NBER WORKING PAPER SERIES

THE PRICING OF SHORT-TERM MARKET RISK: EVIDENCE FROM WEEKLY OPTIONS

Torben G. Andersen
Nicola Fusari
Viktor Todorov

Working Paper 21491
http://www.nber.org/papers/w21491

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
August 2015

Andersen gratefully acknowledges support from CREATES, Center for Research in Econometric Analysis of Time Series (DNRF78), funded by the Danish National Research Foundation. Todorov's work was partially supported by NSF Grant SES-0957330. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

The Pricing of Short-Term market Risk: Evidence from Weekly Options
Torben G. Andersen, Nicola Fusari, and Viktor Todorov
NBER Working Paper No. 21491
August 2015
JEL No. C01,C14,C52,C58,G12,G13,G17,G32

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. Economically, this reflects a desire among investors for actively managing their exposure to very short-term risks. Such short-dated options provide an easy and direct way to study market volatility and jump risks. Unlike longer-dated options, they are largely insensitive to the risk of intertemporal shifts in the economic environment, i.e., changes in the investment opportunity set. Adopting a novel general semi-nonparametric approach, we uncover variation in the shape of the negative market jump tail risk which is not spanned by market volatility. Incidents of such tail shape shifts coincide with serious 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 on the market that are not always "signaled" by the level of market volatility and elude standard asset pricing models.

Torben G. Andersen
Kellogg School of Management
Northwestern University
2001 Sheridan Road
Evanston, IL 60208
and NBER
t-andersen@kellogg.northwestern.edu

Viktor Todorov
Department of Finance
Kellogg School of Management
Northwestern University
2001 Sheridan Road
Evanston, IL
v-todorov@kellogg.northwestern.edu

Nicola Fusari
Department of Finance
Carey School of Business
The Johns Hopkins Univer
Baltimore, MD 21202
nicola.fusari@jhu.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 about one week or less to maturity have seen their share of trading at the Chicago Board of Options Exchange (CBOE) rise steadily from about 12% in 2010 to 25% in 2014. Furthermore, the volume in shorter-dated options is skewed disproportionate towards out-of-the-money (OTM) options relative to the pattern for longer-dated options. Qualitatively 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."

This begs the question of what economic function these newly popular instruments serve. The primary distinguishing feature, relative to regular longer-dated options, is the intimate link between the pricing of options close to expiry and the state of the underlying asset return process.[1] When tenor is short, the volatility and jump intensity are not expected to vary much over the remaining life of the option. This implies, in particular, that the prices of deep OTM options are largely independent of the level of diffusive volatility. They reflect solely the characteristics of the risk-neutral jump process. Likewise, the pricing of short-dated at-the-money (ATM) options depends primarily on spot volatility. These 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 the joint volatility and jump dynamics involve a complex interaction among these distinct components, rendering semi-closed form option pricing and traditional statistical inference feasible only under strong parametric assumptions, e.g., the specifications reside within the affine jump-diffusion model class of Duffie et al. (2000).

Consequently, short-dated OTM options represent exposures tied closely to the possibility of an abrupt shift, or jump, in the near future, i.e., such instruments are well suited for short-term crash (or boom) protection. In contrast, a position in short-dated ATM options reflects an exposure to regular (diffusive) price movements, whose size scale with spot volatility. In other words, short-dated options allow us to decompose the exposure to distinct features of the underlying return process. They enable investors to trade securities representing jump or crash risk on the one hand, and instruments with primary exposure to diffusive risks on the other.

Hence, the introduction of weeklies has improved on the spanning of market jump risk and represents a step towards market completion. This topic has a long history, with Ross (1976)

[1] A 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.

emphasizing the enhanced spanning and the associated 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.[2]

Reversing the above reasoning, we infer that prices for actively traded short-dated options may simplify the task of identifying the concurrent spot volatility and pertinent features of the risk-neutral jump process, subject to only minimal assumptions on the return generating process. Specifically, short-maturity ATM options should help pin down spot volatility while the relative prices of deep OTM options assist in determining the jump intensity and jump distribution. The emergence of the weekly options has moved 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 very 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 all available information in our short-dated option sample, we develop a new asymptotic pricing approximation, operative across all strikes, and not just for ATM and deep OTM options. Exploiting this novel approximation, we proceed semi-nonparametrically and impose only weak parametric restrictions on the jump distribution, but remain 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 the short-dated options. In particular, we generate separate estimates of the underlying state of the volatility and jump intensity process as well as the current jump size distribution at the conclusion of each single 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. Our approach explicitly imposes no-arbitrage constraints in estimation and synthesizes the option price information into consistent estimates for the spot volatility and the key jump characteristics, amenable to direct econometric analysis. In contrast, standard calibration delivers a smoothed risk-neutral density for a specific horizon, but provides no direct guidance for the ex-

[2]The 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.

traction of spot volatility, jump intensities or jump size distributions. Indeed, 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 semi-martingale representing the risk-neutral asset price process. Moreover, 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, a sequence of smoothed risk-neutral densities over specific horizons provides a purely descriptive account of the evolution of the system. 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 on the basis of distinct estimates at the end of each individual trading day.[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 vastly 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.

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 in the range $(-2, 3)$, suggesting no major change in spot volatility or implied upside jump risk across these days. However, on September 20 and 21, the option prices are systematically higher for the deep OTM put options than on the preceding day, September 19. The gap emerges around $m = -2$ and increases steadily for further OTM puts. The finding is clearly robust, with close to 20 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, one may also observe a tendency towards an elevation in the prices of the longer-dated deep OTM put options, but the differences across the days are now

[3]See 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.

smaller and there is considerable noise in the far OTM quotes, which generates non-monotonicity and further blurs the inference. This reflects the fact that longer-dated OTM puts are sensitive to not only the negative jump distribution, but also the pricing of the time-varying volatility and jump intensity as well as associated leverage-type effects. The risks due to time variation in 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.

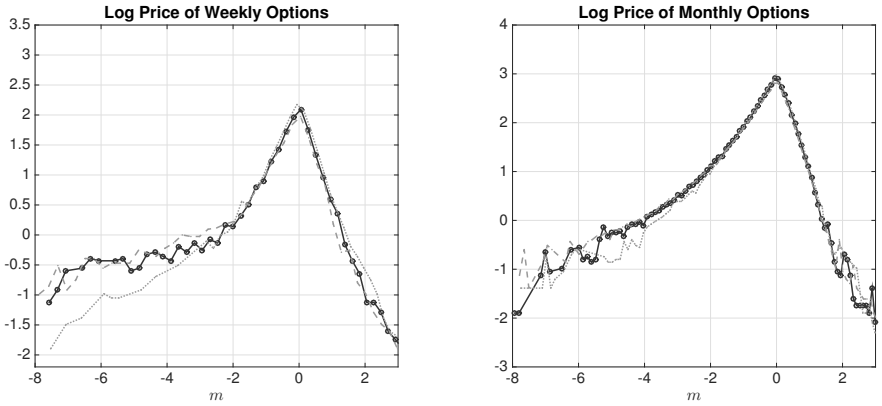

Figure 1: **Log Option Prices on September 20th 2012.** The Log Prices on September 20, 2012, are indicated by the full drawn line. On this date, the tenors were 7 calendar days for the short-dated options and 28 calendar days for the regular options. The dotted and broken lines represent log prices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

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 the 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 decidedly negative, the S&P 500 recovered most initial losses and experienced only a marginal day-over-day decline.[4] Likewise, on September 21, the economic news were at best mixed, and the S&P 500 ended flat.[5] Hence, one may conjecture that the broader outlook for equities remained fairly positive, but an element of uncertainty, or fear, may have entered the minds of some investors. If the monetary policy announcements were able to contain any escalation in the perception of overall risk, the negative news updates may have increased the expected size or pricing of contingencies associated with a sudden market downturn.

The above finding of market jump risk 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 equilibrium consumption-based models, such as Drechsler and Yaron (2011) and Wachter (2013) which, although quite flexibly specified, impose tight (affine) connections between the variation of market jump risks and volatility. Furthermore, to the extent the variation in the priced negative jump tail risk is unmatched by actual variation in negative market jumps under the statistical measure, our option evidence points to substantial nonlinearities in the part of the economy-wide pricing kernel that concerns the pricing of "large" negative jumps.

On the constructive side, our semi-nonparametric approach provides an easy way to extract a measure for perceived downside short-term tail risk. This measure can be used both for modeling and analyzing the latter in a robust way, and hence may also serve as a useful input to policy making.[6] In particular, we demonstrate that periods of heightened tail risk, as measured from the short-dated options, are closely connected with periods in which traditional asset pricing models

[4] According to Zacks Stock Market News, the weekly 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.

[5] The Zacks stock market report 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.

[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).

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.

The rest of the paper is organized as follows. In Section 2, we describe the option data on the S&P 500 index, and in particular its decomposition in terms of tenor. Our general setup and notation is presented in Section 3. In Section 4, using data for standard longer-dated options, we estimate general parametric models nesting many of the models estimated in prior work and assess their ability to price the short-dated options. In Section 5, we describe our semi-nonparametric approach for conducting inference based on the short-dated options and estimate models with time-invariant and time-varying jump distributions. In Section 6, we analyze the implications of the estimated models for the connections between market jump and volatility risks as well as the pricing of longer-dated OTM puts. Section 7 concludes. Some details on the filters applied to the options data and additional results from the estimation are provided 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 date, 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, an actively traded front-maturity option with expiry within nine calendar days is always present.

Figure 2 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 2010-2014 was around 17,000, 71,000, 107,000, 198,000, and 330,000 contracts, respectively. This raised their proportion from below 8% in early 2010 to well above 40% by the end of 2014. Since weeklies have a relatively short maturity upon issuance, this development has contributed to a fundamental shift in the maturity profile: options with tenor below 9 days

were thinly traded up till five years ago – and routinely excluded from empirical studies. Now, they constitute one of the dominant segments of the maturity spectrum.

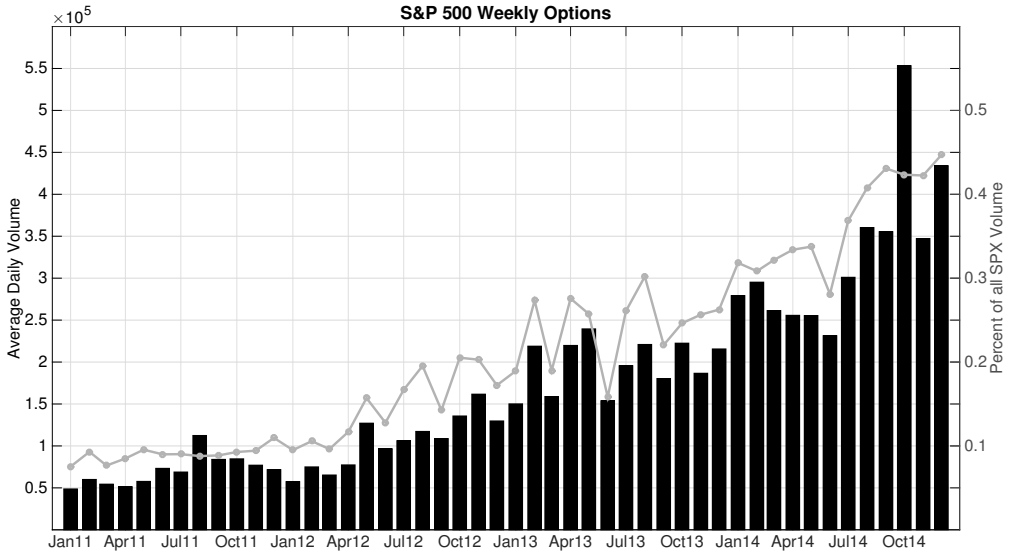

Figure 2: **Weekly (SPXW) Options**. The bar plot (left y-axes) represents the average daily volume in weekly options. The light gray line (right y-axes) depicts the weekly options' volume as a percentage of the total volume in SPX (sum of SPX, SPXQ and SPXW) options.

We focus primarily on extracting information about the current state of the risk-neutral dynamics from short-maturity options. To allow meaningful inference, we require good coverage across the strike range at the front maturity. 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 December 31, 2014. 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 SPXW, SPXQ and SPX options for 2011-2014.[7] 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.

[7]The 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.

We then 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 we can 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 1,505,535 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.[8] These two criteria eliminate about 240,000 additional quotes. The subsequent moneyness and bid-ask quote conditions leave us with slightly more than one million quotes. The OTM coverage condition eliminates five trading days during August 2011. Finally, removing Holidays and partial trading days leaves us with 973,866 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 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 percentage, or relative, error induced by the rounding of option values to the price grid is larger for the 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 A, to remove stale quotes or quotes that constitute clear (no-arbitrage) violations. Since the 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 61,756 to 41,206 contracts, covering 925 trading days. Thus, we have, on average, 54.5 bid-ask quotes at day's end for a single short

[8]The 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.

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 912,110 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.

2.2 Characteristics of the Two Option Samples

Table 1 summarizes the composition of our option sample. The top panel gives the number of contracts that fall in each moneyness-tenor category, while the bottom panel provides the relative trading volume. The shaded column concerns our short-maturity sample, and the next three columns refer to our regular sample.

The top panel reveals that the short-dated options, consisting exclusively of the front-maturity contracts, accounts for about 4.3% of the sample. Moreover, for options with tenor below 60 days, the OTM put options constitute a marginally larger fraction than the ATM options ($-2 \leq m \leq 2$), but this is sharply reversed for the longer maturities. Finally, there are very few quotes for OTM calls. The bottom panel shows that the trading in our short-maturity sample, comprising 19.2% of the total, is much more intense than for the regular options. Furthermore, the volume is higher for ATM than OTM options, with the volume in OTM calls being extraordinarily light in the regular sample. Finally, the volume for tenors beyond 180 days is low, especially for far OTM options.

Hence, 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 definite 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 3 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 3 combines information across tenors in our regular sample. Hence, the extreme positive and negative moneyness is often attained at different maturities.

	$\tau \leq 9$	$9 < \tau \leq 60$	$60 < \tau \leq 180$	$\tau > 180$	Total
			Number of contracts (%)		
$m < -2$	1.95	23.22	12.88	6.73	44.78
$\lvert m \rvert \leq 2$	2.17	21.25	17.40	12.34	53.16
$m > 2$	0.20	1.63	0.17	0.07	2.07
Total	4.32	46.10	30.44	19.13	100
			Volume (%)		
$m \leq -2$	5.97	14.58	3.55	0.66	24.77
$\lvert m \rvert \leq 2$	12.74	40.67	16.34	4.18	73.93
$m > 2$	0.47	0.78	0.04	0.01	1.30
Total	19.18	56.03	19.93	4.85	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 December 31, 2014. **Bottom Panel**: Percentage of contract volume over different combinations of tenor (τ, in days) and moneyness (m) between January 3, 2011, and December 31, 2014.

Moreover, the moneyness is not truncated, as is the case for the short-maturity sample. As such, the two panels are not directly comparable. Instead, the right panel merely serves to convey general information regarding the time variation in coverage offered by our regular option sample. With this caveat in mind, we note the broadly similar coverage on the upside while, up through 2013, there typically are some longer-dated options that 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 third of the dates in the sample. At that point, no significant information is available from additional downside coverage. In summary, our short-maturity sample covers an impressively broad range of strikes on the downside and matches the upside coverage provided by standard option samples.

Finally, we address the concern that the bid-ask spread may be excessive for short-dated options. Figure 4 depicts the average spreads for our two categories. For the range $[-4, 1]$, the relative spread is indeed lower for the longer-dated options, but outside this range the liquidity of the short-dated options – reflecting our sample selection criteria – manifests itself in a relatively narrow spread. This suggests the reliability of the pricing for this critical segment of far OTM short-dated options is comparable to, if not better than, that of the typical OTM options used in prior studies.

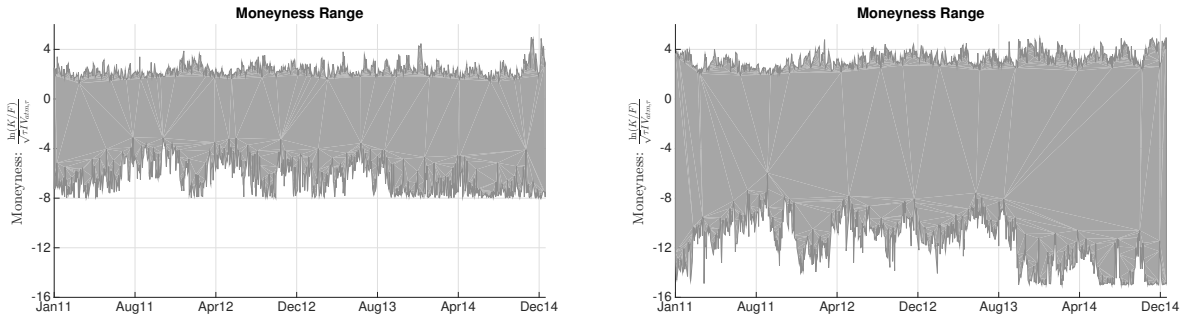

Figure 3: **Moneyness Range**. **Left Panel**: Moneyness range for short-maturity options. **Right Panel**: moneyness range for long-maturity options. The sample period is January 3, 2011 – December 31, 2014. Moneyness is computed as in the main text, $m = \frac{\ln(K/F)}{\sqrt{\tau}IV_{atm,\tau}}$.

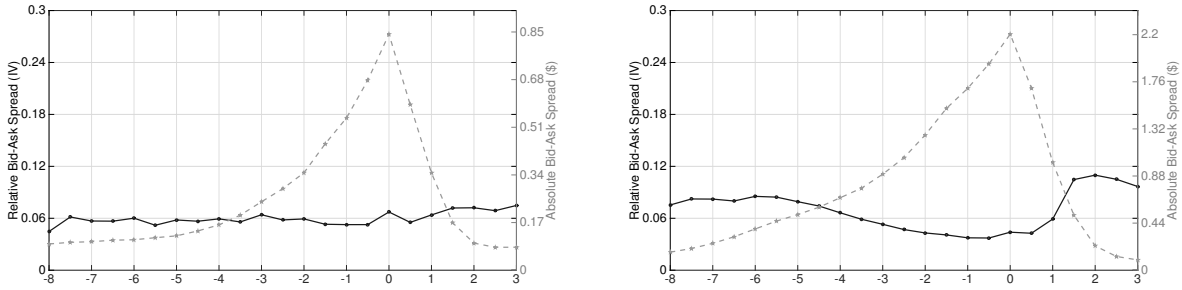

Figure 4: **Bid Ask Spread**. Kernel regression of the Relative (left axes) and absolute (right axes) bid-ask spread as a function of moneyness, m.

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,

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

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 $dt \otimes \nu_t(dx, dy)$ and the difference $\widetilde{\mu}(dt, dx, dy) = \mu(dt, dx, dy) - dt \otimes \nu_t(dx, dy)$ is the associated martingale jump measure. The jump specification involves two separate components, x and y. The former captures the price jumps and the latter the jumps in the diffusive variance. The marginal of the jump measure that counts only the price jumps will be denoted by $\mu(dt, dx) \equiv \mu(dt, dx, \mathbb{R})$ and its associated compensator by $dt \otimes \nu_t(dx)$.

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

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.[10]

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 out-of-the-money (OTM) options on X at time t by $O_{t,k,\tau}$. Assuming frictionless trading, option prices equal the expected discounted payoffs under the risk-neutral measure,

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

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

For each trading day, the option data is split into two groups. The first comprises the "standard" data used in empirical work, namely OTM options with maturities beyond nine calendar days, and the second consists of short-dated options with less than or equal to nine days to expiry. The number of options in these groups on day t are denoted N_t and M_t, respectively.

4 How Well Do Existing Parametric Models Fit the Weeklies?

We start our formal analysis of the short-dated options by exploring the ability of standard parametric asset pricing models to capture both their level and dynamics. We use estimates of these models based on "standard" data, namely S&P 500 index options with tenor beyond nine calendar days. Subsequently, we contrast our findings to a more general semi-nonparametric setup in which certain parts of the pricing model are left unspecified while others are generalized.

The parametric models we consider below parameterize the dynamics of V and its dependence with W and μ, as well as that of the jump compensator $\nu_t(dx, dy)$. In particular, we consider two-factor affine volatility specifications, and we set $V_t = V_{1,t} + V_{2,t}$, where

$$\begin{aligned} dV_{1,t} \;&=\; \kappa_1\,(\overline{v}_1 \,-\, V_{1,t})\,dt \,+\, \sigma_1\,\sqrt{V_{1,t}}\,dB_{1,t} \,+\, \int_{\mathbb{R}^2} y\,\mu(dt, dx, dy), \\ dV_{2,t} \;&=\; \kappa_2\,(\overline{v}_2 \,-\, V_{2,t})\,dt \,+\, \sigma_2\,\sqrt{V_{2,t}}\,dB_{2,t}, \end{aligned} \tag{3}$$

[10] One 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.

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

We consider two alternative specifications for the jump measure that allow for simultaneous jumps, so-called co-jumps, in price and volatility. The first follows the double-jump volatility representation of Duffie et al. (2000),

$$\frac{\nu_t(dx,dy)}{dxdy} \;=\; c_t \;\frac{e^{-\frac{(x-\mu_x-\rho_j y)^2}{2\,\sigma_x^2}}}{\sqrt{2\pi}\;\sigma_x}\;\frac{e^{-\frac{y}{\mu_v}}}{\mu_v}\;1_{\{y>0\}}, \tag{4}$$

where $c_t = \eta_0 + \eta_1 V_{1,t-} + \eta_2 V_{2,t-}$. In words, the log-price jump, conditional on a volatility jump of size y, is Gaussian with mean $\mu_x + \rho_j y$ and variance σ_x^2, while the volatility jump is exponentially distributed with mean μ_v. The jump intensity itself is time-varying and affine in the two volatility factors, as originally proposed by Bates (2000).

Our second specification for the jumps stems from Andersen et al. (2015b) and is given by,

$$\frac{\nu_t(dx,dy)}{dxdy} \;=\; c_t^- \cdot 1_{\{x<0,\; y=\mu_v x^2\}} \cdot \lambda_- \, e^{-\lambda_-|x|} \;+\; c_t^+ \cdot 1_{\{x>0,\; y=0\}} \cdot \lambda_+ \, e^{-\lambda_+ x}, \tag{5}$$

where $c_t^{\pm} = \eta_0^{\pm} + \eta_1^{\pm} V_{1,t-} + \eta_2^{\pm} V_{2,t-}$. Following Kou (2002), the price jumps are exponentially distributed, with separate tail decay parameters, λ_- and λ_+, for negative and positive jumps. The volatility jumps are proportional to the squared negative price jumps, mimicking a discrete asymmetric GARCH specification.[11] As in the previous representation, the jump intensities are time-varying, but we now further allow for separate variation for the positive and negative jumps.

For the models above, 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. Hence, we denote the model-implied BSIV by $\kappa_{k,\tau}(\mathbf{S}_t, \theta)$, where the state vector is $\mathbf{S}_t = (V_{1,t}, V_{2,t})$. They differ from the observed ones due to measurement errors in the option prices and possible model misspecification.

Our first goal is to assess whether standard empirical models from the extant literature perform satisfactorily in pricing short-maturity options. To this end, we initially estimate the parametric models using our regular option sample. As is common in the literature, to limit the computational burden, we exploit Wednesday end-of-day quotes for estimation.[12] The estimation is based on the following objective function,

[11] Option pricing implications of qualitatively similar features of the return distribution are pursued within a discrete-time GARCH setting by Christoffersen et al. (2012).

[12] We switch to a nearby business day in case of a Wednesday market closure.

$$\left(\{\widehat{\mathbf{S}}_t\}_{t\in\mathcal{T}^W},\widehat{\theta}\right) \;=\; \underset{\{\mathbf{S}_t\}_{t\in\mathcal{T}^W},\theta}{\operatorname{argmin}} \sum_{t\in\mathcal{T}^W}\left\{\sum_{j:\tau_j>9}\frac{\left(\kappa_{t,k_j,\tau_j}\;-\;\kappa_{k_j,\tau_j}(\mathbf{S}_t,\theta)\right)^2}{N_t}\;+\;\frac{k_n}{N_t}\frac{\left(\sqrt{\widehat{V}_t^n}\;-\;\sqrt{V_{1,t}+V_{2,t}}\right)^2}{\widehat{V}_t^n/2}\right\}, \tag{6}$$

where $\widehat{V}_t^n$ is a nonparametric estimate of the spot diffusive variance,[13] using k_n high-frequency increments preceding time t, and $\mathcal{T}^W$ denotes the set of time indices corresponding to the Wednesdays in our sample. The main component of the objective function (6) is the mean squared error in fitting the option (BSIV) panel, but we add a term penalizing the deviation of the implied spot volatility from the nonparametric high-frequency estimate.

Given the estimated parameter vector $\widehat{\theta}$ from the "regular" option data set, we estimate the state vector, $\mathbf{S}_t=(V_{1,t},V_{2,t})$, each trading day using the short-dated option sample,

$$\widehat{\mathbf{S}}_t \;=\; \underset{\mathbf{S}_t}{\operatorname{argmin}}\left\{\sum_{j:\tau_j\leq 9}\left(\kappa_{t,k_j,\tau_j}\;-\;\kappa_{k_j,\tau_j}(\mathbf{S}_t,\widehat{\theta})\right)^2+k_n\frac{\left(\sqrt{\widehat{V}_t^n}\;-\;\sqrt{V_{1,t}+V_{2,t}}\right)^2}{\widehat{V}_t^n/2}\right\},\quad t=1,...,T. \tag{7}$$

The estimation results for the two parametric models are reported in Appendix B. For the Gaussian jump model, the mean jump intensity is 0.30, or only three jumps per ten years, while the mean jump size is large and negative at -21.15%. The root-mean-squared-error (RMSE) of the fit to the short-maturity option sample is 2.80%, expressed in annualized BSIV. In comparison, the point estimates for the double-exponential jump model imply about 3.5 jumps per year, roughly equally split between positive and negative jumps. The mean negative jump size is -4.64% and it is 2.05% for positive jumps. Finally, this model has a lower RMSE at 2.39%. Thus, both models feature left-skewed jump distributions, but the Gaussian representation involves rare and large negative jumps, while the double-exponential implies more frequent, moderately sized, negative and positive jumps, with the negative jumps being about double the size of the positive ones.

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,

$$\widehat{T}^{\text{of}}_{\mathcal{K},t} \;=\; \mathcal{Z}_{\mathcal{K},t}/\sqrt{\widehat{\text{Avar}}\left(\mathcal{Z}_{\mathcal{K},t}\right)}, \tag{8}$$

where $\mathcal{Z}_{\mathcal{K},t}\;=\;\sum_{j:k_j\in\mathcal{K},\;\tau_j\leq 9}\left(\kappa_{t,k_j,\tau_j}-\kappa_{k_j,\tau_j}(\widehat{\mathbf{S}}_t,\widehat{\theta})\right)$ is the average pricing error over the relevant

[13] In our application, $\widehat{V}_t^n$ is a standard robust threshold realized volatility estimator, implemented as in Andersen et al. (2015a).

part of the moneyness region, $\mathcal{K}$. Andersen et al. (2015a) show that $\widehat{T}^{\text{of}}_{\mathcal{K},t}$ 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, $\widehat{V}_t = \widehat{V}_{1,t} + \widehat{V}_{2,t}$, and the corresponding nonparametric high-frequency estimate, $\widehat{V}^n_t$. The formal test is based on the normalized "volatility fit" (vf) statistic,

$$\widehat{T}^{\text{vf}}_t = (\widehat{V}^n_t - \widehat{V}_t) / \sqrt{\widehat{\text{Avar}}(\widehat{V}^n_t - \widehat{V}_t)}. \tag{9}$$

Under correct model specification $\widehat{T}^{\text{vf}}_t$ is asymptotically standard normal, and diverges otherwise. We defer 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 (8) and (9).

For conciseness, Figures 5 and 6 only depict the Z-scores for the best performing model, based on the double-exponential jump specification. The diagnostics for the alternative parametric specification are relegated to Appendix B.

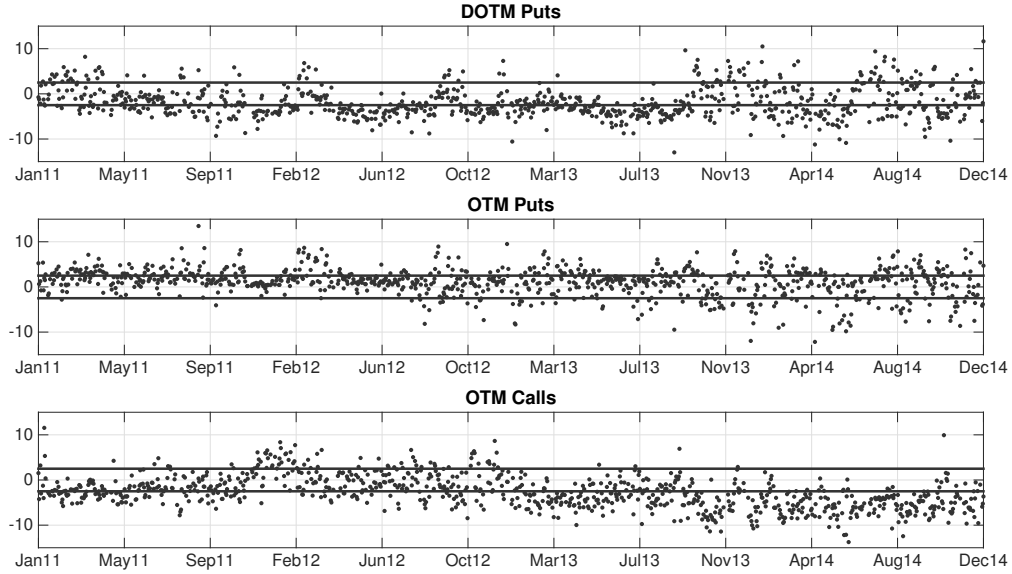

Figure 5: **Short-dated option price fit based on a parametric model estimated from the "regular" option sample.** The figure reports the daily test statistics (8) for the short-dated options based on the parametric model in equations (3) and (5) with parameters estimated from the "regular" option sample and the state vector estimated separately from the short-dated options each trading day. 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.

The panels in Figure 5 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. In all cases, we observe very large pricing errors with t-statistics often exceeding 5 in absolute value and even going beyond 10. In particular, we identify persistent overpricing of deep OTM puts between March 2012 and October 2013 and of OTM calls in the entire second part of the sample. Likewise, there are highly significant deviations between the model-implied spot volatilities and the corresponding high-frequency estimates in Figure 6. In this case, the lower panel reveals that the model-implied estimates quite often exceed the high-frequency estimates by a wide margin, i.e., the model-implied volatility has excessive positive outliers. Overall, the standard models provide a very poor fit to the short-dated option prices. This suggests that these models fail to capture critical features of the underlying risk-neutral return dynamics.

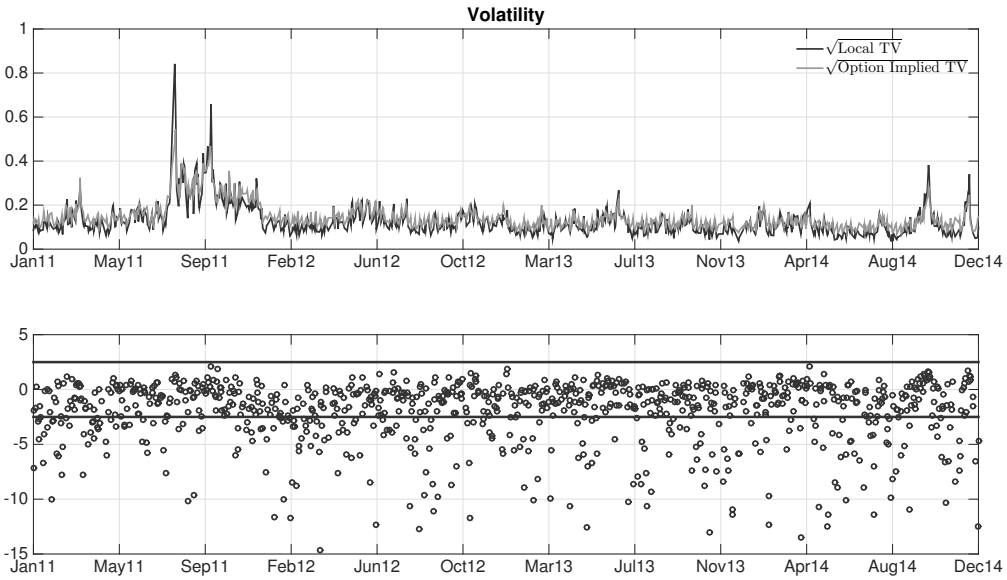

Figure 6: **Volatility fit based on a parametric model estimated using "regular" options.** The top panel reports end-of-day spot volatilities based on nonparametric estimates from high-frequency data (solid line) or option-implied values (grey line). The option-implied volatility estimate is computed using the parametric model defined in equations (3) and (5) with parameters estimated from the "regular" option sample and the state vector estimated separately from the short-dated options on each trading day. The volatility estimates are reported in annualized units. The bottom panel reports the corresponding daily volatility test statistic (9). The solid lines indicate the symmetric 95% confidence band for the test.

These findings are quite unsettling. Standard asset pricing models fail systematically in terms of providing a suitable characterization of the equity return dynamics and the associated short-horizon risk pricing. Is this due to the specific parametric models used here? If so, what features are to blame for the poor performance? More generally, what type of no-arbitrage model is consistent with our short-maturity option sample? Of course, it is possible that short-maturity options are

subject to a great deal of idiosyncratic noise and serve as a poor basis for assessing the performance of asset pricing models. We explore this possibility further in the following sections, as we impose successively weaker restrictions on the risk-neutral dynamics. In this manner, the option data are given prominence in shaping the results, allowing us to more readily distinguish the effects of model misspecification versus excessive noise in the option prices.

5 Semi-nonparametric Models for Short-Maturity Option Pricing

The 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 more robust semi-nonparametric approach in which parts of the pricing kernel are left unspecified, avoiding the unforeseeable consequences of misspecification along particular dimensions. However, these techniques apply only for a restrictive set of strikes and for very short maturities. 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. In turn, this enables us to perform feasible inference on the risk-neutral dynamics from short-dated options alone, while imposing only weak distributional assumptions.

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

Various asymptotic expansions have been used in prior work. For example, Carr and Wu (2003) and Bollerslev and Todorov (2014) exploit the following OTM option price approximation,

$$\frac{O_{t,\tau,k}}{\tau X_{t-}} \longrightarrow \begin{cases} \int_{\mathbb{R}} (e^x - e^k)^+ \, \nu(dx; \mathbf{J}_t, \theta), & \text{if } k > 0 \\ \int_{\mathbb{R}} (e^k - e^x)^+ \, \nu(dx; \mathbf{J}_t, \theta), & \text{if } k < 0 \end{cases}, \qquad \text{as } \tau \downarrow 0 \,. \tag{10}$$

That is, the valuation of short-dated OTM options is largely determined by the prices of "big" jumps. However, in practice, this approximation works well only for very deep OTM options, i.e., for k large in absolute value. Moreover, once we move beyond very short-maturity options, the presence of the diffusion term and the time-variation in volatility and jump intensity have a non-trivial impact on the pricing.

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

$$\kappa_{t,0,\tau} \longrightarrow \sqrt{V_t}, \quad \text{as } \tau \downarrow 0. \tag{11}$$

The approximations (10) and (11) work either for deep OTM or ATM options, but not for the broad range with 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 $\widetilde{X}$. It is defined as,

$$\frac{d\widetilde{X}_s}{\widetilde{X}_{s-}} = (r_t - \delta_t)ds + \sqrt{V_t}dW_s + \int_{\mathbb{R}} (e^x - 1)(\mu(ds, dx) - ds \otimes \nu(dx; \mathbf{J}_t, \theta)), \quad s \geq t, \quad \widetilde{X}_t = X_t\,. \tag{12}$$

$\widetilde{X}$ 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 $\mathcal{F}_t$, $\widetilde{X}$ is a Lévy process. The reason $\widetilde{X}$ may approximate X well over short intervals (like a few days) is that, over such horizons, the volatility and the jump intensity of X are not expected to vary much. In fact, since risk premiums tend to be highly persistent, the stochastic volatility and jump intensity are likely to be even less variable under the pricing measure than under the actual statistical measure.

Pricing options written on $\widetilde{X}$ 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). The theoretical price of an option, in terms of BSIV, at time t written on $\widetilde{X}$ 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 denote it by $\widetilde{\kappa}_{k,\tau}(\mathbf{S}_t, \theta)$. The formal analysis of the approximation error, $\widetilde{\kappa}_{k,\tau}(\mathbf{S}_t, \theta) - \kappa_{t,k,\tau}$, relies on an asymptotic analysis of the price increment $\widetilde{X}_{t+\tau} - X_{t+\tau}$ for $\tau \downarrow 0$. While the approximation $\widetilde{X}$ 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.

Option prices obtained via the approximation (12) improve on equations (10) and (11) 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 (10) includes also approximation errors 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, equation (11) includes an approximation error due to the

presence of jumps in X. Thus, if X has constant volatility and jump intensity, our approximation $\widetilde{\kappa}_{k,\tau}(\mathbf{S}_t, \theta)$ is exact, unlike the ones in equations (10) and (11).

We will use our approximation, exploiting the simplified return dynamics in equation (12), to better 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 $\widetilde{\kappa}_{k,\tau}(\mathbf{S}_t, \theta)$, we need to specify only the jump distribution. Thus, the inference based on $\widetilde{\kappa}_{k,\tau}(\mathbf{S}_t, \theta)$ is far more robust than one based on a full parametric model which hinges also on correctly specified volatility and jump intensity dynamics.

5.1 Models with Fixed Jump Distribution

Using short-dated options and the Lévy-based approximation (12), we now estimate semi-nonparametric models that mirror our two parametric specifications but, importantly, we avoid specifying the dynamics of the volatility and the jump intensity as well as their interdependence. Hence, the success of these models does not depend on correct parametric specification for the dynamics of the volatility and jump intensity. Consequently, the relative quality of fit for the parametric and semi-nonparametric specifications provides an informal diagnostic tool. In particular, it sheds some light on the extent to which the representation for the dynamics of the volatility and jump intensity in equation (3) is to blame for the poor performance of the parametric models in Section 4.

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

$$\nu(dx; \mathbf{J}_t, \theta) \;=\; c_t \, \frac{1}{\sqrt{2\pi}\,\sigma_x} \, e^{-\frac{(x-\mu_x)^2}{2\,\sigma_x^2}} \, dx \,. \tag{13}$$

This formulation involves Gaussian jumps, as in equation (4), but we now allow the jump intensity, c_t, to display (arbitrary) time variation unrelated to volatility. In this specification $\theta = (\mu_x, \sigma_x^2)$ and $\mathbf{J}_t = c_t$. We note that the jump distribution remains time invariant in this specification.

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

$$\nu(dx; \mathbf{J}_t, \theta) \;=\; c_t \left\{ \frac{e^{-\lambda^- |x|}}{|x|^{1+\alpha}} \, 1_{\{x<0\}} \;+\; \frac{e^{-\lambda^+ |x|}}{|x|^{1+\alpha}} \, 1_{\{x>0\}} \right\} dx, \quad \alpha < 2. \tag{14}$$

This jump specification extends the double-exponential model (5) in two dimensions. First, the intensity parameter c_t now varies unrestrictedly. Second, although we retain the double-exponential specification for now, we introduce, for later reference, an additional parameter, α, which controls

the behavior of the jump measure around zero. $\alpha = -1$ corresponds to the original double-exponential and $\alpha = 0$ to the variance gamma jump process. For $\alpha < 0$, the jump process is of finite activity, while $\alpha \in [0, 1)$ implies infinite activity, but finite variation jumps. Finally, $\alpha \in [1, 2)$ generates jumps of infinite variation. For constant c_t, the process (14) belongs to the tempered stable class, introduced by Carr et al. (2002, 2003), and estimated from return data by Bates (2012).[14] We find that α cannot be identified with precision from the data. Therefore, we do not treat α as a parameter, but instead estimate the jump specification (14) for different fixed values of α. Thus, the parameter vector is $\theta = (\lambda^-, \lambda^+)$ and the state vector is $\mathbf{J}_t = c_t$.

On the other hand, there is also a new constraint imposed in equation (14). The time-varying jump intensities now take a similar form for positive and negative jumps, whereas the earlier double-exponential formulation (with $\alpha = -1$) in equation (5) allows for these jump intensities to evolve separately. We impose this additional restriction because separate identification of the jump intensity for each observation date requires reliable observations for deep OTM options. However, on many trading days, the coverage for deep OTM calls is too limited to allow meaningful inference. Thus, effectively, the formulation in equation (17) implies we "borrow" information from the highly liquid OTM put options to pin down the jump intensity, also for the positive jumps. Without such a restriction, minor model misspecification or measurement errors in the option prices induce a large degree of instability and imprecision in the estimates for the right jump intensities. This implies that we should view our fit to the OTM call options with some skepticism.

The estimation of our semi-nonparametric models is done via,

$$\left(\{\widehat{\mathbf{S}}_t\}_{t=1}^T, \widehat{\theta}\right) = \underset{\{\mathbf{S}_t\}_{t=1}^T, \theta}{\operatorname{argmin}} \sum_{t=1}^{T} \left\{ \sum_{j:\tau_j \leq 9} \frac{\left(\kappa_{t,k_j,\tau_j} - \widetilde{\kappa}_{k_j,\tau_j}(\mathbf{S}_t, \theta)\right)^2}{M_t} + \frac{k_n}{M_t} \frac{\left(\sqrt{\widehat{V}_t^n} - \sqrt{V_t}\right)^2}{\widehat{V}_t^n/2} \right\}, \tag{15}$$

which is a direct analogue to equation (6).

We defer the estimation results to Appendix C. Briefly summarizing the findings, however, we note a substantial improvement in the fit to the short-dated options with RMSEs of, respectively, 1.19% and 1.02% for the Gaussian and double-exponential ($\alpha = -1$) jump specifications. Moreover, the jump intensity for the Gaussian representation rises to 2.5 per year with a more moderate jump size of -5%. In comparison, the double-exponential model has mean jump sizes of -3.8% and 1.6%, i.e., it features only slightly smaller jumps than before.

Given the model estimates, we construct Z-scores for the short-maturity option and volatility fit, exactly as in equations (8) and (9). For brevity and ready comparison with Figures 5 and 6, we only provide diagnostic plots for the better fitting double-exponential representation, while the plots

[14] More precisely, (14) can be viewed as tempered stable for $\alpha \in (0, 2)$.

for the Gaussian jumps are deferred to Appendix C. Comparing the fit of the double-exponential parametric model and its semi-nonparametric counterpart in this section, the improvement is evident. The average magnitudes of the Z-scores both for the fit to the option cross-section and the nonparametric volatility estimates are now much smaller, suggesting a lower degree of misspecification. This is consistent with the hypothesis that the parametric modeling of the return dynamics – including the two-factor affine specification as well as the affine dependence of the jump intensity on the volatility factors – is partially responsible for the poor empirical performance of the models in Section 4. This conclusion is also in line with the empirical evidence based on option data and flexible one-factor nonlinear volatility models provided by Christoffersen et al. (2010).

Nevertheless, even for the current model, the fit is problematic. Figure 7 displays the Z-scores for the fit to the option prices. In line with the evidence for the fully parametric model in Section 4, we continue to observe severe overvaluation for the deep OTM puts during March 2012 to 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 overall shape of the option skew, even with a fully flexible specification for the jump intensity. Moreover, Figure 8 documents that large outliers occur systematically on the downside, representing model-implied overestimation of the spot volatilities relative to the ones based on the high-frequency data. Qualitatively similar issues are apparent from the diagnostic tests for the Gaussian specification in Appendix C.

5.2 Models with Time-Varying Jump Distribution

In spite of the flexibility afforded by our semi-nonparametric models based on equations (13) and (14), they still do not account satisfactorily for the time variation in the pricing of short-dated options. What is the reason? The specifications (13) and (14) impose structure on the jump process along two dimensions. First, they assume that the jump distribution, at any point in time, may be captured through a finite-dimensional parameter vector. Second, like essentially all prior parametric models in the extant literature, they imply that the relative intensity of differently sized jumps is invariant over time. That is, $\nu_t(x)/\nu_t(x')$ is constant across all jump sizes x, $x' \in \mathbb{R}$.

It is unlikely that the parametric structure of the specifications (13) or (14) is the main problem. In fact, equation (14) is quite flexible and allows for separate parameters to govern small and big jumps. A more plausible culprit is the constraints on the time variation of the jump characteristics and, in particular, the imposition of a time-invariant shape for the intensity across jump sizes.[15]

[15]Indeed, we can formally test the hypothesis of a time-invariant shape across the jump sizes. Recall the short-

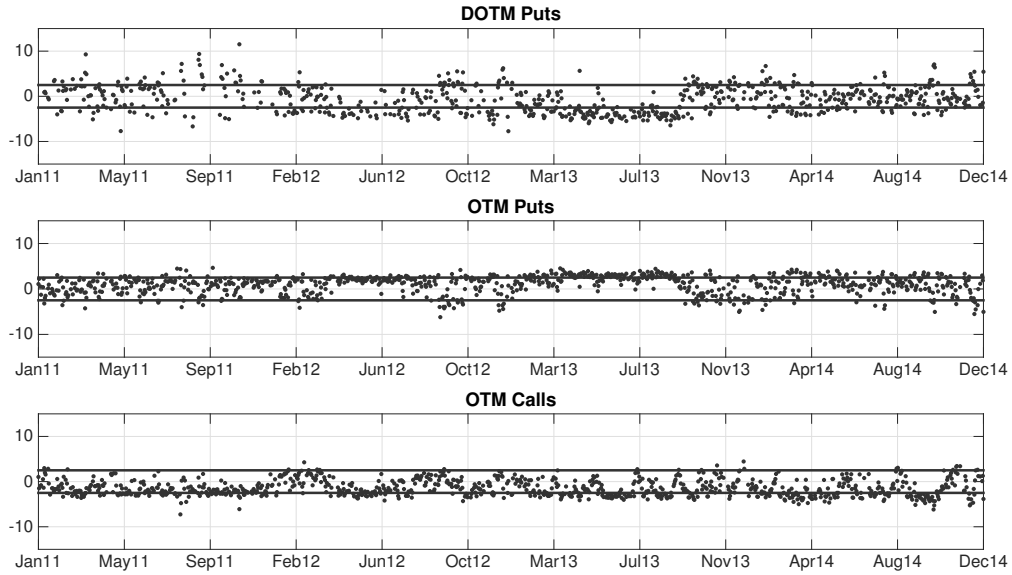

Figure 7: **The fit to the prices of short-dated options based on the semi-nonparametric model (14).** The figure reports the daily test statistics (8) for the short-dated option sample based on model (14) with the parameters and state vector estimated from the short-maturity options. The regions of the cross-section are for tenor of $\tau \leq 9$ 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.

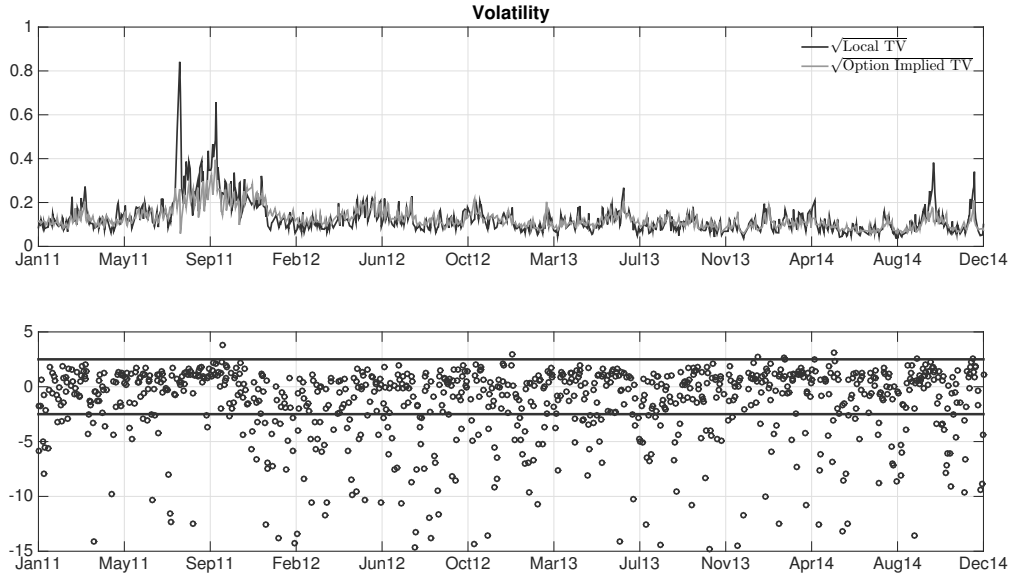

Figure 8: **The fit to spot volatility based on the semi-nonparametric model (14).** The top panel reports end-of-day spot volatilities based on nonparametric estimates from high-frequency data (solid line) or option-implied values (grey line). The option-implied volatility estimate is computed on the basis of model (14) with the parameters and state vector estimated from the short-dated options on each trading day. The volatility estimates are reported in annualized units. The bottom panel reports the corresponding daily volatility test statistics (9). The solid lines indicate the symmetric 95% confidence band for the test.

Hence, we now further generalize the models by allowing the parameters in equations (13) and (14) to vary over time. We consider the following *time-varying* jump distribution models,

$$\nu(dx; \mathbf{J}_t, \theta) \;=\; c_t \, \frac{e^{-\frac{(x-\mu_{x,t})^2}{2\sigma_{x,t}^2}}}{\sqrt{2\pi}\sigma_{x,t}} \, dx \,, \tag{16}$$

and

$$\nu(dx; \mathbf{J}_t, \theta) \;=\; c_t \, \left\{ \frac{e^{-\lambda_t^- |x|}}{|x|^{1+\alpha}} \, 1_{\{x<0\}} \;+\; \frac{e^{-\lambda_t^+ |x|}}{|x|^{1+\alpha}} \, 1_{\{x>0\}} \right\} dx \,. \tag{17}$$

In these specifications, the volatility and jump intensity remain unrestricted, but the coefficients of the jump distributions now also shift freely across the observation dates. Thus, all the jump coefficients are time-varying and no longer constitute part of the parameter vector. Instead, they enter $\mathbf{S}_t$ as additional state variables that are necessary for a full characterization of the prevailing risk-neutral price dynamics. In other words, θ is now the empty set — formally $\theta = \varnothing$ — while $\mathbf{S}_t = (V_t, \mathbf{J}_t)$ with $\mathbf{J}_t = (c_t, \mu_{x,t}, \sigma_{x,t})$ and $\mathbf{J}_t = \left(c_t, \lambda_t^+, \lambda_t^-\right)$ for the two models, respectively. We stress that the two models (16) and (17), by construction, remain arbitrage-free.

As in Section 5.1, we estimate the system via criterion function (15). 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 jump intensity, c_t, enters symmetrically for the positive and negative jumps is now even more critical. On many days, the information attainable through the few available quotes on deep OTM calls is insufficient for reasonable inference. Moreover, the richness 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. As a result, we may occasionally observe a poor fit to the prices of OTM calls, but this is a natural consequence of the lack of identifying information on this dimension. While this is a genuine limitation, we note that, from an economic perspective, most of the 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.

Our modeling via equation (1) and the jump specifications (16) and (17) is reminiscent of the calibration to option surfaces, often adopted by practitioners. However, there are fundamental

maturity asymptotic approximation (10). A specification of $\nu_t(dx) = c_t \times \nu(dx)$ for some intensity process c_t (like our parametric jump specifications thus far), together with the asymptotic result in equation (10), implies that the ratio $O_{t,\tau,k}/O_{t,\tau,k'}$ for some $k, k' > 0$ or $k, k' < 0$ fixed will remain constant as we vary t, i.e., across the days in the sample. This can be easily rejected for our short-maturity data set.

differences. Model (1), complemented by either (16) or (17), is built around the principles of no-arbitrage. In particular, it enforces (statistical) equality of the diffusive volatility under the physical and risk-neutral measures, and we explicitly enforce this restriction in the estimation procedure. There is no counterpart of this no-arbitrage condition for "large" jumps due to the market incompleteness induced by the presence of jumps. Our flexible semi-nonparametric specifications (16) and (17) 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.

Hence, 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. This is not feasible from the standard practitioner calibration of option surfaces via flexible curve-fitting procedures, e.g., smoothing splines. We subsequently employ this characterization to study the jump risk and volatility dynamics as well as the implications for longer-dated options. Again, such applications are precluded if calibration is performed through curve-fitting void of any connection to volatility, jump intensity and jump distribution as well as their no-arbitrage links with the underlying asset returns.

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. Since we are close to a (statistically) acceptable characterization, we provide a more comprehensive review of the diagnostic tests.

First, we consider the Gaussian jump specification, widely adopted in the empirical literature. Of course, our current version (16) is 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 options in our short-maturity sample are depicted in Figure 21 of Appendix D. 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. In other words, the model has difficulty in accommodating the shape of the left tail. Likewise, the diagnostics for the volatility fit in Figure 22 reveal that the model-implied volatility often greatly exceeds the corresponding high-

frequency based estimate of spot volatility. Even so, the overall RMSE drops to 0.85%, confirming the vast improvement in average fit relative to the prior Gaussian specifications.

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 overall RMSE is now 0.68% and the associated Z-score diagnostics, presented in Appendix D, also indicate a superior fit relative to the Gaussian specification. However, the fit 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 jumps – 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 values for this coefficient. We find a substantial improvement for increasing values of α, culminating around $\alpha = 0.5$.[16] We now obtain a much lower RMSE of 0.57%, and 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 the models we entertain, the pricing of the OTM call options is somewhat problematic. This is, of course, not unexpected, given our restriction that the right jump intensity is proportional to the left jump intensity which, in turn reflects the limited number of OTM call options in our sample. As noted earlier, this implies that a much higher weight is assigned to the OTM put option fit in our criterion function. The Z-scores for the option price and volatility fit are depicted in Figures 9 and 10, while more detailed information on the fit of the tempered stable models for alternative values of α is provided in Appendix D.

6 Characterizing the Time-Varying Jump Distributions

Section 5 demonstrates that a suitable fit to the pricing of short-dated equity-index options requires an accommodation of pronounced time-variation in the jump distribution. If this is, indeed, a robust empirical finding, it has serious implications for the proper specification of asset and derivatives pricing models. In this section, we analyze the main features of our final, and statistically satisfactory, semi-nonparametric representation for the risk-neutral asset return process. We focus on the time-variation in the volatility and jump characteristics to gauge if they are economically rational and sensible. Alternatively, one may hypothesize that the results are driven either by a substantial

[16] We note that a value of $\alpha = 0.5$ implies that jumps are of infinite activity, i.e., there are many small ones, under the risk-neutral probability. This property is preserved under equivalent martingale measure, and thus continues to hold for the statistical law of the market index.

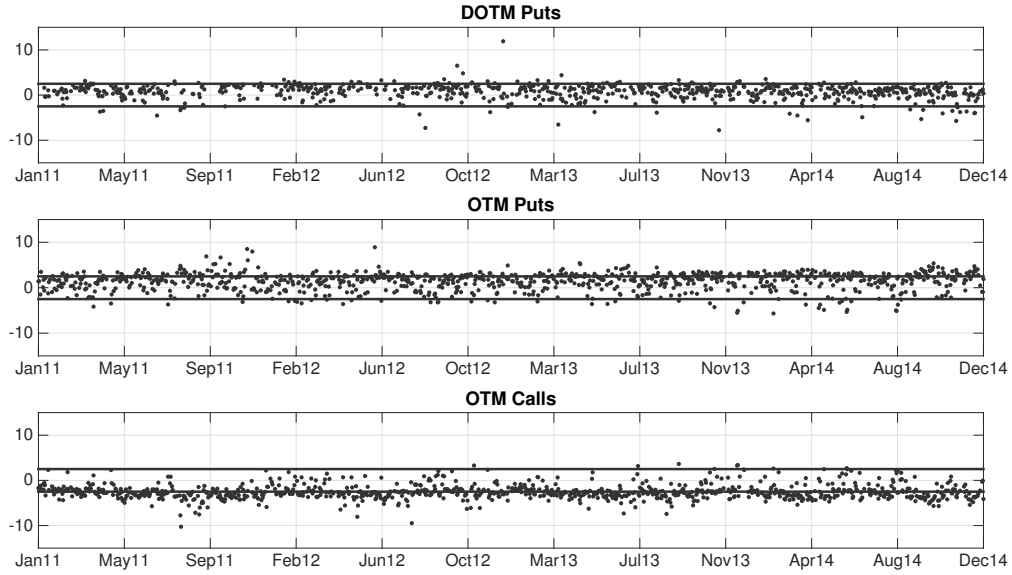

Figure 9: **The fit to short-dated option prices based on the semi-nonparametric model (17).** The figure reports the daily test statistics (8) for the short-maturity options based on the semi-nonparametric model (17) with $\alpha = 0.5$ and the state vector estimated from the short-dated options. The 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.

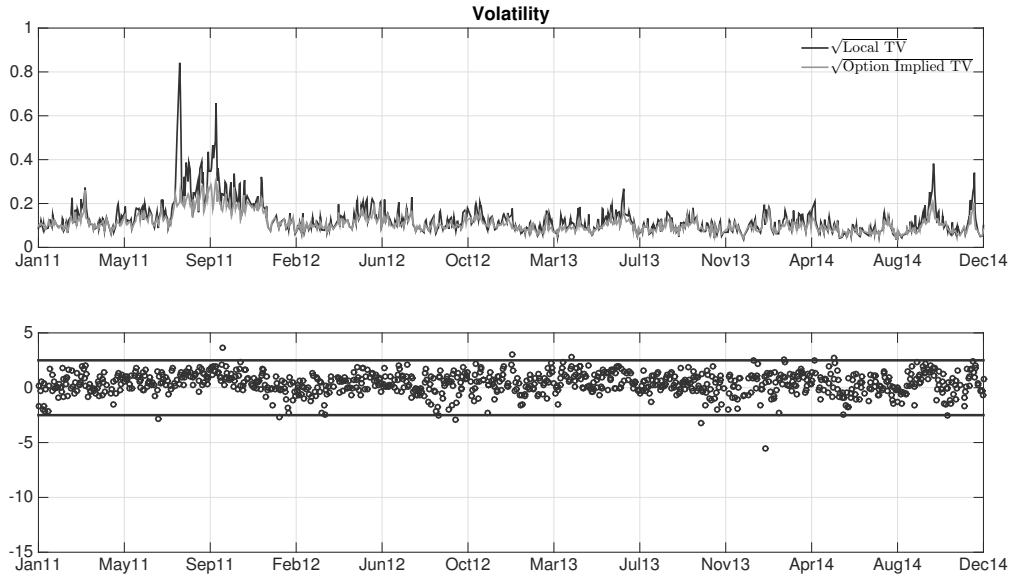

Figure 10: **The volatility fit based on the semi-nonparametric model (17).** The top panel reports end-of-day spot volatilities based on nonparametric estimates from high-frequency data (solid line) or option-implied values (grey line). The option-implied volatility estimate is computed on the basis of the semi-nonparametric model (17) 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 (9). The solid lines indicate the symmetric 95% confidence band for the test.

degree of noise in the quotes for short-dated options or by idiosyncratic pricing in this segment of the derivatives market, which may be effectively separated from mainstream capital markets.

Section 6.1 explores why the Gaussian jump model consistently performs poorly in capturing the implied jump size distribution; Section 6.2 depicts the time-variation in the extracted spot volatility and relate the more significant tail shifts to changes in the prevailing economic environment; Section 6.3 decomposes the fluctuations in the implied spot return variation into the diffusive, left- and right-jump variation components, documenting precise and economically meaningful identification of the overall return variation and its constituents; Section 6.4 shows that large fluctuations in our left jump variation measure systematically spill over into contemporaneous pricing errors for standard parametric models within regular option panels, confirming the existence of strong links between our jump variation measure and the state of the broader option market.

6.1 The Shape of the Jump Distribution

The empirical results in Sections 4 and 5 beg the question of 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 seems to underperform our alternative representations. To illustrate a common shortcoming, Figure 11 displays the observed log-price for our short-dated options on October 15, 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, there is again significant underpricing but this region contains few observations and has a limited impact on the estimation for this trading day. The tempered stable model, on the other hand, provides an impressive fit to the entire left tail, only missing marginally for extreme OTM puts. The point is that the Gaussian specification generates a negative jump mode, which induces 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 model.

We also note that both specifications produce quite substantial pricing errors for OTM calls. This reflects the dominant role of OTM puts in determining the estimation results. In turn, this stems from the relative sparsity and illiquidity of the OTM calls. As noted previously, an extended model specification, involving separately evolving intensities for large positive and negative jumps, is likely to improve the fit to the right tail. Unfortunately, the information embedded in the available OTM call options is insufficient for disentangling the intensity and size distribution for positive jumps with reasonable precision on a day-to-day basis.

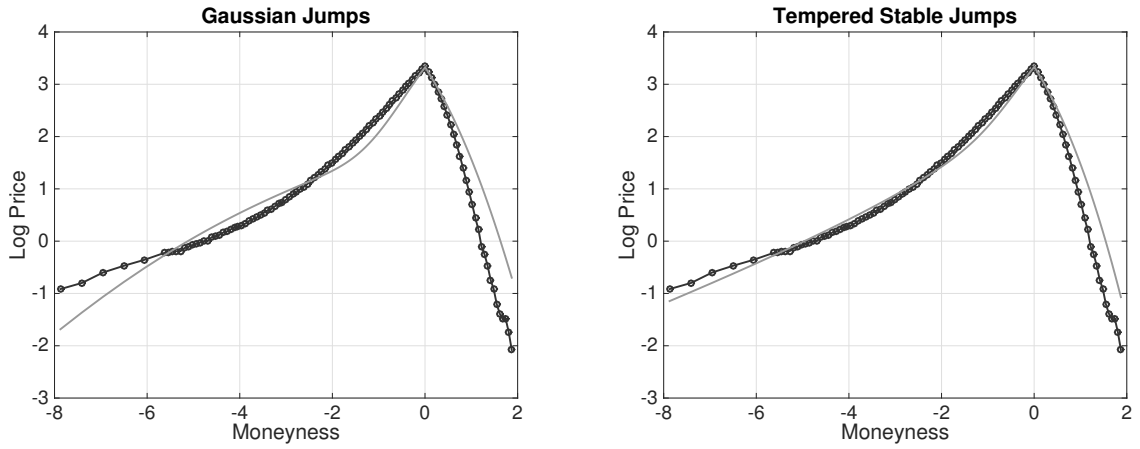

Figure 11: **Option prices and model fit on October 15, 2014**. **Left Panel**: fitted option prices (in log) according to the Gaussian semi-nonparametric model (16). **Right Panel**: fitted option prices (in log) according to the semi-nonparametric Generalized Tempered Stable model (17) with $\alpha = 0.5$. In both panels the dots represent the observed option prices while the light gray line indicates the model fitted prices. The tenor of the options is nine days.

6.2 Time Variation in Diffusive and Jump Risk

We now take a closer look at the time variation in the volatility and jump characteristics implied by the short-dated options. We use the best performing model (17) with $\alpha = 0.5$. As discussed in Section 5, our model diagnostics indicate that this model successfully captures the dynamics in the pricing of short-maturity options, with the exception of the OTM calls. Since we allow both c_t and $\lambda_t^{\pm}$ to vary unrestrictedly and the option bid-ask spread is relatively large, the daily point estimates are invariably noisy. Hence, to reduce the estimation error and better bring out the systematic time series variation in the system, Figure 12 presents 20-day moving averages of the jump parameters and the spot diffusive variance.

Figure 12 reveals some interesting dynamic patterns. Starting in the upper left panel, we see that the estimated diffusive variance is dominated by one major spike associated with the downgrade of the U.S. federal debt in August 2011. In comparison, the elevated (diffusive) volatility during the European sovereign debt crisis from early 2012 until July 2012, the fears of a prolonged government shutdown in the summer of 2013, and the worries about European growth and Ebola outbreaks in October 2014, were all quite moderate.

The estimated jump intensity series in the upper right panel is highly correlated with the diffusive variance. This is consistent with standard asset pricing models and our parametric specifications in Section 4, which attribute variation in the jump intensity to shifts in market volatility. Nevertheless, there are periods for which the two series behave markedly differently. For example,

during the European debt crisis in 2012 and the turbulence surrounding the potential government shutdown in 2013, the response of the intensity was in relative terms much larger than for volatility, while the opposite was the case towards the end of 2014.

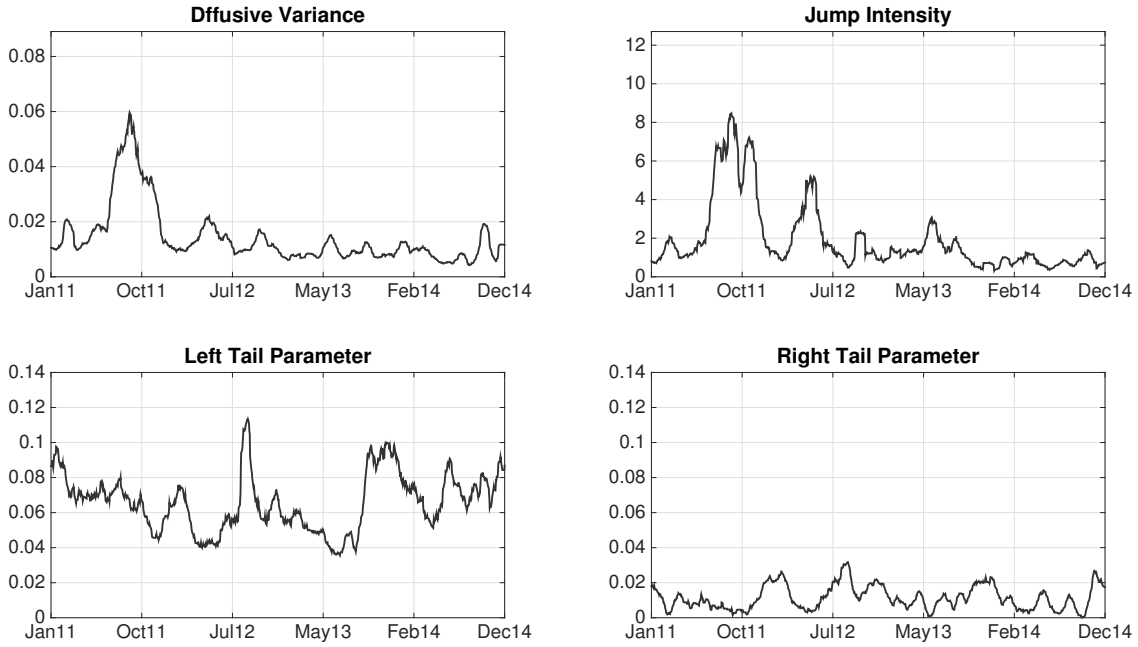

Figure 12: **Estimates for the semi-nonparametric model (17).** We plot the estimates for V_t (top left panel), c_t (top right panel), $1/\lambda_t^-$ (bottom left panel), and $1/\lambda_t^+$ (bottom right panel) based on the model (17) with $\alpha = 0.5$. For each series, we plot a 20-day moving average of the daily estimates.

The strong correlation between the spot variance and jump intensity is compatible with the standard option pricing framework as long as the jump size distribution remains invariant. However, the bottom left panel reveals that the left tail parameter λ_t^- exhibits pronounced time variation. This suggests that the standard paradigm will have difficulty in capturing the type of time variation in the pricing of short-dated options observed across our sample. To further interpret the variation in the negative jump size distribution, we first review the features of the cross-section of option prices that facilitate separate identification of the jump intensity and jump size distribution.

For a fixed value of α, λ_t^- determines the shape of the risk-neutral jump distribution for the given day and, in particular, the slope of the implied volatility curve for "sufficiently" low moneyness, i.e., the far left tail (recall approximation (10) for deep OTM options). On the other hand, the jump intensity c_t governs the overall level of implied volatility for "sufficiently" deep OTM options and it impacts both the level and curvature of the implied volatility curve in the near-ATM region.

In summary, a shift in λ_t^- alters the slope, or thickness, of the far left tail of the implied volatility curve, while c_t moves the curve up or down in a parallel fashion.

The left tail of the risk-neutral jump distribution is thickest during the beginning and end of our sample as well as during July of 2012 and summer through fall of 2013. The initial high level may reflect the aftermath of the first European sovereign debt crisis. The slow decay in $1/\lambda_t^-$ merely paused for the downgrade of U.S. debt in August 2011. That is, while the volatility and jump intensity soared, there was no perception of an increase in the size of a potential market crash. In contrast, a substantial thickening of the tail is evident during the European debt crisis in 2012, the fear of a government shutdown in 2013, and during a prolonged period in 2014, characterized by the Russian annexation of Crimea and the imposition of sanctions, concerns about global economic growth, collapsing oil prices and, finally, the Russian foreign exchange crisis in December 2014. These episodes seem to induce thicker tails rather than enhance the general sense of uncertainty, as the corresponding movements in volatility and jump intensity are muted.

These observations suggest that the variation of λ_t^- is not strongly related to that of the diffusive volatility or c_t. Indeed, the correlation between $1/\lambda_t^-$, on the one hand, and V_t or c_t, on the other, is 0.01 and -0.37, respectively. Since $1/\lambda_t^-$ and c_t have a fairly similar effect on the pricing of deep OTM put options (recall approximation (10)), the latter correlation can be explained, at least in part, by the correlation of the estimation errors in recovering these two parameters from the data.

Finally, the bottom right panel of Figure 12 confirms the difficulty of recovering the shape of the positive jump distribution. The estimates of λ_t^+ suggest the right tail is an order of magnitude thinner than the left tail. The series is mildly correlated with the left tail parameter, but the systematic fluctuations across the sample are modest. This reflects the lower prices, large spreads, and relatively poor liquidity for deep OTM calls, as discussed previously.

To further illustrate the separation of the volatility dynamics from the behavior of the left jump tail, suggested by Figure 12, we display the logarithmic option prices on three select days in Figure 13. They complement Figure 1 by providing additional examples of isolated shifts in the slope of the left tail. The upper panels depict log prices of the short-dated options for the day of the tail shift as well as for a pair of nearby trading days. The lower panels cover the identical days, but display log option prices from our regular sample with about one month to expiry.

We note that the log prices in each panel of the top row are very close across the nearby trading days for $m \in (-2, 3)$. This implies, in particular, that neither the return volatility nor jump intensity shifted significantly across those days. However, in all cases, we observe a striking increase in the (negative) slope of the left tail. This type of movement is not compatible with standard asset

pricing models for which the jump size distribution is invariant and the jump intensity is tied to volatility. The challenge is only magnified by noting that there is a plausible economic rationale for the tail shift on each of these dates, including the one in Figure 1.[17]

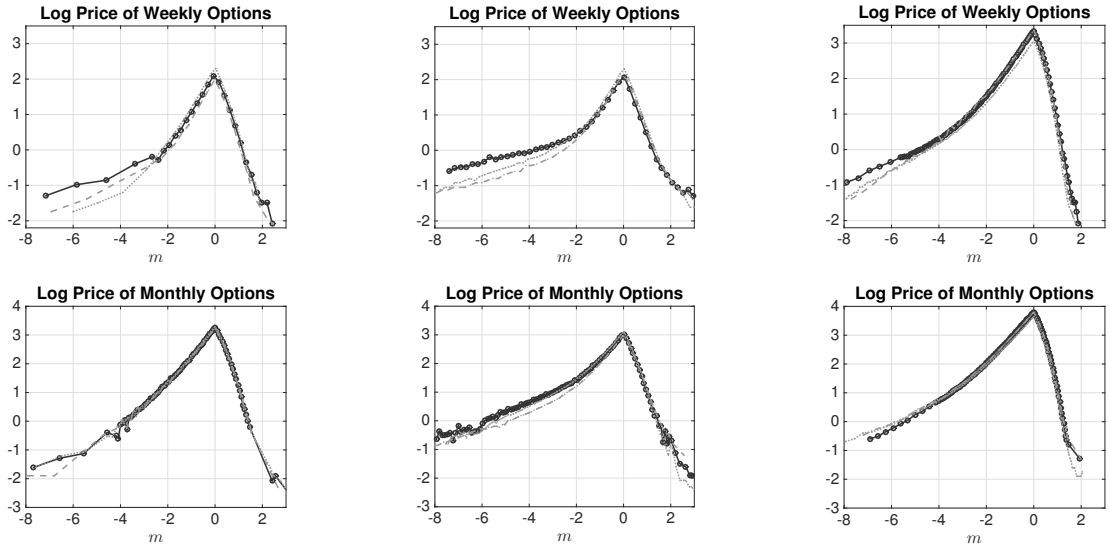

Figure 13: **Log Option Prices**. **First column: January 7, 2011**. The short-dated options in the top panel have 7 calendar days to expiry, while the regular options in the bottom panel have a tenor of 42 calendar days. The surrounding trading days are January 6 and 10, 2011. **Second column: December 20, 2013**. The short-dated options in the top panel have 7 calendar days to expiry, while the regular options in the bottom panel have a tenor of 28 calendar days. The surrounding trading days are December 19 and 26, 2013. **Third column: October 15, 2014**. The short-dated options in the top panel have 9 calendar days to expiry, while the regular options in the bottom panel have a tenor of 30 calendar days. The surrounding trading days are October 17 and 16, 2014.

Comparing the top and bottom panels, we note that the regular options also provide a hint of the tail movements, but the evidence is less clear, often blurred by idiosyncratic noise and, overall, not particularly compelling. This is perhaps not surprising. As the regular options are longer-dated, the effect from a change in the jump size distribution is compounded by time variation in the volatility and jump intensity as well as their interaction. Consequently, the power to detect shifts in the jump distribution should, indeed, be much better for the short-maturity sample.

[17]For example, on January 7, 2011, the employment news were disappointing and a couple of critical rulings concerning the conduct of banks during the financial crisis caused losses for the financial sector. Moreover, this occurred against the backdrop of tension related to the European sovereign debt crisis. Further, we previously discussed the global economic growth concerns and Ebola outbreaks dominating the news on October 15, 2014.

6.3 Implied Variation Measures

We next explore the implications of our estimates for measures of diffusive and jump variance. In particular, we plot model-implied estimates for the spot (diffusive) variance, V_t, and the left and right jump variances, $LJV_t = \int_{x<0} x^2 \nu_t(dx)$ and $RJV_t = \int_{x>0} x^2 \nu_t(dx)$. These quantities are the spot counterparts for the measures that comprise the theoretical valuation of the VIX index.[18]

Figure 14 provides the 95% confidence bounds for the above spot estimates implied by our best performing model (17) with $\alpha = 0.5$. From the top panel of Figure 14, we note that spot diffusive variance is recovered with very good precision. As expected, its confidence intervals become wider during the more volatile periods. The alternative nonparametric estimator of V_t from high-frequency data on X is much noisier. Nonetheless, these two spot volatility series are compatible, as documented in Section 5.2. On the other hand, as discussed above, estimates for V_t based on regular longer-dated options are more prone to model misspecification. Hence, the estimates for V_t based on the semi-nonparametric specification (17) and short-maturity options strike a good balance between precision and robustness to model misspecification. As such, they arguably provide the most reliable spot variance estimates available in the literature.

The last two panels of Figure 14 depict the risk-neutral spot jump variation measures. Unlike 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 even tighter confidence bounds than for V_t. In contrast, the relative precision in recovering the right jump variation is poor, reflecting the scarcity of deep OTM call options. However, this does not impinge on the overall accuracy of our risk-neutral variation measure, as the magnitude of the right jump variation is inconsequential relative to the other components.

Figure 14 reveals interesting connections between the diffusive variance and the left jump variance. On average, they are about equal in magnitude. This provides a sharp contrast to the much more modest contribution of jumps to the return variation under the physical 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. The first is early August 2011, when there is a surge in both the diffusive and left jump variation, but the latter is almost three times the size of the former. As noted previously, this period was dominated by fears linked to the S&P downgrade of the US credit rating. While the level of general uncertainty

[18] Formally, the 30-day VIX measure is given by the expected value of the cumulative sum of these measures over the coming calendar month, i.e., $VIX_t \approx \mathbb{E}_t^{\mathbb{Q}} \left[\int_t^{t+\tau} (V_s + LJV_s + RJV_s)\, ds \right]$ with τ equal to one calendar month.

rose sharply, the event disproportionally impacted the priced tail risk. The second episode is the first half of October 2014 where, again, the diffusive and the left jump variance increase sharply.[19] 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.

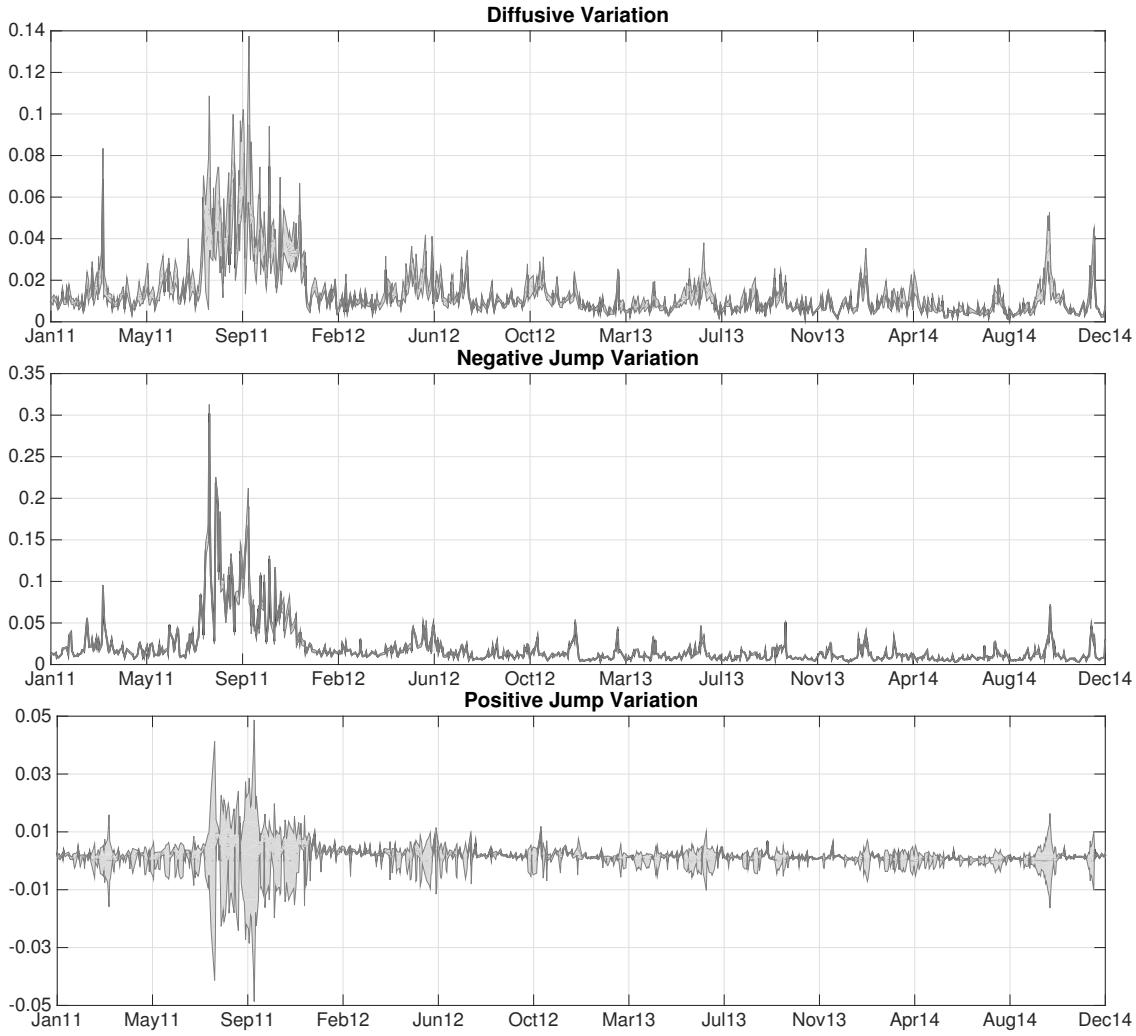

Figure 14: **Measures of Return Variation**. The figure plots 95% confidence intervals for the daily V_t (top panel), $\int_{x<0} x^2 \nu_t(dx)$ (middle panel) and $\int_{x>0} x^2 \nu_t(dx)$ (bottom panel) implied by estimates for the semi-nonparametric model (17) with $\alpha = 0.5$ from the short-dated options. All quantities are reported in annualized units.

[19] As mentioned earlier, the main sources of concern at that time were signs of weakness in the European economy and its potential global impact as well as the spread of incidents linked to the Ebola virus.

6.4 Tail Jump Variation and Standard Option Pricing Models

The preceding sections document a significant amount of variation in the pricing of negative market jumps and show that these fluctuations are not directly linked to market volatility. Moreover, they demonstrate that such time-variation in the pricing of jump risk is most readily identified from short-dated options. In line with this reasoning, the bulk of our empirical evidence hinges of information extracted exclusively from options with very short tenor.

It is natural to ask whether our evidence concerning the pricing of jump risk is strictly limited to the short-dated options. That is, can we identify implications for the type of longer-dated options that are used in standard asset pricing and derivatives studies? In theory, this should be feasible if the pricing of jump risk is economically meaningful and persistent as opposed to purely idiosyncratic. In this case, periods characterized, via short-dated options, as containing elevated jump variation should manifest themselves as distinct in terms of the pricing of longer-dated options as well. Of course, for longer-dated options, other sources of risks 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 critical ingredients in the computation of the 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 (17) with $\alpha = 0.5$, along with the one implied by estimates from the parametric model (3) and (5), based on our longer-dated (regular) option sample. 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 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, going beyond 20 several times, and even topping 30 around the downgrade of US debt. The truly striking feature, however, is the remarkable correlation between these episodes of blatant parametric mispricing for OTM put options and the periods for which we observe very large deviations between the jump

variation estimates across the two models, indicated in the top panel.

In summary, the jump variation measures, extracted exclusively from the short-dated option sample, are highly informative regarding qualitatively similar features of the regular option panel. Whenever the parametric models – estimated from and evaluated against regular longer-dated options only – fail to capture elevation in these jump variation measures, they also fail, 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 "regular" option panels. The liquid trading of short-dated options affords us with a simple and efficient procedure for identifying such priced jump tail risks.

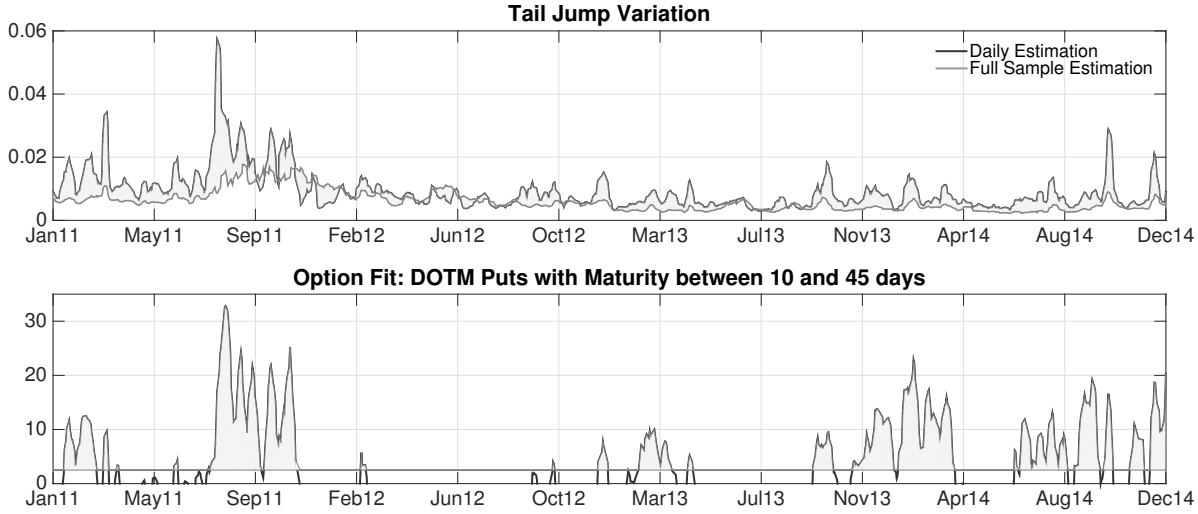

Figure 15: **Tail Jump Variation**. **Top Panel**: Jump Variation implied by the parameters and state vector from the estimates of the parametric model defined by equations (3) and (5) based on the regular option panel and the semi-nonparametric model (17) with $\alpha = 0.5$ based on the short-dated options. For each model, we report the quantities $\int_{-\infty}^{b_t} x^2 \nu_t(dx)$, using the corresponding parameter and state vector estimates, where b_t equals the negative of three times the daily ATM implied volatility. **Bottom Panel**: Option Z-scores for the fit of the parametric model defined by equations (3) and (5) to options with moneyness [-7,-2] and tenor between 10 and 45 days. The shaded area on the plot corresponds to Z-scores in excess of 2.5 (i.e., when the model severely underestimates the OTM Puts). The correlation between the shaded areas on the two plots equals 0.54. All series are reported as five-day moving averages.

6.5 Interpretation and Implications

We conclude by briefly outlining some potential implications of our findings. The top panel of Figure 15 depicts a highly volatile priced negative jump tail risk whose time variation is distinct

from that of market volatility. This suggests that it may be useful 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 (17) 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 enable us to incorporate information about tail risks and volatility embedded in short-maturity options into the pricing of regular longer-dated options. The evidence from Figure 17 strongly suggests this will boost the performance of such models in capturing the price dynamics of broader option panels. This, in turn, will allow for more robust extraction of information regarding the pricing of the 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 suggest 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 pricing of market jump risk. This, in turn, suggests time-varying risk aversion towards tail risk and/or heterogeneity in the attitude towards tail risk among investors. To pin down the economic forces that can rationalize the empirical evidence, obtained from our short-dated option sample, requires an extensive study exploiting data on the full option panel and the underlying asset as well as, potentially, volume data for option trading.

7 Conclusion

In this paper we study short-dated options written on the S&P 500 market index. The trading in this segment of the market for equity-index options has increased significantly over the last five years with the introduction and increased popularity of the so-called "weeklies". In this paper we seek to understand the dynamic behavior of this sizable and rapidly growing part of the option market which hitherto has been essentially unexplored. We confirm 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 the 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 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 further show that the dynamic behavior of priced tail risk is reflected in the prices of longer-dated options. Moreover, we document that this feature is a major source of misspecification of 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 investors concerns about this critical component of risk and for direct exploration of its dynamic behavior as well as its interaction with other observed economic quantities.

References

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. and D. Madan (2000). Spanning and Derivative Security Valuation. *Journal of Financial Economics 55*, 205–238.

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.

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). Tails Risk Premia and Return Predictability. *Journal of Financial Economics, forthcoming*.

Breeden, D. and R. Litzenberger (1978). Prices of State Contingent Claims Implicit in Option Prices. *Journal of Business 51*, 621–652.

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 K. Mimouni (2010). Volatility Dynamics for the S&P 500: Evidence from Realized Volatility, Daily Returns, and Option Prices. *Review of Financial Studies 23*, 3141–3189.

Christoffersen, P., K. Jacobs, and C. Ornthanalai (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.

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.

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, forthcoming*.

Kou, S. (2002). A Jump Diffusion Model for Option Pricing. *Management Science 48*, 1086–1101.

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 Data 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 16 illustrates how the filtering procedure works for three trading days in our sample.

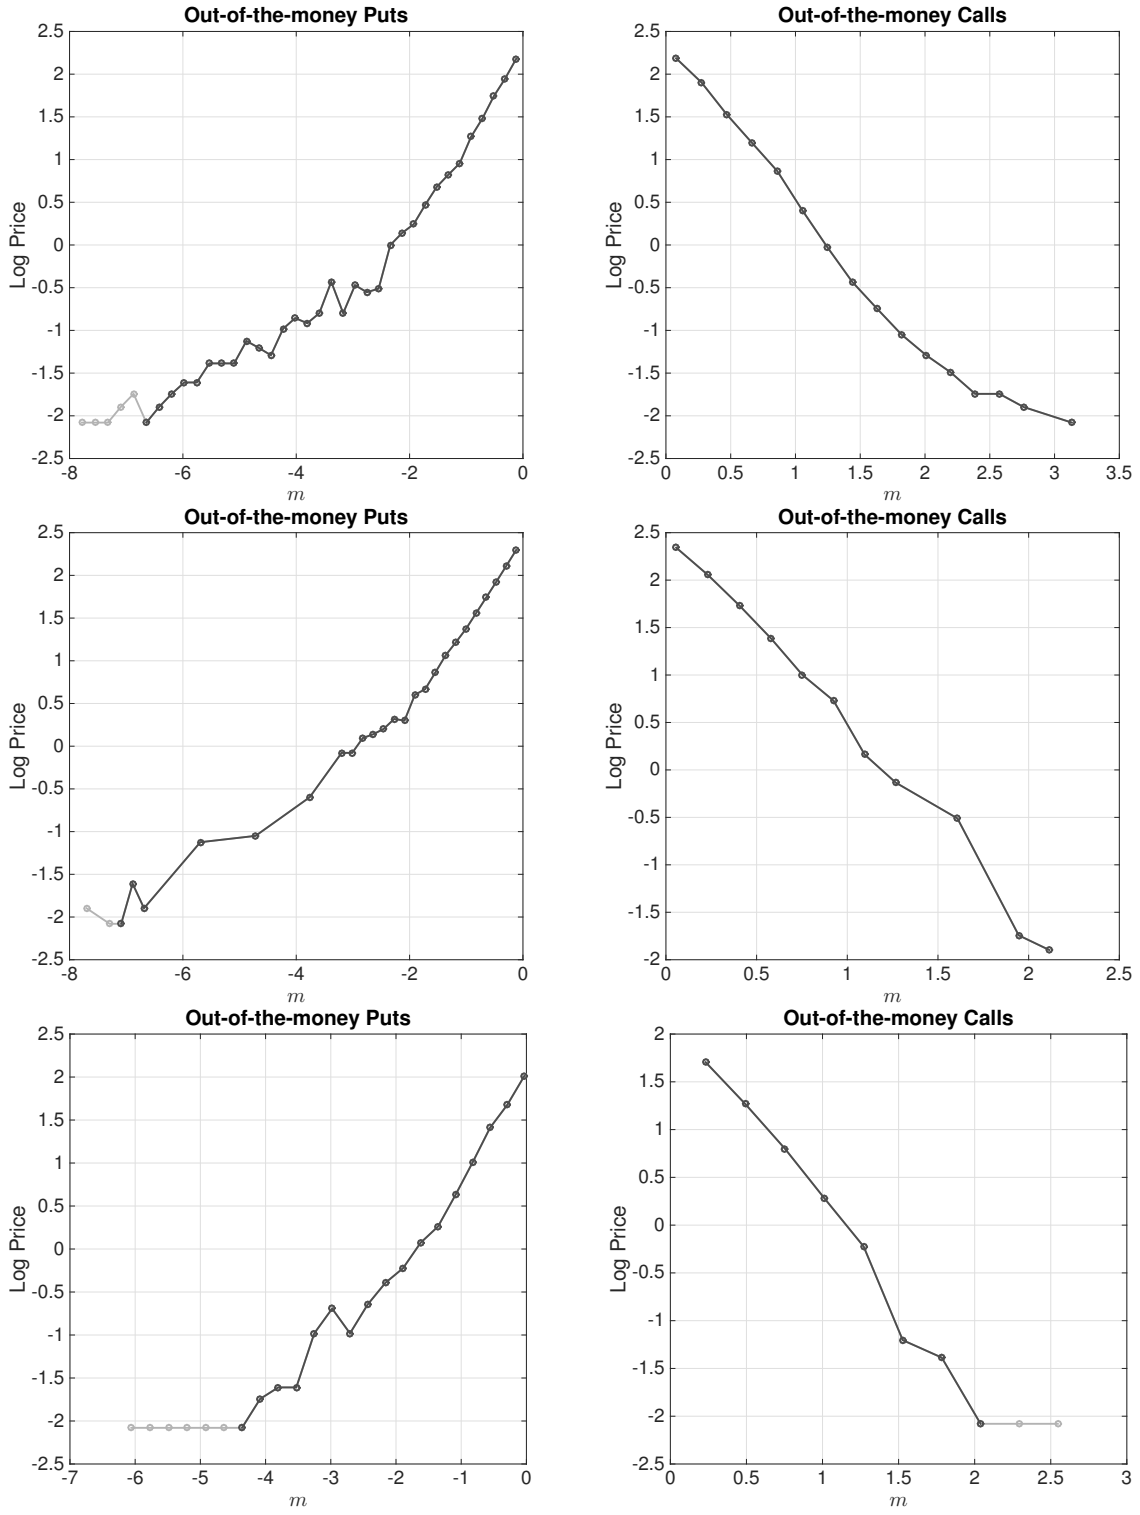

Figure 16: **Data Filters**. Illustration of the filtering procedure applied to the option data. Top panel: January 13th 2012, middle panel: March 21st 2012 and bottom panel: May 16th 2012. The dots on the plots correspond to the observed option prices (in logs). The solid line connects the data points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

B Additional Results: Parametric Estimation based on the "Regular" Option Sample

Panel A: Parameter Estimates

Parameter	Estimate	Std.	Parameter	Estimate	Std.	Parameter	Estimate	Std.
ρ_1	−0.9998	0.0122	$\overline{v}_2$	0.0088	0.0001	η_2	5.0739	0.2159
$\overline{v}_1$	0.0173	0.0002	κ_2	11.5603	0.2371	μ_x	−0.1241	0.0022
κ_1	2.0769	0.0219	σ_2	0.3708	0.0052	σ_x	0.1141	0.0005
σ_1	0.2678	0.0060	η_0	0.0033	0.0029	μ_v	0.0850	0.0011
ρ_2	−0.7185	0.0049	η_1	20.2664	0.3938	ρ_j	−0.3260	0.0185

Panel B: Summary Statistics

RMSE	1.5179%
Mean jump intensity (yearly)	0.2963
Mean jump size	−0.2115
Mean diffusive variance	0.0234
Mean negative jump variance	0.0171
Mean postive jump variance	0.0000

Table 2: **Estimation results for the parametric model defined by equations (3) and (4).** 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.

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
$\overline{v}_1$	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
$\overline{v}_2$	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 postive jump variance	0.0014

Table 3: **Estimation results for the parametric model defined by equations (3) and (5).** 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.

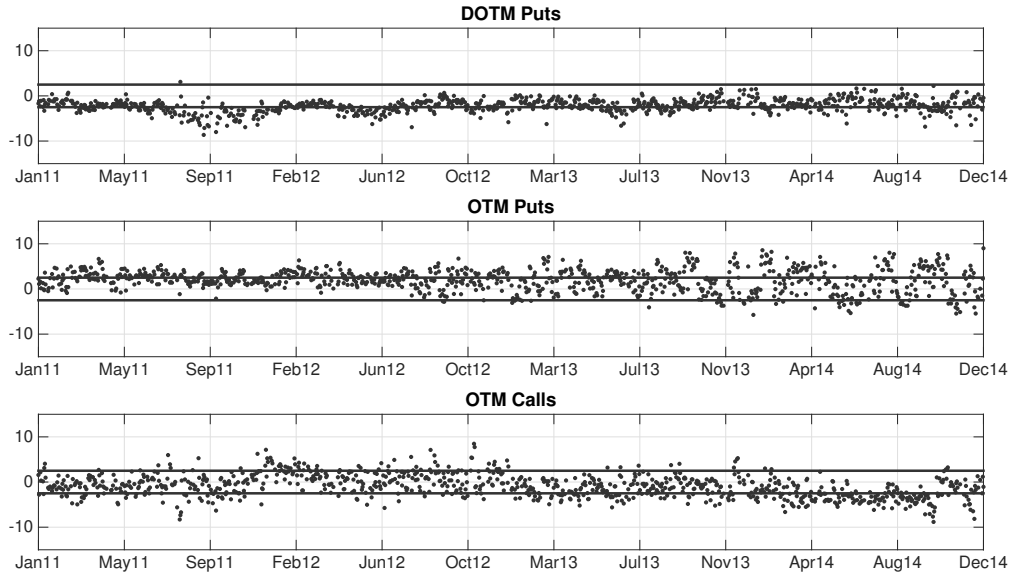

Figure 17: **Short-dated option price fit based on the parametric model with Gaussian jumps estimated from the "regular" option sample.** The figure reports the daily option fit tests in (8) for the short-dated options based on the parametric model in equations (3) and (4) with parameters estimated from the "regular" option sample and the state vector estimated separately from the short-dated options each trading day. 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.

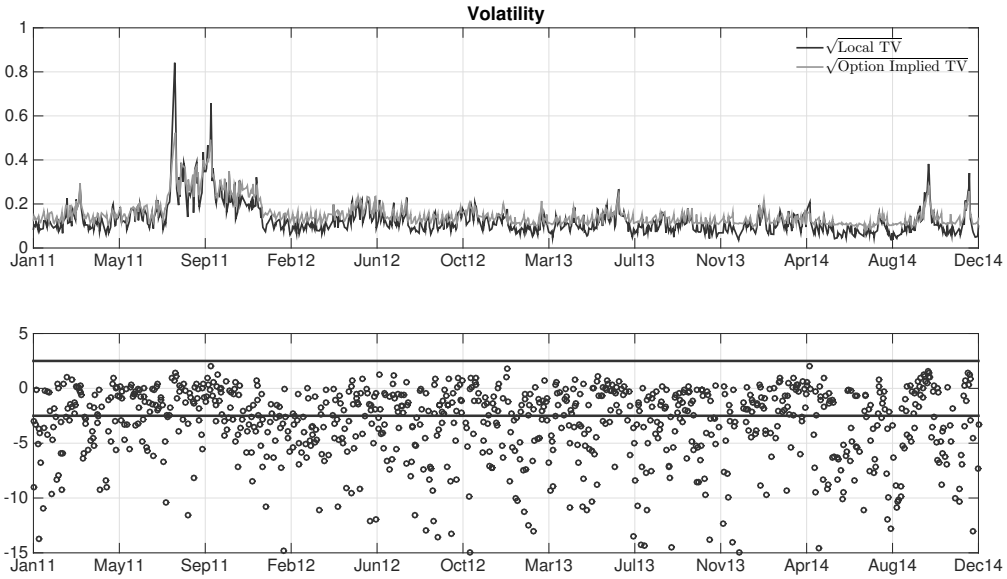

Figure 18: **Volatility fit based on a parametric model with Gaussian jumps estimated using "regular" options.** The top panel reports end-of-day spot volatilities based on nonparametric estimates from high-frequency data (solid line) or option-implied values (grey line). The option-implied volatility estimate is computed using the parametric model defined in equations (3) and (4) with parameters estimated from the "regular" option sample and the state vector estimated separately from the short-dated options on each trading day. The volatility estimates are reported in annualized units. The bottom panel reports the corresponding daily volatility test statistic (9). The solid lines indicate the symmetric 95% confidence band for the test.

C Additional Results: Semi-nonparametric Estimation with Fixed Jump Distribution Parameters using Short-Dated Options

	Mean	Std	Q05	Q50	Q95
Gaussian Jumps					
RMSE	1.189	0.619	0.568	1.048	2.322
Diffusive Variance	0.019	0.017	0.006	0.014	0.053
Negative Jump Variance	0.014	0.023	0.003	0.008	0.049
Positive Jump Variance	0.001	0.002	0.000	0.001	0.003
Double Exponential Jumps					
RMSE	1.023	0.528	0.488	0.913	1.909
Diffusive Variance	0.018	0.015	0.005	0.013	0.050
Negative Jump Variance	0.016	0.024	0.004	0.009	0.052
Positive Jump Variance	0.001	0.002	0.000	0.001	0.004

Table 4: **Estimation results for the semi-nonparametric models (13) and (14).** 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 – December 2014. The Gaussian jumps refer to model (13) and the double exponential jumps to model (14) 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.

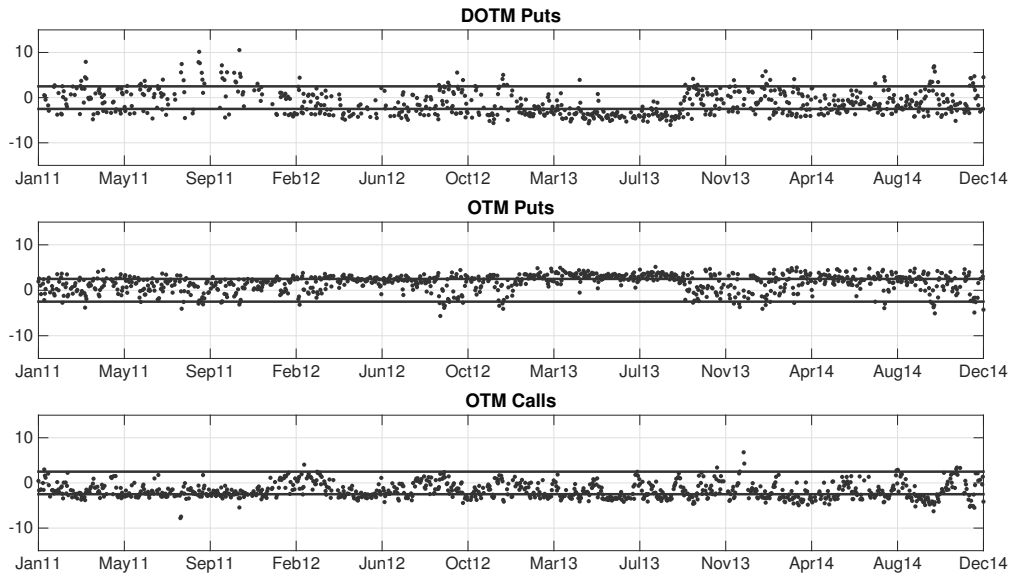

Figure 19: **The fit to short-dated option prices based on the semi-nonparametric model with Gaussian jumps estimated from the short-maturity sample.** The figure reports the daily test statistics (8) for the short-maturity options based on the semi-nonparametric model (13) with the parameters and state 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.

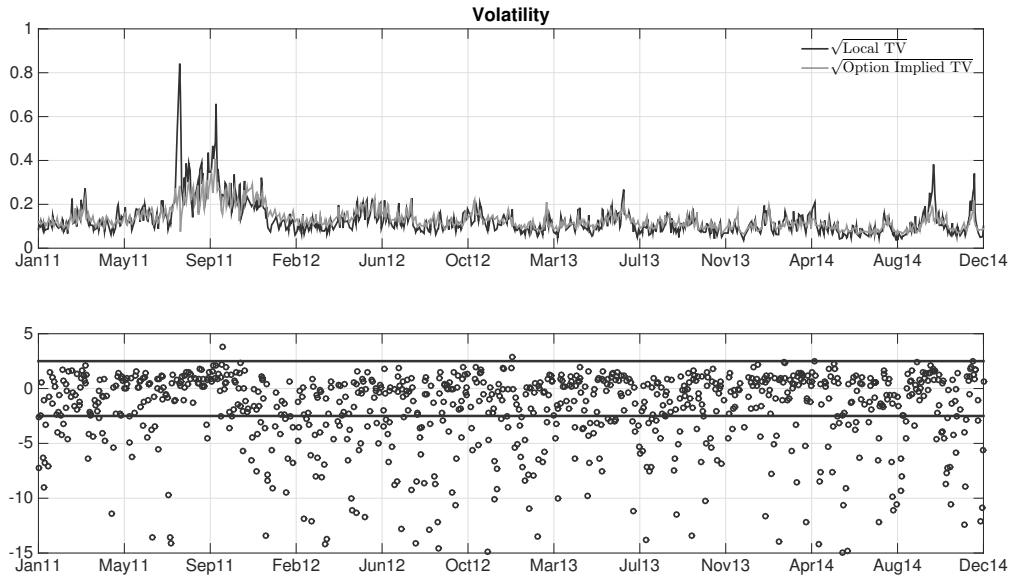

Figure 20: **The volatility fit based on the semi-nonparametric model with Gaussian jumps estimated from the short-dated options.** The top panel reports end-of-day spot volatility estimates based on nonparametric techniques and high-frequency data (solid line) or option-implied values (grey line). The option-implied volatility estimate is computed on the basis of the semi-nonparametric model (13) 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 (9). The solid lines indicate the symmetric 95% confidence band for the test.

D Additional Results: Semi-nonparametric Estimation from Short-Dated Options; Time-Varying Jump Distribution Parameters

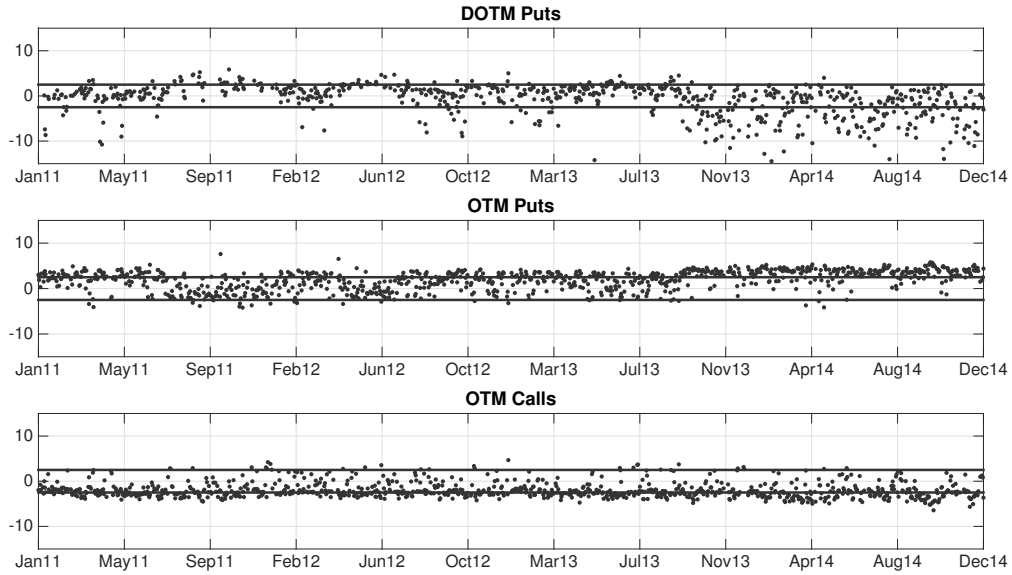

Figure 21: **The fit to short-dated option prices based on the semi-nonparametric model (16) with Gaussian jumps estimated from the short-maturity sample.** The figure reports the daily test statistics (8) for the short-maturity options based on the semi-nonparametric model (16) with the state vector 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	0.855	0.402	0.404	0.754	1.717
Diffusive Variance	0.018	0.016	0.005	0.013	0.054
Negative Jump Variance	0.015	0.021	0.004	0.009	0.051
Positive Jump Variance	0.001	0.001	0.000	0.001	0.003
Double Exponential Jumps ($\alpha = -1$)					
RMSE	0.677	0.299	0.346	0.590	1.309
Diffusive Variance	0.016	0.014	0.005	0.012	0.047
Negative Jump Variance	0.018	0.023	0.004	0.011	0.061
Positive Jump Variance	0.001	0.001	0.000	0.001	0.003
Tempered Stable Jumps with $\alpha = 0$					
RMSE	0.606	0.273	0.312	0.531	1.173
Diffusive Variance	0.015	0.013	0.005	0.011	0.043
Negative Jump Variance	0.019	0.025	0.005	0.012	0.065
Positive Jump Variance	0.001	0.001	0.000	0.001	0.004
Tempered Stable Jumps with $\alpha = 0.5$					
RMSE	0.571	0.270	0.292	0.495	1.114
Diffusive Variance	0.014	0.012	0.004	0.010	0.040
Negative Jump Variance	0.020	0.026	0.005	0.012	0.069
Positive Jump Variance	0.001	0.002	0.000	0.001	0.004

Table 5: **Estimation results for the semi-nonparametric models (16) and (17).** 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 – December 2014. The Gaussian jumps refer to model (16) and those with specific α values refer to model (17). 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.

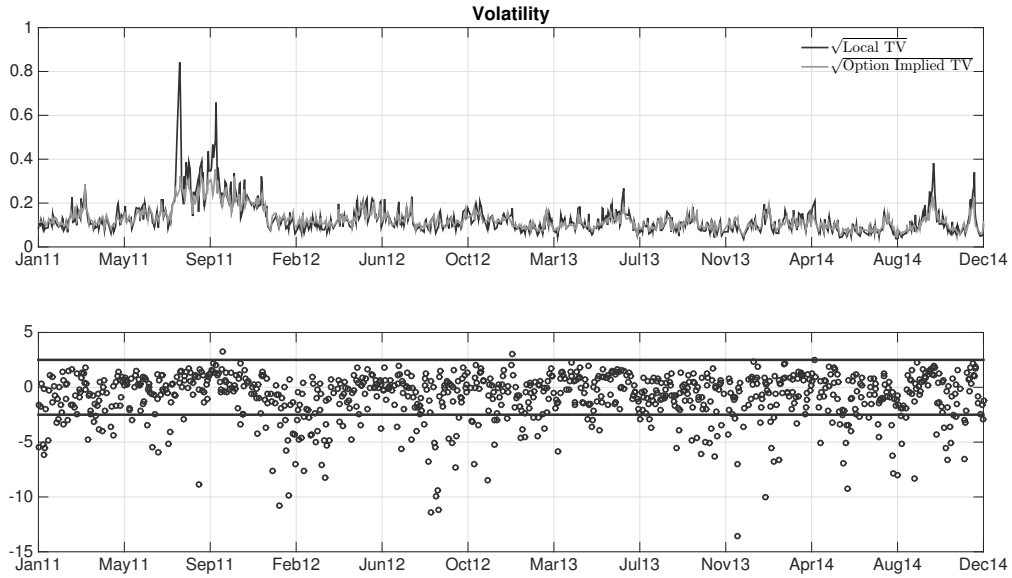

Figure 22: **The volatility fit based on the semi-nonparametric model (16) with Gaussian jumps estimated from the short-dated options.** The top panel reports end-of-day spot volatilities based on nonparametric estimates from high-frequency data (solid line) or option-implied values (grey line). The option-implied volatility estimate is computed on the basis of the semi-nonparametric model (16) 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 (9). The solid lines indicate the symmetric 95% confidence band for the test.

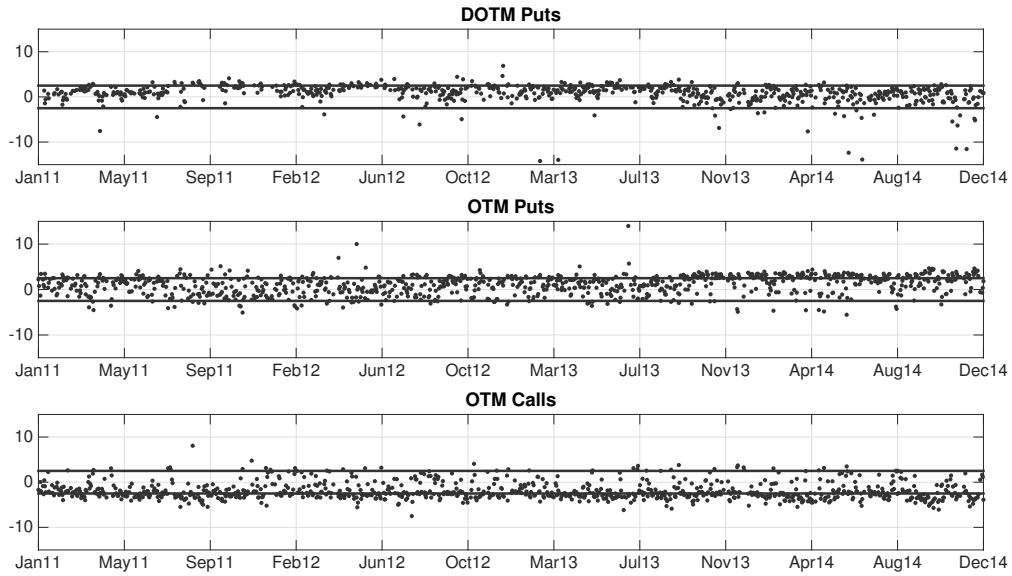

Figure 23: **The fit to short-dated option prices based on the semi-nonparametric model (17) with double-exponential jumps estimated from the short-maturity sample.** The figure reports the daily test statistics (8) for the short-maturity options based on the semi-nonparametric model (17) with $\alpha = -1$ 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.

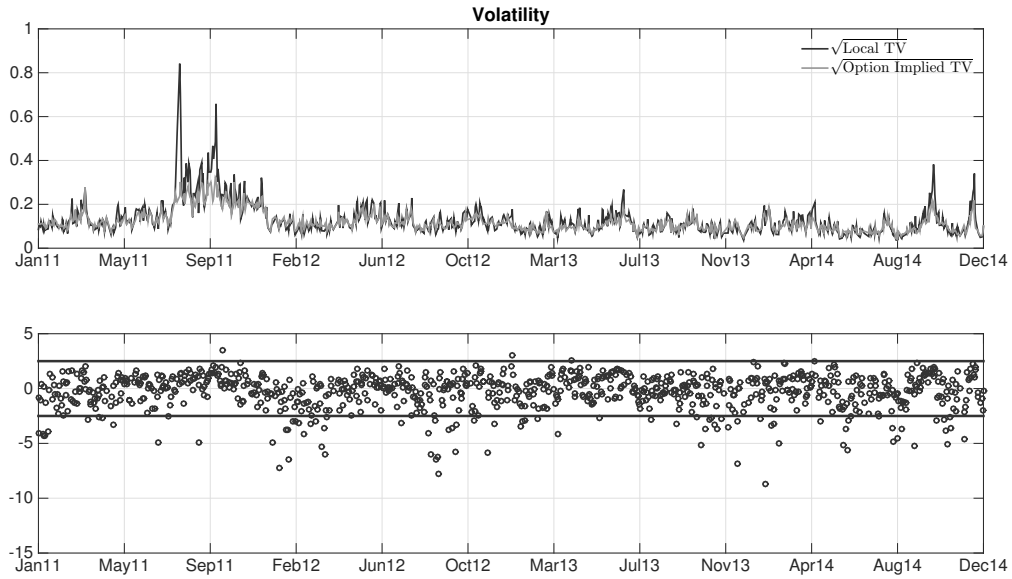

Figure 24: **The volatility fit based on the semi-nonparametric model (17) with double-exponential jumps estimated from the short-maturity sample.** The top panel reports end-of-day spot volatilities based on nonparametric estimates from high-frequency data (solid line) or option-implied values (grey line). The option-implied volatility estimate is computed on the basis of the semi-nonparametric model (17) 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 (9). The solid lines indicate the symmetric 95% confidence band for the test.