



The fine structure of equity-index option dynamics[☆]



Torben G. Andersen^{a,b,c,*}, Oleg Bondarenko^d, Viktor Todorov^a, George Tauchen^e

^a Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208, United States

^b NBER, Cambridge, MA, United States

^c CREATES, Aarhus, Denmark

^d Department of Finance, University of Illinois at Chicago, Chicago, IL 60607, United States

^e Department of Economics, Duke University, Durham, NC 27708, United States

ARTICLE INFO

Article history:

Available online 10 March 2015

JEL classification:

C51

C52

G12

Keywords:

High-frequency data

Implied volatility

Jump activity

Kolmogorov–Smirnov test

Stable process

Stochastic volatility

VIX index

ABSTRACT

We analyze the high-frequency dynamics of S&P 500 equity-index option prices by constructing an assortment of implied volatility measures. This allows us to infer the underlying fine structure behind the innovations in the latent state variables driving the evolution of the volatility surface. In particular, we focus attention on implied volatilities covering a wide range of moneyness (strike/underlying stock price), which load differentially on the different latent state variables. We conduct a similar analysis for high-frequency observations on the VIX volatility index as well as on futures written on it. We find that the innovations over small time scales in the risk-neutral intensity of the negative jumps in the S&P 500 index, which is the dominant component of the short-maturity out-of-the-money put implied volatility dynamics, are best described via non-Gaussian shocks, i.e., jumps. On the other hand, the innovations over small time scales of the diffusive volatility, which is the dominant component in the short-maturity at-the-money option implied volatility dynamics, are best modeled as Gaussian with occasional jumps.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

Volatility risk is a major concern for investors and they require compensation for bearing it. Over the last two decades trading in derivatives, allowing for speculation and hedging vis-a-vis volatility risk, has grown dramatically. These instruments include plain vanilla options but also, more directly, so-called variance swaps, which are forward contracts on realized volatility (which in turn are nonparametric estimates for the unobserved quadratic variation). The price of a variance swap can be recovered in model-free fashion from the price of a portfolio of out-of-the-money (OTM) options on the underlying asset. The Chicago Board Options

Exchange (CBOE) relies on this methodology in computing the well-known VIX volatility index based on S&P 500 index options. In recent years, the VIX index itself has become the underlying instrument for futures and options, further expanding the opportunities for managing exposures to equity market volatility risk.

The abundance of reliable data on volatility-related derivative contracts enables us to take a closer look at the properties of the process driving the innovations to spot (stochastic) volatility and jump intensity, which are otherwise latent, or “hidden”, within the stock returns. Todorov and Tauchen (2011b) show that, although the VIX is a risk-neutral expectation of future realized volatility, under conventional model settings, the VIX index preserves important information about the behavior of the latent stochastic volatility over small time scales. In particular, the presence of VIX jumps can be traced back to discontinuities in the spot volatility process itself, even if the magnitudes of the jumps in the two series generally differ. Similarly, the “degree of concentration” of small jumps, known as the jump activity or Blumenthal–Gettoor index, in (spot) volatility and the VIX index coincides. Finally, the presence of a diffusive component is preserved when going from spot volatility to VIX. Based on this line of reasoning, using high-frequency VIX data, Todorov and Tauchen (2011b) conclude that the VIX index – and by extension spot volatility – contains jumps

[☆] 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 National Science Foundation grant SES-0957330. We are also grateful for support from a grant by the CME Group. We thank Makoto Takahashi for providing assistance with collecting the VIX futures data.

* Corresponding author at: Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208, United States.

E-mail addresses: t-andersen@northwestern.edu (T.G. Andersen), olegb@uic.edu (O. Bondarenko), v-todorov@northwestern.edu (V. Todorov), george.tauchen@duke.edu (G. Tauchen).

and is best characterized as a pure-jump process with infinite variation jumps.

The existing literature on the activity level of spot volatility is extremely limited, with the main contribution being [Todorov and Tauchen \(2011b\)](#), which is based solely on high-frequency data for the VIX index and relies on ad hoc procedures in dealing with the confounding effects of market microstructure noise. This state-of-affairs reflects the fact that, until recently, we neither had access to high-frequency data for alternative volatility-sensitive derivatives nor suitable theoretical tools for drawing inference regarding the activity index for volatility series.¹ The goal of the current paper is to generate new and robust empirical evidence concerning the properties of the latent spot volatility and jump intensity over small time scales by exploiting high-frequency data across a greatly expanded set of derivative contracts relative to prior studies in the literature.

Under the common assumption that the risk-neutral jump intensity is a sole function of (components of) spot volatility, the corresponding Black–Scholes implied volatility (BSIV) measures, extracted from S&P 500 index options, are also functions of volatility alone (along with the characteristics of the option contracts). Hence, they likewise typically “inherit” the behavior of volatility over small scales. Furthermore, the identical logic applies to derivatives on the VIX index such as VIX futures. Therefore, data for an extended set of volatility derivatives will enhance efficiency and robustness in evaluating existing findings based strictly on high-frequency VIX series.

Importantly, the additional derivatives data also allow us to gain qualitatively new insights. For example, many studies conclude that the return variation is governed by multiple factors. Further, recent evidence suggests that the dynamics of the risk-neutral negative jump intensity for the equity index cannot be captured fully by (components of) spot volatility, see, e.g., [Bollerslev and Todorov \(2011\)](#) and [Andersen et al. \(forthcoming\)](#). In this case, VIX is governed by the factors driving both the jump intensity and the spot volatility. The use of derivative data loading differentially on such factors (or state variables) helps us discern the properties of those factors and thus fosters a deeper understanding of the fine structure of the VIX dynamics. For example, the BSIV of short-maturity OTM put options load primarily on the risk-neutral intensity of negative jumps, so their high-frequency increments reflect the small scale behavior of the factors driving the jump intensity. Similarly, the BSIV of short-maturity at-the-money (ATM) options is mostly determined by spot volatility and, hence, provides more direct evidence on the fine structure of spot volatility.

The issue of microstructure noise in high-frequency volatility indices is also not trivial. [Andersen et al. \(2013\)](#) document problems associated with the construction of the VIX at high frequencies. These are mostly related to the rules for truncating deep OTM options in the computation of the index.² Using high-frequency S&P 500 index options data directly allows us to construct implied volatility measures whose increments are much less sensitive to such features.

In summary, our empirical analysis exploits the following high-frequency series. First, we use short-maturity S&P 500 futures and futures options traded on the Chicago Mercantile Exchange (CME). Using the option prices, we construct one-month BSIV series with fixed log moneyness (strike/futures price) relative to the ATM BSIV. Second, we use data on the S&P 500 index futures and the VIX index. Finally, we use data on the two nearby VIX futures. Combined, our data cover January 2007 till May 2012, but we filter out problematic observations and there is a bit of mismatch in the our high-frequency data across the different types of derivatives contracts.

We apply two very different procedures in our investigation of the fine structure of the asset price and volatility dynamics at small time scales. The first technique relies on the ratio of power variations at two different frequencies, which allows us to estimate the activity index of the given process. Estimators of this type have been studied by [Ait-Sahalia and Jacod \(2010\)](#) and [Todorov and Tauchen \(2010, 2011a\)](#). The activity index for a process is two if it has a diffusion component. On the other hand, if no diffusion term is present, the activity index is smaller and equals the jump activity index for the (alternative) pure-jump model. Our second econometric tool is the empirical cumulative distribution function (cdf) of the nonparametrically de-volatilized high-frequency increments of the process and was developed by [Todorov and Tauchen \(2014\)](#). The de-volatilized high-frequency increments should be approximately Gaussian if the process is jump-diffusive and, alternatively, follow a stable distribution if the process is of pure-jump type.

Our empirical results display an intriguing pattern. The lowest point estimates for the activity index, of around 1.6, are obtained for the BSIV of deep OTM put options. As we move from OTM puts to ATM options, the estimates gradually increase to the maximal value of two, indicative of a jump-diffusive process. Our additional test based on the empirical cdf of the de-volatilized increments corroborates this finding, i.e., it also suggests the BSIV for deep OTM puts are pure-jump processes, while pointing towards the opposite conclusion for the BSIV of near-the-money options. These results are consistent with the risk-neutral intensity of the negative price jumps being driven, solely or predominantly, by state variable(s) of pure-jump type, and the spot volatility being governed by a jump-diffusion. Moreover, the value of volatility related indices and contracts, such as the VIX and VIX futures, are functions of the state variables driving the volatility as well as those determining the jump intensity. Hence, in finite samples, we would expect point estimates for the activity index of such series to be close to, or fall within, the range of the activity estimates obtained across the moneyness dimension of the S&P 500 implied volatility surface. And this is what we find: the estimated activity indices for VIX fall just below the values obtained for the deep OTM IV measures while those for the VIX futures are around the values attained for IV4.

The rest of the paper is organized as follows. Section 2 introduces notation and formally defines the option-based quantities that we study in the paper. In Section 3, we describe the separate datasets used in our analysis and we conduct an initial analysis regarding the liquidity of the individual instruments. Section 4 reviews the econometric tools we use, while Section 5 contains the main empirical results, and Section 6 concludes.

2. Setup

We assume the underlying asset price process, X , is an Itô semimartingale under the statistical measure \mathbb{P} , characterized through the following return dynamics,

$$\frac{dX_t}{X_t} = \alpha_t dt + \sqrt{V_t} dW_t + \int_{\mathbb{R}} x \tilde{\mu}^{\mathbb{P}}(ds, dx), \quad \text{under } \mathbb{P}, \quad (1)$$

¹ Studying the properties of the underlying S&P 500 index, by contrast, is much easier as it is directly observable and high-frequency data are readily available. By now, there is ample evidence that the index contains both a diffusion ([Todorov and Tauchen, 2011b](#)) and jump components ([Ait-Sahalia and Jacod, 2009](#); [Barndorff-Nielsen and Shephard, 2006](#); [Lee and Mykland, 2008](#)). In addition, [Carr and Wu \(2003\)](#) find support for a jump-diffusive characterization via nonparametric analysis of the time decay of short-maturity options.

² Theoretically, the VIX index involves option prices for log-moneyness across the whole real line. In practice, we have a limited set of option prices available and this inevitably induces (time-varying) approximation errors.

where W_t is a Brownian motion and μ is an integer-valued measure with compensator $\nu^\mathbb{P}$ and $\tilde{\mu}^\mathbb{P} = \mu - \nu^\mathbb{P}$ is the associated martingale measure.

Under mild regularity conditions, no arbitrage implies the existence of a locally equivalent probability measure \mathbb{Q} , labeled risk-neutral, under which X evolves as follows,

$$\frac{dX_t}{X_t} = (r_t - \delta_t) dt + \sqrt{V_t} dW_t^\mathbb{Q} + \int_{\mathbb{R}} x \tilde{\mu}^\mathbb{Q}(ds, dx), \quad \text{under } \mathbb{Q}, \quad (2)$$

where r_t is the instantaneous risk-free rate and δ_t the dividend yield; $W_t^\mathbb{Q}$ is a Brownian motion with respect to \mathbb{Q} , $\tilde{\mu}^\mathbb{Q} = \mu - \nu^\mathbb{Q}$, where μ is the integer-valued measure in Eq. (1) and $\nu^\mathbb{Q}$ is a jump compensator of the form $\nu^\mathbb{Q}(dt, dx) = a_t dt \otimes \nu^\mathbb{Q}(dx)$ for some Lévy measure $\nu^\mathbb{Q}(dx)$.

The primary latent state variables in Eq. (2) are the spot volatility and jump intensity. We assume that $V_t = f(\mathbf{S}_t)$ and $a_t = g(\mathbf{S}_t)$ for some $n \times 1$ dimensional stochastic process \mathbf{S}_t governed by a Lévy-driven SDE, while $f(\cdot)$ and $g(\cdot)$ are arbitrary smooth functions: $\mathbb{R}^n \rightarrow \mathbb{R}_+$. Our task is to gain insights into the dynamic features of this latent state vector through high-frequency observations on derivatives that are highly sensitive to variation in volatility and jump intensity.

We denote European-style OTM option prices for the asset X at time t by $O_{t,k,\tau}$. Assuming frictionless trading in the options market, the option prices are given as,

$$O_{t,k,\tau} = \begin{cases} \mathbb{E}_t^\mathbb{Q} \left[e^{-\int_t^{t+\tau} r_s ds} (X_{t+\tau} - K)^+ \right], & \text{if } K > F_{t,t+\tau}, \\ \mathbb{E}_t^\mathbb{Q} \left[e^{-\int_t^{t+\tau} r_s ds} (K - X_{t+\tau})^+ \right], & \text{if } K \leq F_{t,t+\tau}, \end{cases} \quad (3)$$

where τ is time-to-maturity, K is the strike price, $F_{t,t+\tau}$ is the forward price for X at t with respect to date $t + \tau$, and $k = \ln(K/F_{t,t+\tau})$ is log-moneyness. We further denote the BSIV associated with the option price $O_{t,k,\tau}$ (plus risk-free rate r_t and dividend yield δ_t) by $\kappa_{t,k,\tau}$.

The volatility VIX index, computed by the CBOE from short-maturity OTM options on the S&P 500 index, has a theoretical value of,

$$VIX_{t,\tau} = \sqrt{\frac{1}{\tau} \mathbb{E}_t^\mathbb{Q} \left(\int_t^{t+\tau} V_s ds + \int_t^{t+\tau} \int_{\mathbb{R}} (e^x - 1 - x) \mu(ds, dx) \right)}. \quad (4)$$

Finally, in this no-arbitrage setting, the futures price for the one-month VIX index is,

$$VF_{t,\tau} = \mathbb{E}_t^\mathbb{Q} [VIX_{t+\tau,30}]. \quad (5)$$

We assume that the joint conditional distribution of $(\log(X_{t+\tau}/X_t), \int_t^{t+\tau} r_s ds, \int_t^{t+\tau} \delta_s ds)$ is determined uniquely by the state vector \mathbf{S}_t . Provided the log-moneyness k is non-random or, more generally, \mathcal{F}_t -adapted, the above then implies $\frac{e^{r_{t,t+\tau} O_{t,k,\tau}}}{F_{t,t+\tau}}$ is a function only of the tenor, moneyness, and the state vector (and t , if \mathbf{S}_t is non-stationary under \mathbb{Q}), where $r_{t,t+\tau}$ is the risk-free rate for $[t, t + \tau]$. Hence, BSIV, evaluated for a suitably normalized moneyness, is immune to the underlying asset price. That is, given the assumed general structure for the process X , all such option-based measures are (different) functions solely of the state vector \mathbf{S}_t (as well as tenor and moneyness). Thus, while \mathbf{S}_t is not directly observed – it is “hidden” in the characteristics of X – we can analyze the derivatives-based volatility measures by nonparametric means to gain insight into certain features of the state vector dynamics. In particular, we exploit the increments to these measures over short intervals to learn about the small scale, or fine structure, of the volatility and jump intensity processes. We provide much more details regarding this issue later on.

3. Data

We draw on multiple sources for information about high-frequency S&P 500 volatility. If our findings are consistent across different securities, traded on distinct exchanges, and over longer periods of time, they are less likely to be driven by microstructure effects or other idiosyncrasies specific to a given contract, exchange, or sample. We exploit tick data for the CME E-mini S&P 500 futures and options. CBOE provides us with high-frequency data for the VIX and VIX futures. Finally, we rely on Treasury bill rates from the Federal Reserve to proxy for the risk-free rate.

The payoff, and pricing, of the various CME and CBOE derivative contracts are linked to the underlying S&P 500 index. As such, they are closely interconnected. Fig. 1 provides a flowchart depicting the relation among the different products within the two exchanges. The contracts in the oval gray boxes are utilized in this paper. At the top of the diagram we have individual stocks listed on the NYSE or NASDAQ, providing the constituents of the broad S&P equity indices. In the second row, we have the S&P 500 (cash) index on the left, obtained as a value-weighted average of the stock prices. This is the first derivative, or synthetic index, in the chain. There are two branches which originate from the S&P 500 index. On the left branch, we first encounter the CME E-mini S&P 500 futures contract. The E-mini futures price is tied closely to the (non-traded) S&P 500 cash index by strong arbitrage forces. The E-mini futures are routinely viewed as the primary location for price discovery in the U.S. equity market. There is also a very active market for options written on the E-mini S&P 500 futures. We exploit these options to construct BSIV measures across the moneyness spectrum. Since the E-mini futures options represent the third link in the chain, our derived BSIV measures represent the fourth level derivative generated from the underlying portfolio of individual stocks.

On the right branch originating from the S&P 500 index, we find the CBOE SPX options. The CBOE computes the (model-free) VIX index from the cross-section of SPX option prices across the strikes at a fixed 30 day maturity. While the VIX is not directly traded, the VIX futures and options are very liquid. We construct a synthetic 30 day fixed maturity VIX futures price from the two shortest, and most liquid, VIX futures maturities at a given point in time. Our analysis exploits both the VIX index itself and this 30 day VIX futures price. Since the VIX futures represent the fourth layer of derivatives for the S&P 500 stocks, our fixed maturity VIX futures price is also a high level derivative relative to the original stock portfolio.

Finally, for completeness, we display the S&P 100 index and the associated options and old VIX (now labeled VXO) index in the most right branch of the figure.

Fig. 2 plots the S&P 500 index along with the VIX, 30 day VIX futures, and three representative BSIV series. We discuss the construction of the individual series in detail below. The figure is included at this stage to provide a sense of the complementarity and potential incremental information obtained by exploring the full set of measures. As expected, the volatility series are highly correlated, but there are noteworthy differences. For example, the VIX futures series is less erratic than the VIX index, reflecting the longer effective maturity of the futures contracts (the 30 days maturity plus the payoff reflecting the one-month forward-looking feature of the VIX). Likewise, the IV6 series for ATM options is much less volatile than the corresponding IV2 series for deep OTM put options. We demonstrate below that these distinct features are aligned with corresponding discrepancies in the fine structure of the price dynamics for the different measures.

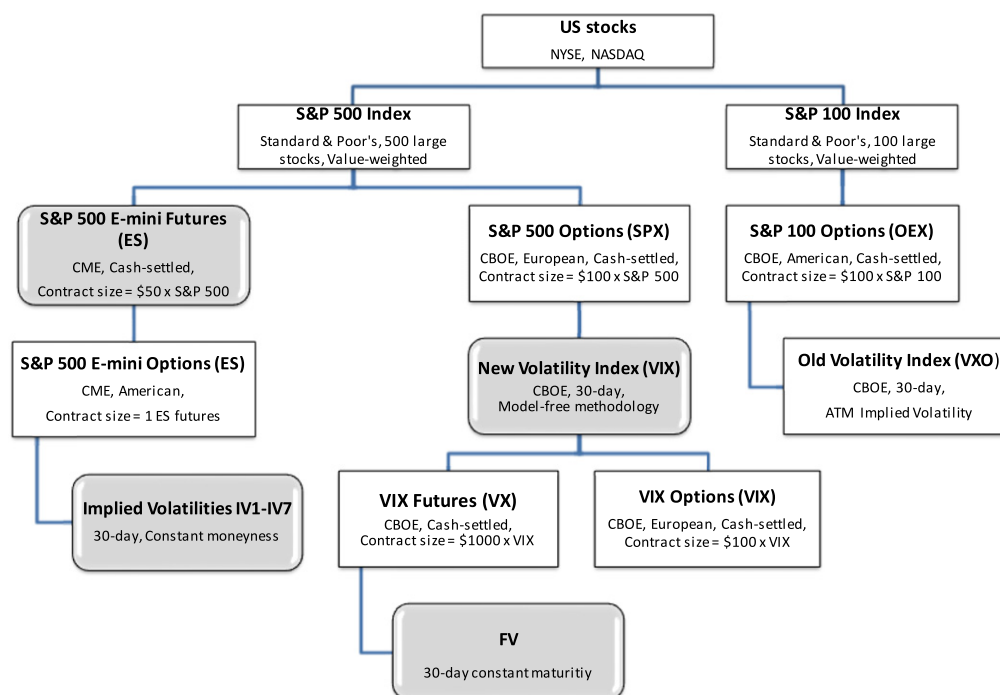


Fig. 1. This flow chart illustrates the relationship between various S&P 500 volatility derivatives. The shaded boxes with rounded corners indicate series used in this study.

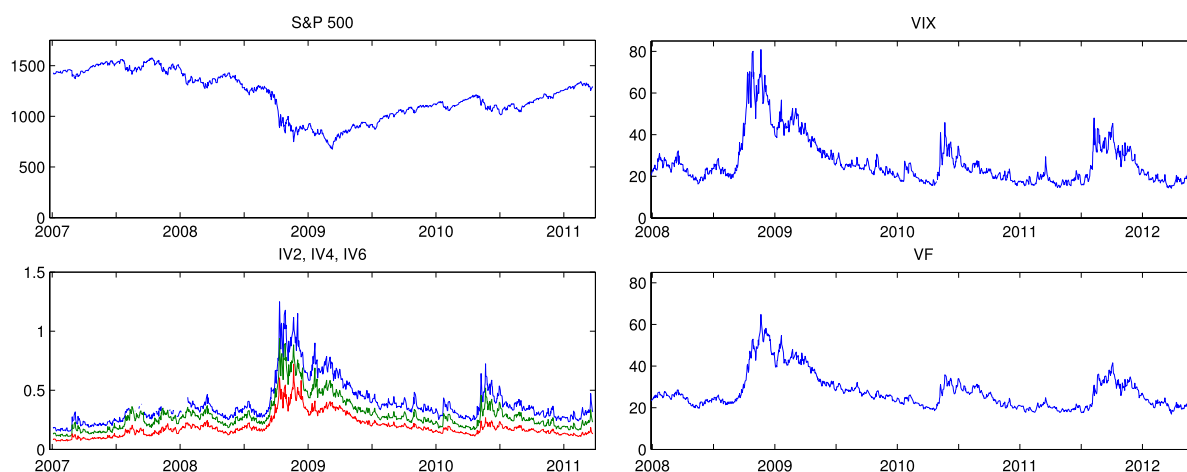


Fig. 2. This figure plots daily values of S&P 500 in the top left panel, the implied volatilities IV2 (blue), IV4 (green), and IV6 (red) in the bottom left panel, the VIX index in the top right panel, and the 30 day VIX futures price, VF, in the bottom right panel. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

3.1. E-mini S&P 500 futures and options

The E-mini S&P 500 futures and options (commodity ticker ES) are traded exclusively on the CME GLOBEX electronic platform. They trade essentially 24 hours a day, five days a week and are among the deepest and most liquid worldwide. Our sample covers January 3, 2007–March 22, 2011, or 1062 trading days. [Appendix A](#) provides additional details about these markets.

We construct 15-second series for the futures and options using the “previous tick” method, i.e., we retain the last quotes prior to the end of each 15 seconds interval. We focus on the regular trading hours from 8:45 to 15:15 CT, yielding 1560 15 seconds intervals per trading day.³ While the E-mini futures market is extremely liquid,

less is known, especially about the deep OTM futures options. To address this issue, we require a consistent notion of moneyness across time. Hence, letting σ_{BS} denote the ATM BSIV and F the forward price at tenor τ , we define the normalized moneyness measure m , where

$$m = \frac{\ln(K/F)}{\sigma_{BS} \sqrt{\tau}}. \quad (6)$$

In [Fig. 3](#), we plot the average values of various option liquidity measures versus m . The top left panel depicts the average dollar option bid–ask spread, while the bottom left panel depicts the normalized mid quote. Evidently, the bid–ask spread grows as the option prices increase, but much less than proportionally. In fact, for the deep OTM options, the spreads flatten so that low priced options carry a non-trivial dollar spread. As a consequence, the upper right panel indicates a dramatic increase in the relative spread as we move away from ATM options. Once we go beyond the moneyness of -5 to the left or 2 to the right, the spread exceeds 70%

³ Since the market occasionally faces irregularities in the quoting or transmission process following the cash market opening, we are conservative and remove the first fifteen minutes of regular trading.

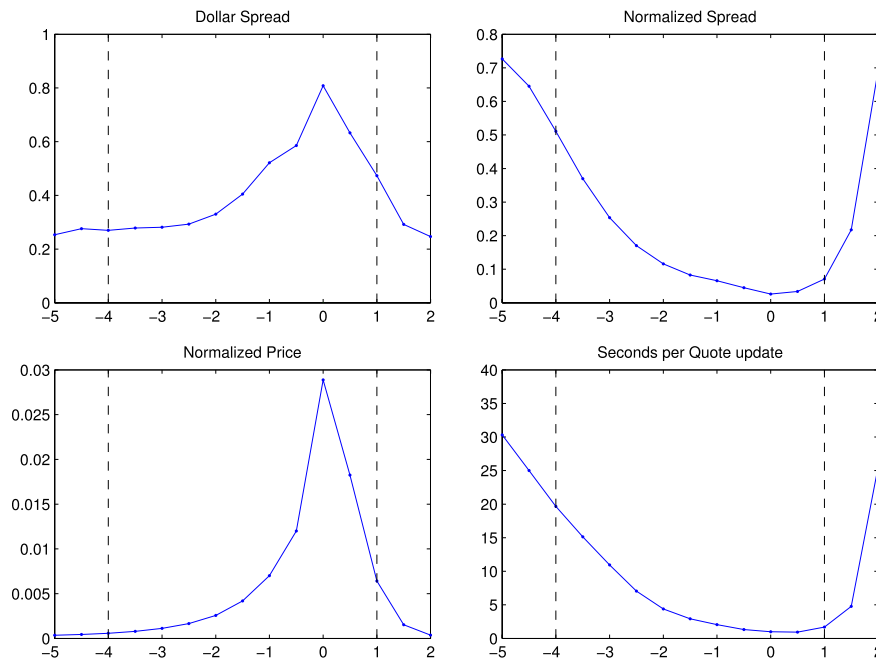


Fig. 3. This figure plots average values of various option liquidity measures for different levels of the normalized moneyness m . For $m < 0$, puts are used; for $m > 0$, calls are used; and for $m = 0$, the average of puts and calls is computed. The dollar spread is $A - B$, where A and B are the dollar ask and bid prices; the normalized spread is $(A - B)/M$, where $M = \frac{1}{2}(A + B)$ is mid quote price; the normalized price is M/F . The last panel shows the average number of seconds per quote update. The dashed vertical lines indicate the range of moneyness m used in our analysis.

of the mid-quote, implying that transactions at the bid or ask occur at prices that, on average, deviate from the quote midpoint by 35% or more. For $m = -4$, the half-spread drops to around 25%, and it is much less elevated at approximately 12% for $m = -3$ and $m = 1.5$. Finally, in the bottom right panel, we note a corresponding pattern for liquidity expressed in terms of the average number of seconds between each quote update.⁴ For $m = -3.5$ or $m = 2$, the quote frequency falls short of the 15 second mark, while for $m \in [-3, 1.5]$ it is about 10 seconds or less, and around $m = 0$ it is about one per second. Finally, we face the issue that the strike range occasionally is quite limited. For example, for $m = -4$, there is no strike available at or below this point for about 16% of our observations, so this region is not covered during parts of our sample. Likewise, the corresponding non-coverage rates for $m = -3$ and $m = -2$ are 7.3% and 1.8%, respectively. For ATM options this problem vanishes, with only a fraction of a percent of observations missing for $m = -1$ or $m = 0.5$, and none for $m = 0$. However, for $m = 2$, non-coverage jumps to fully 47% of the observations, leading to some concern about the integrity of our procedures due to potential endogenous sampling issues for this region. The coverage deteriorates even more for further OTM call options.

In summary, we have timely quotes across a broad range of moneyness for the futures options. However, at the 15-second frequency, the price series are inevitably noisy due to staleness and the bid–ask spread. Thus, we do not use 15-second observations directly, but rely on quote mid-points from a lower frequency such as five minutes. We do this in a couple of ways. The direct approach is to use only every 20th observation in the 15-second series. However, inspired by Podolskij and Vetter (2009), we can exploit additional information and simultaneously alleviate lingering microstructure effects by pre-averaging the 15-second quotes across five-minute blocks and use the resulting series as input to the subsequent analysis. We exploit this approach in the empirical analysis in Section 5. Since our theory is

developed without consideration of microstructure noise, this implementation scheme is attractive. We provide details on this procedure in Appendix A.1.

Given the above evidence, we construct seven separate 30 day BSIV measures IV1–IV7 for each 15-second interval, corresponding to moneyness $m = -4, -3, -2, -1, 0, 0.5$, and 1 .⁵ However, before including a specific observation in our analysis, we impose a number of checks regarding the validity of the underlying quotes. We defer the details of our construction to Appendix A.2. Lastly, when interpreting our results, we discount conclusions that stem from the deepest OTM options, as these are more prone to noise and measurement error.

3.2. VIX and VIX futures

The volatility index VIX is disseminated in real time by the CBOE. It is based on the model-free methodology and equals the square root of the par 30 day variance swap. The index is calculated from option prices across a range of strikes on the S&P 500 index and quoted in terms of the annualized standard deviation. While the VIX index is not directly traded, the VIX futures (ticker symbol VX) are traded on the CBOE Command platform from 7:00 am to 3:15 pm CT every weekday. Appendix B provides additional details.

The VIX futures contract started trading on March 26, 2004. Initially, the volume was low but it has since increased dramatically, and the contract is now considered one of the most successful CBOE product launches of all time. Our tick data for the VIX index and VIX futures were obtained from the CBOE Market Data Express (MDX). Table 1 provides daily summary statistics for the two nearby (most liquid) VIX futures contracts, FV1 and FV2, with tenors spanning 0–30 and 30–60 days, respectively. Between 2006 and 2012, the

⁴ A quote update implies that the current bid and ask are “actionable”.

⁵ The measure for a given m is obtained by linearly interpolating the BSIVs for strikes just below and above m . This is identical to the procedure used by the CBOE in computing the ATM BSIV.

Table 1
Daily statistics for VIX futures.

	First maturity FV1						
	2006	2007	2008	2009	2010	2011	2012
Trading volume, 000 s	0.58	1.98	2.10	1.97	8.17	18.89	27.37
Number of quotes, 000 s	0.66	1.40	6.75	11.15	70.90	217.72	109.13
Normalized spread, %	0.82	0.55	0.42	0.34	0.26	0.24	0.26
	Second maturity FV2						
	2006	2007	2008	2009	2010	2011	2012
Trading volume, 000 s	0.41	0.79	1.13	1.46	5.49	13.20	21.00
Number of quotes, 000 s	0.20	0.72	6.09	9.85	57.69	164.52	94.33
normalized Spread, %	0.83	0.63	0.55	0.36	0.26	0.23	0.24

Note: FV1 and FV2 denote the two nearby maturities. The normalized spread is $(A - B)/M$, where A and B are the ask and bid prices and $M = \frac{1}{2}(A + B)$ is mid-quote price.

volume grew more than 45-fold, the number of quotes increased more than one hundred-fold and the percentage spreads declined to less than a third of the original size. Certainly, by 2008, the liquidity is sufficient to ensure that new quotes arrive almost every 15 second during regular trading hours. Thus, we rely on the period January 2, 2008–May 31, 2012, covering 1112 trading days, and we construct 15-second series using the “previous tick” method from 8:45 to 15:15 CT.⁶ Moreover, we pre-average the 15-second quotes for the VIX and VIX futures within non-overlapping five-minute blocks.

Finally, we convert the VIX futures series into a fixed 30 day maturity series, denoted FV, via a weighted linear interpolation of the prices for the two nearby futures contracts.

4. Econometric tools for the empirical analysis

4.1. The formal framework

This section reviews the various estimators and tests we apply later in our empirical analysis. We denote the generic process under investigation by Z . This will be one of the option based measures or the S&P 500 index futures price. Our econometric analysis proceeds under the assumption that the dynamic evolution of Z may be captured by the following general specification,

$$dZ_t = \alpha_t dt + \sigma_{t-} dS_t + dY_t, \quad 0 \leq t \leq T, \quad (7)$$

where α_t and σ_t are processes with càdlàg paths; Y_t is a process whose second characteristic is identically zero, i.e., it is a process without a continuous local martingale component; S_t is a (strictly) stable process with characteristic function (see, e.g. Sato, 1999),⁷

$$\log[\mathbb{E}(e^{iuS_t})] = -t|cu|^\beta. \quad (8)$$

Note that Y_t need not be independent of S_t (or α_t and σ_t). Hence, Z_t does not necessarily “inherit” the tail properties of the stable process S_t . This implies, for example, that Z_t can be driven by a tempered stable process whose tail behavior is very different from that of the stable process. The model in Eq. (7) covers most, if not all, models routinely used for modeling financial asset prices. Key examples include the affine jump–diffusion model class of Duffie et al. (2000), the time-changed Lévy models of Carr et al. (2003), and the Lévy-driven SDEs, provided the Lévy density is locally stable (which is almost always true for parametric specifications of Lévy densities).

For $\beta = 2$, the S_t process is a Brownian motion, and it constitutes the (first order) leading term driving Z_t in Eq. (7) at high frequencies. For $1 < \beta < 2$, S_t is a pure jump process, but it continues to (first order) dominate the drift term at high frequencies. Throughout, we restrict attention to this empirically relevant case, i.e., $1 < \beta \leq 2$. The self-similarity of the strictly stable process implies the following scaling property,

$$S_t - S_s \stackrel{d}{=} |t - s|^{1/\beta} S_1, \quad \forall 0 \leq s < t. \quad (9)$$

Finally, to retain the interpretation of S_t as the dominant (high frequency) term – and separately identify S_t from Y_t in this regard – we require Y_t to satisfy a scaling bound for high frequencies of the following type. There exists a sequence of stopping times increasing to infinity, T_p , and $\beta' < \beta$, such that for $\forall 0 \leq s < t$, we have,

$$\mathbb{E}|Y_{t \wedge T_p} - Y_{s \wedge T_p}|^q \leq K_p |t - s|, \quad \forall q > \beta', \quad (10)$$

where K_p is some constant depending on the sequence of stopping times.

These conditions imply that Y_t , in Eq. (7), plays the role of a high-frequency “residual” component. Combining the self-scaling result in Eq. (9) and the scaling bound on Y in Eq. (10), we obtain, under suitable regularity, for $s \in [0, 1]$ and $\forall t \geq 0$ with $\sigma_t > 0$,

$$h^{-1/\beta} \frac{Z_{t+sh} - Z_t}{\sigma_t} \xrightarrow{\mathcal{L}} S'_{t+s} - S'_t, \quad \text{as } h \downarrow 0, \quad (11)$$

where S'_t is a Lévy process with a distribution identical to that of S_t , and the convergence is defined in the space of càdlàg functions equipped with the Skorokhod topology.

Intuitively, the result in Eq. (11) implies that each high frequency increment of Z behaves like that of a stable process with a constant scale σ_t known at the onset of the increment. If we directly observe the stochastic process σ_t , we can scale the increments of Z accordingly, and the limit result in (11) states that the scaled increments should be i.i.d. stable (in the limit for ever more frequent sampling). Separating the case $\beta = 2$ from $\beta < 2$ is of central importance as it is equivalent to distinguishing between Z being a jump–diffusion or of the pure jump type.

4.2. Estimation and inference procedures

The goal of the empirical analysis in Section 5 is to test, using high-frequency observations, whether Z is a jump–diffusion or a pure jump process and, furthermore, to estimate the activity index β . This section presents the econometric tools we apply for that task. We assume the process Z is observed at the equidistant grid points $0, \frac{1}{n}, \dots, T$ with a fine mesh, $\frac{1}{n}$, that asymptotically approaches 0, and T is a fixed positive integer. We will use either the increments $\Delta_i^n Z = Z_{\frac{i}{n}} - Z_{\frac{i-1}{n}}$ or their so-called pre-averaged analogue (see (Jacod et al., 2009) and (Jacod et al., 2010)) constructed via additional sampling within each interval. The

⁶ For the VIX series we apply a very mild filter by removing observations that fall outside the daily high–low range, as reported at the end of trading. This correction reflects errors that have been recognized by the exchange.

⁷ This setting can be extended to allow for S_t being an asymmetric stable process. This is accommodated by differencing the increments in the statistics below, as proposed in Todorov (2013).

pre-averaged increments are defined as,

$$\widetilde{\Delta}_i^n Z = \sum_{\ell=1}^L \left\{ \left(Z_{\frac{i-1}{n} + \frac{\ell}{nL}} - Z_{\frac{i-1}{n} + \frac{\ell-1}{nL}} \right) \left(\frac{\ell}{L} \wedge \frac{L-\ell}{L} \right) \right\},$$

$$i = 1, \dots, Tn. \quad (12)$$

Given the observation frequency of our data, we restrict attention to the case of L fixed. In this setting, the asymptotic behavior of our statistics is identical whether $\widetilde{\Delta}_i^n Z$ or $\Delta_i^n Z$ is used.⁸ The purpose of pre-averaging is simply to mitigate the potential impact of moderate microstructure noise on the estimation, as the asymptotic analysis below does not account for observation noise.

We consider two types of tests for the presence of a diffusion component, exploiting different aspects of the limiting result in Eq. (11). The first test is based on direct estimates of the index β for S_t . It is constructed from the ratio of power variations at two distinct frequencies. Formally, letting $1\{A\}$ be one for A true, and zero otherwise, our estimator of β is,

$$\widehat{\beta}(p) = \frac{p \log(2)}{\log \left(V_T^{n,2}(p, Z) \right) - \log \left(V_T^{n,1}(p, Z) \right)} \times 1 \left\{ V_T^{n,2}(p, Z) \neq V_T^{n,1}(p, Z) \right\}, \quad (13)$$

where,

$$V_T^{n,1}(p, Z) = \sum_{i=1}^{nT} \left| \widetilde{\Delta}_i^n Z \right|^p,$$

$$V_T^{n,2}(p, Z) = \sum_{i=2}^{nT} \left| \widetilde{\Delta}_{i-1}^n Z + \widetilde{\Delta}_i^n Z \right|^p. \quad (14)$$

This represents an extension of the activity estimator of [Todorov and Tauchen \(2011a,b\)](#), allowing for overlapping observations in the construction of $V_T^{n,2}(p, Z)$, as in [Todorov \(2013\)](#), to enhance efficiency.⁹ The estimator is consistent for β , when $p < \beta$. Moreover, under suitable conditions on Z and, importantly, $p < \beta/2$, we obtain the following asymptotic limit result,

$$\sqrt{n} \left(\widehat{\beta}(p) - \beta \right) \xrightarrow{\mathcal{L}-s} \frac{\sqrt{\int_0^T |\sigma_s|^{2p} ds}}{\int_0^T |\sigma_s|^p ds} \times \frac{\beta^2}{\mu_p(\beta) p \log(2)} \sqrt{\widetilde{\mathcal{E}}} \times \mathcal{N}. \quad (15)$$

\mathcal{N} denotes a standard normal random variate which is defined on an extension of, and is independent from, the original probability space. Meanwhile, $\widetilde{\mathcal{E}}$ is given by,

$$\widetilde{\mathcal{E}} = \mathcal{E}^{(1,1)} - 2^{1-p/\beta} \mathcal{E}^{(1,2)} + 2^{-2p/\beta} \mathcal{E}^{(2,2)},$$

where $\mathcal{E} = \Sigma_0(p, \beta) + \Sigma_1(p, \beta) + \Sigma_1'(p, \beta)$ and $\Sigma_j(p, \beta) = \mathbb{E}(\mathbf{Z}_j \mathbf{Z}_{1+j}') for $j = 0, 1$ with $\mathbf{Z}_i = (|S_i|^p - \mu_p(\beta), |S_i + S_{i+1}|^p - 2^{p/\beta} \mu_p(\beta))'$, $\mu_p(\beta) = \mathbb{E}|S_1|^p$, and S_0, S_1, \dots are i.i.d. β -stable random variables, defined through the characteristic function in Eq. (8) for $t = 1$. Consistent estimators for the integrated power variation terms that appear in the limiting distribution are readily obtained from corresponding realized power variation statistics; we refer to [Todorov and Tauchen \(2011a\)](#) for details on this as well as an enumeration of the requisite regularity conditions.$

⁸ This is because both statistics introduced below are scale-free and a weighted sum of i.i.d. stable random variables continues to be a stable random variable, albeit with a different scale.

⁹ See [Ait-Sahalia and Jacod \(2010\)](#) for an estimator based on truncated power variation.

Our second test is based on the distributional implications of the small scale result in Eq. (11) and follows from the developments in [Todorov and Tauchen \(2014\)](#). In particular, we split the data into blocks of asymptotically shrinking length and form local volatility estimates for each block based only on observations from within the block. Next, we use these estimates to rescale (de-volatilize) the high-frequency increments within the block, in order to “annihilate” the stochastic volatility. Finally, we estimate the empirical cumulative distribution function (cdf) of the de-volatilized high-frequency increments. The test is then based on the distance between this empirical cdf and that of the standard normal, which constitutes the limiting distribution in Eq. (11) under the null of Z being a jump-diffusive process over the given time interval.

To formally define the procedure, we require some notation. Each block contains k_n high-frequency increments, each of length $1/n$, with $k_n \rightarrow \infty$ and $k_n/\sqrt{n} \rightarrow 0$. There are a total of $J_n = \lfloor Tn/k_n \rfloor$ separate blocks, $j = 1, \dots, J_n$. The set of increments within block j is $\mathcal{I}_j = \{(j-1)k_n + 1, \dots, jk_n\}$, while the same index set excluding element i is denoted $\mathcal{I}_j(i) = \mathcal{I}_j \setminus \{i\}$. We exploit the so-called Truncated Variation statistic of [Mancini \(2009\)](#) for estimating the quadratic variation of the diffusion component of Z . For each increment we define the following (noisy) volatility estimator, scaled to express the return variation per unit time interval,

$$\widehat{c}_i^n = n \cdot |\widetilde{\Delta}_i^n Z|^2 \cdot 1 \left(|\widetilde{\Delta}_i^n Z| \leq \alpha n^{-\varpi} \right),$$

$$i = 1, \dots, Tn, \alpha > 0, \varpi \in (0, 1/2). \quad (16)$$

Our local volatility estimator for block j improves precision by averaging over the individual contributions and, again, scaling to reflect the return variation per time unit,

$$\widehat{C}_j^n = \frac{1}{k_n} \sum_{i \in \mathcal{I}_j} \widehat{c}_i^n, \quad j = 1, \dots, J_n. \quad (17)$$

However, feasible inference requires we use a smaller number of increments within the block, $m_n \ll k_n$, in the construction of the empirical cdf. Furthermore, as we scale each high-frequency increment during the computation of the empirical cdf, we must ensure that the effect of that specific increment is asymptotically vanishing. Thus, we introduce an alternative estimator for volatility (per unit time) over block j that excludes the impact from increment i ,

$$\widehat{C}_{j(i)}^n = \frac{1}{k_n - 1} \sum_{i \in \mathcal{I}_j(i)} \widehat{c}_i^n, \quad j = 1, \dots, J_n. \quad (18)$$

The block volatility estimator in Eq. (17) is then modified as follows,

$$\begin{aligned} \widehat{C}_j^n(i) &= \widehat{C}_j^n + \frac{\eta_n}{k_n - 1} (\widehat{C}_j^n - \widehat{c}_i^n) \\ &= \widehat{C}_{j(i)}^n + \frac{1 - \eta_n}{k_n} (\widehat{c}_i^n - \widehat{C}_{j(i)}^n), \end{aligned} \quad (19)$$

where $\eta_n = 1 - \frac{m_n}{k_n}$. Since $0 < \eta_n < 1$ and $\eta_n \rightarrow 1$, the middle expression in Eq. (19) indicates the dampening of the contribution from the i th increment, while the expression to the right shows how the impact of \widehat{c}_i^n is eliminated asymptotically, as all weight shifts to $\widehat{C}_{j(i)}^n$ in the limit.

Normalizing the increments by the modified local volatility estimator and exploiting $m_n (< k_n)$ terms in each block, the empirical cdf of the de-volatilized high-frequency increments becomes,

$$\begin{aligned} \widehat{F}_n(\tau) &= \frac{1}{N^n(\alpha, \varpi)} \sum_{j=1}^{J_n} \sum_{i=(j-1)k_n+1}^{(j-1)k_n+m_n} 1 \left\{ \frac{\sqrt{n} \widetilde{\Delta}_i^n Z}{\sqrt{\widehat{C}_j^n(i)}} \leq \tau \right\} \\ &\times 1 \left\{ |\widetilde{\Delta}_i^n Z| \leq \alpha n^{-\varpi} \right\}, \end{aligned} \quad (20)$$

where, for the identical $\alpha > 0$ and $\varpi \in (0, 1/2)$ as in (17), we have defined,

$$N^n(\alpha, \varpi) = \sum_{j=1}^{\lfloor n/k_n \rfloor} \sum_{i=(j-1)k_n+1}^{(j-1)k_n+m_n} 1(|\widetilde{\Delta}_i^n Z| \leq \alpha n^{-\varpi}), \quad (21)$$

so $N^n(\alpha, \varpi)$ denotes the total number of increments used in the formation of the empirical cdf.

Under appropriate regularity, Todorov and Tauchen (2014) show that if Z is a jump–diffusion, i.e., S_t in Eq. (7) is a Brownian motion, we have,

$$\widehat{F}_n(\tau) \xrightarrow{\mathbb{P}} F(\tau), \quad (22)$$

where $F(\tau)$ is the cdf of a standard normal random variate and the convergence is uniform in $\tau \in \mathbb{R}$. Again, there is a CLT associated with this convergence in probability result, enabling us to construct a distribution-based test for whether Z is a jump–diffusion via the Kolmogorov–Smirnov distance between the empirical cdf $\widehat{F}_n(\tau)$ and the cdf $F(\tau)$ of a standard normal variate. To conserve space, we refer to Todorov and Tauchen (2014) for details.

Finally, our choice of tuning parameters for constructing $\widehat{F}_n(\tau)$ in Eq. (22) reflects standard practice as well as the underlying observation frequency. Letting one trading day represent our unit interval, we have $n = 78$ five-minute increments for most of our series. Moreover, we set $k_n = 26$, so that we have three blocks per day, and $m_n = 20$, i.e., our (finite sample) implementation exploits about 75% of the increments in constructing the empirical cdf. Finally, following prior work, we set the truncation parameters to $\varpi = 0.49$ and $\alpha = 3\sqrt{\widehat{B}_j^n}$, where,

$$\widehat{B}_j^n = \frac{\pi}{2} \frac{n}{k_n - 1} \sum_{i=(j-1)k_n+2}^{jk_n} |\widetilde{\Delta}_{i-1}^n Z| |\widetilde{\Delta}_i^n Z|, \quad (23)$$

$$j = 1, \dots, \lfloor n/k_n \rfloor,$$

is the Bipower Variation of Barndorff-Nielsen and Shephard (2006). Under our null that Z is a jump–diffusion, it consistently estimates the diffusive quadratic variation. That is, we apply a time-varying threshold to ensure a better finite sample separation of the “big” jumps from the purely continuous increments. Lastly, to account for the well-known diurnal pattern in volatility, the raw high-frequency returns are standardized by a time-of-day scaling factor prior to analysis.¹⁰

5. Empirical findings

This section implements the tests introduced in the previous section for the alternative data series introduced in Section 3. Before progressing to the empirical results, we briefly discuss how the activity indices, i.e., the “fine structure” of the return dynamics, should be interconnected, in theory, for the different derivatives instruments and the underlying asset.

5.1. Linkages in the fine structure of returns across assets

For clarity and reference, we initially assume that the latent state vector, S_t , governing the volatility and (risk-neutral) jump intensity dynamics of the underlying index, X_t , introduced in Section 2, is a univariate Markov process. In this case, we may readily show that $VIX_{t,\tau}$ is a (smooth) function of the latent state

variable at time t , S_t , as well t and τ . Similar reasoning implies that $VF_{t,\tau}$ are smooth functions of S_t and the tenor of the contracts. In addition, the Markovian assumption implies that the ratio $\frac{e^{f_{t,t+\tau} O_{t,k,\tau}}}{F_{t,t+\tau}}$ is a function only of S_t (and the specific characteristics of the option contract), as long as the measure of log-moneyness, k , is either a fixed constant or a function solely of the state vector, i.e., no extraneous variables enter the expression. Moreover, these properties carry over to the BSIV corresponding to $O_{t,k,\tau}$, as the latter also is a smooth transform of the ratio $\frac{e^{f_{t,t+\tau} O_{t,k,\tau}}}{F_{t,t+\tau}}$. Hence, the implied volatility measures, constructed from OTM S&P 500 options in Section 3, henceforth abbreviated IV1–IV7, are functions of S_t alone as well.

In summary, in the above scenario, all option-based derivative prices are related to the state variable, S_t , through a smooth functional. It follows, via Itô calculus, that the local behavior of these series “inherits” the qualitative features of the underlying process, i.e., they share the “fine structure” of the increments which, in turn, is determined by the activity index. In other words, they all contain a diffusive component or they are all pure (infinite activity) jump processes. Likewise, the jump activity of all our option-based quantities equals that of S_t . We stress, however, that the actual jump sizes will differ as the valuation of the various derivative securities arise from very different (smooth) transformations of S_t . What (typically) coincides is the timing of jumps across the series. In conclusion, under the univariate Markovian assumption, the ability to identify and estimate the activity index is greatly enhanced by exploiting information across the series.

Matters are less straightforward when the state vector, S_t , is multivariate. While the option-based derivatives prices continue to be smooth functions of the vector S_t , the features dominating the local behavior of the derivatives price process now depend critically on the relative sensitivity to distinct components of S_t . For example, recent studies find that the equity-index option surface dynamics is governed by different factors, with some including a diffusive term and others being governed by pure jump processes.¹¹ In this case, even if the stochastic volatility is a jump–diffusion, the price path for certain derivative securities will be dominated by the exposure to the pure jump component of the latent state vector S_t . Under such circumstances, the activity index for the derivative security is determined by the largest such index across the subcomponents of S_t for which it has a meaningful exposure. In short, the fine structure of the processes governing the option-based derivative prices and jump activity indices will hinge on their specific type of exposure to the underlying asset dynamics. They will only share the identical activity index if they all remain sensitive to the full set of elements in the state vector S_t . Since derivatives contracts, deliberately, are designed to offer unique exposures to, and thus span, distinct features of the risks driving the underlying asset, it is plausible we will encounter systematic differences in this regard across alternative derivatives with diverging exposures to underlying asset risks.

5.2. Activity estimates for our equity-index related securities

Our estimators and test statistics are constructed from increments obtained by pre-averaging the series over five-minute blocks with an underlying frequency of 15 seconds, as described in Section 3. The lower frequency increments are then obtained by simply adding two consecutive increments based on the five-minute blocks. Table 2 reports estimates of the activity index, along

¹⁰ The scaling factor is computed by dividing the high-frequency increments by a simple sample average of the volatility over the corresponding time of day in the sample.

¹¹ See, e.g., Andersen et al. (forthcoming) for empirical support for this type of model.




Fig. 4. Scatter plots for $\hat{\beta}(0.7)$. RV stands for the realized volatility, based on 5-minute pre-averaged returns, over the period over which $\hat{\beta}(p)$ is computed. To account for the effect of pre-averaging, we rescaled RV by $\sum_{l=1}^L (\frac{l}{L} \wedge \frac{L-l}{L})^2$.

Table 2
Activity index estimates and diffusion tests for the implied volatility measures.

	p	$\text{med}(\hat{\beta}(p))$	MAD	5%	1%
IV1	0.50	1.58	0.07	92.11	78.95
	0.70	1.64	0.07	84.21	57.89
	0.90	1.69	0.08	57.89	26.32
IV2	0.50	1.57	0.12	90.70	79.07
	0.70	1.60	0.11	81.40	62.79
	0.90	1.64	0.12	67.44	58.14
IV3	0.50	1.60	0.07	89.13	84.78
	0.70	1.63	0.06	89.13	71.74
	0.90	1.65	0.06	84.78	54.35
IV4	0.50	1.69	0.09	68.09	48.94
	0.70	1.72	0.08	63.83	40.43
	0.90	1.75	0.08	57.45	25.53
IV5	0.50	1.74	0.08	65.96	38.30
	0.70	1.75	0.08	63.83	38.30
	0.90	1.75	0.08	59.57	27.66
IV6	0.50	1.87	0.11	27.66	10.64
	0.70	1.87	0.13	27.66	6.38
	0.90	1.87	0.13	27.66	6.38
IV7	0.50	2.02	0.14	0.00	0.00
	0.70	2.02	0.12	0.00	0.00
	0.90	2.03	0.11	0.00	0.00
F	0.50	1.97	0.07	2.08	0.00
	0.70	2.00	0.07	0.00	0.00
	0.90	2.02	0.06	0.00	0.00
VIX	0.50	1.49	0.13	96.00	88.00
	0.70	1.53	0.12	94.00	84.00
	0.90	1.58	0.10	86.00	72.00
VF	0.50	1.62	0.11	72.00	62.00
	0.70	1.71	0.14	62.00	40.00
	0.90	1.79	0.12	42.00	14.00

Note: The estimator, $\hat{\beta}(p)$, is defined in Eq. (13). The rejection rates in the last two columns are based on a one-sided test for $\beta = 2$, as well as the point estimate $\hat{\beta}(p)$ and the corresponding estimate for the asymptotic variance provided in Eq. (15). The estimation and tests are performed over periods of 22 trading days.

with a one-sided test for the null hypothesis $\beta = 2$, i.e., a diffusion component is present in the return dynamics, against the alternative that the process is of the pure jump type. The point estimates for the various securities are obtained via Eq. (13), exploiting the ratio of the p th order power variations for a few different values of p . We obtain separate estimates across the sample for 22 non-overlapping trading day periods, i.e., months. The tables summarize the point estimates through robust statistics, indicating the median monthly estimate, $\text{med}(\hat{\beta}(p))$, as well as the median deviation from this median value, MAD. Likewise, we provide the rejection frequencies for the one-sided test of $\beta = 2$ at the five and one percent levels.

In Table 2 there is a sharp contrast between our findings for the equity (futures) returns and the VIX index. The point estimates

suggest that the S&P 500 index is a jump–diffusion while the VIX measure is best characterized as a pure jump process with jumps of infinite variation and an activity index in the range 1.5–1.6. Moreover, for the VIX series, the rejection rates obtained using $p = 0.90$ are consistent with those we expect if the null hypothesis is false, but the driving jump process is highly active with an index between 1 and 2; see the simulation evidence in Todorov and Tauchen (2010). We also note that the estimates for the S&P 500 and VIX indices are similar to those reported in Todorov and Tauchen (2011b) although the sample periods, 2007–2012 versus 2003–2008, are quite different, thus providing a useful robustness check.¹² Turning to the activity estimates for our fixed 30 day maturity VIX futures, we find the evidence supportive of our findings regarding the VIX index, as it also seems to be of the pure-jump type. Nonetheless, the estimates for the activity index are now somewhat larger, falling in the 1.62–1.79 range.

Moving on to the results for the implied volatility measures based on S&P 500 options in Table 2, we observe an interesting pattern. The implied volatilities corresponding to the OTM put options all have activity index estimates close to those for the VIX index, although the point estimates increase as we move from the far OTM towards the ATM measures. On the other hand, the implied volatilities for normalized log-moneyness $m = 1$ yield estimates slightly in excess of 2, and the test is consistent with the presence of a diffusion term for this series. Obviously, this runs counter to the estimates for many of our other option-based derivatives contracts.¹³ We defer the discussion of these features to Section 5.3, after we have presented the full range of evidence.

Finally, we note that our estimator for the activity index in Eq. (13) depends solely on the ratio of power variation measures obtained at two distinct sampling frequencies. As such, the performance of the estimator hinges on our ability to estimate this ratio with reasonable precision. While we do not observe actual realizations of the power variation statistics, we may check some of the implied properties of our procedures. Specifically, the fine structure of the increments for the various processes should be invariant to the level of volatility. On the other hand, if microstructure effects or other issues associated with our empirical power variation measures induce distortions, we may observe systematic biases related to the level of volatility and/or noise. Thus, in Fig. 4, we plot our activity estimates $\hat{\beta}(0.7)$ against the corresponding (nonparametric) pre-averaged realized

¹² Another distinction between the current activity estimates and Todorov and Tauchen (2011b) is that we now use overlapping observations in the construction of $V_p^{n,2}(p, Z)$ in (14), enhancing efficiency.

¹³ The estimated activity index may exceed 2 due to either regular sampling error or the presence of microstructure noise. In fact, the latter causes our activity estimator (for L fixed) to diverge as we approach continuous sampling.




Fig. 5. Scatter plots for $\hat{\beta}(0.7)$. RV_1 and RV_2 stand for the realized volatility, based on 5-minute pre-averaged and raw returns respectively, over the period over which $\hat{\beta}(p)$ is computed. Scale adjustment for pre-averaged RV is done as for Fig. 4.

volatility measure, $\log(RV)$, for each month in the sample. Fig. 5 plots $\hat{\beta}(0.7)$ against the logarithm of the ratio of the pre-averaged and raw realized volatility, which is a measure of the noise. Figs. 4 and 5 display only the estimates for the S&P 500 futures and the IV3 series, as they are representative of measures generating starkly different implications for the activity index. In both cases, the plots reveal no significant dependence of the activity estimates on the level of RV or noise, thus confirming that our empirical activity measures are largely invariant with respect to the fluctuations in volatility and noise. The corresponding plots for the remaining series (not displayed) produce identical conclusions.

As a final robustness check of the activity estimation results, we present estimates of $\hat{\beta}(p)$ based on 10 min pre-averaged returns in Table 4 of Appendix C. Comparing Tables 2 and 4, we see that the estimates are relatively close to each other, with perhaps the exception being the results for the VIX futures which increase moderately at the coarser frequency. Thus, noise related issues might affect the empirical evidence for this particular series to some degree.

We next present results from the alternative distribution-based test for a jump–diffusion based on the empirical cdf of the de-volatilized high-frequency increments given in Eq. (20). As discussed in the previous section, this test is based on “distributional implications” stemming from the presence of a diffusion term in the high-frequency increments. As such, it serves as an additional robustness check for our findings above. This is useful as the simulation evidence in Todorov and Tauchen (2014) suggests that the activity estimator, unlike the distribution-based test, can be somewhat sensitive to microstructure noise and thus, erroneously, may signal the existence of a diffusive term in the price process. The findings from our distribution-based Kolmogorov–Smirnov test are reported in Table 3. To highlight potential deviations between the empirical distribution of the de-volatilized increments and the standard normal, we split the domain of the test into four regions, reflecting the quantiles of the standard normal distribution. A key advantage of the test is that the limiting distribution of the difference $\hat{F}_n(\tau) - F(\tau)$ is independent of volatility, σ_t , and hence the critical values for the test are identical across assets and time periods. These critical values are reported in the bottom row of the table.

The results in Table 3 are largely in agreement with the point estimates for the activity index reported in Table 2. In particular, the test statistics for the S&P 500 futures and the implied volatility series IV7 are below the critical values. Hence, this test also is consistent with the presence of a diffusion term in these two series. Moreover, the Kolmogorov–Smirnov test cannot reject a null hypothesis that IV5 and IV6 contain a diffusion, while the tests based on the activity index provide small, yet non-trivial, rejection rates for these series. For all remaining series, however, the test statistics uniformly exceed the critical values, thus rejecting the null of a jump–diffusion. Another interesting pattern is that the rejections

Table 3
Kolmogorov–Smirnov tests for local Gaussianity.

Series	Range of test				Sample Size (in days)
	$Q_{0.01}-Q_{0.2}$	$Q_{0.2}-Q_{0.4}$	$Q_{0.6}-Q_{0.8}$	$Q_{0.8}-Q_{0.99}$	
IV1	6.18	10.30	10.41	5.99	842
IV2	5.13	8.63	9.05	5.26	953
IV3	3.90	6.39	6.76	3.90	1021
IV4	2.21	4.54	4.52	1.96	1046
IV5	1.71	1.24	1.73	1.67	1047
IV6	1.26	1.05	1.22	1.43	1047
IV7	1.24	1.04	1.04	1.16	1046
F	1.21	0.47	0.45	1.41	1062
VIX	4.60	7.94	7.42	5.04	1112
VF	2.94	8.59	7.96	2.64	1112
Critical value of test					
$\alpha = 5\%$	2.52	1.99	1.99	2.52	
$\alpha = 1\%$	2.73	2.29	2.29	2.73	

Note: The Kolmogorov–Smirnov test is based on comparing $\hat{F}_n(\tau)$ in (20) with $F(\tau)$ using the asymptotic distribution of $\hat{F}_n(\tau) - F(\tau)$ derived in Todorov and Tauchen (2014). Q_p denotes the p th quantile of the standard normal distribution and the range of the test indicates the range of values of τ over which the supremum of the difference $\hat{F}_n(\tau) - F(\tau)$ is taken. The critical values are for $T = 4 \times 252$ days.

of the null of a jump–diffusion model invariably are most significant for the de-volatilized increments falling within the 20–40 and 60–80 quantiles of the normal distribution. To visualize this observation, Fig. 6 provides kernel estimates for the log-density of the de-volatilized high-frequency increments against the log-density of the standard normal for our representative series, the S&P 500 futures and IV3. Consistent with the value of the Kolmogorov–Smirnov test, as well as the activity estimates in Table 2, the empirical log-density for the adjusted intraday S&P 500 returns approximates the log-density of the standard normal closely. In contrast, for IV3, the empirical log-density exceeds that of the standard normal significantly in the center of the distribution as well as in the tails beyond -2 and 2 , while in the intermediate ranges of $[-2, -0.5]$ and $[0.5, 2]$, the relative size of the two log-densities is reversed. This highlights the stark deviations from Gaussianity in the middle of the standardized distribution for the IV3 increments. Qualitatively similar displays (not provided) arise for the other series for which the Kolmogorov–Smirnov test signals an absence of a diffusive term.

Can we rationalize the rather striking evidence in Fig. 6? In fact, these findings are exactly what we should expect to observe for a very closely related scenario. Todorov and Tauchen (2014) document that if the underlying process Z is either of pure-jump type or a jump–diffusion then, regardless of the value of β , $\hat{F}_n(\tau)$ converges in probability to the cdf of the β -stable random variate with first absolute moment equal to $\sqrt{2/\pi}$, as long as we normalize the high-frequency increments by the bipower variation statistic, $\hat{B}_j^n(i)$, rather than the truncated variation $\hat{C}_j^n(i)$. However, in practice, the choice among these alternative techniques is immaterial,




Fig. 6. Log-density of de-volatilized returns. Solid lines on the plots are kernel-based (normal kernel) log-density estimates of the de-volatilized high-frequency increments $\frac{\sqrt{n}\Delta_i^n Z}{\sqrt{\hat{C}_j^n(i)}}$, for which $|\Delta_i^n Z| \leq \alpha n^{-\varpi}$, and the dashed lines correspond to the log-density of the standard normal distribution. Left panel: S&P 500 Index, Right panel: IV3.




Fig. 7. Log-density of stable processes. Solid line: log-density of 1.7-stable process with first absolute moment equal to $\sqrt{2/\pi}$; Dashed line: log-density of standard normal.

and normalization via the bipower statistic produces a plot very similar to Fig. 6. Thus, on Fig. 7, we compare the log-densities for the standard normal ($\beta = 2$) and the 1.7-stable distribution, but both scaled to have a first absolute moment of $\sqrt{2/\pi}$.¹⁴ It is evident that Fig. 7 portrays a qualitatively similar pattern to that observed for the log-density of the de-volatilized increments of IV3. Intuitively, since a stable distribution with index less than two is fat-tailed, it can only generate the same first absolute moment as the normal if the density around zero also surpasses that of the normal, inducing a typical leptokurtic shape on the density for the scaled non-Gaussian stable increments.

In summary, we find all our series – apart from the S&P 500 futures, IV7 and, partially, IV5 and IV6 – to display properties that are consistent with them being of the pure jump type. We conclude this section with a few informal robustness checks regarding the finite sample behavior of the tests. In particular, the key idea behind our test based on the empirical cdf $\hat{F}_n(\tau)$ is that normalization by the nonparametric volatility estimators $\hat{C}_j^n(i)$ or $\hat{B}_j^n(i)$ annihilate the effect of time-varying volatility which, otherwise, provides an alternative explanation for the presence of fat-tails in the (unconditional) return distribution. In essence, the test based on $\hat{F}_n(\tau)$ seeks to separate this source of fat-tailedness from the impact of jumps which are the focus in the present work. Of course, volatility is known to be a highly persistent process. Hence, we can gauge whether we have been successful in “removing” the

volatility from the high-frequency returns by checking for significant autocorrelation in the squared de-volatilized high-frequency

increments, $\left(\frac{\sqrt{n}\Delta_i^n Z}{\sqrt{\hat{C}_j^n(i)}}\right)^2$. Fig. 8 plots the serial correlation in the

squared raw and de-volatilized high-frequency increments. It is evident that our scaling by the nonparametric volatility estimates $\hat{C}_j^n(i)$ effectively annihilates the pronounced persistence in the second-moment of the raw returns. To further illustrate this, Fig. 9 depicts the time series of raw and de-volatilized high-frequency increments. The figure indicates that the latter, unlike the former, appear consistent with the realization of an i.i.d. series, further corroborating the conclusion that the finite-sample performance of the Kolmogorov–Smirnov test based on $\hat{F}_n(\tau)$ is satisfactory in the current context.

5.3. Interpreting the results

As discussed in Section 5.2, if the stochastic volatility V_t is Markov relative to its natural filtration and the jump intensity is governed solely by volatility, then all our volatility-based measures must be of the same type as V_t , i.e., a jump–diffusion or of pure-jump type. Moreover, the jump activity should coincide for all the measures. However, the evidence in Tables 2 and 3 is at odds with this prediction. While most of the measures appear to be of pure jump type, the IV7 series along with the S&P 500 futures is best described via a jump–diffusive model. Consequently, we need a multivariate (latent) state vector for the volatility and jump intensity of the underlying stock price X to potentially explain the empirical findings. Specifically, we may have different state variables driving different components of the volatility-based derivative measures. This is feasible if the implied volatility measures differ in their exposures to the jump and volatility risks in the underlying stock price process, X .

For example, short-maturity deep OTM put options load mostly on the negative jump intensity and have hardly any exposure to the diffusive spot volatility, and analogous results hold for short-maturity deep OTM call options. More formally, from the limiting arguments of Bollerslev and Todorov (2011) and Carr and Wu (2003), we have

$$\frac{O_{t,k,\tau}}{\tau F_{t,t+\tau}} \rightarrow a_t \begin{cases} \int_{\mathbb{R}} (e^x - e^k)^+ \nu^Q(dx), & \text{if } k > 0 \\ \int_{\mathbb{R}} (e^k - e^x)^+ \nu^Q(dx), & \text{if } k < 0 \end{cases}, \quad \text{as } \tau \downarrow 0 \text{ and } \forall t \geq 0. \quad (24)$$

Moreover, simulation evidence in Bollerslev and Todorov (2011) suggests that for the range of τ used here (1-month options), this approximation works relatively well.

¹⁴ $\beta = 1.7$ is roughly compatible with our point estimates for IV3 and IV4.




Fig. 8. Serial correlation of squared increments. Solid lines on the plot correspond to the de-volatilized high-frequency increments $\frac{\sqrt{n}\Delta_t^n Z}{\sqrt{\hat{c}_t^n(0)}}$, for which $|\Delta_t^n Z| \leq \alpha n^{-\varpi}$, and the dashed lines to the raw high-frequency increments $\Delta_t^n Z$.

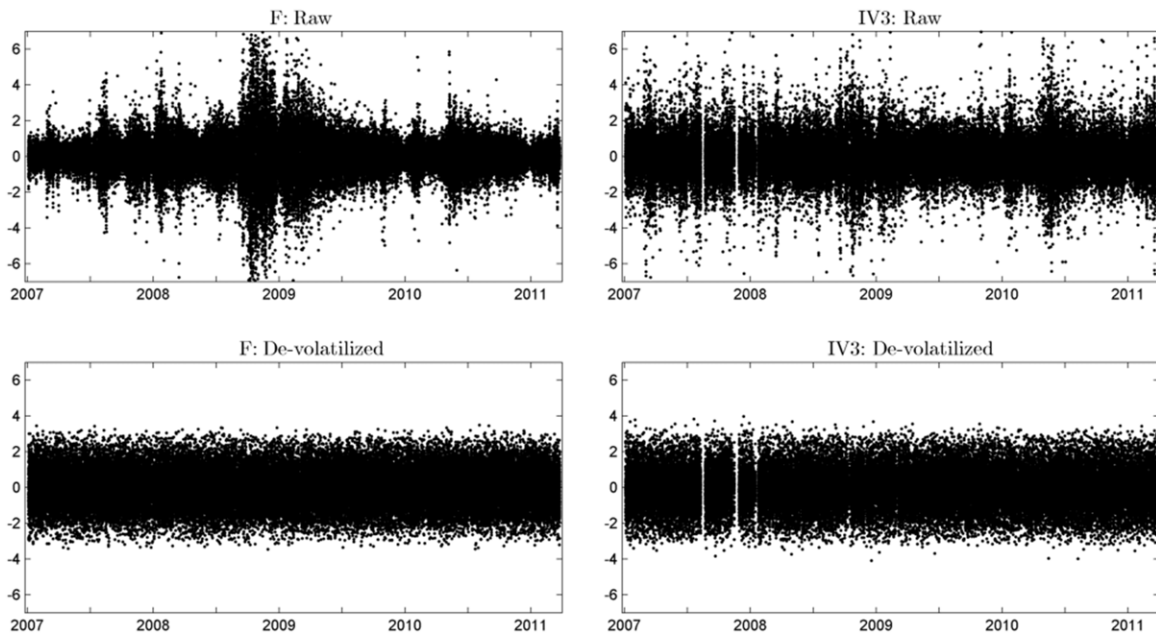


Fig. 9. Time series of high frequency increments. Top panels: raw high-frequency increments $\Delta_t^n Z$, standardized to have unit sample variance, Bottom panels: the de-volatilized high-frequency increments $\frac{\sqrt{n}\Delta_t^n Z}{\sqrt{\hat{c}_t^n(0)}}$, for which $|\Delta_t^n Z| \leq \alpha n^{-\varpi}$. Left panels: S&P 500 Index, Right panels: IV3.

At the same time, the BSIV of the close-to-maturity ATM options are largely determined by the current spot value of the stochastic volatility. More formally, it can be shown, see, e.g., Muhle-Karbe and Nutz (2011),

$$\kappa_{t,0,\tau}^{ATM} \longrightarrow \sqrt{V_t}, \quad \text{as } \tau \downarrow 0 \text{ and } \forall t \geq 0. \quad (25)$$

Thus, the deep OTM measures IV1 and IV2 load predominantly on the left jump intensity, while IV5, IV6 and IV7 load heavily on spot volatility and only marginally on the left jump intensity. For the remaining measures, IV3 and IV4, we expect a relative loading on the volatility and jump intensity factors to fall between that for the deep OTM and the ATM implied option measures. As such, their fine structure, as characterized through the presence or absence of a diffusive term, will be determined by the empirically dominant high-frequency component.¹⁵

¹⁵ Since our log-moneyness measure is scaled by ATM BSIV, the expression on the right-hand side of (24) will be affected by the latter. As a robustness check, we also estimated the activity indices for fixed (unscaled) log-moneyness of $-0.2, -0.15, -0.10, -0.05$. As for the IV1–IV4 series used here, we found strong evidence against the presence of a diffusion driving their dynamics.

For illustration, consider the case where the observed process is of the form $\alpha W_t + (1 - \alpha)S_t$, where W_t is a Brownian motion and S_t is a β -stable process with $\beta < 2$ and independent from W_t , so that $\alpha \in [0, 1]$ captures the relative weight of the two processes. In this scenario, for $\alpha = 0$, the observed process is of the pure jump type and its activity index coincides with the parameter of the stable process β . In contrast, as soon as $\alpha > 0$, the asymptotic small scale behavior of the process is determined by W_t , see, e.g., Eq. (11). However, this is an asymptotic result. In the presence of microstructure noise and sampling errors, the size of α will have a significant impact on the finite sample behavior of our estimators and test statistics. We illustrate this through Fig. 10, which plots the mean of $\hat{\beta}(0.7)$, obtained via a Monte Carlo simulation, for different values of α . As expected, if the Brownian motion receives considerable weight, the mean of $\hat{\beta}(0.7)$ is close to the asymptotic limit of 2. On the other hand, for low values of α , the mean of $\hat{\beta}(0.7)$ is quite close to β , rather than the asymptotic limit of 2, indicating that the finite sample behavior of the high-frequency increments is dominated by the β -stable process.

Clearly, the same type of factors will be operative when the diffusive and jump components display time-varying volatility and intensity, respectively. Specifically, the implied volatility measures are nonlinear functions of the state variables driving the volatility




Fig. 10. Small sample behavior of $\hat{\beta}(0.7)$. The estimator is based on 22 days, each containing 100 high-frequency returns simulated from the model $\alpha W_t + (1 - \alpha)S_t$ for S_t being 1.75-stable with scale normalized such that $E|S_t| = \sqrt{\frac{2}{\pi}}t$. The figure plots the mean of $\hat{\beta}(0.7)$ over 1000 Monte Carlo replications.

and jump intensity, so the weight of the different factors in a locally linear expansion will fluctuate over time, depending on the factor realizations. Thus, if, e.g., the negative jump intensity factors are of pure jump type and the diffusive volatility evolves diffusively then, depending on the loadings of the implied volatility measures on the jump intensity and diffusive volatility factors, we should observe the estimates for the activity index fall between β and 2. This is consistent with our findings in Table 2, where we noted the gradual increase in the estimates as we move across the moneyness spectrum from IV1 to IV7. A similar line of reasoning applies to the VIX index and the VIX futures. Since the VIX index may be viewed as a weighted average over the BSIVs for the OTM options, it too can be approximated by a weighted sum of the diffusive volatility and jump intensity factors.

The discussion also underscores the usefulness of including a diverse set of S&P 500 implied volatility measures in the analysis. The VIX index data, by construction, only speak to a certain weighted sum of the state variables driving the jump intensity and diffusive volatility. In contrast, by forming option portfolios that differ systematically in their degree of moneyness, we create a significant discrepancy in the portfolio loadings on the underlying factors which, in turn, should generate a sizable spread in the estimates for the associated activity measures if, indeed, the underlying factors also differ substantially in this regard.

6. Conclusion

In this paper we conduct econometric analysis of the behavior over small time scales of various derivatives written on the S&P 500 index and the CBOE VIX volatility index. In particular, in our empirical analysis we study high-frequency intraday data on a wide range of options written on the S&P 500 index, the VIX index as well as futures written on the latter. We conduct formal tests for the presence of a diffusive component in the dynamics of the series and further estimate the degree of jump activity in the absence of a diffusion. We analyze how the results regarding the different derivatives are theoretically connected within general no-arbitrage models for the underlying S&P 500 index. Our joint evidence suggests that the diffusive volatility and the (risk-neutral) intensity of the negative jumps in the S&P 500 index display different behavior over small time scales: the increments of the

former look locally Gaussian unlike those of the latter which are best described as locally stable with an activity index strictly less than two.

Appendix A. E-mini S&P 500 futures and options

The E-mini S&P 500 futures and options trade almost continuously for five days a week. Specifically, from Monday to Thursday, the trading is from 3:30 pm to 3:15 pm of the following day, with a one-hour maintenance shutdown from 4:30 pm to 5:30 pm. On Sunday, trading is from 5:00 pm to 3:15 pm of the following day. We exploit the so-called best bid-offer (BBO) files, which, among other variables, provide the best bid, bid depth, ask, ask depth, last trade price, and last trade volume. Quotes and trades are time stamped to the second. However, the files also contain a sequence indicator that identifies the order in which quotes and trades arrive to the exchange. Thus, we know the exact sequence of orders within each second. The BBO files record every change in the best bid or offer. Importantly, the quote updates arrive in pairs, one for the bid and one for the ask, synchronized by the sequence variable.

The ES futures contract expires quarterly, on the March expiration cycle. The notional value of one contract is \$50 times the value of the S&P 500 stock index.¹⁶ The ES futures contract has a tick size of 0.25 index points, or \$12.50. On any trading day, the ES options are available for seven maturity months: four months from the March quarterly cycle and three additional nearby months (“serial” options). The options expire on a third Friday of a contract month. However, the quarterly options expire at the market open, while the serial options expire at the market close. The option contract size is one S&P 500 E-mini futures. The minimum price movement is 0.05 index points. The strikes are multiples of 5 for near-term months and multiples of 25 for far months. If at any time the S&P 500 futures contract trades through the highest or lowest strike available, additional strikes are usually introduced. A disadvantage of the ES options is their American-style feature. However, we conduct our empirical analysis in such a way that the effect of the early exercise is minimal.

To guard against data errors, we apply the so-called “non-convexity” filter. The filter imposes a maximum threshold for the degree of no-arbitrage violations implied by the option mid-quotes. Apparent arbitrage violations could arise from staleness of quotes due to a temporary malfunction of the dissemination system and other issues. The procedures for implementing this filter is detailed in Appendix A.3. Whenever the “non-convexity” threshold is exceeded, we deem a cross-section too noisy and unreliable and do not use it in our analysis. Furthermore, we also additional “bounce-back” filter described in Appendix A.4. This filter detects situations where an index experiences a very large move, which is immediately offset by a jump of a similar magnitude, but in the opposite direction.

A.1. Pre-Averaging of the option IV increments

For each trading day, we start with 1620 15-second return from 8:30 am–3:15 pm and discard NaN values.¹⁷

¹⁶ The notional value of the original (“big”) S&P 500 futures contract is \$250 times the values of the index.

¹⁷ An IV observation can be labeled NaN due to a violation of the non-convexity or bounce-back filters, as explained above, or due to a lack of coverage for this strike range by the exchange. Since the moneyness varies with the ATM BSIV, an option can move from being covered to not being covered and vice versa on a given day, depending on the shifts in volatility as well as the existing coverage of the strike range by the exchange for that day. The lack of coverage is particularly problematic for the OTM call options.

Next, if all 1620 returns are non-NaN, we drop the first 60, corresponding to the first 15-minutes of trading, when most of the problems with the option quotes occur. If no more than 60 are NaNs, we drop the missing returns and possibly more from the start, so that we again are left with 1560 returns. If there are more than 60 NaN returns, the entire day is deemed unusable and is dropped from our sample.

For the remaining days, we use 78 non-overlapping intervals of 20 15-second returns ($1560 = 78 \times 20$) and average them using the tent-shaped kernel as in (12).

A.2. Constructing implied volatility measures

There are 1620 15-second intervals per trading day. For each interval, we construct seven implied volatility measures, IV1–IV7, corresponding to a fixed maturity of 30 calendar days. They are constructed as follows.

1. In the dataset, we match all puts and calls by trading date t , time to maturity τ , and strike K . We take the mid-quote as the observed option prices and convert them into forward prices, where we approximate the risk-free rate r_f over $[t, t + \tau]$ by the rate of T-bills. The implied forward price $F_{t,t+\tau}$ is obtained via the put–call parity relationship for the ATM put and call options.

2. Because the ES options are American type, their prices $P_t^A(K, \tau)$ and $C_t^A(K, \tau)$ could be slightly higher than prices of the corresponding European options $P_t(K, \tau)$ and $C_t(K, \tau)$. The difference, however, is very small for short maturities that we focus on. This is particularly true for OTM and ATM options.¹⁸

To infer prices of European options $P_t(K, \tau)$ and $C_t(K, \tau)$ we proceed as follows. First, we discard all ITM options. Prices of OTM and ATM options are more reliable, have smaller bid–ask spreads, and are less affected by the early exercise feature. Second, we correct American option prices $P_t^A(K, \tau)$ and $C_t^A(K, \tau)$ for the value of the early exercise feature by using the Barone-Adesi and Whaley (1987) approximation.¹⁹ Third, we compute prices of ITM options through the put–call parity relationship.

3. From generated prices of European options, we compute BSIVs $IV(K, \tau)$ and linearly interpolate them to obtain implied volatilities for seven values of the normalized moneyness $m = -4, -3, -2, -1, 0, 0.5$, and 1, defined in (6).

4. We repeat the above steps for the two option maturities with $\tau_j, j = 1, 2$ closest to 30 days, but excluding options with less than seven calendar days to expire. Finally, given two available maturities, τ_j , the seven constant moneyness implied volatilities IV1–IV7 are obtained computed for a fixed maturity of 30 calendar days via a linear interpolation.

A.3. Non-convexity filter

To preclude arbitrage opportunities, theoretical call and put prices must be monotonic and convex functions of the strike. In particular, the call prices must satisfy the following convexity restriction:

$$D_i = \frac{C(K_{i+1}) - C(K_i)}{K_{i+1} - K_i} - \frac{C(K_i) - C(K_{i-1}))}{K_i - K_{i-1}} \geq 0,$$

and a similar restriction for the put prices. We obtain the option prices as the average of the bid and ask quotes and use the

¹⁸ As shown in Whaley (1986), the early exercise premium increases with the level of the risk-free rate, volatility, time to maturity, and degree to which an option is in-the-money.

¹⁹ It is important to point out that this correction is always substantially smaller than typical bid–ask spreads. In particular, the correction generally does not exceed 0.2% of an option price.

Table 4

Activity index estimates and diffusion tests for the implied volatility measures based on 10-minute pre-averaged return.

	p	$\text{med}(\hat{\beta}(p))$	MAD	5%	1%
IV1	0.50	1.55	0.10	84.21	52.63
	0.70	1.62	0.09	65.79	34.21
	0.90	1.66	0.10	44.74	7.89
IV2	0.50	1.57	0.07	81.40	41.86
	0.70	1.62	0.09	72.09	32.56
	0.90	1.65	0.10	60.47	20.93
IV3	0.50	1.61	0.08	73.91	30.43
	0.70	1.64	0.05	71.74	17.39
	0.90	1.67	0.05	54.35	10.87
IV4	0.50	1.70	0.06	40.43	10.64
	0.70	1.71	0.07	42.55	6.38
	0.90	1.72	0.08	36.17	4.26
IV5	0.50	1.71	0.10	44.68	12.77
	0.70	1.72	0.10	44.68	6.38
	0.90	1.72	0.09	40.43	6.38
IV6	0.50	1.77	0.09	29.79	4.26
	0.70	1.77	0.08	25.53	4.26
	0.90	1.76	0.08	23.40	2.13
IV7	0.50	1.93	0.15	10.64	2.13
	0.70	1.95	0.14	6.38	2.13
	0.90	1.96	0.13	4.26	2.13
F	0.50	1.99	0.11	0.00	0.00
	0.70	1.99	0.10	0.00	0.00
	0.90	2.01	0.09	0.00	0.00
VIX	0.50	1.53	0.10	86.00	62.00
	0.70	1.55	0.08	84.00	56.00
	0.90	1.58	0.08	76.00	30.00
VF	0.50	1.85	0.12	18.00	8.00
	0.70	1.88	0.11	12.00	2.00
	0.90	1.92	0.10	4.00	2.00

Note: The estimator, $\hat{\beta}(p)$, is defined in Eq. (13). The rejection rates in the last two columns are based on a one-sided test for $\beta = 2$, as well as the point estimate $\hat{\beta}(p)$ and the corresponding estimate for the asymptotic variance provided in Eq. (15). The estimation and tests are performed over periods of 22 trading days.

above restriction to identify “suspect” cross-section with apparent arbitrage violations, which could arise from recording errors, staleness, and other issues. Specifically, for each strike K_i , we compute the following measure of local non-convexity:

$$NC_i = -\min\{D_i, 0\}.$$

For low strikes ($K_i \leq F$), we compute NC_i from OTM puts and, for high strikes ($K_i > F$), we use OTM calls. We then average NC_i across all strikes to obtain the aggregate measure of non-convexity NC .

When $NC > 0.1$, we deem a cross-section unreliable and do not use it in our econometric analysis. In practice, we declare all IV series at that point in time to be NaN. For those cross-sections, the option prices indicate sizeable apparent arbitrage opportunities. Andersen et al. (2013) uses a similar non-convexity filter.

A.4. Bounce-back filter

We detect situations where a volatility index experiences a very large move, which is immediately offset by a jump of a similar magnitude, but in the opposite direction. Such behavior is usually indicative of a data error which impacts only a few recorded quotes. Such an error induces a jump in returns when it occurs, and then another one when the error is eliminated and the recorded quotes reverse to their appropriate level. Specifically, we label such an incident a “bounce-back” if the index jumps by at least 15 “sigmas” and either (i) more than 75% of this jump is reversed with the next 15-second return, or (ii) more than 80% of the jump is reversed within the next two 15-second returns. This bounce-back filter is similar to the one used in Andersen

et al. (2013), except now we apply it at 15-second frequency and use a more conservative cutoff of 15 sigmas. Sigma denotes the robust standard deviation and all index returns are adjusted for the intraday pattern. See Andersen et al. (2013) for more details.

Appendix B. The VIX index and VIX futures

The contract multiplier for the VIX futures contract is \$1000. The contract has a tick size of 0.05 points, or \$50.00. CBOE usually lists up to nine near-term serial months and five months on the February quarterly cycle. Normally, the settlement date is the Wednesday that is thirty days prior to the third Friday of the calendar month immediately following the month in which the contract expires. The VIX futures are cash settled. The settlement value is for the VIX futures is determined via Special Opening Quotation (SOQ), which calculates VIX from the sequence of opening prices of the options on the settlement date. The opening price for any series in which there is no trade is equal to the average of that option's bid price and ask price as determined at the opening of trading. The final settlement value is rounded to the nearest \$0.01.

We apply the pre-averaging and bounce-back filters described in Appendices A.1 and A.4 to the VIX and VIX futures series as well.

Appendix C. Activity estimates on a Coarser frequency

(See Table 4).

References

- Ait-Sahalia, Y., Jacod, J., 2009. Testing for jumps in a discretely observed process. *Ann. Statist.* 37, 184–222.
- Ait-Sahalia, Y., Jacod, J., 2010. Is Brownian motion necessary to model high frequency data? *Ann. Statist.* 38, 3093–3128.
- Andersen, T.G., Bondarenko, O., Gonzalez-Perez, M., 2013. Uncovering Novel Features of Equity-Index Return Dynamics via Corridor Implied Volatility. Working Paper, Northwestern University and University of Illinois.
- Andersen, T.G., Fusari, N., Todorov, V., 2013a. Parametric inference and dynamic state recovery from option panels. *Econometrica*, forthcoming.
- Andersen, T.G., Fusari, N., Todorov, V., 2013b. The risk premia embedded in index options. *J. Financ. Econom.*, forthcoming.
- Barndorff-Nielsen, O., Shephard, N., 2006. Econometrics of testing for jumps in financial economics using bipower variation. *J. Financ. Econom.* 4, 1–30.
- Barone-Adesi, Giovanni, Whaley, Robert E., 1987. Efficient analytical approximation of American option values. *J. Finance* 41, 301–320.
- Bollerslev, T., Todorov, V., 2011. Tails, fears and risk premia. *J. Finance* 66, 2165–2211.
- Carr, P., Geman, H., Madan, D., Yor, M., 2003. Stochastic volatility for Lévy processes. *Math. Finance* 13, 345–382.
- Carr, P., Wu, L., 2003. What type of process underlies options? A simple Robust test. *J. Finance* 58, 2581–2610.
- Duffie, D., Pan, J., Singleton, K., 2000. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68, 1343–1376.
- Jacod, J., Li, Y., Mykland, P., Podolskij, M., Vetter, M., 2009. Microstructure noise in the continuous case: The pre-averaging approach. *Stochastic Process. Appl.* 119, 2249–2276.
- Jacod, J., Podolskij, M., Vetter, M., 2010. Limit theorems for moving averages of discretized processes plus noise. *Ann. Statist.* 38, 1478–1545.
- Lee, S., Mykland, P., 2008. Jumps in financial markets: A new nonparametric test and jump dynamics. *Rev. Financ. Stud.* 21, 2535–2563.
- Mancini, C., 2009. Non-parametric threshold estimation for models with Stochastic diffusion coefficient and jumps. *Scand. J. Stat.* 36, 270–296.
- Muhle-Karbe, J., Nutz, M., 2011. Small-time asymptotics of option prices and first absolute moments. *J. Appl. Probab.* 48, 1003–1020.
- Podolskij, M., Vetter, M., 2009. Estimation of volatility functionals in the simultaneous presence of microstructure noise and jumps. *Bernoulli* 15, 634–658.
- Sato, K., 1999. Lévy Processes and Infinitely Divisible Distributions. Cambridge University Press, Cambridge, UK.
- Todorov, V., 2013. Power variation from second order differences for pure jump semimartingales. *Stochastic Process. Appl.* 123, 2829–2850.
- Todorov, V., Tauchen, G., 2010. Activity signature functions for high-frequency data analysis. *Journal of Econometrics* 154, 125–138.
- Todorov, V., Tauchen, G., 2011a. Limit theorems for power variations of pure-jump processes with application to activity estimation. *Ann. Appl. Probab.* 21, 546–588.
- Todorov, V., Tauchen, G., 2011b. Volatility jumps. *J. Bus. Econom. Statist.* 29, 356–371.
- Todorov, V., Tauchen, G., 2014. Limit theorems for the empirical distribution function of scaled increments of Ito semimartingales at high frequencies. *Ann. Appl. Probab.* 24, 1850–1888.
- Whaley, Robert E., 1986. Valuation of American futures options: Theory and empirical tests. *J. Finance* 41, 127–150.