa}_{k,\tau}(\mathbf{S}_t, \theta)$ applies far more broadly than one based on a full parametric model which hinges also on correctly specified volatility and jump intensity dynamics.¹⁶

5 Estimation and Inference

We estimate our semi-nonparametric models, using the M_t options available at a given point in time during the trading day, via the criterion,

$$\left(\{\hat{\mathbf{S}}_t\}_{t=1}^T, \hat{\theta} \right) = \underset{\{\mathbf{S}_t\}_{t=1}^T, \theta}{\operatorname{argmin}} \sum_{t=1}^T \left\{ \sum_{j: \tau_j \leq 9} \frac{(\kappa_{t,k_j,\tau_j} - \tilde{\kappa}_{k_j,\tau_j}(\mathbf{S}_t, \theta))^2}{M_t} + \frac{\hat{\eta}_t k_n}{M_t} \frac{(\sqrt{\hat{V}_t^n} - \sqrt{V_t})^2}{\hat{V}_t^n/2} \right\}, \quad (6)$$

where \hat{V}_t^n is a nonparametric estimate of the spot diffusive variance, using k_n high-frequency increments preceding time t , and $\hat{\eta}_t$ is a consistent estimator of the observation error variance for the option BSIV at time t .¹⁷ The main component of the objective function (6) is the mean squared error in fitting the option (BSIV) panel, but we add a (scaled) term penalizing the deviation of the option-implied spot volatility from the nonparametric high-frequency estimate. The estimation is done via concentration of the state variables. On each day, given a parameter vector θ , we first optimize over the state vector on this day (a low-dimensional optimization), and then we optimize the concentrated objective function over the parameter θ . Further implementation details may be found in Andersen et al. (2015a).

Our primary focus here is the quality of fit for the short-dated options. To assess model performance, we make use of the diagnostic tests developed in Andersen et al. (2015a). Our first test is a Z-score that reflects the fit to a specific part of the option surface over a given period. We focus on the day-by-day fit. The test is given by the normalized “option-fit” (of) statistic,

$$\hat{T}_{\mathcal{K},t}^{\text{of}} = \mathcal{Z}_{\mathcal{K},t} / \sqrt{\widehat{\text{Avar}}(\mathcal{Z}_{\mathcal{K},t})}, \quad (7)$$

¹⁶Internet Appendix I documents superior performance of our approximation based on the simplified model (5) relative to the commonly used procedure in equation (3) for a set of empirically relevant specifications.

¹⁷We use a standard robust threshold realized volatility estimator, \hat{V}_t^n , implemented as in Andersen et al. (2015a). The scaling factor, $\hat{\eta}_t$, representing the observation error, is obtained via a first-stage estimator of the error variance, $\hat{\eta}_t = \frac{1}{M_t} \sum_j \left(\kappa_{t,k_j,\tau_j} - \tilde{\kappa}_{k_j,\tau_j}(\tilde{\mathbf{S}}_t, \tilde{\theta}) \right)^2$, where $(\{\tilde{\mathbf{S}}_t\}_{t=1}^T, \tilde{\theta})$ is a preliminary consistent (unpenalized) estimator.

where $\mathcal{Z}_{\mathcal{K},t} = \sum_{j:k_j \in \mathcal{K}, \tau_j \leq 9} (\kappa_{t,k_j,\tau_j} - \tilde{\kappa}_{k_j,\tau_j}(\hat{\mathbf{S}}_t, \hat{\theta}))$ is the average pricing error over the relevant part of the moneyness region, \mathcal{K} . Andersen et al. (2015a) show that $\hat{T}_{\mathcal{K},t}^{\text{of}}$ is asymptotically standard normal under the null of correct model specification, and it diverges towards infinity otherwise.

Our second diagnostic is based on the discrepancy between the option-implied model-based estimate for spot diffusive volatility, \hat{V}_t , and the corresponding nonparametric high-frequency estimate, \hat{V}_t^n . The formal test is based on the normalized “volatility fit” (vf) statistic,

$$\hat{T}_t^{\text{vf}} = (\hat{V}_t^n - \hat{V}_t) / \sqrt{\widehat{\text{Avar}}(\hat{V}_t^n - \hat{V}_t)}. \quad (8)$$

Under correct model specification \hat{T}_t^{vf} is asymptotically standard normal, and diverges otherwise. We refer to Andersen et al. (2015a) for details regarding the construction of feasible estimates for the two asymptotic variance terms involved in the tests based on equations (7) and (8).

6 Empirical Results for Alternative Jump Size Distributions

This section explores the ability of different jump specifications to accommodate the short-dated option prices. By using the Lévy-based approximation (5), we develop semi-nonparametric representations that mirror parametric jump specifications explored in earlier work. Importantly, however, we avoid reliance on any particular model for the volatility dynamics and jump intensity as well as their interdependence. Hence, our empirical success hinges exclusively on the specification for the jump size distribution.

6.1 Models with Fixed Jump Distribution

Our first model for $\nu(dx; \mathbf{J}_t, \theta)$ is

$$\nu(dx; \mathbf{J}_t, \theta) = c_t \frac{1}{\sqrt{2\pi} \sigma_x} e^{-\frac{(x-\mu_x)^2}{2\sigma_x^2}} dx. \quad (9)$$

This formulation involves Gaussian jumps, following the seminal work of Merton (1976). In most prior parametric models, e.g., Bates (2000), c_t is an affine function of V_t (or the factors which comprise it). In our semi-nonparametric model we allow the jump intensity, c_t , to display (arbitrary) time variation unrelated to volatility, and we later assess this issue empirically. We have $\theta = (\mu_x, \sigma_x^2)$ and $\mathbf{J}_t = c_t$. We note that the jump distribution remains time-invariant.

Our second model for $\nu(dx; \mathbf{J}_t, \theta)$ is given by

$$\nu(dx; \mathbf{J}_t, \theta) = c_t \left\{ \frac{e^{-\lambda^-|x|}}{|x|^{1+\alpha}} 1_{\{x < 0\}} + \frac{e^{-\lambda^+|x|}}{|x|^{1+\alpha}} 1_{\{x > 0\}} \right\} dx, \quad \alpha < 2. \quad (10)$$

Similar to the specification (9), the time-variation is governed by c_t and left entirely unrestricted. Since c_t scales the intensity of jumps of all sizes, the jump distribution of the specification (10) is time-invariant. This representation of $\nu(dx; \mathbf{J}_t, \theta)$ is very general, as it allows separate parameters to control different features of the jump distribution. In particular, the parameters λ^- and λ^+ govern the behavior of the large negative and positive jumps, respectively, while the parameter α controls the behavior of small jumps of either sign. For $\alpha < 0$, the jump process is of finite activity, whereas $\alpha \in [0, 1)$ implies infinite activity, but finite variation jumps, while, finally, $\alpha \in [1, 2)$ generates paths of infinite variation. For constant c_t , the process (10) belongs to the tempered stable class, introduced by Carr et al. (2002, 2003), and estimated from return data by Bates (2012).¹⁸ When c_t is constant and $\alpha = -1$, we recover the double-exponential model of Kou (2002) adopted in many applications. Another popular variant is obtained for c_t constant and $\alpha = 0$, as this reduces to the variance gamma model of Madan et al. (1998).

We find that α cannot be identified with precision from the data. Hence, we estimate the jump specification (10) for different fixed values of α . Therefore, the parameter vector is $\theta = (\lambda^-, \lambda^+)$ and the state vector is $\mathbf{J}_t = c_t$.

We reiterate that our semi-nonparametric approach provides a more flexible representation of the risk-neutral return process than adopted in studies for which tractable pricing formulas are required for longer-dated options. This implies that we obtain a superior fit to the short-dated option prices and spot volatility than can be achieved with traditional specifications, as the latter impose additional (strong) parametric restrictions on the volatility and jump intensity processes. In other words, the evidence below provides an upper bound on the performance of standard models involving the same family of jump size distributions.

Turning to the empirical results, we invariably find the double-exponential jump representation, across the full range of relevant values for α , to outperform the Gaussian specification.¹⁹ For brevity, we only provide the diagnostic plots for the option pricing fit for the better performing double-exponential representation here, while the plot for the fit to spot volatility and the corresponding plots for the Gaussian model are deferred to Appendix E. Even for the former model, the fit is problematic. The panels in Figure 7 plot the Z-scores for each trading day across the short-dated option sample for three separate segments of the option cross-section, corresponding to the deep OTM put options in the top panel, the OTM puts in the middle panel and the OTM call options in the bottom panel. We observe severe overvaluation for the deep OTM puts during March 2012 to

¹⁸More precisely, (10) can be viewed as tempered stable for $\alpha \in (0, 2)$.

¹⁹The estimated average jump intensity and jump size are in line with estimates from earlier studies using a Gaussian model for S&P 500 index price jumps, e.g., Bakshi et al. (1997), Pan (2002), Andersen et al. (2002), Eraker (2004) and Broadie et al. (2007).

October 2013 and for the OTM calls in the second part of the sample. In addition, the moderately OTM puts are, in contrast to the deep OTM puts, systematically underpriced in the period March 2012 to October 2013. This reveals a tension in jointly fitting the shape of the option skew, even with a fully flexible specification for the jump intensity. Qualitatively similar issues are apparent from the diagnostic option pricing tests for the Gaussian specification reported in Appendix E.




Figure 7: Fit to short-dated option prices using the semi-nonparametric model (10). The figure reports the daily test statistics (7) for our short-dated options based on model (10) with the parameters and state vector estimated from the short-maturity option sample. The regions cover options with tenor $\tau \leq 9$ and deep OTM puts (top panel, $-8 \leq m < -4$), OTM puts (middle panel, $-4 \leq m < 0$), and OTM calls (bottom panel, $0 \leq m < 5$). The solid lines indicate the symmetric 95% confidence band for the test.

In summary, it is evident that – despite the flexibility afforded by our semi-nonparametric models exploiting the jump specifications (9) and (10) – such representations cannot account satisfactorily for the time variation in the pricing of short-dated options. This implies that traditional parametric option pricing models using the same type of jump size distributions will fare even more poorly when confronted with the full cross-section of short-maturity option prices.

There are two potential explanations. One, the jump distribution is fixed, but our parametric representations are misspecified. Alternatively, the (risk-neutral) jump size distribution displays substantial variation over time and we cannot obtain an adequate fit to the option prices with an invariant distribution. To discriminate between these competing hypotheses, we turn to a nonparametric diagnostic for the left tail of the distribution. The point is that, under the null hypothesis of an invariant jump size distribution, the *relative* intensity of differently sized jumps is also invariant over time, i.e., $\nu_t(x)/\nu_t(x')$ is constant across all jump sizes $x, x' \in \mathbb{R}$.

To develop a generally applicable diagnostic, we exploit our generic semi-nonparametric specification for processes with time-invariant jump size distribution, $\nu_t(dx) = c_t \cdot \nu(dx)$, where c_t is an arbitrary intensity process. Combining this representation with the asymptotic result in the short-maturity approximation (3), it follows that the ratio $O_{t,\tau,k}/O_{t,\tau,k'}$, for sufficiently large fixed $k, k' > 0$ or $k, k' < 0$, is constant across all days t in the sample. To ensure that we isolate the pricing of jumps, we rely on very short-dated and far OTM options. Specifically, we consider the ratio $O_{t,\tau,-0.07}/O_{t,\tau,-0.06}$, for options with time to expiry of at most 4 days.




Figure 8: **The ratio of OTM puts with moneyness of -0.07 over OTM puts with moneyness of -0.06 .** The ratio is computed for each day in the sample where we have options with time-to-maturity less than or equal to four days (a total of 309 days in the sample). We first interpolate the BSIV for the desired level of moneyness using a fourth-order spline, and then convert back the implied volatility into a put option price. We report the five-day backward looking moving average of the ratio.

From Figure 8, it is evident that the relative expensiveness of -7% OTM puts over -6% OTM puts varies widely and quite persistently over the sample, contradicting the invariance of the jump distribution. For example, from October 2012 till around October 2013, the ratio gradually declines, implying a “thinning” of the negative jump tail. Similar developments are observed over March-July 2012 and January-March 2015. We further note that all these episodes are consistent with our diagnostic in Figure 7, which indicates systematic overpricing of deep OTM puts and underpricing of OTM puts over the latter part of these time intervals. Since the shape of the jump tail is invariant in equation (10), the model simply cannot accommodate the pronounced variation in the relative pricing of these options.

6.2 Models with Time-Varying Jump Distribution

Given the evidence in Section 6.1, we now generalize the semi-nonparametric models by allowing the parameters in equations (9) and (10) to vary over time. Our *time-varying* jump distribution

models take the form,

$$\nu(dx; \mathbf{J}_t, \theta) = c_t \frac{e^{-\frac{(x-\mu_{x,t})^2}{2\sigma_{x,t}^2}}}{\sqrt{2\pi}\sigma_{x,t}} dx, \quad (11)$$

and

$$\nu(dx; \mathbf{J}_t, \theta) = c_t \left\{ \frac{e^{-\lambda_t^- |x|}}{|x|^{1+\alpha}} 1_{\{x < 0\}} + \frac{e^{-\lambda_t^+ |x|}}{|x|^{1+\alpha}} 1_{\{x > 0\}} \right\} dx. \quad (12)$$

In these specifications, the volatility and jump intensity remain unrestricted, but the coefficients of the jump distributions now also shift freely across the observation times. Thus, all the jump coefficients are time-varying and no longer constitute part of the (fixed) parameter vector. Instead, they enter \mathbf{S}_t as additional state variables, necessary for a full characterization of the prevailing risk-neutral price dynamics. In other words, θ is now the empty set — formally $\theta = \emptyset$ — while $\mathbf{S}_t = (V_t, \mathbf{J}_t)$ with $\mathbf{J}_t = (c_t, \mu_{x,t}, \sigma_{x,t})$ and $\mathbf{J}_t = (c_t, \lambda_t^+, \lambda_t^-)$ for the two models, respectively. We stress that the two models (11) and (12), by construction, remain arbitrage-free.

As in Section 6.1, we estimate the system via criterion function (6). Since the parameter vector is absent (the empty set), this reduces to a sequence of optimization problems with the spot volatility, jump intensity, and jump distribution coefficients estimated independently for each observation date. Clearly, this provides a tremendous amount of flexibility in fitting the option cross-section. At the same time, the constraint that the intensity, c_t , enters symmetrically for positive and negative jumps is now even more critical. The abundance of quotes for deep OTM puts implies that the coefficients typically adjust to allow a good fit to the left tail of the implied volatility skew, while we occasionally observe a poor fit for the OTM calls as a consequence of the lack of identifying information for the right tail. Although this is a genuine limitation, we note that, from an economic perspective, the most interesting variation is associated with the negative jump tail. Specifically, the latter has been linked to risk premiums that account for a large component of the time-varying equity and variance risk premiums in the aggregate U.S. stock market.

This modeling approach is somewhat reminiscent of the calibration to option surfaces, often adopted by practitioners. However, there are fundamental differences. Model (1), complemented by the specification (11) or (12), is built on the principles of no arbitrage. We explicitly enforce the (statistical) equality of the diffusive volatility under the physical and risk-neutral measures in our estimation procedure. There is no counterpart to this no-arbitrage condition for “large” jumps due to the market incompleteness induced by jumps. Our flexible semi-nonparametric specifications (11) and (12) reflect this basic fact. In short, we rely on the short-dated options to inform us about time-variation in the pricing of jump risk as well as the state of spot volatility. Although longer-dated options also impose some “discipline” on the risk-neutral jump intensity, this information

cannot be extracted in a model-free manner.

Thus, in essence, we identify the features of the semimartingale model for X , i.e., the spot volatility and characteristics of the jump process, from the short-dated options written on X . This endows our day-by-day estimates with a direct “structural” interpretation vis-a-vis the state of the underlying risk-neutral asset price process, which is not possible on the basis of standard calibrations of the option surface through flexible curve-fitting procedures, e.g., smoothing splines. We subsequently employ our characterization to study the jump and volatility dynamics as well as the implications for longer-dated options and its implications for the equity risk premium. Again, such applications are precluded if calibration is performed through curve-fitting that is not linked directly to the underlying volatility, jump intensity and jump distribution.

The ability to accommodate each individual implied volatility skew through variation in the coefficients of the jump distribution across observation dates via “structural calibration” generates a substantial improvement in fit to the short-dated options. While we defer most of the detailed estimation results to Appendix F, we provide an overview of the diagnostic tests, as we now are close to a (statistically) acceptable characterization.

First, we consider the Gaussian jump specification, widely adopted in the empirical literature. Model (11) is obviously much more flexible than the standard representation which invokes fixed jump parameters and relies on a particular parametric representation for the spot volatility and jump intensity dynamics. The Z-scores associated with our pricing approximation for the short-maturity option sample are depicted in Appendix F. The rejections are far less frequent and less dramatic than observed previously. Nonetheless, the model is strongly rejected, as there are lengthy periods with pronounced clustering of Z-scores outside the standard error bands. Specifically, over the last fifteen months of the sample, the model-implied estimates overprice deep OTM puts, while close-to-ATM puts are underpriced, indicating the model has difficulty accommodating the shape of the left tail. Likewise, the diagnostics for the volatility fit in Appendix F reveal that the model-implied volatility often greatly exceeds the corresponding high-frequency based estimate.

Second, we consider the class of tempered stable distributions with time-varying tail parameters, λ_t^- and λ_t^+ , across observation dates. Consistent with our previous findings, the double-exponential jump representation ($\alpha = -1$) improves on the Gaussian model. The associated Z-score diagnostics, presented in Appendix F, indicate a superior fit relative to the Gaussian specification. However, it is still not stellar. There are indications of mispricing for close-to-ATM puts towards the end of the sample – less dramatic but qualitatively in line with the evidence for the Gaussian model – and there is an excessive number of instances with model-implied overestimation of spot volatility

relative to the nonparametric estimator. Since the shape of the implied volatility skew close to the origin is significantly affected by α , we explored alternative specifications. We find a substantial improvement for increasing values, culminating around $\alpha = 0.5$.²⁰ The Z-scores for the option price and volatility fit are depicted in Figures 9 and 10. The diagnostic checks are no longer indicative of systematic pricing errors for the OTM put options or any persistent errors in the model-implied volatility estimate. However, as for all models entertained, the pricing of OTM call options is somewhat problematic. This is, of course, not unexpected, given our prior observation that the quality of the option quotes in the right tail does not allow for precise day-by-day identification of the positive jump process.




Figure 9: **Fit to short-dated option prices using the semi-nonparametric model (12).** The figure reports the daily test statistics (7) for the short-maturity options based on the semi-nonparametric model (12) with $\alpha = 0.5$ and the state vector estimated from the short-dated options. The regions cover options with tenor $\tau \leq 9$ and deep OTM puts (top panel, $-8 \leq m < -4$), OTM puts (middle panel, $-4 \leq m < 0$), and OTM calls (bottom panel, $0 \leq m < 5$). The solid lines indicate the symmetric 95% confidence band for the test.

7 Characterizing the Time-Varying Jump Distributions

Our findings in Section 6 imply that a reasonable fit to the pricing of short-dated equity-index options requires accommodation of pronounced time-variation in the jump distribution. If this is,

²⁰We note that $\alpha = 0.5$ implies that jumps are of infinite activity under the risk-neutral probability. This property is preserved under equivalent martingale measure transformations, so this continues to hold for the statistical law. The model with $\alpha = 0$ represents a generalization of the variance gamma model in Madan et al. (1998) to allow for time-varying jump size and intensity parameters. It produces a comparable RMSE to the model with $\alpha = 0.5$, but results in slightly worse option and volatility Z-scores, so we retain $\alpha = 0.5$ as our main specification.




Figure 10: **Volatility fit based on the semi-nonparametric model (12)**. The top panel reports end-of-day spot volatilities based on nonparametric estimates from high-frequency data (solid line) or option-implied values (grey line). The option-implied volatility estimate is based on the semi-nonparametric model (12) with $\alpha = 0.5$ and the state vector estimated from the short-maturity option sample. The volatility estimates are reported in annualized units. The bottom panel reports the corresponding daily volatility test statistics (8). The solid lines indicate the symmetric 95% confidence band for the test.

indeed, a robust empirical finding, it has important implications for the proper specification of asset and derivatives pricing models. In this and the following section, we explore some of the main features of our final, and statistically satisfactory, semi-nonparametric representation for the risk-neutral asset return process. We focus on time-variation in the volatility and jump characteristics to gauge the extent to which they are economically rational and sensible.

7.1 The Shape of the Jump Distribution

The empirical results in Section 6 beg the question why the popular Gaussian specification, even with the extreme degree of flexibility afforded by unrestricted time-variation in the volatility, jump intensity and jump parameters, invariably underperforms our alternative representations. To illustrate a common shortcoming, Figure 11 displays the observed log-price for our short-dated options on a representative day, October 17, 2014, along with the corresponding model-implied ones generated by the Gaussian and tempered stable jump specifications. In the moneyness range $(-2, 0)$, the Gaussian model systematically underprices OTM put options, while it overprices them over the range $(-5.5, -2.5)$. Finally, for the extreme left tail, the options are again underpriced but this region contains relatively few observations and has a limited impact on the estimation. The tempered stable model, on the other hand, provides an impressive fit to the entire left tail. The

Gaussian specification generates a negative jump mode, inducing a pronounced non-convexity in the left tail as a function of m . In contrast, the actual log option prices consistently display a convex pattern in the left tail, signifying systematic mispricing by the Gaussian jump model.




Figure 11: **Option prices and model fit on October 17, 2014.** **Left Panel:** fitted log option prices according to the Gaussian model (11). **Right Panel:** fitted log option prices implied by the Generalized Tempered Stable model (12) with $\alpha = 0.5$. In both panels the dots represent observed option prices while the light gray line indicates model implied prices. The option tenor is nine days.

7.2 Time Variation in Diffusive and Jump Risk

We now look more closely at the time variation in the volatility and jump characteristics implied by the short-dated options. We rely on our best performing model (12) with $\alpha = 0.5$, in which the spot variance V_t , the jump intensity factor c_t , and the left and right tail shape parameters λ_t^\pm vary unrestrictedly. The top left panel in Figure 12 depicts the end-of-trading-day estimate of the spot variance. The volatility is mildly elevated at the start of the sample in January 2011 reflecting the aftermath of the initial phase of the European sovereign debt crisis. The temporary spike in mid March reflects the market reaction to the Fukushima nuclear disaster. A prolonged period of sharply higher volatility is observed during the late summer of 2011 through the end of the year. The U.S. federal debt was downgraded in August and the viability of the Italian and Spanish sovereign debt and banks became a major concern, with markets only calming down after the European Central Bank (ECB) announced the long-term refinancing operation (LTRO) for Euro-zone financial institutions. However, concerns resurfaced in 2012, prompting ECB President Mario Draghi to declare, in August 2012, that the central bank “was ready to do whatever it takes to preserve the single currency.” Subsequently, volatility was subdued, though interrupted in the second part of 2014 by the imposition of sanctions in the wake of the Russian annexation of Crimea, the weakening prospects for economic growth in Europe, and collapsing oil prices. Further, there

is a spike during the contentious negotiations about a resolution to the Greek debt crisis in June 2015, sparking fears of a “Grexit.” Finally, there is an explosion in volatility associated with the downturn in US stock markets, triggered by a collapse of Chinese equity prices and concerns about global economic growth at the very end of our sample in August 2015.

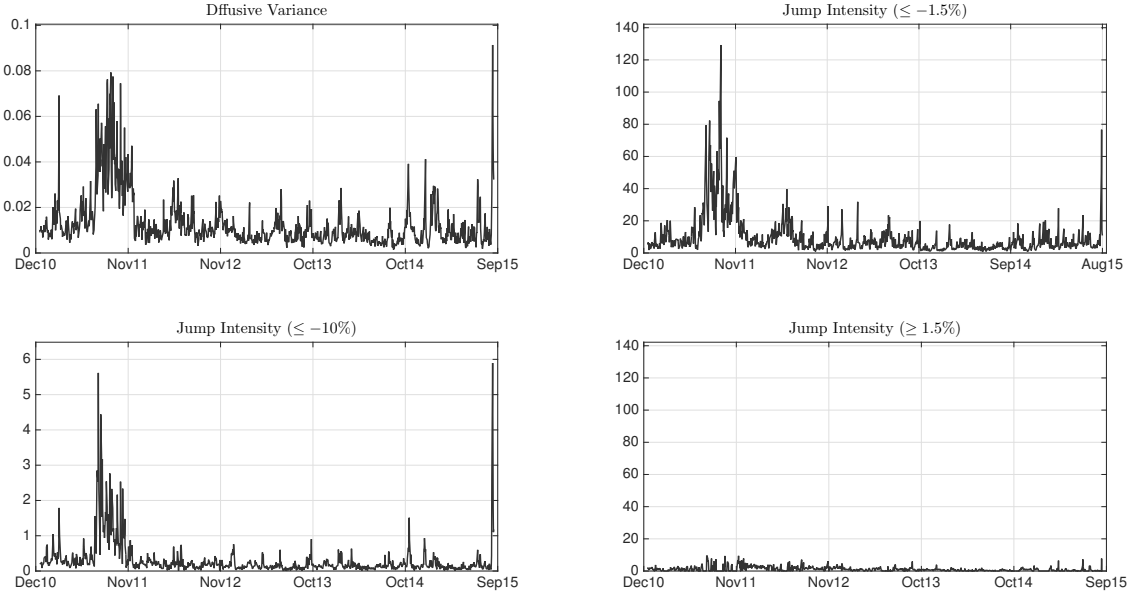


Figure 12: **Return Characteristics Extracted from the semi-nonparametric model (12).** We plot the day-by-day estimates for V_t (top left panel), $\int_{x < -0.015} \nu_t(dx)$ (top right panel), $\int_{x < -0.10} \nu_t(dx)$ (bottom left panel), and $\int_{x > 0.010} \nu_t(dx)$ (bottom right panel) expressed in annualized terms, based on the model (12) with $\alpha = 0.5$. The sample period is January 2011 - August 2015.

The remaining three panels of Figure 12 capture different features of the time-varying jump process. While the risk-neutral jump intensity is infinite in our generalized tempered stable model, the probability of jumps beyond any specific size is finite and readily interpretable. The top right panel provides the (annualized) expected number of jumps below -1.5% (under \mathbb{Q}). For most of the days, the probability is fairly low, yet consistent with many existing studies we find negative jumps to occur about one to five times per year. However, there are also several periods in which the conditional intensity rises to between 10 and 20 expected jumps, and during the stressed market conditions surrounding the sovereign debt crisis in 2011, the level typically exceeded 20 annual jumps for several months, topping out at a reading above 100. Likewise, there is a sharp spike in August 2015, where it reaches 70. These extreme observations are associated with turbulent market conditions and there is clearly some degree of cohesion between volatility and the intensity of moderately sized negative price jumps. However, there also are distinct differences. Foremost,

the higher degree of skewness in the jump series is evident, with the jump intensity often spiking proportionally more than the spot variance. Likewise, there are several instances of divergent behavior, including the absence of a large outlier in the jump series around the Fukushima incident.

The lower right panel provides a corresponding series for the intensity of positive jumps in excess of 1.5%. The vertical axis is scaled identically to the panel above, rendering the lower intensity of positive jumps (under \mathbb{Q}) striking visually.

Finally, the lower left panel depicts the probability of very large negative price jumps, beyond -10%, reflecting the time-varying fears of a genuine market crash implied by the pricing of deep OTM put options. While the intensity of moderate jumps is governed primarily by the intensity factor c_t , these crash probabilities also load heavily on the shape of the negative jump size distribution, controlled by the λ_t^- parameter. On most days, the expected number of crashes per year is very small, although the series clearly displays persistent fluctuations, even during times of subdued risk pricing. More significantly, however, are the dramatic spikes during the sovereign debt crisis and in August 2015 where the expected number of such jumps momentarily exceeds five annually. In addition, we observe smaller, but distinct, spikes in the crash probability during the stressed conditions following the Fukushima episode and in October 2014. In short, these tail probabilities do not appear to be tied closely to the evolution of spot variance, although extreme crash probabilities invariably are associated with volatile markets. The persistent variation in the jump intensity for (large) jumps relative to the spot variance is also evident from Figure 13. As we document later, these extreme tail pricing events also signify a pronounced variation in the equity risk premium.




Figure 13: Jump Intensity vs Diffusive Variance. The figure depicts the ratio of the estimated intensity of jumps smaller than -10%, i.e., $\int_{x < -0.10} \nu_t(dx)$, over the estimated spot variance, using model (12) with $\alpha = 0.5$. We take a 21-day backward looking moving average which is normalized to have mean zero and unit variance.

7.3 Implied Return Variation Measures

We next explore how the fluctuations in the features governing the intensity of the price movements, explored in Section 7.2, translate into standard measures of market risk such as the diffusive and jump components of the quadratic return variation. Exploiting the inference techniques developed in Andersen et al. (2015a), we provide standard error bands for several return variation series, thus conveying the level of precision with which we are able to extract the different measures from short-dated options. Initially, we focus on the spot (diffusive) variance, V_t , and the left and right jump variances, $LJV_t = \int_{x<0} x^2 \nu_t(dx)$ and $RJV_t = \int_{x>0} x^2 \nu_t(dx)$. Since the two latter quantities involve both very small (infinite activity) and larger jumps, we also present evidence regarding the return variation generated by large negative jumps, which we subsequently find particularly informative regarding the pricing of market-wide equity risk.

The diffusive, left and right jump variation constitute spot counterparts for the measures that comprise the theoretical valuation of the VIX index.²¹ For the jump specification in equation (12), the model-implied expressions for LJV_t and RJV_t are,

$$LJV_t = c_t (\lambda_t^-)^{\alpha-2} \Gamma(2-\alpha) \quad \text{and} \quad RJV_t = c_t (\lambda_t^+)^{\alpha-2} \Gamma(2-\alpha).$$

In Figure 14, the gray area in each panel indicates a 95% confidence region for the given spot return variation estimate, implied by our best performing model (12) with $\alpha = 0.5$. From the top panel, we note that the spot diffusive variance generally is recovered with good precision although, as expected, its confidence intervals become wider during volatile periods. In comparison, the alternative nonparametric estimator of V_t based on high-frequency return data is much less accurate.²² Nonetheless, these two spot volatility series are compatible, as documented in Section 6.2. Moreover, as noted previously, estimates for V_t based on regular (longer-dated) options are more prone to distortions due to model misspecification. Hence, the implied V_t values obtained from the semi-nonparametric specification (12) and short-dated options are both accurate and robust, arguably representing the most reliable spot variance estimates available in the literature.

The next two panels of Figure 14 depict the risk-neutral spot jump variation measures. Unlike for V_t , high-frequency data on X do not deliver model-free estimates for these quantities, underscoring the lack of market completeness in the presence of jumps. The left jump variation is estimated very accurately, with tighter confidence bounds than for V_t . In contrast, the relative precision

²¹Formally, the VIX measure equals the expected value of the sum of these three measures, cumulated over the coming calendar month, i.e., $VIX_t \approx \mathbb{E}_t^Q \left[\int_t^{t+\tau} (V_s + LJV_s + RJV_s) ds \right]$ with τ equal to one calendar month.

²²In separate analysis, we have confirmed that the asymptotic variances for the high-frequency based spot variance estimator are much larger than the option-implied ones.

in recovering the right jump variation is poor, reflecting the scarcity of deep OTM call options. However, this does not seriously impinge on the accuracy of our overall variation measure, since the magnitude of the right jump variation is inconsequential relative to the other components and the associated estimation errors are strongly negatively correlated with those for the spot variance, implying a large degree of cancelation of errors in the sum of these two measures.

Finally, the bottom panel portrays the negative tail variation, induced by jumps exceeding -10% on the downside. In our semi-nonparametric model, it takes the explicit form,

$$NTJV_t = \int_{x < -0.10} x^2 \nu_t(dx) = c_t (\lambda_t^-)^{\alpha-2} \Gamma(2 - \alpha, 0.1 \cdot \lambda_t^-),$$

where $\Gamma(x, y)$ denotes the upper incomplete gamma function. This is effectively a measure of the (risk-neutral) expected return variation associated with crash events. For the vast majority of days, this measure is essentially negligible. However, when the probability of such large negative jumps spikes, as displayed in Figure 12, we observe a dramatic rise in the proportion of the return variation stemming from this component. In fact, for the two extreme events, associated with the onset of the sovereign debt crisis in the summer of 2011 and the Chinese stock meltdown in August 2015, this tail component dominates the diffusive return variation. Likewise, for the two other notable spikes in the “crash probability” – the Fukushima incident and October 2014 – the left tail return variation measure is again the largest component of the overall return variation. In other words, during periods of market turmoil, fear of crash scenarios often turn into a prominent driver of the option-implied risk measures such as the VIX.

More generally, Figure 14 reveals interesting connections between the diffusive and the left jump variance. On average, they are roughly of equal magnitude. This provides a sharp contrast to the much more modest contribution of jumps to the return variation under the statistical measure, identified in numerous studies using high-frequency data on X . In other words, our results suggest a very rich pricing of jump risk. At the same time, Figure 14 also shows that the relative importance of the diffusive and left jump variation varies greatly over time. We highlight two episodes. In early August 2011, when there is a surge in both the diffusive and left jump variation, the latter is almost three times the size of the former. This period was dominated by fears linked to the S&P downgrade of the US credit rating and the sovereign debt crisis. While the level of general uncertainty rose sharply, the event disproportionately impacted tail risk. The second episode is the first half of October 2014 where, again, the diffusive and the left jump variance increase sharply reflecting, among other factors, concerns about weakness in the European economy and its potential global impact. However, unlike the previous episode, the diffusive variance remains roughly on par with the left jump variance throughout the period. Hence, in this instance, we observe an increased level

of general uncertainty, but no relative elevation in the importance of tail risk.

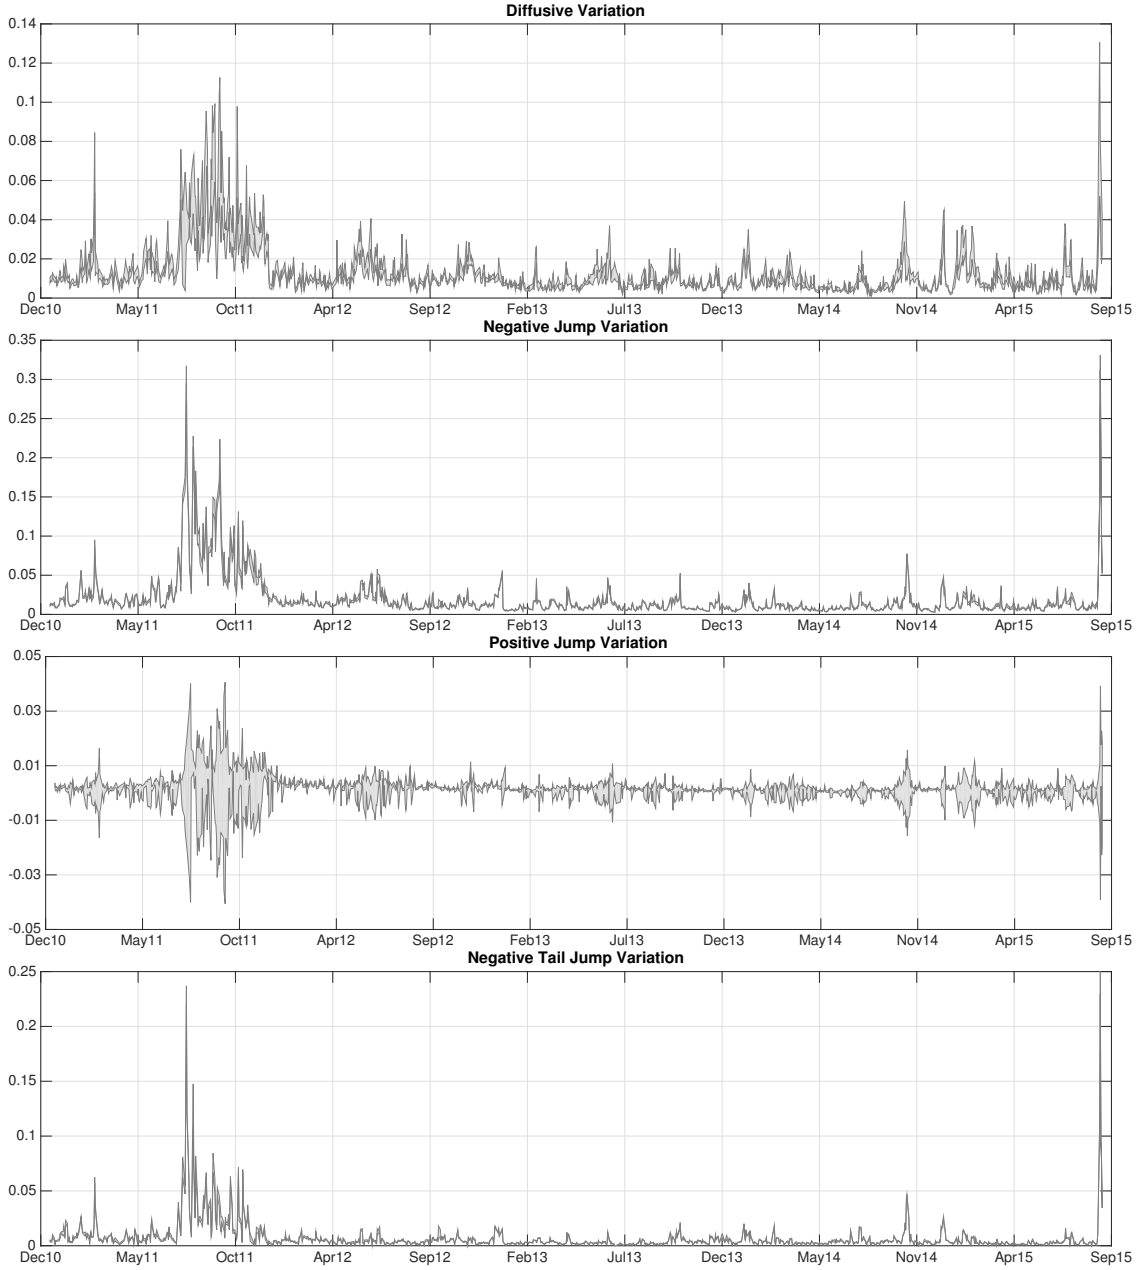


Figure 14: **Measures of Return Variation.** The figure plots 95% confidence intervals for the daily V_t (top panel), $\int_{x<0} x^2 \nu_t(dx)$ (second panel), $\int_{x>0} x^2 \nu_t(dx)$ (third panel) and $\int_{x<-0.10} x^2 \nu_t(dx)$ (bottom panel) implied by estimates for the semi-nonparametric model (12) with $\alpha = 0.5$ from the short-dated options. All quantities are reported in annualized units.

8 Asset Pricing Implications

We now exemplify how the time-variation in the short-dated option-implied left jump tail measure is associated with broader financial market developments. Section 8.1 documents that large fluctuations in our left jump variation measure systematically spill over into contemporaneous pricing errors for standard parametric models applied to regular option panels, confirming the existence of strong links between our jump variation measure and the state of the broader option market. Section 8.2 explores whether the left jump tail possesses explanatory power vis-a-vis the short-horizon equity risk premium, extending prior work based on longer-dated option measures.

8.1 Implication for Standard Option Pricing Models

Sections 1.1 and 1.2 illustrate how ongoing economic developments may induce significant shifts in the pricing of options that solely reflect negative tail events. This is consistent with the evidence in Section 7 documenting a substantial amount of variation in the pricing of negative market jumps that are not directly linked to market volatility.

It is natural to ask if our evidence concerning the pricing of jump risk is limited to short-dated options, or whether it has implications for the longer-dated options used in standard asset pricing and derivatives studies.²³ In theory, this is feasible if the pricing of jump risk is economically meaningful and persistent, as opposed to purely idiosyncratic. In this case, the parts of the sample characterized, via short-dated options, as representing elevated (risk-neutral) jump variation episodes should manifest themselves as distinct in terms of the pricing of longer-dated options as well. Of course, for longer horizons, other sources of risk start playing an important role. Hence, to better isolate the effect of jump risk pricing, we focus on moderately-dated OTM options, i.e., tenors of 10–45 days and moneyness $m \in [-7, -2]$. The SPX options in this category are highly liquid and constitute a critical component in the computation of the CBOE VIX index.

In the top panel of Figure 15, we display the tail jump variation extracted from short-dated options, via our semi-nonparametric representation (12) with $\alpha = 0.5$, along with the one implied by estimates from a general parametric two-factor jump-diffusive option pricing model and our longer-dated (regular) option sample. The parametric model is given in Appendix G.²⁴ It is evident that our tail estimates, extracted from the short-dated sample, are far more volatile than those obtained from the standard parametric model and commonly used option data. In particular, there are several instances where the former estimates for the tail jump variation spike sharply relative to

²³Early work experimenting with time-varying parameter models for standard longer-dated options include Bates (1991) and Bakshi et al. (1997).

²⁴This specification performs well relative to existing parametric models, see, e.g., Andersen et al. (2015b).

the latter. If these periods truly feature elevated tail risk and this, to a large extent, is missed by standard empirical option pricing procedures, then this should manifest itself in systematic pricing errors. Specifically, the parametric model should severely underprice deep OTM put options during such episodes. The lower panel of Figure 15 provides strong confirmatory evidence for this hypothesis. The panel displays formal Z-scores, indicating the degree to which the parametric model underprices the OTM put options within the regular sample. Notice that these option prices are part of the data used for estimation of the parametric model itself. Nonetheless, dramatic mispricing is evident, with Z-scores exceeding 10 on many occasions, and even topping 20 around the downgrade of US debt. The truly striking feature, however, is the high correlation of 0.46 between these episodes of blatant parametric mispricing for OTM puts and the periods for which we observe very large deviations between the jump variation estimates across the two models, indicated in the top panel.




Figure 15: **Tail Jump Variation.** **Top Panel:** Jump Variation implied by the parameters and state vector from the estimates of the parametric model defined by equations (19) and (20) based on the regular option panel and the semi-nonparametric model (12) with $\alpha = 0.5$ based on the short-dated options. For each model, we report the quantities $\int_{x < -0.10} x^2 \nu_t(dx)$, using the corresponding parameter and state vector estimates. **Bottom Panel:** Option Z-scores for the fit of the parametric model defined by equations (19) and (20) to options with moneyness $[-7, -2]$ and tenor between 10 and 45 days. The shaded area corresponds to Z-scores in excess of 2.5 (i.e., when the model severely underestimates OTM Puts). The correlation between the shaded areas on the two plots equals 0.46. All series are reported as five-day moving averages.

In summary, the jump variation measures, extracted exclusively from the short-dated option sample, are highly informative regarding qualitatively similar features in the regular option panel. Whenever the parametric model – estimated from and evaluated against regular longer-dated options only – fails to capture elevation in these jump variation measures, it also fails, often

spectacularly, in pricing in-sample observations on OTM put options. Hence, our flexible semi-nonparametric jump variation measures, constructed from short-dated options, identify large and systematic pricing errors in standard parametric models across critical segments of the implied volatility surface. Put differently, the tail risks priced in the very short-dated options are closely linked to those priced in standard option panels. Hence, the liquid trading of short-dated options affords us a simple and efficient procedure for identifying the risk-neutral jump tail risks.

8.2 The Short-Horizon Equity Risk Premium

The preceding sections have documented substantial independent variation in the negative risk-neutral tail and shown direct implications for the entire option surface. We now explore potential links between this tail factor and the underlying equity-index market. We already noted, in Section 7.3, that the negative tail jump variation (*NTJV*) is an order of magnitude larger than the corresponding return variation estimates obtained from actual equity return data. It suggests that the pricing of short-dated deep OTM put options captures a risk premium that is separate from the compensation of standard (diffusive) volatility risk. This feature is qualitatively consistent with the evidence in Andersen et al. (2015b), yet it is obtained through a nonparametric approach which cannot be directly reconciled with their parametric (affine) specification.

These considerations motivate exploring whether this tail risk measure contains useful information regarding future equity returns. A standard state variable in this context is, of course, the spot variance. Another celebrated factor is the so-called variance risk premium (*VRP*), which is defined as the gap between the VIX measure and the expected 30-day return variation, see, e.g., Bollerslev et al. (2009). We construct the latter as the difference between the CBOE VIX and a standard forecast of the expected one-month-ahead return variation for the S&P 500 index based on the HAR model of Corsi (2009) and high-frequency returns for the underlying S&P 500 index.²⁵

We first note that the model-implied spot variance V from our preferred nonparametric representation with $\alpha = 0.5$ is correlated with each of the supplementary measures, although the latter are weakly correlated only. The sample correlations of the spot variance with *VRP* and *NTJV* are 0.51 and 0.57, respectively, while the *VRP-NTJV* correlation is a mere 0.12. Turning to the predictive evidence, we now regress the future excess equity returns on the current values of the alternative explanatory variables.

Table 2 provides results for univariate regressions of the realized equity risk premium on each variable of interest as well as bivariate regressions involving the *NTJV* as one of the regressors.

²⁵See Appendix H for the construction of these forecasts.

Given the short sample for which we have the requisite observations on the short-dated options, we limit the evidence to the one- and two-week horizon.

| | $R_{t,t+5}$ | $R_{t,t+5}$ | $R_{t,t+5}$ | $R_{t,t+5}$ | $R_{t,t+5}$ |
|-------------------------|--------------|--------------|--------------|--------------|--------------|
| Spot Variance (V_t) | 8.123 | | | | 5.829 |
| | (3.531) | | | | (2.312) |
| VRP (30 days) | | 3.971 | | 2.446 | |
| | | (1.790) | | (1.435) | |
| Negative Tail JV | | | 6.254 | 5.539 | 3.261 |
| | | | (4.507) | (4.345) | (2.413) |
| R^2 | 0.054 | 0.019 | 0.044 | 0.050 | 0.062 |
| | $R_{t,t+10}$ | $R_{t,t+10}$ | $R_{t,t+10}$ | $R_{t,t+10}$ | $R_{t,t+10}$ |
| Spot Variance (V_t) | 5.643 | | | | 3.884 |
| | (3.047) | | | | (2.033) |
| VRP (30 days) | | 1.561 | | 0.413 | |
| | | (0.865) | | (0.308) | |
| Negative Tail JV | | | 4.626 | 4.508 | 2.536 |
| | | | (4.160) | (4.752) | (2.321) |
| R^2 | 0.066 | 0.007 | 0.057 | 0.057 | 0.077 |

Table 2: **Predictive regression on future excess market returns.** Data on excess market returns are from Kenneth R. French website. The sample covers the period January 2011 – August 2015. For each regression we report the value of the regression coefficients with the corresponding Newey-West robust t-statistics (we include a number of lags equal to the number of forecasted days) in parenthesis; the regression R^2 is reported in the last row of each panel.

The top panel reports on the five (trading) day regressions. Both the spot variance and NTJV are highly significant predictors, with the former achieving a slightly higher degree of explained variation, while the latter enters more significantly. In contrast, VRP obtains the correct sign, but is insignificant. When combining NTJV with V , NJTV remains the most significant predictor and the regression R^2 rises to 6%, a substantial value given the sizable and highly idiosyncratic shocks

that dominate the return realizations over short horizons. The bottom panel reports on the ten-day regressions. This horizon allows for better diversification of idiosyncratic return innovations but the effective sample size shrinks, limiting our statistical power. Nonetheless, we now find NTJV to enter even more significantly, while the relative explanatory power of spot variance drops, and the explanatory power of VRP largely vanishes. For example, when added to the univariate NTJV regression, VRP adds no auxiliary forecast power. Likewise, in the last column, the NTJV measure enters more significantly than the spot variance.

Overall, the evidence is consistent with both the spot variance and negative jump tail variation being informative regarding the short-term equity risk premium, with the tail variation providing superior forecast power. This points to the importance of negative tail risk pricing in understanding the overall equity risk premium. It suggests that the pronounced fluctuations in the risk-neutral left jump tail are of broad economic significance, reflecting a joint and consistent variation in the reward for risk across the equity and short- and longer-maturity option markets.

9 Concluding Remarks

In this paper we study short-dated options written on the S&P 500 market index. The trading in this segment of the equity-index option market has increased significantly over the last five years with the introduction and increased popularity of the so-called “weeklies.” We seek to understand the dynamic behavior of this sizable and rapidly growing component of the option market which hitherto is largely unexplored. We note that these instruments provide an easy way to separately manage market jumps and diffusive risks. The flip side is that they provide for a simple and robust way to identify market volatility and priced tail risk. We exploit this informational role of the short-dated options to identify and study the dynamic properties of priced jump risk. Our analysis shows that the latter exhibits significant variation which cannot be directly associated with market volatility. This runs counter to the modeling approach adopted by the vast majority of no-arbitrage reduced-form and equilibrium-based models. We also show that the dynamic behavior of priced tail risk is reflected in the pricing of longer-dated options. Moreover, we document that this feature is a major source of misspecification for traditional asset pricing models, designed to capture the variation in option-implied volatility surfaces. Our analysis yields an easy-to-construct measure of priced negative jump tail risk. It can be used to capture investor concerns about this critical component of risk and for direct exploration of its dynamic behavior as well as its interaction with other economic quantities.

Our results suggest that it will be useful in future research to include the priced tail risk as a

genuine state variable, distinct from market volatility, into parametric models. It would reflect the current concerns regarding negative tail events, priced in the economy separately from volatility. Our semi-nonparametric model (12) indicates that, in order to capture such variation, it may be necessary to disentangle level shifts in the pricing of tail risk from changes in the shape of the tail, i.e., identify c_t separately from λ_t^- . In the current study, we deliberately do not impose any restriction on the variation of these processes. Our goal is to let the short-dated options convey information directly, unencumbered by potentially misspecified parametric representations. However, in light of the above evidence, a natural next step is to model the dynamic evolution of the priced tail risk by parametric means. This will allow us to incorporate information about tail risks and volatility embedded in short-maturity options into the pricing of regular longer-dated options. It should boost the performance of such models in capturing the price dynamics of broader option panels and, in turn, allow for a more robust extraction of information regarding the pricing of other risks that are critical for the valuation of longer-dated options, like the term structure of volatility and the pricing of volatility and jump intensity shocks, i.e., the main risks driving changes in the investment opportunity set.

From an economic perspective, it is of interest to explore the origin behind the pronounced fluctuation in the price of jump tail risk. The large variation in the shape of the risk-neutral jump distribution, implied by the short-dated options, is notoriously difficult to reconcile with the dynamics of market jump risks obtained from actual return data, see, e.g., the nonparametric analysis in Bollerslev and Todorov (2011). This may be explained, in part, by the limited number of actual tail risk realizations. Nevertheless, the current evidence suggests that a large component of the observed variation in the risk-neutral jump variation stems from shifts in the pricing of negative jump tail risk. That is, there is a sharp separation between the dynamics of the actual jump risk and its pricing. This is indicative of nonlinearities in the aggregate pricing kernel along the dimensions associated with the pricing of market jump risk. This may reflect time-varying aversion towards tail risk and/or heterogeneity in the attitude towards tail risk among investors. Identifying the economic forces that can rationalize our empirical evidence requires an extensive study exploiting data on the full option panel and the underlying asset as well as, potentially, volume data for option trading.

References

- Andersen, T. G., L. Benzoni, and J. Lund (2002). An Empirical Investigation of Continuous-Time Equity Return Models. *Journal of Finance* 57, 1239–1284.
- Andersen, T. G., N. Fusari, and V. Todorov (2015a). Parametric Inference and Dynamic State Recovery from Option Panels. *Econometrica* 83, 1081–1145.
- Andersen, T. G., N. Fusari, and V. Todorov (2015b). The Risk Premia Embedded in Index Options. *Journal of Financial Economics* 117, 558–584.
- Bakshi, G., C. Cao, and Z. Chen (1997). Empirical Performance of Alternative Option Pricing Models. *Journal of Finance* 52, 2003–2049.
- Bakshi, G. and D. Madan (2000). Spanning and Derivative Security Valuation. *Journal of Financial Economics* 55, 205–238.
- Bates, D. S. (1991). The Crash of '87 – Was It Expected? The Evidence from Options Markets. *Journal of Finance* 46, 1009–1044.
- Bates, D. S. (2000). Post-'87 Crash Fears in S&P 500 Future Options. *Journal of Econometrics* 94, 181–238.
- Bates, D. S. (2012). U.S. Stock Market Crash Risk, 1926 - 2010. *Journal of Financial Economics* 105, 229–259.
- Bentata, A. and R. Cont (2012). Short-Time Asymptotics for Marginal Distributions of Semimartingales. Working paper.
- Bollerslev, T., G. Tauchen, and H. Zhou (2009). Expected Stock Returns and Variance Risk Premia. *Review of Financial Studies* 22, 4463–4492.
- Bollerslev, T. and V. Todorov (2011). Estimation of Jump Tails. *Econometrica* 79, 1727–1783.
- Bollerslev, T. and V. Todorov (2014). Time Varying Jump Tails. *Journal of Econometrics* 183, 168–180.
- Bollerslev, T., V. Todorov, and L. Xu (2015). Tail Risk Premia and Return Predictability. *Journal of Financial Economics* 118, 113–134.
- Breeden, D. and R. Litzenberger (1978). Prices of State Contingent Claims Implicit in Option Prices. *Journal of Business* 51, 621–652.
- Broadie, M., M. Chernov, and M. Johannes (2007). Specification and Risk Premiums: The Information in S&P 500 Futures Options. *Journal of Finance* 62, 1453–1490.
- Carr, P., H. Geman, D. Madan, and M. Yor (2002). The Fine Structure of Asset Returns: An Empirical Investigation. *Journal of Business* 75, 305–332.
- Carr, P., H. Geman, D. Madan, and M. Yor (2003). Stochastic Volatility for Lévy Processes. *Mathematical Finance* 13, 345–382.
- Carr, P. and L. Wu (2003). What Type of Process Underlies Options? A Simple Robust Test. *Journal of Finance* 58, 2581–2610.
- Christoffersen, P., K. Jacobs, and C. Ornathanalai (2012). Dynamic Jump Intensities and Risk Premiums: Evidence from S&P 500 Returns and Options. *Journal of Financial Economics* 106, 447–472.

- Collin-Dufresne, P., R. Goldstein, and C. Jones (2008). Identification of Maximal Affine Term Structure Models. *Journal of Finance* 63, 743–795.
- Cont, R. and P. Tankov (2004). *Financial Modelling with Jump Processes*. Boca Raton, Florida, U.S.A.: Chapman and Hall.
- Corsi, F. (2009). A Simple Approximate Long-Memory Model of Realized Volatility. *Journal of Financial Econometrics* 7, 174–196.
- Drechsler, I. and A. Yaron (2011). What’s Vol Got to Do with It? *Review of Financial Studies* 24, 1–45.
- Duffie, D., J. Pan, and K. Singleton (2000). Transform Analysis and Asset Pricing for Affine Jump-Diffusions. *Econometrica* 68, 1343–1376.
- Durrleman, V. (2008). Convergence of at-the-money Implied Volatilities to the Spot Volatility. *Journal of Applied Probability* 45, 542–550.
- Eraker, B. (2004). Do Stock Prices and Volatility Jump? Reconciling Evidence from Spot and Option Prices. *Journal of Finance* 59, 1367–1403.
- Fang, F. and C. Oosterlee (2008). A Novel Pricing Method for European Options Based on Fourier-Cosine Series Expansions. *SIAM Journal on Scientific Computing* 31(2), 826–848.
- Green, R. and R. Jarrow (1987). Spanning and Completeness in Markets with Contingent Claims. *Journal of Economic Theory* 41, 202–210.
- Jacod, J. and P. Protter (2012). *Discretization of Processes*. Berlin: Springer-Verlag.
- Jarrow, R. and S. Kwok (2015). Specification Tests of Calibrated Option Pricing Models. *Journal of Econometrics* 189, 397–414.
- Kou, S. (2002). A Jump Diffusion Model for Option Pricing. *Management Science* 48, 1086–1101.
- Madan, D. B., P. Carr, and E. C. Chang (1998). The Variance Gamma and Option Pricing. *European Finance Review* 2, 79–105.
- Medvedev, A. and O. Scaillet (2007). Approximation and Calibration of Short-Term Implied Volatilities under Jump-Diffusion Stochastic Volatility. *Review of Financial Studies* 20, 427–459.
- Merton, R. (1976). Option Pricing when Underlying Asset Returns are Discontinuous. *Journal of Financial Economics* 3, 125–144.
- Nachman, D. (1988). Spanning and Completeness with Options. *Review of Financial Studies* 1, 311–328.
- Pan, J. (2002). The Jump-Risk Premia Implicit in Options: Evidence from an Integrated Time-Series Study. *Journal of Financial Economics* 63, 3–50.
- Pan, J. and J. Liu (2003). Dynamic Derivative Strategies. *Journal of Financial Economics* 69, 401–430.
- Ross, S. (1976). Options and efficiency. *Quarterly Journal of Economics* 90, 75–89.
- Wachter, J. A. (2013). Can Time-Varying Risk of Rare Disasters Explain Aggregate Stock Market Volatility? *Journal of Finance* 68, 987–1035.

Appendix

A Illustrative Shifts in the Left Option Tail

Figure 1 provides an example of an isolated left tail shift in the risk-neutral distribution during the year 2012. This tail event is readily identified from the changes in the prices of short-dated options across the strike range, but it is hard to discern from prices of options with tenors akin to those used in regular option pricing studies.

Below, we provide illustrations of similar isolated left tail shifts during the other full calendar years covered by our sample period. The depiction of the lines indicating the option prices for the day of the tail shift (dark and full drawn) and for the surrounding trading days (lighter gray and dashed or dotted) follow the convention of Figure 1.

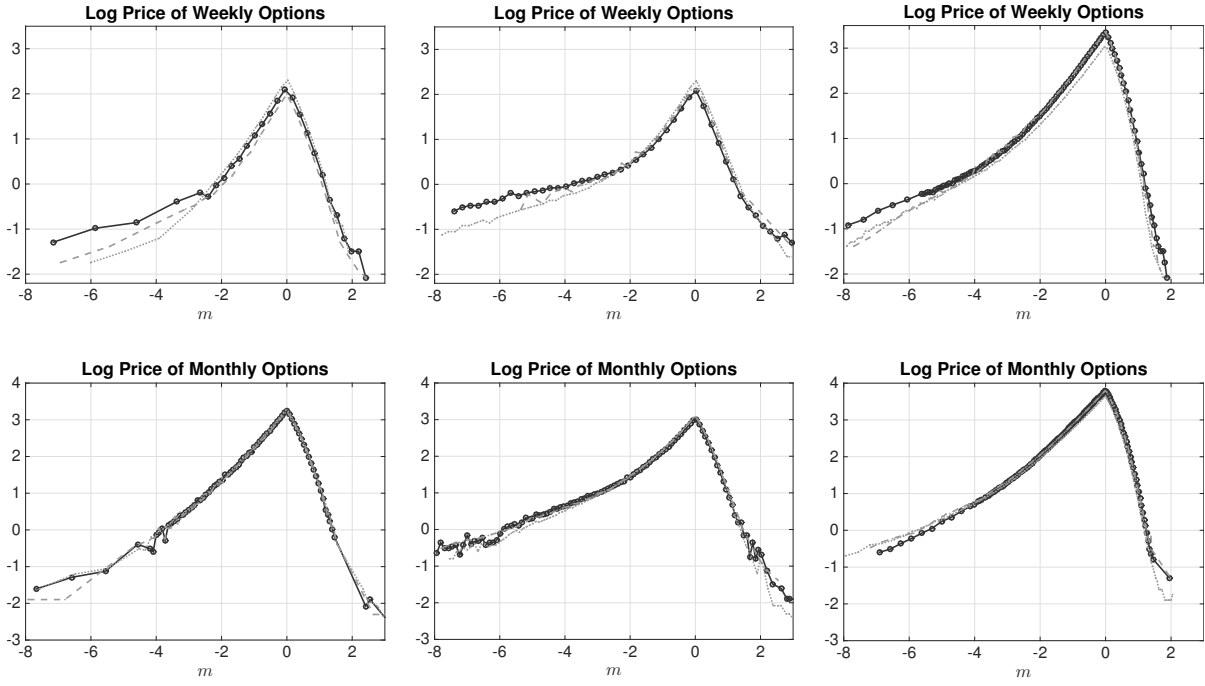


Figure 16: **Log Option Prices.** **First column: January 7, 2011.** The short-dated options in the top and regular options in the bottom panel have 7 and 42 calendar days to expiry. The surrounding trading days are January 6 and 10, 2011. **Second column: December 20, 2013.** The short-dated options in the top and regular options in the bottom panel have 7 and 28 calendar days to expiry. The surrounding trading days are December 19 and 23, 2013. **Third column: October 15, 2014.** The short-dated options in the top and regular options in the bottom panel have 9 and 30 calendar days to expiry. The surrounding trading days are October 17 and 16, 2014.

B Quality of the CME High-Frequency Option Data

Below, we depict measures of the bid-ask spread and quote staleness around the relevant FOMC statements. For some events, we observe a slight and fairly uniform elevation in the spread. However, the shifts in Figure 2 are clearly not linked directly to the size of the spreads. First, the option price patterns are much smoother than those for the spreads, and, second, they shift in a different manner, with ATM prices declining and deep OTM put prices remaining stable or increasing. In short, there is no sign of a distortion in the option price pattern due to a change in the spreads.

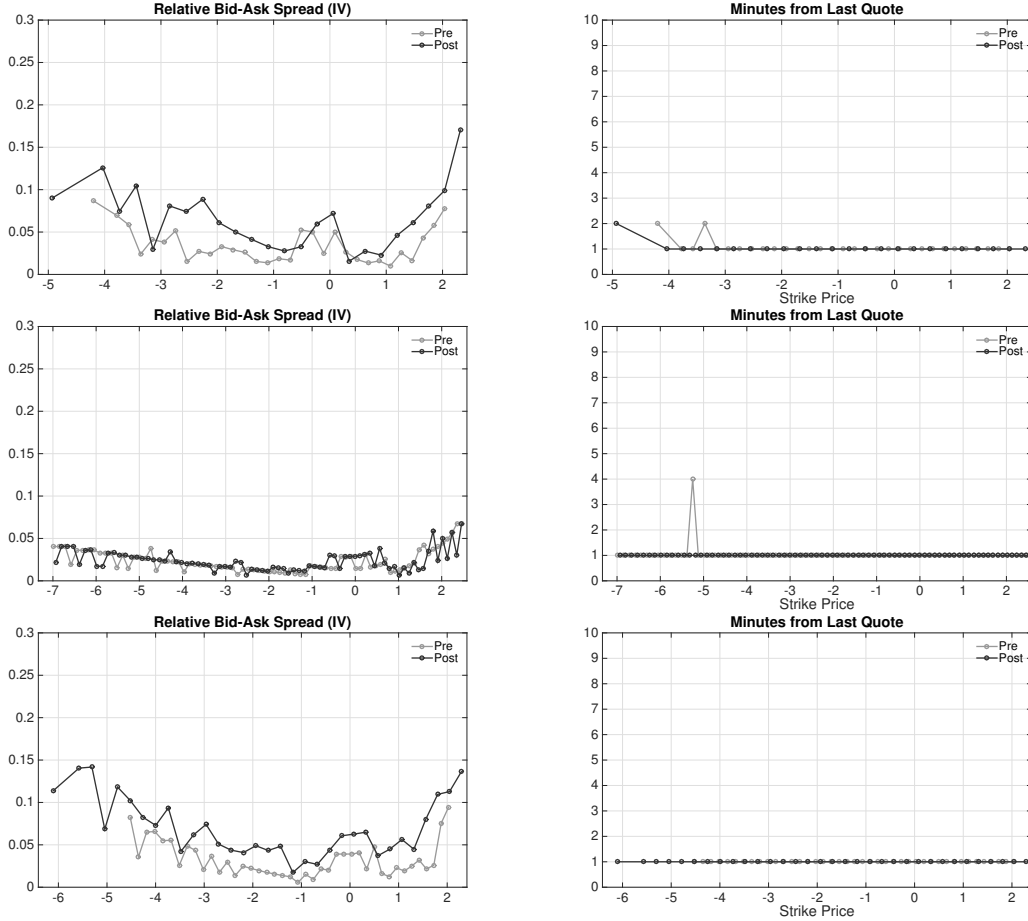


Figure 17: **Option Prices Pre- and Post-FOMC Announcements.** The rows correspond to the shortest tenor available for the CME e-mini S&P 500 futures options forty-five minutes before and after the FOMC announcements on September 18th 2013, October 30th 2013, and December 18th 2013. The announcements occur at 2 p.m. ET (1 p.m. CT). The left panels show the relative bid-ask option spread measured in implied volatility as $(IV_{ask} - IV_{bid})/IV_{mid}$ where IV_{ask} , IV_{ask} , and IV_{ask} are the implied volatilities corresponding to the ask, bid, and mid price, respectively. The right panels show, for each option contract, how many minutes have passed since either the bid or the ask quotes have been updated. On October 30th 2013 the tenors are 16 and 51 days. For the two remaining days, the tenors are 2 and 30 days.

C Data Filters

We apply a first set of filters, retaining only trading days and option quotes satisfying the following criteria: (i) For each included tenor, there are at least ten distinct option quotes across the strike range; (ii) The front maturity contract has at most 9 calendar days to maturity; (iii) The maturity is less than or equal to 365 calendar days; (iv) The moneyness is not extreme: $-15 \leq m \leq 5$; (v) The ratio of ask to bid price is less than five: $\frac{Ask}{Bid} < 5$. Notice this also implies that the bid quotes must be strictly positive; (vi) The front-maturity contract has valid quotes for OTM put options beyond moneyness $m = -3.5$; (vii) It is not an abbreviated trading day, a U.S. holiday, or a low-activity trading day just prior to a U.S. holiday.

The first criterion ensures that we mitigate noise by diversifying the measurement errors across multiple contracts at each maturity, enabling more accurate inference. The availability of multiple contracts is also important for robust identification of the forward price, which is a critical input to our analysis. We have 2,131,436 quotes satisfying this initial condition. The second criterion guarantees we have short-dated options available throughout the sample, while the exclusion of maturities beyond one year mirrors standard practice.²⁶ These two criteria eliminate about 280,000 additional quotes. The subsequent moneyness and bid-ask quote conditions leave us with slightly less than 1.5 million quotes. The OTM coverage condition eliminates five trading days during August 2011. Finally, removing holidays and partial trading days leaves us with 1,453,152 quotes.

We next split these “eligible” quotes into two distinct maturity categories: “short-dated” options with tenor less than or equal to nine calendar days versus “regular” options with strictly more than nine days to expiry. The sorting brings out another source of heterogeneity. The short-dated options, almost tautologically, are cheaper than longer-dated ones. Given the tick size of \$0.05, the relative error induced by the rounding of option values to the price grid is larger for short-dated options. Likewise, the effect of illiquidity can be exacerbated for the cheaper options. Hence, we impose additional filters on the short-maturity sample: (a) for each trading day, we retain only the (most liquid) cross-section with the shortest maturity; (b) The moneyness is further restricted to $-8 \leq m \leq 5$; (c) at least five units of each included contract are traded during the day; (d) we impose a final condition, detailed in Appendix C.1, to remove stale quotes or quotes that constitute clear (no-arbitrage) violations. Since short-maturity options are critical for our inference, these filters provide an additional safeguard against biases stemming from excessively noisy quotes.

These final conditions reduce our short-maturity sample from 82,528 to 51,670 contracts, cov-

²⁶The availability of quotes for actively traded options with tenor beyond one year varies over time and is correlated with market conditions. Hence, their inclusion will bias our comparison of the characteristics across our short- and long-dated option samples.

ering 1,105 trading days. Thus, we have, on average, 46.7 bid-ask quotes at day’s end for a single short maturity across a broad strike range, with a minimum of five units traded for each contract during the course of the day. Consistent with standard practice, we do not impose active trading conditions on the longer-dated options. The latter are subject to relatively less distortion from discreteness or illiquidity. Our set of “regular” options comprises 1,370,624 contracts with a tenor-strike composition roughly matching that of the extant literature. If anything, it provides better coverage of relatively shorter maturities due to the growth in the trading of weeklies. These contracts are introduced in a staggered fashion with an initial maturity of about six weeks. Hence, these contracts typically first enter our regular option sample and only later our short-dated sample.

C.1 Additional Filters

It can be problematic to infer the underlying market value of very deep OTM options from the corresponding bid and ask quotes. The tick size becomes large relative to the option value and the options may be quite illiquid for some of the strikes in the extreme tail, as market makers seek to concentrate trading in a few contracts. Both factors tend to inflate the percentage spread and induce non-monotonicity in the spread midpoint (as a function of strike price). Consequently, the relative pricing of the options in this region is particularly noisy. In particular, one will often encounter a sequence of identical quotes in the tail end of the strike range. Such “flat” pricing induces very thick tails in the extracted risk-neutral density. To mitigate the impact of rounding and potential illiquidity or quote staleness, we impose a filter that eliminates flat pricing and non-monotonicity for the quote midpoints in the extreme tail from our short-dated option sample. We stress that this, all else equal, will induce thinner tails in our estimates for the size of the jump distribution.

Specifically, for OTM puts, we remove all quotes at the end of the moneyness spectrum until the extreme quote midpoint is smaller than all other quote midpoints for put options positioned closer to the ATM strike. A similar procedure is applied separately for the call options each trading day. Figure 18 illustrates how the filtering procedure works for three trading days in our sample.

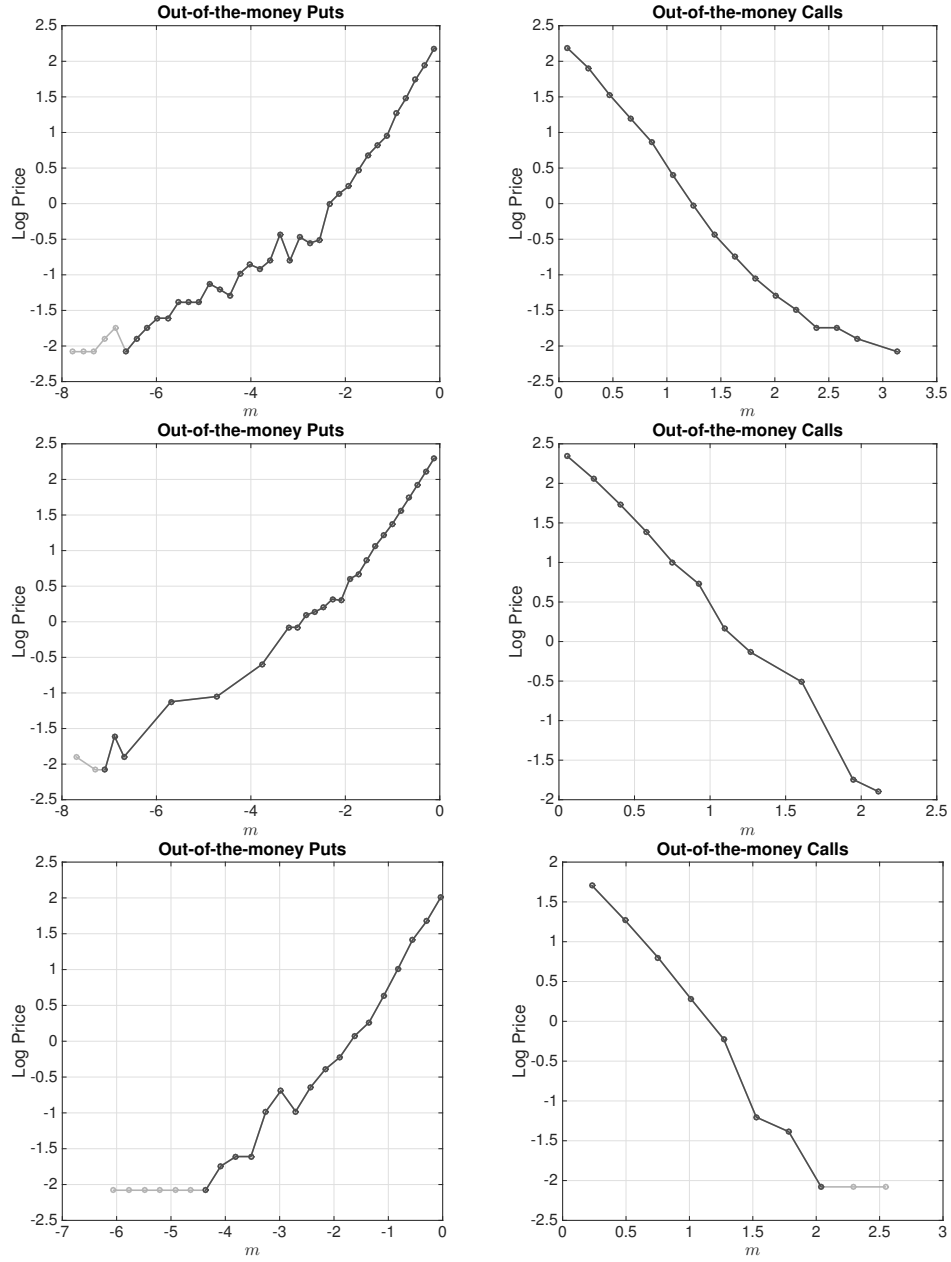


Figure 18: **Data Filters.** Illustration of the additional filtering procedure applied to the option data. Top panel: January 13th 2012, middle panel: March 21st 2012 and bottom panel: May 16th 2012. The dots on the plots correspond to the observed option prices (in logs). The solid line connects the data points that are preserved after applying the filter, while the grey line connects the data points which are removed after applying the filter.

D On the Computation of Option Prices

This section provides explicit formulas for the conditional characteristic function (CCF) of the log-return for our alternative jump size distributions. From the CCF, one can price options according to one of the many Fourier method proposed in the literature. We found the Fourier-Cosine series approximation method by Fang and Oosterlee (2008) to be fast and accurate.

For both our parametric representations, the CCF takes the form

$$\mathbb{E}_t[e^{uy_{t+\tau}}] = e^{\alpha_t(\tau, u)}, \quad u \in \mathbb{C},$$

where $y_{t+\tau} = \ln(X_{t+\tau}/X_t)$. Specifically we have,

- for the Model with Gaussian jumps in Equation (11), and $\tilde{\mu} = e^{\mu_{x,t} + \frac{\sigma_{x,t}^2}{2}} - 1$:

$$\alpha_t(\tau, u) = \tau \left[u \left((r_t - \delta_t) - \frac{V_t}{2} \right) + \frac{u^2}{2} V_t + c_t \left(e^{u\mu_{x,t} + \frac{u^2}{2}\sigma_{x,t}^2} - 1 - u\tilde{\mu} \right) \right].$$

- for the model with generalized tempered stable jumps in equation (12):

$$\alpha_t(\tau, u) = \tau \left[u \left((r_t - \delta_t) - \Psi(1) - \frac{V_t}{2} \right) + \frac{u^2}{2} V_t + \Psi(u) \right]$$

where $\Psi(u) = \Psi^-(u) + \Psi^+(u)$. Specifically, from Cont and Tankov (2004), p. 122, we have:

- if $\alpha_t^- = \alpha_t^+ \neq 0$ and $\alpha_t^- = \alpha_t^+ \neq 1$,

$$\Psi^-(u) = \Gamma(-\alpha^-)(\lambda_t^{-\alpha^-})c_t^- \left\{ \left(1 + \frac{u}{\lambda_t^-} \right)^{\alpha^-} - 1 - \frac{u\alpha^-}{\lambda_t^-} \right\} \quad (13)$$

$$\Psi^+(u) = \Gamma(-\alpha^+)(\lambda_t^{+\alpha^+})c_t^+ \left\{ \left(1 - \frac{u}{\lambda_t^+} \right)^{\alpha^+} - 1 + \frac{u\alpha^+}{\lambda_t^+} \right\} \quad (14)$$

- if $\alpha_t^- = \alpha_t^+ = 0$,

$$\Psi^-(u) = -c_t^- \left\{ -\frac{u}{\lambda_t^-} + \ln \left(1 + \frac{u}{\lambda_t^-} \right) \right\} \quad (15)$$

$$\Psi^+(u) = -c_t^+ \left\{ -\frac{u}{\lambda_t^+} + \ln \left(1 - \frac{u}{\lambda_t^+} \right) \right\} \quad (16)$$

- if $\alpha_t^- = \alpha_t^+ = 1$,

$$\Psi^-(u) = -c_t^- u + c_t^- (\lambda_t^- + u) \ln \left(1 + \frac{u}{\lambda_t^-} \right) \quad (17)$$

$$\Psi^+(u) = +c_t^+ u + c_t^+ (\lambda_t^+ - u) \ln \left(1 - \frac{u}{\lambda_t^+} \right) \quad (18)$$

Cases with $\alpha^- \neq \alpha^+$ can be handled readily, but in all our model specifications, we have $\alpha^- = \alpha^+ = \alpha$ and $c_t^- = c_t^+ = c_t$.

E Semi-Nonparametric Estimation with Fixed Jump Parameters

| | Mean | Std | Q05 | Q50 | Q95 |
|---------------------------------|-------|-------|-------|-------|-------|
| Gaussian Jumps | | | | | |
| RMSE | 1.238 | 0.732 | 0.561 | 1.072 | 2.400 |
| Diffusive Variance | 0.018 | 0.015 | 0.006 | 0.014 | 0.047 |
| Negative Jump Variance | 0.015 | 0.027 | 0.003 | 0.007 | 0.050 |
| Positive Jump Variance | 0.001 | 0.002 | 0.000 | 0.000 | 0.003 |
| Double Exponential Jumps | | | | | |
| RMSE | 1.063 | 0.645 | 0.491 | 0.929 | 1.966 |
| Diffusive Variance | 0.016 | 0.014 | 0.005 | 0.012 | 0.042 |
| Negative Jump Variance | 0.016 | 0.028 | 0.004 | 0.009 | 0.053 |
| Positive Jump Variance | 0.001 | 0.002 | 0.000 | 0.001 | 0.004 |

Table 3: **Estimation results for the semi-nonparametric models (9) and (10).** The columns provide the mean, standard deviation and 5th, 50th and 95th quantiles for the sequence of daily estimates obtained for the series indicated in the rows. The estimates are obtained from our short-dated option sample, covering January 2011 – August 2015. The Gaussian jumps refer to model (9) and the double exponential jumps to model (10) with $\alpha = -1$. The Root-Mean-Squared Error (RMSE) summarizes the fit to the end-of-day cross-section of annualized Black-Scholes implied volatilities (BSIVs). All variances are also reported in annualized units.

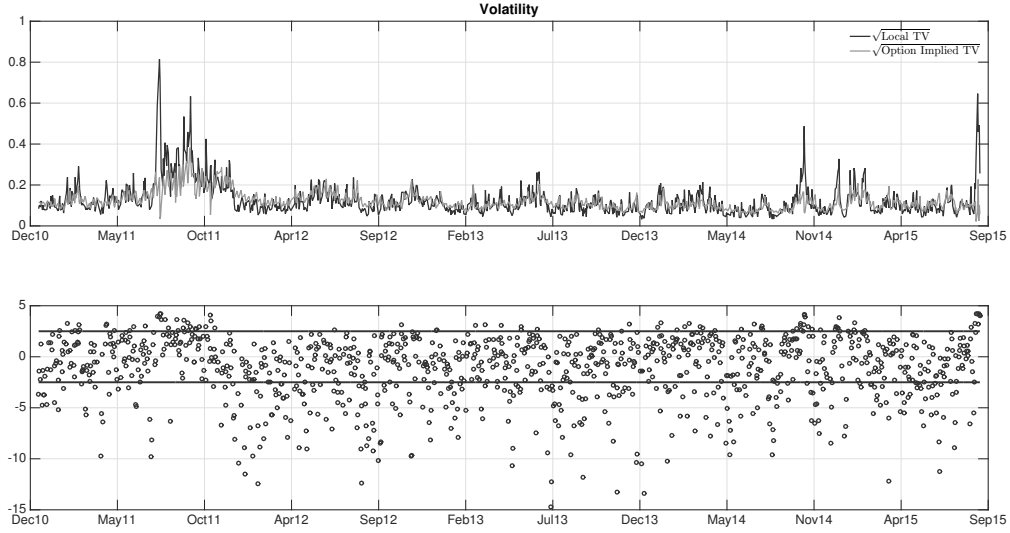


Figure 19: **Fit to spot volatility using the semi-nonparametric model (10).** The top panel reports end-of-day spot volatilities based on nonparametric estimates from high-frequency data (solid line) or option-implied values (grey line). The option-implied volatility estimate is based on model (10) with the parameters and state vector estimated from our short-dated options each trading day. The volatility estimates are reported in annualized units. The bottom panel reports the corresponding daily volatility test statistics (8). The solid lines indicate the symmetric 95% confidence band for the test.




Figure 20: **The fit to short-dated option prices based on the semi-nonparametric model with Gaussian jumps estimated from the short-maturity sample.** The figure reports the daily test statistics (7) for the short-maturity options based on the semi-nonparametric model (9) with the parameters and state vector estimated from the short-dated options. The regions of the option cross-section are for tenor $\tau \leq 9$ and deep OTM puts (top panel, $-8 \leq m < -4$), OTM puts (middle panel, $-4 \leq m < 0$), and OTM calls (bottom panel, $0 \leq m < 5$). The solid lines indicate the symmetric 95% confidence band for the test.

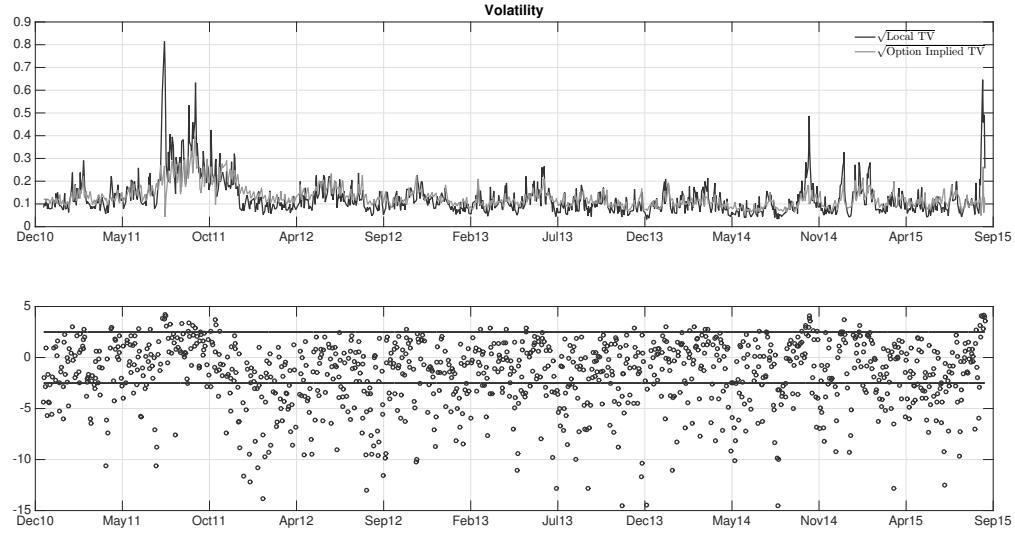


Figure 21: **The volatility fit based on the semi-nonparametric model with Gaussian jumps estimated from the short-dated options.** The top panel reports end-of-day spot volatility estimates based on nonparametric techniques and high-frequency data (solid line) or option-implied values (grey line). The option-implied volatility estimate is computed on the basis of the semi-nonparametric model (9) with the parameters and state vector estimated from the short-maturity option sample. The volatility estimates are reported in annualized units. The bottom panel reports the corresponding daily volatility test statistics (8). The solid lines indicate the symmetric 95% confidence band for the test.

F Semi-Nonparametric Estimation with Time-Varying Jump Parameters




Figure 22: **The fit to short-dated option prices based on the semi-nonparametric model (11) with Gaussian jumps estimated from the short-maturity sample.** The figure reports the daily test statistics (7) for the short-maturity options based on the semi-nonparametric model (11) with the state vector estimated from the short-dated options. The regions of the option cross-section are for tenor $\tau \leq 9$ and deep OTM puts (top panel, $-8 \leq m < -4$), OTM puts (middle panel, $-4 \leq m < 0$), and OTM calls (bottom panel, $0 \leq m < 5$). The solid lines indicate the symmetric 95% confidence band for the test.

| | Mean | Std | Q05 | Q50 | Q95 |
|---|-------|-------|-------|-------|-------|
| Gaussian Jumps | | | | | |
| RMSE | 0.886 | 0.437 | 0.397 | 0.782 | 1.744 |
| Diffusive Variance | 0.018 | 0.016 | 0.006 | 0.013 | 0.051 |
| Negative Jump Variance | 0.015 | 0.022 | 0.004 | 0.009 | 0.051 |
| Positive Jump Variance | 0.001 | 0.001 | 0.000 | 0.000 | 0.002 |
| Tempered Stable Jumps, $\alpha = -1$ | | | | | |
| RMSE | 0.701 | 0.326 | 0.331 | 0.612 | 1.370 |
| Diffusive Variance | 0.016 | 0.014 | 0.005 | 0.012 | 0.045 |
| Negative Jump Variance | 0.017 | 0.025 | 0.004 | 0.010 | 0.057 |
| Positive Jump Variance | 0.001 | 0.001 | 0.000 | 0.000 | 0.003 |
| Tempered Stable Jumps, $\alpha = 0$ | | | | | |
| RMSE | 0.624 | 0.289 | 0.303 | 0.544 | 1.206 |
| Diffusive Variance | 0.014 | 0.012 | 0.004 | 0.010 | 0.039 |
| Negative Jump Variance | 0.019 | 0.026 | 0.005 | 0.012 | 0.061 |
| Positive Jump Variance | 0.001 | 0.001 | 0.000 | 0.001 | 0.004 |
| Tempered Stable Jumps, $\alpha = 0.5$ | | | | | |
| RMSE | 0.581 | 0.274 | 0.283 | 0.509 | 1.154 |
| Diffusive Variance | 0.013 | 0.011 | 0.004 | 0.009 | 0.035 |
| Negative Jump Variance | 0.020 | 0.028 | 0.005 | 0.012 | 0.064 |
| Positive Jump Variance | 0.001 | 0.002 | 0.000 | 0.001 | 0.005 |

Table 4: **Estimation results for the semi-nonparametric models (11) and (12).** The columns provide the mean, standard deviation and 5th, 50th and 95th quantiles for the sequence of daily estimates obtained for the series indicated in the rows. The Gaussian jumps refer to model (11) and those with specific α values refer to model (12). The Root-Mean-Squared Error (RMSE) summarizes the fit to the end-of-day cross-section of annualized Black-Scholes implied volatilities (BSIVs). All variances are also reported in annualized units.

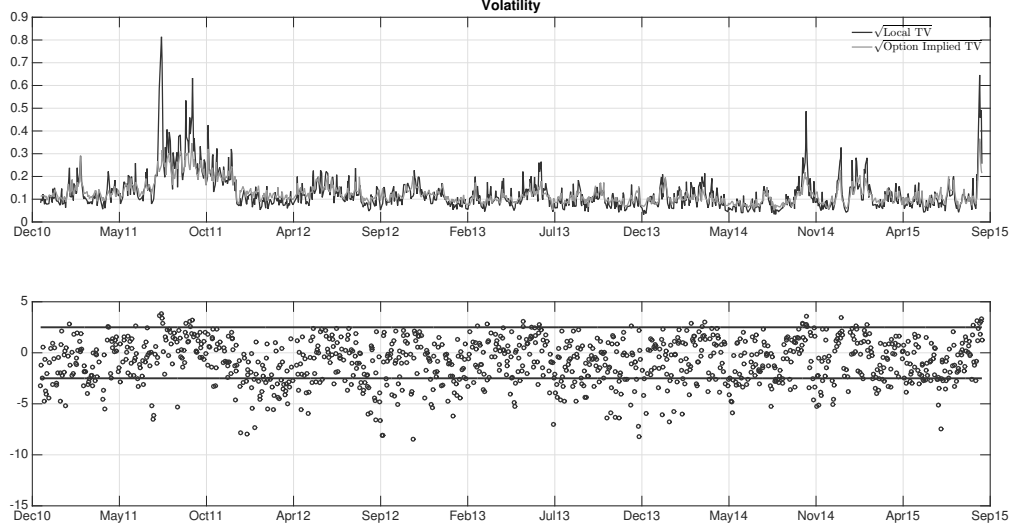


Figure 23: **The volatility fit based on the semi-nonparametric model (11) with Gaussian jumps estimated from the short-dated options.** The top panel reports end-of-day spot volatilities based on nonparametric estimates from high-frequency data (solid line) or option-implied values (grey line). The option-implied volatility estimate is computed on the basis of the semi-nonparametric model (11) with the state vector estimated from the short-maturity option sample. The volatility estimates are reported in annualized units. The bottom panel reports the corresponding daily volatility test statistics (8). The solid lines indicate the symmetric 95% confidence band for the test.




Figure 24: **The fit to short-dated option prices based on the semi-nonparametric model (12) with double-exponential jumps estimated from the short-maturity sample.** The figure reports the daily test statistics (7) for the short-maturity options based on the semi-nonparametric model (12) with $\alpha = -1$ and the state vector estimated from the short-dated options. The regions of the option cross-section are for tenor $\tau \leq 9$ and deep OTM puts (top panel, $-8 \leq m < -4$), OTM puts (middle panel, $-4 \leq m < 0$), and OTM calls (bottom panel, $0 \leq m < 5$). The solid lines indicate the symmetric 95% confidence band for the test.




Figure 25: **The volatility fit based on the semi-nonparametric model (12) with double-exponential jumps estimated from the short-maturity sample.** The top panel reports end-of-day spot volatilities based on nonparametric estimates from high-frequency data (solid line) or option-implied values (grey line). The option-implied volatility estimate is computed on the basis of the semi-nonparametric model (12) with $\alpha = -1$ and the state vector estimated from the short-maturity option sample. The volatility estimates are reported in annualized units. The bottom panel reports the corresponding daily volatility test statistics (8). The solid lines indicate the symmetric 95% confidence band for the test.

G Results for a Standard Parametric Model and Regular Options

In this section we introduce the parametric model and report its estimates and diagnostics based on “regular” options. The “regular” sample consists of S&P 500 index options with tenor beyond nine calendar days. The parametric model estimates are used to construct the tail jump variation which is compared with the one from our best-performing semi-nonparametric model in Figure 15.

The model explicitly parameterizes the dynamics of V and its dependence with W and μ , as well as that of the jump compensator. Since the model features jumps in volatility, we extend the jump measure μ to include these jumps, i.e., $\mu(dt, dx, dy)$ counts the jumps over the interval dt of size dx in the price and of size dy in V . The associated jump compensator is extended to $\nu_t(dx, dy)$. We consider a two-factor affine volatility specification, and we set $V_t = V_{1,t} + V_{2,t}$, where

$$\begin{aligned} dV_{1,t} &= \kappa_1 (\bar{v}_1 - V_{1,t}) dt + \sigma_1 \sqrt{V_{1,t}} dB_{1,t} + \int_{\mathbb{R}^2} y \mu(dt, dx, dy), \\ dV_{2,t} &= \kappa_2 (\bar{v}_2 - V_{2,t}) dt + \sigma_2 \sqrt{V_{2,t}} dB_{2,t}, \end{aligned} \quad (19)$$

and $(B_{1,t}, B_{2,t})$ is a two-dimensional Brownian motion with independent increments, $\text{corr}(dW_t, dB_{1,t}) = \rho_1 \sqrt{V_{1,t}/V_t}$, and $\text{corr}(dW_t, dB_{2,t}) = \rho_2 \sqrt{V_{2,t}/V_t}$.

Our jump specification is from Andersen et al. (2015b) and is given by,

$$\frac{\nu_t(dx, dy)}{dxdy} = c_t^- \cdot 1_{\{x < 0, y = \mu_v x^2\}} \cdot \lambda_- e^{-\lambda_- |x|} + c_t^+ \cdot 1_{\{x > 0, y = 0\}} \cdot \lambda_+ e^{-\lambda_+ x}, \quad (20)$$

where $c_t^\pm = \eta_0^\pm + \eta_1^\pm V_{1,t-} + \eta_2^\pm V_{2,t-}$. The volatility jumps are proportional to the squared negative price jumps, mimicking a discrete asymmetric GARCH specification. Option pricing implications of qualitatively similar features of the return distribution are pursued within a discrete-time GARCH setting by Christoffersen et al. (2012). The jump intensities are time-varying and we allow for separate variation for the positive and negative jumps.

For the above model, the option prices are known in semi-analytic form and are solely functions of the parameter vector θ , the volatility states $V_{1,t}$ and $V_{2,t}$, and the option characteristics τ and k . The estimation now follows the procedure outlined in Andersen et al. (2015b). The estimation results are reported in Table 5.

| Panel A: Parameter Estimates | | | | | | | | |
|---------------------------------------|----------|--------|------------|----------|--------|-------------|----------|--------|
| Parameter | Estimate | Std. | Parameter | Estimate | Std. | Parameter | Estimate | Std. |
| ρ_1 | -0.9998 | 0.0159 | κ_2 | 2.1994 | 0.0263 | η_2^- | 109.1400 | 2.2914 |
| \bar{v}_1 | 0.0040 | 0.0001 | σ_2 | 0.2720 | 0.0043 | η_2^+ | 0.0034 | 3.2853 |
| κ_1 | 10.6500 | 0.1045 | η_0^- | 0.0016 | 0.0116 | λ^- | 21.5550 | 0.1285 |
| σ_1 | 0.2905 | 0.0058 | η_0^+ | 1.6494 | 0.0678 | λ^+ | 48.7747 | 0.5458 |
| ρ_2 | -0.9990 | 0.0075 | η_1^- | 65.9215 | 1.5383 | μ_v | 16.5960 | 0.2541 |
| \bar{v}_2 | 0.0169 | 0.0002 | η_1^+ | 0.0078 | 4.9065 | | | |
| Panel B: Summary Statistics | | | | | | | | |
| RMSE | | | | 1.4210% | | | | |
| Mean positive jump intensity (yearly) | | | | 1.8784 | | | | |
| Mean negative jump intensity (yearly) | | | | 1.6495 | | | | |
| Mean negative jump size | | | | -0.0464 | | | | |
| Mean positive jump size | | | | 0.0205 | | | | |
| Mean diffusive variance | | | | 0.0217 | | | | |
| Mean negative jump variance | | | | 0.0081 | | | | |
| Mean positive jump variance | | | | 0.0014 | | | | |

Table 5: **Estimation results for the parametric model defined by equations (19) and (20).** The results are based on our regular option sample with tenor $10 \leq \tau \leq 365$, covering January 2011 – December 2014. The parameter estimates are obtained using weekly observations at day’s end on Wednesday, or Tuesday in case of a market closure on Wednesday. The state vector is estimated daily from the short-dated option sample, given the estimated parameter vector. **Panel A** provides the point estimates for the parameters and the associated asymptotic standard errors. **Panel B** reports summary statistics for the daily series of model-implied jump and variance estimates. All variances are given in annualized units.

H Construction of the Variance Risk Premium (VRP) Measure

Our variance risk premium (VRP) measure is computed as the difference between the end-of-day CBOE VIX measure and a twenty-one trading day (one month) forecast of the realized volatility. The latter is obtained via the HAR model of Corsi (2009). We construct the predictions of RV from estimates based on five-minute returns on the e-mini S&P 500 futures from a longer sample, comprising January 2007 through August 2015. Specifically, we let $RV_{s,t}$, $s \leq t$, denote the realized volatility, i.e., the cumulative sum of the 5-minute squared returns over the full electronic (Globex) trading session (but excluding the overnight squared returns) from the end-of-trading day s to the end of trading day t , and we denote the corresponding RV, inclusive of the overnight (close-to-open) squared returns, by $\overline{RV}_{s,t}$.

Our forecasts are now generated from the following linear regression (with a regression R^2 of 0.58 and t-statistics in parentheses),

$$\begin{aligned} \overline{RV}_{t,t+21} &= 0.0130 + 0.1638 RV_{t-1,t} + 0.2651 RV_{t-5,t} + 0.3054 RV_{t-21,t}. \\ &\quad (9.52) \qquad (8.25) \qquad (8.49) \qquad (11.36) \end{aligned}$$

INTERNET APPENDIX

I The Quality of the Approximate Short-Maturity Option Pricing

In this section we assess the quality of the approximation (5) which is used for estimation of the semi-nonparametric models in Sections 6.1 and 6.2.

Option prices obtained via the approximation (5) improve on the ones in (3) and (4) along several dimensions. First, our new approximation applies for any degree of moneyness. Second, our approximation error stems only from the time-variation of volatility and jump intensity over the life of the option. In contrast, equation (3) includes additional sources of error due to the presence of a diffusive component in X , the presence of “small” jumps, and the possible occurrence of multiple “big” jumps prior to expiry. Similarly, the approximation in (4) contains a non-trivial error because it ignores the presence of jumps in X . Thus, if X has constant volatility and jump intensity, our approximation $\tilde{\kappa}_{k,\tau}(\mathbf{S}_t, \theta)$ is exact, unlike the ones in equations (3) and (4).

We confirm these observations, and more generally the good quality of our approximation, for two popular parametric specifications used in prior empirical option pricing work. First, we consider the so-called double-jump volatility model of Duffie et al. (2000). The volatility follows a one-factor affine jump-diffusion which (co-)jumps simultaneously with the price at a constant intensity. We employ the parameter estimates from Broadie et al. (2007). On Figure 26 we report the true values of the options prices (known in semi-closed form) along with the ones based on the semi-nonparametric approximation in (5) and the nonparametric approximation in (3). We experiment with different volatility levels and tenors. We find the quality of our semi-nonparametric approximation to be very good, both for low and high levels of spot volatility, and, importantly, the approximation works well across the full range of moneyness, including close-to-the-money, where the alternative nonparametric approximation is ineffective. Moreover, the latter method displays serious downward biases when volatility is high. This is intuitive: the nonparametric method approximates the OTM option price, ignoring the diffusive component of the stock price. When volatility is high, therefore, this approximation deteriorates.

Second, we test the quality of the approximation methods for the two-factor stochastic volatility model of Bates (2000). In this model, the diffusive volatility is a sum of two square-root diffusions and the price jumps have a time-varying intensity which is an affine function of the two volatility factors. The results are presented in Figure 27. They are largely consistent with those for the one-factor double-jump volatility model discussed above. Our semi-nonparametric approximation

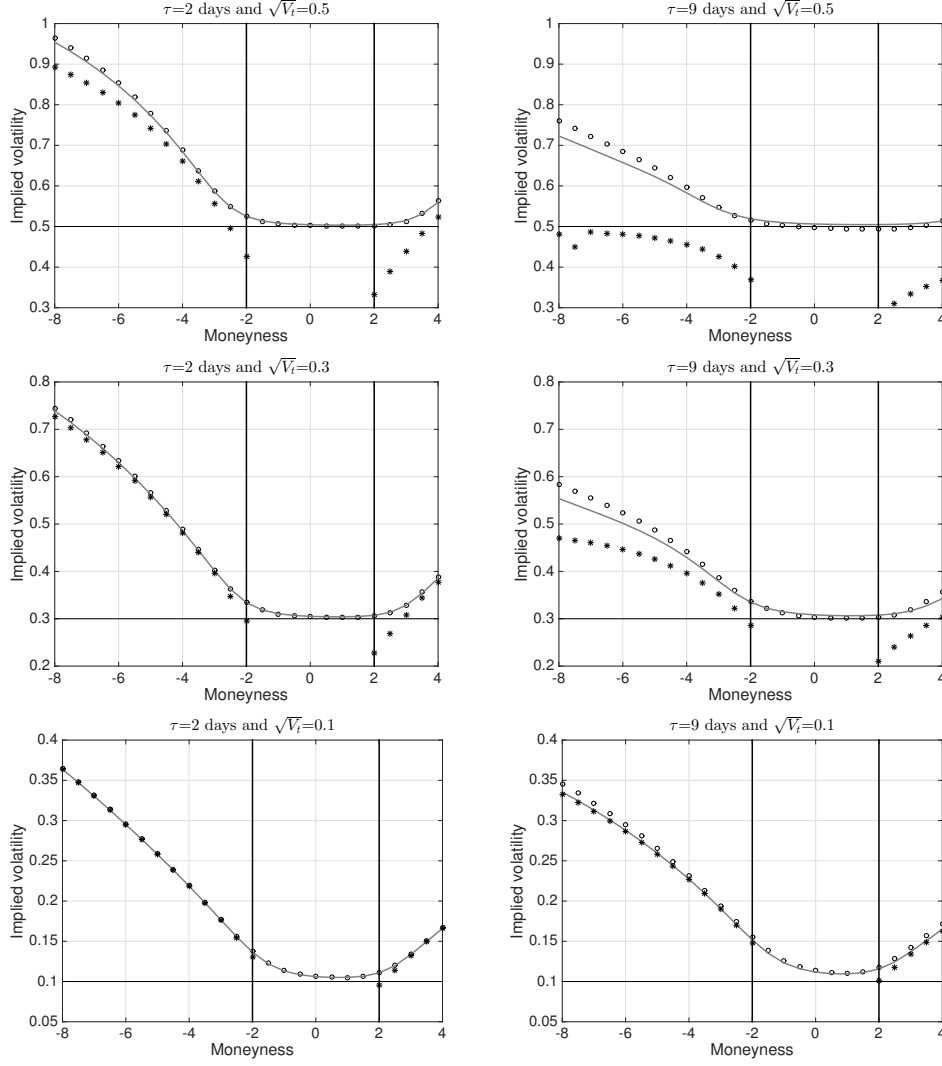


Figure 26: **Option price approximation: one-factor model with volatility jumps.** The figure reports the implied volatility generated by the SVCJ one-factor model in Broadie et al. (2007) (dots), the implied volatility generated by the semi-nonparametric approximation in equation (5) (solid line), and the implied volatility obtained using the nonparametric approximation in equation (3) (stars). The parameter values are taken from Tables I and IV (last row) of Broadie et al. (2007). Left and right panels show option cross-sections with two and nine, respectively, days to maturity. The plots in the three different rows are for three different level of spot volatility of 0.1, 0.3 and 0.5 (expressed in annualized terms).

continues to work well for the relevant range of tenors and moneyness. We only encounter minor underpricing for moderately OTM calls for our longest “short” maturity, $\tau = 9$. In fact, this bias is small compared to the inevitable measurement error associated with features like the bid-ask spread. The alternative nonparametric approximation has the same drawbacks as identified above. It performs worse and fares particularly badly for high levels of volatility and long tenors.

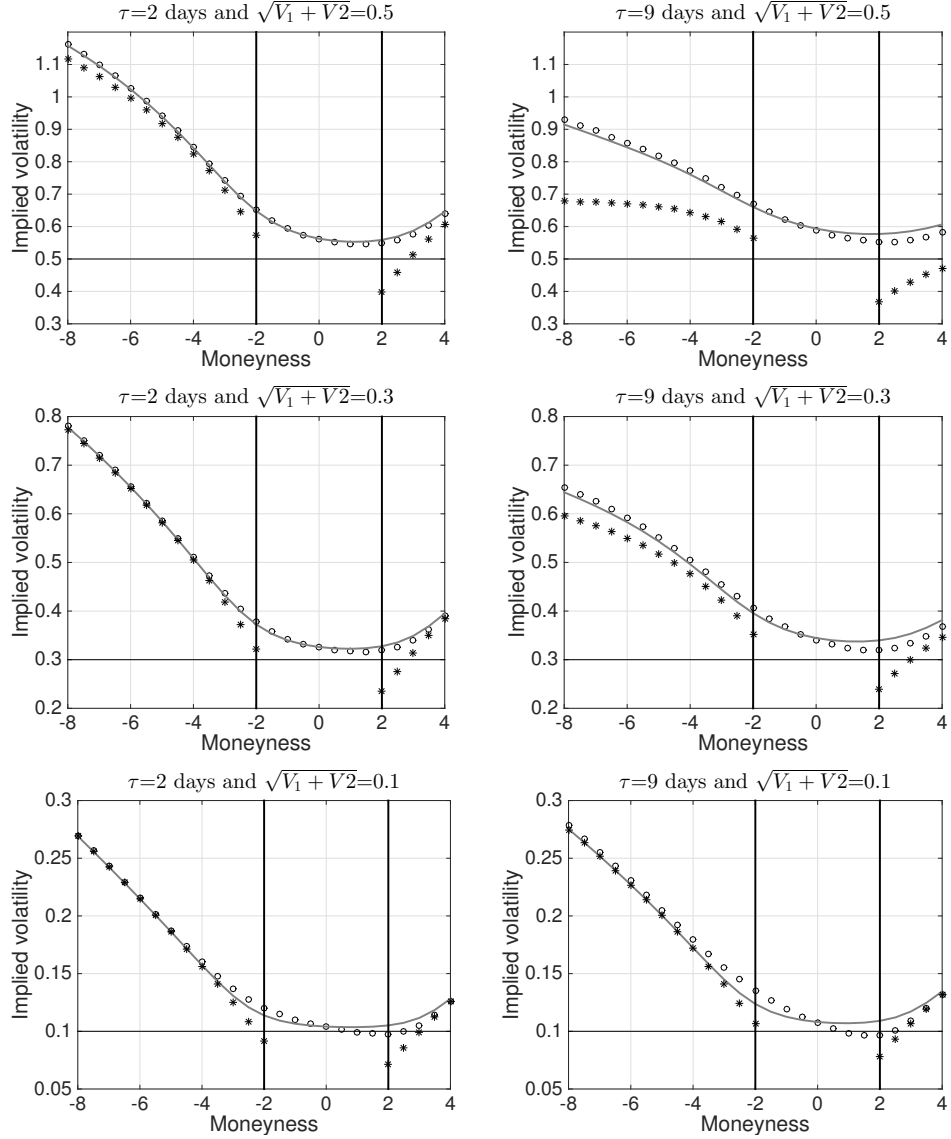


Figure 27: **Option price approximation: two-factor model.** The figure reports the implied volatility generated by the SVJD2 two-factor model of Bates (2000) (dots), the implied volatility generated by the semi-nonparametric approximation in equation (5) (solid line), and the implied volatility obtained using the nonparametric approximation in equation (3) (stars). The parameter values are taken from unconstrained estimates in Table 2 of Bates (2000). Left and right panels show option cross-sections with two and nine, respectively, days to maturity. The plots in the three different rows are for three different level of spot volatility of 0.1, 0.3 and 0.5 (expressed in annualized terms).

Overall, our evidence in this section shows that the newly-proposed semi-nonparametric option approximation, apart from being easy to compute and robust (to model misspecification of the volatility and jump intensity dynamics), is accurate in realistically calibrated settings.