

Volatility Jumps

Viktor Todorov ^{*} and George Tauchen [†]

April 12, 2010

Abstract

The paper undertakes a non-parametric analysis of the high frequency movements in stock market volatility using very finely sampled data on the VIX volatility index compiled from options data by the CBOE. We derive theoretically the link between pathwise properties of the latent spot volatility and the VIX index, such as presence of continuous martingale and/or jumps, and further show how to make statistical inference about them from the observed data. Our empirical results suggest that volatility is a pure jump process with jumps of infinite variation and activity close to that of a continuous martingale. Additional empirical work shows that jumps in volatility and price level in most cases occur together, are strongly dependent, and have opposite sign. The latter suggests that jumps are an important channel for generating leverage effect.

Keywords: Stochastic volatility, activity index, jumps, jump risk premium, leverage effect, VIX index.

^{*}Department of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208; e-mail: v-todorov@kellogg.northwestern.edu

[†]Department of Economics, Duke University, Durham, NC 27708; e-mail: george.tauchen@duke.edu.

1 Introduction

Jumps are intrinsically a continuous time concept that can be defined only relative to a theoretical stochastic process satisfying mild regularity conditions. Models for such processes are convenient paradigms that should, of course, provide close approximations to the dynamics of discretely observed data. Models without jumps, i.e., models with continuous sample paths are especially convenient because then asset prices respond in a locally linear manner, hedging arguments work, and convenient, easy to manipulate closed-form expressions for the reduced forms of economic models are available. In the presence of jumps, however, markets are fundamentally incomplete and the analysis far less tractable. A fairly complete discussions of the complications induced by jumps is (Cont and Tankov, 2004, Chapter 10, pp. 319–351). Technical issues aside, jumps are important because they represent a significant source of non-diversifiable risk as discussed at length in (Bollerslev et al., 2008) and the references therein. Policy makers must make decisions in real time during times of jump-inducing chaotic conditions in financial markets, and it is thereby economically important to develop a statistical understanding of the time series behavior of jumps.

There is currently fairly compelling empirical evidence for jumps in the level of financial prices. The most convincing evidence comes from recent nonparametric work using high-frequency data as in Barndorff-Nielsen and Shephard (2007) and Ait-Sahalia and Jacod (2009a) among others. Preceding that evidence are the findings from parametric studies using daily data such as Chernov et al. (2003), Andersen et al. (2002), Eraker et al. (2003), which are strongly suggestive but arguably not overwhelming evidence for price jumps in the daily record.

A very prominent model that underlies much empirical work for continuous time processes with jumps is the setup of Duffie et al. (2000), which we call here the *affine double-jump model*. Since the double-jump model is in the affine class, just as in Heston (1993), it admits reduced form solutions for asset prices and derivatives that are closed form in the sense that they can be

readily computed on modern computing equipment using straightforward numerical techniques for Fourier series and ordinary differential equations. The double-jump model presumes rare jumps, e.g., compound Poisson process, for both asset prices and their variances. It has been applied empirically by Broadie et al. (2007), Chernov et al. (2003), Eraker et al. (2003), Eraker (2004) among others (see also empirical work by Wu (2010) who allows for the jumps to be of infinite activity). It is especially useful for specification and estimation of continuous time models that use data on both the underlying security and derivatives written on it. These studies generally find evidence for both jumps in the price level and its volatility.

In this paper we aim to understand better the nature of changes, both small and big, in the market volatility, which have important implications for volatility modeling, developing hedging strategies and specification of market risk premia. In particular we answer the following questions. Is the market volatility moving through occasional and relatively infrequent changes like in a model driven by a compound Poisson process, or it involves a lot of small moves, which over short intervals look like Gaussian as in the Heston (1993) model? Are there “sufficiently big” moves in the volatility to justify inclusion of jumps in its modeling? Are volatility and price jumps related?

To date, the answers to these questions come predominantly through estimation of parametric models built around the affine double-jump model (the recent nonparametric results of Bandi and Reno (2008, 2009) are an exception). However, these questions are intrinsically nonparametric and importantly they are related with the properties of the observed paths of the volatility for which we do not need long-span asymptotics. Therefore, the persistence in volatility, e.g. how many autoregressive factors are needed for its modeling, is a completely separate issue from the type of changes through which the volatility evolves over time. Here we are interested in the latter. Goodness-of-fit type tests for parametric volatility models would inevitably be joint type hypothesis and therefore they should always be interpreted with caution when making conclusions about pathwise properties of volatility. Here, we separate the pathwise properties of volatility from its long-span ones (like persistence) by using high-frequency data and resorting to fill-in

asymptptotics. The analysis is fully nonparametric and thus the evidence we provide here is robust.

Our estimation is based on inferring from the data the value of a generalized activity index (Ait-Sahalia and Jacod (2009a) and Todorov and Tauchen (2010a)), which is a generalization of the classical Blumenthal-Gettoor index of Blumenthal and Gettoor (1961). The generalized activity index is defined for an arbitrary stochastic process unlike the Blumenthal-Gettoor index which is defined for jump processes only. It lies in the interval $[0, 2]$ and measures the vibrancy of the process. The index divides the stochastic processes used in the volatility modeling into equivalence classes. For example the compound Poisson jump process, which is a building block in the affine double-jump model has an activity level of 0. On the other extreme is the Brownian motion (and any diffusion process), whose activity is 2. Values of the index in $(0, 1)$ correspond to jump processes of finite variation, i.e., processes whose trajectories over finite intervals are finite. Values of the index in $(1, 2)$ correspond to jump processes of infinite variation, i.e., their trajectories over finite intervals have infinite length.

We estimate the activity index of high-frequency data on the VIX volatility index computed by the Chicago Board of Options Exchange (CBOE), which is based on close-to-maturity S&P 500 index options, and then make inferences about activity level of the unobserved market volatility. Our estimation of the activity is based on constructing from the high-frequency data an activity signature function, a diagnostic tool proposed in Todorov and Tauchen (2010a). The latter provides also evidence whether the “big” moves in volatility should be modeled as jumps. Finally, to explore the link between the discontinuities on market level and market volatility we use co-jumping statistics proposed in Jacod and Todorov (2009).

The nonparametric evidence regarding the types of moves in the market volatility provides empirical information on the plausibility of the various parametric volatility models that have been proposed in the literature. The set of parametric models includes the double-jump model discussed above along with many others reviewed in Section 3 below. In some of these models volatility is

continuous, and in others it is a pure jump process. Of course there are also models with both continuous and jump components. The various parametric models have different implications for the activity level of the VIX index, the presence of jumps in it and their relationship with the ones in the price level. We find that our nonparametric evidence identifies with reasonable accuracy the most plausible class of parametric models and rules out many others.

There are certain advantages and also some notable pitfalls entailed with using the VIX data. High-frequency data, of course, provide far more information about jumps, both large and small, than do daily data, which is a major plus. Furthermore, since the VIX index is computed from quoted options prices, which are highly sensitive to volatility, it provides far more information on volatility than does the financial price series itself. Some care is needed, however, because the VIX index is not a direct measure of volatility, but rather it is actually the forward price, and thus a risk-neutral expectation, of future variance. The issues are discussed in more detail below. Finally, volatility is known to be a long memory process, and this interacts with the VIX index in some subtle ways regarding traded securities, semimartingales, and lack of arbitrage. As discussed below, it turns out that use of the general activity index permits us to separate jumps from long memory, and therefore we can make statements about the characteristics of volatility jumps without having to account for the long memory.

Turning to our main empirical findings, we can summarize them as follows. First, we find that market volatility is a very vibrant process - it involves many small changes as well as occasional big moves. The presence of big moves justifies the use of jumps in volatility modeling. In terms of modeling the small moves in volatility we find some evidence against using Brownian motion because it is somewhat more “active” than what the data implies for the volatility. On the other hand, the “activity” of the small volatility changes cannot be captured by a compound Poisson process or even a process of finite variation like a Lévy subordinator, (i.e. a jump process with non-negative increments as in the non-Gaussian OU model of Barndorff-Nielsen and Shephard (2001)). The reason for this is that a finite variation jump process would imply too “little” activity

in volatility than what is observed. We conclude that a model for the volatility that can reconcile the empirical evidence is a pure jump model, where the driving jump process is far more active than a process of finite variation, but on the other hand not as active as a continuous martingale (though the jump activity that we estimate from the data is nevertheless relatively close to that of a continuous martingale). This is to be contrasted with our findings about the market level where we find that we need a continuous martingale to capture the small changes and jumps to capture the big ones.

Second, using both high-frequency data on the VIX index as well as the S&P 500 index, we find strong evidence that the jumps in the volatility and the price level occur at the same time. We also find that these jumps exhibit strong negative dependence. These findings suggest that the underlying risks behind the occurrence of stock market discontinuities and the spikes in market volatility (and the corresponding risk premia) are similar if not the same. Therefore plausible equilibrium-based models for the market risk premia should be able to generate endogenously such links between volatility and jump risk (and their compensation).

The paper is organized as follows. In Section 2 we define the measures of stochastic volatility and in particular the VIX index, data on which is used in the empirical part. In Section 3 we present some popular stochastic volatility models and analyze their implications for the VIX index. Section 4 introduces our measure of activity of a continuous-time process and proposes methods for its inference from discrete observations. Section 5 applies the estimation technique to simulated data and Section 6 contains the empirical part. Finally Section 7 concludes.

2 The VIX Index

Let $\{S_t\}_{t \geq 0}$ denote the log of a financial price evolving in continuous time. We are interested in the high frequency dynamics of the so-called volatility index (VIX) pertaining to S_t . The VIX index is computed by the CBOE for the S&P 500 index using written options on it, but the methodology for its computation can be applied to other assets as well. Theoretically, the VIX

index is based on a portfolio of out-of-the-money options written on S_t over a continuum of strike prices whose value equals that of a variance swap, see e.g. Britten-Jones and Neuberger (2000), Jiang and Tian (2005) and Carr and Wu (2009). The latter is defined as a forward contract on the total quadratic variation of the log-price of the underlying asset over a fixed interval into the future. Following (Protter, 2004, pp. 66–76), let $[S, S]$ denote the quadratic variation process associated with S_t . Hence the VIX index is given by

$$v_t \equiv \mathbb{E}^{\mathbb{Q}}([S, S]_{t+N} - [S, S]_t \mid \mathcal{F}_t), \quad (2.1)$$

where $N > 0$ is fixed, $\{\mathcal{F}_t\}$ is the filtration on the probability space on which $\{S_t\}_{t \geq 0}$ is defined, and the expectation is taken under the risk-neutral distribution \mathbb{Q} . Note that in practice the volatility index is typically quoted in terms of annualized volatility, which is easier to interpret, but the form (2.1) is much simpler to work with theoretically so we stick with that. The quadratic variation process $[S, S]$ is adapted, increasing, càdlàg (i.e. with paths that are a.s. right continuous with left limits), and it can be split into continuous and discontinuous components

$$[S, S]_t = [S, S]_t^c + [S, S]_t^d, \quad (2.2)$$

corresponding respectively to the quadratic variation of the continuous and discontinuous parts of the price process S_t . We make a standard assumption in finance and impose absolute continuity of $[S, S]_t^c$ i.e.

$$[S, S]_t^c = \int_0^t \sigma_s^2 ds, \quad (2.3)$$

where σ_t^2 is the *spot variance* of S_t , also referred to as the instantaneous variance by Andersen et al. (2008). The spot variance σ_t^2 is the instantaneous increment to the quadratic variation of the continuous martingale component of S_t . Thus the VIX can be written as

$$v_t = \mathbb{E}^{\mathbb{Q}}\left(\int_t^{t+N} \sigma_s^2 ds \mid \mathcal{F}_t\right) + \mathbb{E}^{\mathbb{Q}}\left([S, S]_{t+N}^d - [S, S]_t^d \mid \mathcal{F}_t\right). \quad (2.4)$$

The first term is the familiar risk-neutral expectation of the forward integrated variance while the second is the risk-neutral expected contribution of the price jumps.

In practice, to generate an empirical measure of v_t in (2.4) the CBOE uses a portfolio of short-maturity out-of-the-money options on the S&P 500 Index over a discrete grid of strike prices. The details of the computation are available at <http://www.cboe.com/micro/vix/vixwhite.pdf>. In practice, there are two very small errors in replicating the price of a variance swap, which is the value on the right hand side of (2.4). The first comes from the fact that a finite number of options is used in the calculation of the VIX index, while the theoretical variance swap rate is equal to the price of continuum portfolio of options. The second error arises when there are jumps in S_t . It is equal to $-2 \int_t^{t+N} \int_{\mathbb{R}_0} (e^x - 1 - x - x^2/2) dt\nu_t^{\mathbb{Q}}(dx)$, where $dt\nu_t^{\mathbb{Q}}(dx)$ is the risk-neutral measure of the jumps. Nonetheless, the measurement is considered to be very accurate, as documented by extensive theoretical and Monte Carlo analysis in Jiang and Tian (2005) and Carr and Wu (2009), and the second error does not influence any of our subsequent results; see Theorem 1 and its proof. Hence in what follows, we treat the CBOE measured VIX as coinciding directly to v_t .

It is always important to keep in mind the distinction between the *observed* VIX and the *unobserved* spot variance. The observed VIX is the CBOE measurement of v_t in (2.4). We use these observations to make inferences about important characteristics of the random process $\{\sigma_t^2\}_{t \geq 0}$ for the spot variance. The inference is complicated because the VIX is forward looking, and its increments are generated by movements in variables that influence the conditional expectations on the right hand side of (2.4). Furthermore, we only observe discretely-sampled observations on the VIX index which also complicates estimation and inference.

To the extent possible, we follow the convention of using the term “variance” for quantities that are squares and measures of variance and the term “volatility” to refer to measures of standard deviation. Variance measures are easier to work with mathematically because they add, while volatility measures are easier to interpret because they are expressed in the same units as the data itself.

As indicated by the many papers reprinted in Shephard (2005b) and the references therein, the dynamics of the spot variance σ_t^2 are extremely important for modeling financial series. However,

the spot variance itself is not directly observed. Our plan here is to adduce nonparametric evidence from high-frequency VIX data on the empirical plausibility of various models for the spot variance. The spot variance itself can also be split into continuous and discontinuous parts

$$\sigma_t^2 = \sigma_{c,t}^2 + \sigma_{d,t}^2. \quad (2.5)$$

Note that a jump discontinuity $\sigma_{d,t}^2 - \sigma_{d,t-}^2$ influences the entire trajectory $\mathbb{E}(\sigma_{t+s}^2 | \mathcal{F}_t)$, $s \geq 0$, and thereby (in general) induces a jump discontinuity in v_t .

Historically, stochastic volatility models have assumed that the spot variance is continuous, i.e., $\sigma_t^2 \equiv \sigma_{c,t}^2$. However, more recently there has been interest in pure jump stochastic volatility models, $\sigma_t^2 \equiv \sigma_{d,t}^2$; see, e.g., Barndorff-Nielsen and Shephard (2001). Of course, the models can be combined, as in the double-jump model of Duffie et al. (2000). Two recent comprehensive reviews of stochastic volatility are Shephard (2005a) and Andersen and Benzoni (2007). In the subsequent section we highlight the more relevant models and their implications for the VIX index.

3 Parametric Models for the Spot Variance

Our objective is to use nonparametric-type evidence from high frequency VIX and returns data to cast light on the empirical plausibility of the various parametric volatility models for the spot variance σ_t^2 that have been proposed in the literature. We briefly review the extant parametric models in this section and then proceed to the nonparametric analysis in the application section farther below. We leave unspecified whether the model pertains to the risk-neutral distribution or the objective distribution, because common practice is to assume a risk premium structure that preserves the basic form of the model across the two distributions. We also suppress here for simplicity the presence of price jumps, since we are considering parametric spot volatility models, and, for example, allowing for price jumps with intensity that is linear in the spot volatility factors will simply lead to affine transformations of the expressions for the VIX index below.

The most widely used model in finance is probably the affine jump diffusion model written in its most general form as:

Affine Jump Diffusion

$$d\sigma_t^2 = \rho(\sigma_t^2 - \psi_0) dt + \psi_1 \sigma_t dB_t + dL_t, \quad \rho < 0, \psi > 0, \quad (3.1)$$

where L_t is a Lévy process of finite variation with non-negative jumps. The model has been widely used in both equilibrium and reduced-form asset pricing modeling, important examples include Merton (1976), Duffie et al. (2000), and Duffie et al. (2003). In this case VIX index v_t is simply an affine function of the volatility process.

In most applications of (3.1), e.g. the affine double-jump model of Duffie et al. (2000), L_t is a compound Poisson process and B_t is present. More recently, an important special case of the affine jump diffusion (3.1) is the non-Gaussian OU model of Barndorff-Nielsen and Shephard (2001) in which the diffusive component is absent:

Non-Gaussian OU

$$d\sigma_t^2 = \rho\sigma_t^2 dt + dL_t, \quad \rho < 0, \quad (3.2)$$

where L_t is a pure jump Lévy process with non-negative increments, also called a subordinator. As for (3.1), the VIX index v_t is an affine function of the volatility process under this specification.

Exponential-type stochastic volatility models have been also widely used in financial econometrics:

EXP-OU- Γ

$$\begin{aligned} \sigma_t^2 &= \exp(\alpha_0 + \alpha_1 f_t), \\ df_t &= \rho f_t dt + d\Gamma_t, \quad \rho < 0, \end{aligned} \quad (3.3)$$

where Γ_t is a generic process. When Γ_t is a Brownian motion, the model is a continuous-time limit of the discrete EGARCH model of Nelson (1991). Many papers relevant for this model are conveniently reprinted in Shephard (2005b). Standard calculations imply that in this case (when Γ_t is a Brownian motion), the VIX index is

$$v_t = \int_0^N \exp\left(\alpha_0 + \alpha_1 e^{\rho u} f_t + \alpha_1^2 \frac{1 - e^{2\rho u}}{4\rho}\right) du. \quad (3.4)$$

Versions of the model (3.3) with Γ_t being a Lévy process with compound Poisson jumps are estimated in Andersen et al. (2002) and Chernov et al. (2003).

More generally, when Γ_t is an arbitrary Lévy process, using (Sato, 1999, Theorem 25.17) the formula for the VIX index generalizes to

$$v_t = \int_0^N \exp[\alpha_0 + \alpha_1 e^{\rho u} f_t + C(u)] du, \quad (3.5)$$

where $C(u)$ is some function of u determined by the characteristic exponent of the driving Lévy process (and the constants α_1 , ρ and N). A very important feature of (3.3) is that the driving Lévy process can be infinitely active and of infinite variation, see Haug and Czado (2007), which is unlike (3.2) where the driving process must be of finite variation.

A common feature of the above models for the volatility and their multi-factor extensions is that they are Markovian (up to augmenting the state space). Given the well-documented long range dependence in volatility (Baillie et al., 1996; Comte and Renault, 1998; Shephard, 2005a), some researchers have alternatively applied models that are non-Markovian, e.g. a fractionally integrated model such as that of Comte and Renault (1998). The latter is an exponential stochastic volatility model with $\Gamma_t = B_{\delta,t}$, where $B_{\delta,t}$ is fractionally integrated Brownian motion with fractional integration parameter $0 < \delta < \frac{1}{2}$. The factor f_t has the following stationary representation

$$f_t = \int_{-\infty}^t a(t-s) dB_s, \quad (3.6)$$

where B_t is standard Brownian motion, and the function $a(\cdot)$ is given by

$$a(u) = \frac{1}{\Gamma(1+\delta)} (u^\delta + \rho e^{\rho u} \int_0^u e^{-\rho x} x^\delta dx). \quad (3.7)$$

The VIX index for this model takes the following form

$$v_t = \int_0^N \exp\left(\alpha_0 + \alpha_1 \int_{-\infty}^t a(t+u-s) dB_s + \frac{\alpha_1^2}{2} \int_0^u a^2(z) dz\right) du. \quad (3.8)$$

As noted in Comte and Renault (1998), the spot variance σ_t^2 in this model is not a semimartingale. Nevertheless there are no arbitrage opportunities of the type discussed in Rogers (1997) because

the spot variance is not a traded security. The observed VIX index, however, is a portfolio of traded securities and it should be a semimartingale to rule out arbitrage, and this indeed is the case for the model-implied VIX index in (3.8); we illustrate the point below.

4 The Activity Level of Volatility

The volatility models in the previous section have been all used in various applications, and our aim is to provide nonparametric evidence on their empirical plausibility using high-frequency observations on the VIX index. Towards this end, we now show in this section how to associate with each continuous-time process an index of its so-called activity and present methods to estimate the index. Since we actually estimate the activity index on VIX data, but we are interested in the spot volatility σ_t^2 process, we end this section by deriving results linking the activity index for the spot variance and the VIX index under mild regularity conditions.

4.1 Activity Index

We start with consideration of a measure of activity for an arbitrary continuous-time process. Intuitively, by activity level we mean the “degree” of vibrancy of the process, i.e. the “roughness” of its trajectories. Formally, the statistical setup is as follows. We observe a generic scalar process X over a long span $[0, T]$. During each subperiod $(t - 1, t]$, where now t is an integer, we have high-frequency observations on X with a sampling interval of length Δ_n . That is, we observe X at times $t - 1, t - 1 + \Delta_n, \dots, t - 1 + [1/\Delta_n]\Delta_n$ during the subperiod. Think of the subperiod as being either a day, week, or month. Following Ait-Sahalia and Jacod (2009a) and Todorov and Tauchen (2010a), we can define the activity of X during an arbitrary interval $(t - 1, t]$ as

$$\beta_{X,t} := \inf \{p > 0 : \text{plim}_{\Delta_n \rightarrow 0} V_t(X, p, \Delta_n) < \infty\}, \quad (4.1)$$

where $V_t(X, p, \Delta_n)$ is the power variation of X over the interval $(t - 1, t]$ given by

$$V_t(X, p, \Delta_n) = \begin{cases} \sum_{i=1}^{[1/\Delta_n]} |x_{t,i}|^p I(|x_{t,i}| \leq c), & p < 2 \\ \sum_{i=1}^{[1/\Delta_n]} |x_{t,i}|^p, & p \geq 2 \end{cases} \quad (4.2)$$

where $x_{t,i} = X_{t-1+i\Delta_n} - X_{t-1+(i-1)\Delta_n}$ and $I(\cdot)$ is the 0-1 indicator function. The importance of the power variation (4.2) for financial econometrics was pointed out in Barndorff-Nielsen and Shephard (2003, 2004) and follow-up papers.

The truncation at $c > 0$ in (4.2) for powers r less than 2 has no effect asymptotically, because the value of the activity index $\beta_{X,t}$ in (4.1) is determined by the “small” price moves (the sample functions of X are càdlàg, and therefore jumps bigger in absolute value than a fixed positive number are always finite over a finite period of time). However, the truncation provides robustness in finite samples of our estimator of the index, constructed from the power variation in the next subsection, to very extreme price movements. For $p > 2$ the truncation is unnecessary because in that region only the “large” moves matter for the asymptotic behavior of the power variation. In practice, we use a very large value of c implying very mild truncation so that typically up to two of the summands are truncated, and we check sensitivity of the findings to the choice of c . We are extremely grateful to an anonymous referee who pointed out the potential lack of finite sample robustness, which can be fixed by this simple expedient of truncation.

Our interest is mainly in the case when X is a semimartingale, because to avoid arbitrage, any traded security, and thus the VIX index as well, needs to be a semimartingale, see Delbaen and Schachermayer (1994). For a semimartingale the activity index takes values in the interval $[0, 2]$. Each semimartingale can be decomposed into drift term along with continuous and discontinuous local martingale parts (Jacod and Shiryaev, 2003). These components of the semimartingale process can be naturally ranked in terms of their activity in the following order from least to most active: finite activity (e.g. compound Poisson) jumps (activity of 0), infinite activity but finite variation jumps (activity in $[0, 1]$), drift (activity of 1), infinite variation jumps (activity of $(1, 2]$), continuous martingales (activity of 2). The activity of the semimartingale is determined by the activity of its most active component. Thus, for example, if X is driven by both a Brownian motion and jumps, the continuous martingale dominates and the activity of X is equal to that of its continuous martingale component, which is 2.

Evidently, the jumps are the most interesting component of a semimartingale in terms of their activity. Our measure of activity for pure jump Lévy processes coincides with their so-called (generalized) Blumenthal-Gettoor index (Blumenthal and Gettoor (1961) and Ait-Sahalia and Jacod (2009a)). This index is analogous to the parameter α of the α -stable distribution.

If X is not a semimartingale things are different. For example, when X is the OU process driven by fractional Brownian motion given in (3.6), then its activity index is determined by its degree of fractional integration, δ , and is equal to $\frac{1}{\delta+0.5}$, see Corcuera et al. (2006).

Finally, note that activity index is defined over the subinterval $(t-1, t]$, instead of the whole sample. This approach allows for the possibility that the process X can change its activity over time. While the activity of most parametric continuous-time models, including those for the spot variance of the previous section, is sometimes presumed constant over long segments of time, we do not make this assumption a priori. Breaking the estimation up across a number of subintervals provides some bootstrap-type indication on sampling fluctuations and more importantly provides protection against parameter shifts.

4.2 Estimation of Activity Index

The estimators of the activity index are really quite simple to compute. The key results from Todorov and Tauchen (2010a) that generate the estimators are as follows. With $V_t(X, p, k\Delta_n)$ denoting the power variation as defined in (4.2) and $\beta_{X,t}$ the activity index in (4.1) to be estimated for period t (think of t as a day or a month), then

$$\begin{aligned}
 (a) \quad & (k\Delta_n)^{1-\frac{p}{\beta_{X,t}}} V_t(X, p, k\Delta_n) \xrightarrow{\mathbb{P}} \Phi_t(p), \quad 0 < p < \beta_{X,t}, \quad \text{for all } k \geq 1 \\
 (b) \quad & V_t(X, p, k\Delta_n) \xrightarrow{\mathbb{P}} \sum_{t-1 \leq s < t} |\Delta X_s|^p 1_{(|\Delta X_s| \leq c \cap p < 2) \cup p \geq 2}, \quad p > \beta_{X,t},
 \end{aligned}
 \tag{4.3}$$

as $\Delta_n \downarrow 0$. In (a) the limit on the right depends only on the power p ; the result in (a) holds for any p in the case when X is continuous (i.e., it does not contain jumps); in (b) the limit on the right is the sum of the absolute jumps raised to the p^{th} power (and possibly truncated).

To get estimators of the activity index, evaluate (a) above at $k = 1, 2$, take the ratio, and

note

$$(\Delta_n)^{1-\frac{p}{\beta_{X,t}}} V_t(X, p, \Delta_n) / (2\Delta_n)^{1-\frac{p}{\beta_{X,t}}} V_t(X, p, 2\Delta_n) \xrightarrow{\mathbb{P}} 1, \quad p < \beta_{X,t}.$$

Now take the log of this ratio, set it equal to its asymptotic value of zero, and solve for the implied “ β ” to get

$$b_{X,t}(p) = \frac{\ln(2)p}{\ln(2) + \ln[V_t(X, p, 2\Delta_n)] - \ln[V_t(X, p, \Delta_n)]}, \quad p > 0. \quad (4.4)$$

The expression $b_{X,t}(p)$ above is termed the *activity signature function* (ASF), and we find that essentially all relevant information about the activity of the X process is contained in the ASF for $p \in (0, 4]$.

As shown in Todorov and Tauchen (2010a), as we sample more frequently, i.e. $\Delta_n \rightarrow 0$ on any fixed interval $(t-1, t]$, the activity signature function $b_{X,t}(p)$ behaves as follows

- A.** $b_{X,t}(p) \rightarrow 2, \quad \forall p > 0$ if X contains continuous martingale,
- B.** $b_{X,t}(p) \xrightarrow{\mathbb{P}} \max(p, 2), \quad \forall p > 0$ if X contains continuous martingale plus jumps,
- C.** $b_{X,t}(p) \xrightarrow{\mathbb{P}} \max(p, \beta_{X,t}), \quad \forall p \neq \beta_{X,t}$ if X is driven by a pure jump process,

where the convergence is locally uniform in p . The right-hand sides above describe the asymptotic behavior of the ASF. Intuitively, to get the result in A for example, one can apply (4.3) to write $V_t(X, p, 2\Delta_n) \approx (2\Delta_n)^{p/2-1} \Phi_t(p)$ and $V_t(X, p, \Delta_n) \approx (\Delta_n)^{p/2-1} \Phi_t(p)$. Simple algebra then leads to the flatness at 2 of the asymptotic limit of $b_{X,t}(p)$ in this case. Similar analysis leads to the limits in the other two cases.

In finite samples the realization of $b_{X,t}(p)$ is a smooth infinitely differentiable function of p . From the asymptotics, we can thus expect it to display a bend around $p \approx \beta_{X,t}$, the population value, with the sharpness of the bend providing an indication of precision of estimation. On the other hand, the behavior of the activity signature function for $p > 2$ can reveal us whether there are jumps, large or small, in the process X in the interval $(t-1, t]$ even in the case when they are dominated by a continuous semimartingale.

Todorov and Tauchen (2010a) suggest a graphical method to use the ASF to get an indication of the value of $\beta_{X,t}$ and the presence of jumps. Specifically, let

$$\mathcal{B}_q(p) = q^{\text{th}} \text{ quantile of } \{b_{X,t}(p)\}_{t=1,2,\dots,N}, \quad q \in (0, 1) \quad (4.5)$$

denote the q^{th} quantile of the $b_{X,t}$ for each power p . $\mathcal{B}_q(p)$ is called the *quantile activity signature function* (**QASF**). The most informative plots are obtained from the lower and upper quartiles, $\mathcal{B}_{0.25}(p)$, $\mathcal{B}_{0.75}(p)$, and the median $\mathcal{B}_{0.50}(p)$ over the range $0 \leq p \leq 4$. Robust methods such as using quantiles are crucial because we are dealing with data sets containing extreme observations.

Inspecting graphs can be helpful for a rough indication, but we can directly estimate the activity index over each subinterval $(t-1, t]$ as

$$\widehat{\beta}_{X,t} = b_{X,t}(p), \quad \text{for some fixed value of } p. \quad (4.6)$$

The estimator (4.6) is consistent for the activity index provided $p < \beta_{X,t}$, and there are several considerations that guide the choice of p . First, obviously we need to pick p lower than the lowest possible activity $\beta_{X,t}$ for the process X . In our case, we can assume the activity level is at least 1, because the volatility process is mean-reverting and the drift term has an activity of 1. This was also illustrated with the different parametric models of Section 3. Second, it can be shown that for very small powers the estimator is relatively inefficient and thus higher values of p are preferred. Based on this discussion, an appropriate choice for p in (4.6) is in the range 0.50 to just under 1.00. Values in this range are nearly optimal for the levels of activity our data suggest, but the formal optimality analysis is technically very demanding and well beyond the scope of the present paper (Todorov and Tauchen, 2010b).

A significant issue is whether the process contains a Brownian component, and we therefore conduct a test whether $\widehat{\beta}_{X,t}$ is statistically less than two. We do the test in logs and construct a one-sided critical region for $\ln(b_{X,t}(p))$ using its asymptotic distribution under the null, which is normal with estimated standard error of $\widehat{AsySE}(\ln(b_{X,t}(p)))$. Details on the asymptotic result and the calculation of the feasible standard error can be found in Todorov and Tauchen (2010a).

4.3 Linking VIX and Spot Variance Activity Indexes

We undertake the estimation described above using the high-frequency VIX data but the interest is on the unobserved spot variance process. Therefore, for the estimation to be meaningful it is important to investigate the relationship between the activity level of the observed VIX index to that of the spot variance, and it turns out they are the same under very weak regularity conditions. Indeed, in a Markov setting, which is most often adopted in parametric volatility modeling, the agreement is established by the following theorem:

Theorem 1 *For the setting in (2.2)-(2.4) suppose in addition the following*

- A.** $\sigma_t^2 = G^{(c)}(\mathbf{f}_t)$ for some twice differentiable function $G^{(c)} : \mathbb{R}^k \rightarrow \mathbb{R}^+$ with non-vanishing first derivatives on the support of \mathbf{f}_t ,
- B.** the compensator of the jumps in $\{S_t\}$ under the measure \mathbb{Q} is of the form $G^{(d)}(\mathbf{f}_t)dt \otimes \eta(dx)$ for some twice differentiable function $G^{(d)} : \mathbb{R}^k \rightarrow \mathbb{R}^+$ and a measure η on \mathbb{R} satisfying $\int_{\mathbb{R}} (|x|^2 \wedge 1)\eta(dx) < \infty$,
- C.** \mathbf{f}_t is a vector with independent elements each of which solves under \mathbb{Q}

$$df_t^{(i)} = \sum_{j=1}^{d_i} g_j^{(i)}(f_{t-}^{(i)})dZ_{tj}^{(i)}, \quad j = 1, \dots, d_i, \quad i = 1, \dots, k, \quad (4.7)$$

where the functions $g_j^{(i)}(\cdot)$ are twice differentiable and $Z_{tj}^{(i)}$ are independent Lévy processes.

Assume further that the expectation $v_t = \mathbb{E}^{\mathbb{Q}}([S, S]_{t+N} - [S, S]_t | \mathcal{F}_t)$ is well defined. Then, $\nu_t = F(\mathbf{f}_t)$ for some continuously differentiable function $F : \mathbb{R}^k \rightarrow \mathbb{R}^+$. Furthermore, if $\frac{\partial F}{\partial f_i}(\cdot) \neq 0$ on the support of \mathbf{f}_t for $i = 1, \dots, k$, then we have for an arbitrary $t > 0$

(a)

$$\beta_{\sigma^2, t} \equiv \beta_{\nu, t} \quad a.s., \quad (4.8)$$

where $\beta_{X, t}$ for an arbitrary semimartingale X_t is defined in (4.1).

- (b) the set of jump times of $\{\sigma_s^2\}_{[0, t]}$ coincides with the set of jump times of $\{\nu_s\}_{[0, t]}$ almost surely provided F is monotone in each of its arguments.

The result of the theorem follows essentially from the fact that under its assumptions, both the spot variance and the associated VIX index are continuously differentiable functions of the volatility factors. Section 3 contains particular parametric examples of this. Such transformations preserve the activity index, and therefore to determine the volatility activity we need only determine the activity of the most active volatility factor.

The theorem does not say that the VIX, which is a market-implied quantity, and the spot variance are the same. Indeed, their dynamics, including persistence and level, might be quite different as is often found in empirical work. The theorem does say, however, that key features of the stochastic processes, i.e., their activity levels and sets of jump times, must agree.

Based on Theorem 1 and the discussion in the previous section (see Todorov and Tauchen (2010a) for formal results), we have the following levels of volatility activity for the Markov models of the previous section

- (a) Affine Jump Diffusion and EXP-OU-Gaussian: $\beta_{\sigma^2,t} = \beta_{v,t} = 2$.
- (b) Non-Gaussian OU: $\beta_{\sigma^2,t} = \beta_{v,t} = 1$.
- (c) EXP-OU-Lévy: $\beta_{\sigma^2,t} = \beta_{v,t} = \max\{\beta_L, 1\}$, where β_L is the Blumenthal-Gettoor index of the driving Lévy process.

Note that for the non-Gaussian OU model the volatility activity is determined by the drift term in (3.2), since the driving jump process is a Lévy subordinator and thus of finite variation. On the other hand, for the affine jump diffusion model (and all models in which Brownian motion is used in the specification of the stochastic volatility), the volatility activity is driven by the continuous martingale part which “dominates” drift and arbitrary jump factors. Thus, the most tractable (and hence most used) stochastic volatility specifications from the different classes of models of the previous section, the affine jump diffusion model and the non-Gaussian model, have very different implications for the activity level of the stochastic volatility.

Finally, the relationship between the activity of the spot variance and the associated VIX index is non-trivial outside of the Markov setting. For the EXP-OU-FI model, we have $\beta_{\sigma^2,t} = \frac{1}{\delta+0.5}$ while $\beta_{v,t} = 2$. As already mentioned, the activity of the spot variance is driven by the degree of the fractional integration. The transformation implied by the VIX index restores the semimartingale property, and, since the model is driven by Brownian motion, we have that the activity of the VIX index is 2.0 as determined by this most active component. Therefore for the EXP-OU-FI model, the activity of the VIX index will not be informative about that of the spot variance. However, this is not a drawback of our analysis. For this model it is the activity of the VIX index that we are most interested in, since it tells us that the volatility process is modeled via (fractionally-integrated) Brownian motion and not jumps, and this is exactly what we are after.

5 Monte Carlo

We now consider the finite sample properties of the estimators and tests developed in the preceding Section 4. Table 1 contains a complete list of the various scenarios considered: affine jump diffusions with a wide range of jump intensities (**Case A**), long memory volatility models (**Case B**), non-Gaussian OU models (**Case C**), and Lévy-driven pure jump models (**Case D**). The parameter values for the first two affine jump diffusion specifications were taken from the estimation results of Eraker et al. (2003), while the other two cases are used as a check for the robustness of the results against various “extreme” scenarios. In case AJD-HJ, we kept the mean of the jumps the same but increased approximately ten times the jump intensity. In case AJD-E, we further increased the mean jump size while keeping the high jump intensity of case AJD-HJ.

For each considered model we calculate the corresponding value of the VIX index using the expressions in Section 3 and do all the calculations of Section 4 using the simulated high-frequency data on it. For the same reasons as in Section 2, we ignore price jumps and risk premia. The main question that we seek to answer is whether the transformations involved in the calculation

Table 1: Parameter Setting for the Monte Carlo

Case	Parameters						β_{σ^2}	β_{ν}
A. Affine Jump Diffusion								
	ψ_0	$-\rho$	ψ_1	λ	μ			
AJD-no jumps	0.8136	0.0128	0.0954			2.0	2.0	
AJD-LJ	0.5585	0.0250	0.0896	0.0055	1.7980	2.0	2.0	
AJD-HJ	0.5585	0.0250	0.0896	0.0500	0.1978	2.0	2.0	
AJD-E	0.5585	0.0250	0.0896	0.0500	0.9889	2.0	2.0	
B. EXP-OU-FI								
	α_0	α_1	δ	$-\rho$				
	0.00	1.00	0.40	1.00		1.1	2.0	
C. Non-Gaussian-OU								
	$-\rho$	β	λ	c				
	0.03	0.50	5.00	0.05		1.0	1.0	
D. EXP-OU-Lévy								
	α_0	α_1	$-\rho$	β	λ	c		
	-0.70	1.00	0.07	1.50	2.50	0.10	1.5	1.5

Note: **Affine Jump Diffusion** model is given in (3.1) with Lévy density of the jump process equal to $\lambda \frac{e^{-x/\mu}}{\mu} 1_{\{x>0\}}$ (compound Poisson process with exponentially distributed jumps). **EXP-OU-FI** model is given in (3.3) with driving process being the fractional Brownian motion $B_{\delta,t}$. **Non-Gaussian-OU** model is given in (3.2) with Lévy subordinator having Lévy density given by $c \frac{e^{-\lambda x}}{x^{\beta+1}} 1_{\{x>0\}}$ (tempered stable process). **EXP-OU-Lévy** model is given in (3.3) with Lévy density of the pure-jump driving Lévy process equal to $c \frac{e^{-\lambda|x|}}{|x|^{\beta+1}}$ (tempered stable process).

of the VIX index have any finite sample effect on our inference for the spot variance activity. For general assessment of finite sample properties of activity estimation and testing we refer to the web-appendix to Todorov and Tauchen (2010a).

5.1 Illustrating the Basic Computations on Simulated Data

We start by summarizing the basic aspects of the computations associated with the theory of Subsection 4.1 and display the outcome on a simulated realizations for a few representative scenarios from Table 1. In this presentation we let β denote $\beta_{X,t}$ for simplicity. In each day we sample 78 times, which corresponds to a 5-minute sampling frequency in a standard 6.5 hours trading day, and this also is the frequency of our high-frequency data that we use in the empirical analysis of the next section. The interval $(t-1, t]$ corresponds to 22 trading days, i.e., a calendar month, so the unit of time is thereby 1 = one month in all calculations that follow. There are $78 \times 22 = 1716$ high-frequency intervals per month. The use of a month as the subinterval is a compromise in the tradeoff between the presumption of constant activity over the subinterval and the associated reduction in sampling error inference with more data points per interval.

To begin, compute the power variation $V_t(X, p, \Delta_n)$ in (4.2), where X is the process and the power $p \in (0, 4]$ ranges over a fine grid in the t^{th} time interval, here a month, with $\Delta_n = 1/(78 \times 22)$. Next compute $V_t(X, p, 2\Delta_n)$ using the coarser 10-minute sampling. When computing the power variation over this coarser frequency and for powers $p < 2$, we first remove the 5-minute price increments bigger in absolute value than the truncation level ($c = 1.50$), and then aggregate to 10-minutes and compute the power variation from them (without any further truncation). Finally, using the power variations over the two frequencies, we compute the *activity signature function* for interval t using (4.4). Since it is impossible to report in any sensible manner each of the activity functions $b_{X,t}(p)$, a summary measure based on robust methods needs to be computed: the quantile activity signature function defined in (4.5) and the quartiles $q = 0.25, 0.50, 0.75$, commonly used in statistics, prove informative.

Recall in presence of jumps

$$b_{X,t}(p) \xrightarrow{\mathbb{P}} \max(\beta, p),$$

from the asymptotic analysis. So, in finite samples we expect the median QASF, $\mathcal{B}_{0.50}(p)$, to be close to β for powers $p < \beta$, close to p for $p > \beta$, and curvilinear for p in a neighborhood of β . The upper and lower QASFs $\mathcal{B}_{0.75}(p)$ and $\mathcal{B}_{0.25}(p)$ provide an indication of sampling dispersion.

As a check, we compute the QASFs on simulated realizations for a few well-known volatility models where the value of β is given. These simulated realizations follow standard conventions with annualized volatility based on 252 trading days per year. We simulate the different volatility models over a total of 4400 days, which corresponds to 200 months. Of course the activity level, which recall we just denote β here, is the same for all simulated months, but that need not be the case with observed data.

We start with an affine jump diffusion where the QASFs are shown in the top left- and right-hand rows of Figure 1. In the top left, jumps are suppressed (Case AJD-no-jumps), the process is continuous, and the QASFs are flat around $\beta = 2$ as expected. In the top right (Case AJD-E) large rare jumps are added in to the Brownian diffusion. Now the QASFs are flat around $\beta = 2$ for $p \leq 2$, since the continuous component dominates here, while for $p > 2$ the curves slope upwards to the asymptotic value p . The sharp break in slope at $p = 2$ in the top right plot in Figure 1 is due to the dominance of the large jumps; this behavior might be unlikely in practice where only few months can have such big jumps, and the plots therefore should be regarded as a robustness check.

The plots in the second two rows in Figure 1 pertain to a Brownian long memory stochastic volatility with parameter settings \mathbf{B} in Table 1. To contrast the different activity of the spot variance and the VIX index in this model, we calculate also the QASFs of the unobservable spot variance. In the second row left-side are the QASFs for the simulated spot variance process, which are flat, reflecting continuity of the process, but around a value well less than 2.0. The reason is that the spot variance is not a semi-martingale so there is no constraint that its QASF

asymptotically pass through the point $(2,2)$. The height of the asymptotic value of the activity signature function is determined by the fractional difference parameter d . Interestingly, for the second row right-hand side the QASFs for the VIX index associated with this spot volatility process are flat lines around 2.0, which has to be the case asymptotically because the VIX is a portfolio of traded securities and thereby must be a semimartingale. Finally, the two plots in

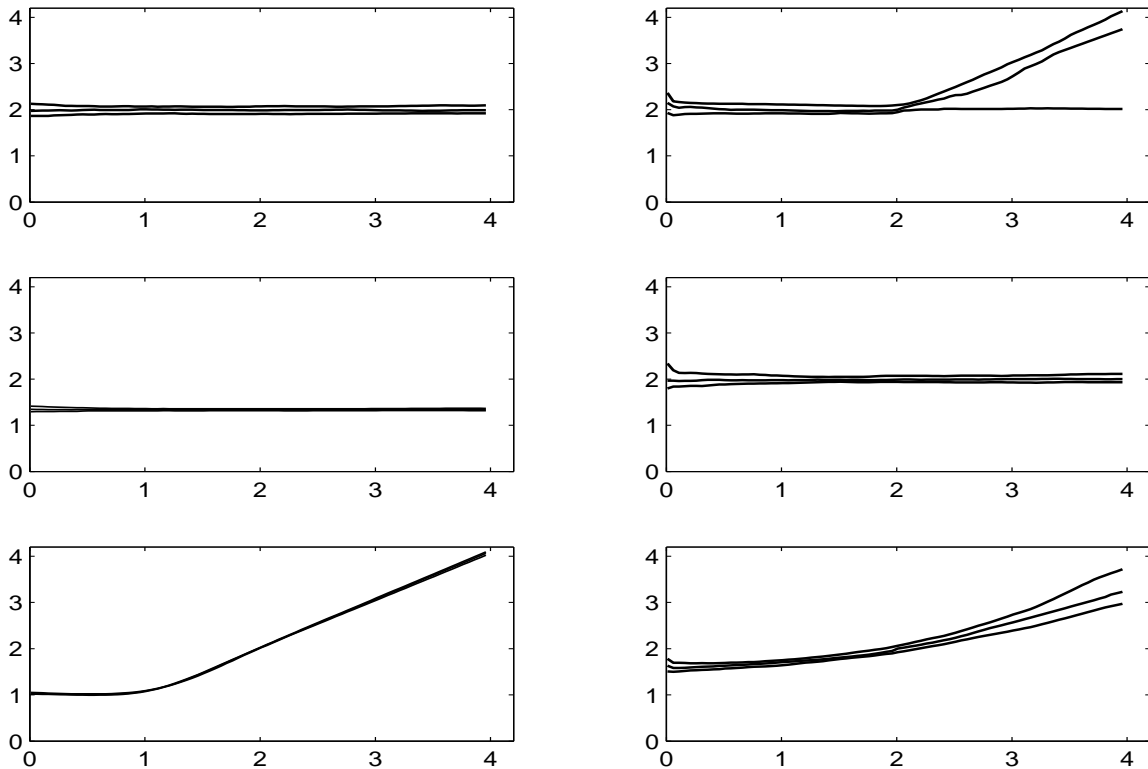


Figure 1: *QASF*-s for various stochastic volatility models. In each panel the three quantiles that are displayed are the 25-th, 50-th and 75-th, and are computed on the basis of 200 months of simulated data. The top left and right panels correspond to the AJD-no jumps and AJD-E respectively models. The middle panels correspond to the EXP-OU-FI model. The bottom left and right panels correspond to the Non-Gaussian-OU and EXP-OU-Lévy respectively specifications. All model specifications are given in Table 1. In all cases but the middle right panel, the *QASF*s are based on the VIX index. *QASF*s for the middle left panel are for the spot variance series. The truncation level in all cases is $c = 1.5$.

the bottom row pertain to models where volatility is a pure jump process with no continuous component. The plot in the lower left of third row, pertains to the non-Gaussian OU model **Case C** in Table 1. The value of β of the driving Lévy process is 0.50, but the bend occurs around $p = 1.0$. The reason is that the non-Gaussian OU model has a drift component, which must have

an activity index of 1.0, and the approach taken here always reveals the index of the dominant component. The plot in the lower right row pertains to the Lévy-driven OU process, **Case D** in Table 1 where $\beta = 1.50$. There is a soft bend around the true value of $p = 1.50$ and the jumps are quite apparent for $p \geq 2.00$. The softness bend around $p = 1.50$ indicates that for higher values of the index the plots are just indicative and will not reveal the actual value with high precision.

5.2 Assessment of the Activity Estimator

The Monte Carlo assessment of the accuracy of the estimator (4.6) for each of the cases is shown in Table 2. We computed the estimator for 5-minute returns for a 6.5 hour day, pooled over a period of a “month” (comprised of 22 trading days) and replicated 1,000 times. The power parameter is $p = 0.95$, but the results are quite insensitive to the choice of p of the range 0.50 to 1.00. Table 2 shows the median and the median absolute deviation about the median as measures of central tendency and variability, respectively. The reported results include no truncation (NT) and truncation (T) at level c in (4.1).

Table 2: Small Sample Behavior of $\hat{\beta}$

Case	β	med($\hat{\beta}$)		MAD	
		NT	T	NT	T
AJD-no jumps	2.00	2.01	2.01	0.081	0.079
AJD-LJ	2.00	2.00	2.01	0.086	0.079
AJD-HJ	2.00	1.98	2.00	0.081	0.081
AJD-E	2.00	1.92	2.00	0.093	0.082
AJD-E-JS	2.00	1.91	2.00	0.080	0.076
EXP-OU-FI	2.00	1.99	1.99	0.088	0.088
NOU	1.00	1.06	1.06	0.010	0.010
EXP-OU-Lévy	1.50	1.69	1.71	0.052	0.052

Note: med is the median function; $\text{MAD} = \text{med}|\hat{\beta} - \text{med}(\hat{\beta})|$; NT indicates no truncation; T indicates truncation with $c = 1.5$; case AJD-E-JS is the same as AJD-E but we keep only simulations in which the estimation period contains at least one jump. There are 1,000 replications of one month’s worth of 5-minute observations. The estimator $\hat{\beta}$ is given in (4.6) for $p = 0.95$.

Results for the affine jump diffusion are in the first four rows of Table 2. The estimator

without truncation (NT) is unbiased and reasonably accurate, except in the cases **AJD-E** and **AJD-E-JS**, where rather large jumps have been added to the diffusion. The case **AJD-E-JS** always contains at least one large jump in each simulated month. We are very grateful to a referee for pointing out that such large jumps could impart a finite sample downward bias. The truncation point $c = 1.50$ (recall VIX index is quoted in annualized percentage units) is very mild, as it eliminates only one or two large moves per period, but as seen from the table in the (T) column it properly corrects for the downward bias.

Overall, Table 2 suggests the estimator is quite well behaved, regardless of whether the jumps are finitely or infinitely active and of bounded or unbounded variation. The truncation has no essential effect in any of the infinite activity cases, and it is really needed only in finite samples to guard against huge large rare jumps (which asymptotically do not matter). The dispersion measure suggests the estimator is accurate to within a range between ± 0.05 to ± 0.10 .

5.3 Assessment of the Test for a Brownian Component

We also evaluated the test for a Brownian component over the same set of replications and summarize the findings in Table 3. For the first five cases of an affine jump diffusion, the null hypothesis is true, so the rejection rates represent the size of the test. Now it is seen that the truncation (T) is much more important for the actual size to agree closely with the nominal size. In the long memory model, the null is also true but the truncation is irrelevant for this case. In the last two cases of pure jump volatility models the test is seen to have very high power.

6 Empirical Application

We use high-frequency data on the VIX index computed by the CBOE along with S&P 500 Index futures returns. The data set spans the period from September 22, 2003 until December 31, 2008, for a total of 1,212 trading days which corresponds to 64 calendar months. Within each day, we use 5-minute records of the VIX index and the S&P 500 futures contract from 9.35 till 16.00

Table 3: Size and Power of the Test for a Brownian Component

	REJECTION RATES (Percent)			
	NT		T	
	$\alpha = 5\%$	$\alpha = 10\%$	$\alpha = 5\%$	$\alpha = 10\%$
SIZE				
AJD-no jumps	3.6	8.1	3.6	8.1
AJD-LJ	8.8	13.1	4.0	7.9
AJD-HJ	7.4	12.9	4.5	9.2
AJD-E	20.8	29.4	3.7	9.6
AJD-E-JS	21.3	32.7	4.3	9.9
EXP-OU-FI	3.5	10.5	3.5	10.5
POWER				
NOU	100.0	100.0	100.0	100.0
EXP-OU-Lévy	92.9	97.0	87.9	94.1

Note: case AJD-E-JS is the same as AJD-E but we keep only simulations in which each estimation period contains at least one jump. The rejection rates are based on 1,000 replications of one month's worth of 5-minute observations. In the construction of the test $p = 0.95$. NT indicates no truncation; T indicates truncation with $c = 1.5$.

(EST) corresponding to 78 price observations per day.

Table 4 shows simple summary statistics and the top two panels of Figure 2 show plots of the high-frequency series. The sample moments of the series as shown in Table 4 are not surprising in view of the fact that the VIX is nonnegative, positively autocorrelated, and right-skewed, together with the fact that the sample includes the very volatile year 2008. The statistics on the ratio of the daily realized variance (RV) at the 5-minute and 10-minute levels are a check on possible microstructure noise, since RV should be invariant to the sampling frequency in the absence of noise. These statistics suggest that noise is unlikely to be much of a problem but we need to be just a little guarded in interpreting the results for the S&P futures returns.

The paths of both VIX and S&P 500 index series exhibit discontinuities. We tested the null hypothesis that in each month there is at least one jump using the test of Ait-Sahalia and Jacod (2009b), where we stress our alternative is of no jumps. At the 5 percent level of significance we can reject the null of the presence of jumps in only 14 and 23 months, respectively for the VIX and the S&P 500 index.

Table 4: Summary Statistics for the Data

Statistics	VIX Index	S&P 500 Index
mean	18.26	-2.94
std	10.52	20.84
skewness	3.28	-0.89
kurtosis	15.02	28.47
5-min Autocorrelation	0.07	-0.03
quant _{.25} (RV_{10}/RV_5)	0.87	0.82
quant _{.50} (RV_{10}/RV_5)	1.00	0.94
quant _{.75} (RV_{10}/RV_5)	1.13	1.04

Note: The mean and standard deviation of the S&P Index daily returns are annualized by multiplying by 252, respectively $\sqrt{252}$, and are reported in percentage terms. The statistics on realized variation (RV) are the quartiles of the ratios of daily RV at the 10- and 5-minute frequencies.

6.1 How Active are Stock Market Volatility and Returns?

To address these questions we start by displaying the Quantile Activity Signature Function (QASF) for each series, computed as developed in Todorov and Tauchen (2010a) for the 25-th, 50-th and 75-th quantiles. The unit interval used in the computation of the ASFs, as well as the rest of the statistics based on them, is a calendar month. The QASFs for 5-minute sampling are shown in the middle panels of Figure 2 with the VIX on the left and the S&P Futures Index on the right.

The contrasts between the VIX and the S&P index QASFs are small but quite noteworthy. The median and 75-th QASFs for the VIX series on the left are just below 2.00 for powers p up to about 1.90, which would be expected for a pure jump process with a relatively high activity level around in the range 1.60–1.90 or so. On the other hand, for the S&P Index the QASF is centered right on 2.00 for powers up to 2.00, which would be expected of a process comprised of a Brownian diffusion plus jumps. These indications appear to be consistent over sampling interval, since the plots in the lower two panels of Figure 2 for the 10-minute frequency appear similar to

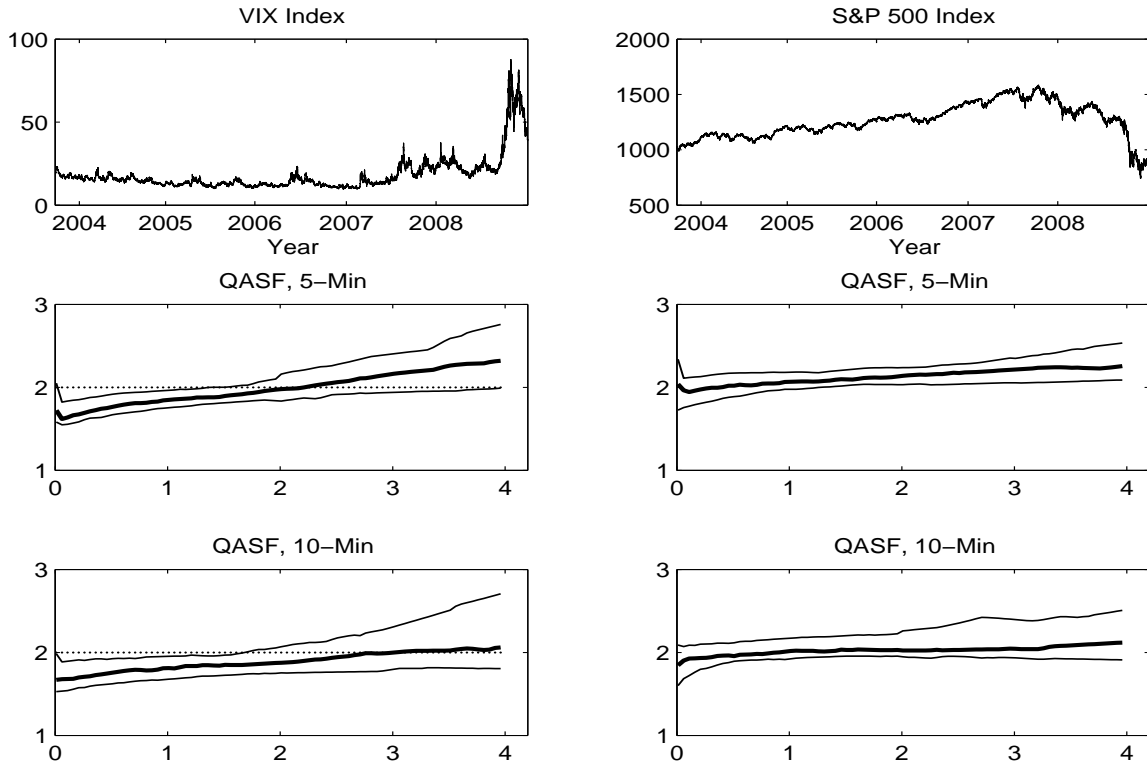


Figure 2: Activity estimation results. The left panels correspond to the VIX index and the right ones to the S&P 500 index. The top two panels plot the high-frequency data. The middle panels report $QASF$ s for 5-minute sampling frequency and the bottom panels for 10-minute sampling frequency. The $QASF$ s are computed using 64 monthly ASF estimates for the sample period September 2003 till December 2008. The quantiles that are displayed are the 25-th, 50-th and 75-th. The truncation level for both series is $c = 1.5$. The dashed lines in the two left bottom panels are straight lines at 2.

the two middle panels.

Visual impressions notwithstanding, we need to examine both the point estimates of the activity levels and the formal test for the presence or absence of a continuous component. We do this across the range of powers $p = 0.50, 0.70, 0.95$. On Figure 3 we also plot a scatter of the activity estimates, corresponding to $p = 0.95$, for the two series and all months in the sample. The left-




Figure 3: Scatter plot of the activity estimates. The estimates of the activity index correspond to $p = 0.95$ and truncation $c = 1.5$.

hand sides of Table 5 show the medians of the monthly point estimates along with the median absolute deviation about the median (MAD). The estimates indicate that the activity index for the VIX is in the range 1.73–1.83 and essentially exactly 2.00 for the S&P index; interestingly, the precision level of ± 0.10 is consistent with that found in the Monte Carlo work for this sampling frequency. The right-hand side of Table 5 shows the outcomes, i.e., the rejection rates, for the formal test for the absence of a continuous component, which is derived in Todorov and Tauchen (2010a) and based on our estimator of the activity index. The rejection rates are for three values of p between 0.50 and 1.00. The null hypothesis of the test is that the underlying process contains continuous martingale plus possibly jumps, where perforce the index is 2.00. The alternative is that the underlying process lacks a continuous martingale and the index is thereby less than 2.00, so the test is one sided. Small values of the log of the estimator relative to $\log(2.00)$

Table 5: Estimates of β_X and Tests for a Brownian Component

	p	Statistics		Rejections(%)	
		$\text{med}(\hat{\beta})$	MAD	5%	10%
VIX	0.50	1.73	0.112	57.8	62.5
	0.70	1.77	0.105	50.0	60.9
	0.95	1.83	0.107	45.3	51.6
S&P 500 Index	0.50	2.02	0.141	3.1	6.3
	0.70	2.04	0.120	3.1	3.1
	0.95	2.06	0.107	1.6	3.1

Note: The median, $\text{MAD} = \text{med}|\hat{\beta} - \text{med}(\hat{\beta})|$, and the rejection rates for the test are computed using 64 monthly estimates and tests for the sample period September 2003 till December 2008. The truncation used for both series is $c = 1.5$.

discredit the null hypothesis. In Table 5, for the VIX the test rejection indicates no continuous component in half of the periods at $p = 0.70$ with similar rejection rates for the other values of p , while for the S&P 500 Index the rejection rates always lie below the nominal significance level of the test.

Since the truncation level c used in computing $b_{X,t}(p)$ is a tuning parameter, it is essential to assess the sensitivity of our key finding regarding the activity level of the VIX index with respect to the choice of the truncation point. Until now in the empirical analysis, as in the Monte Carlo study, we have used very mild truncation corresponding to removing on average only one high-frequency increment per month. In Table 6 we report also estimation results for other choices of c that result in a much more severe truncation. As seen from the table, our findings regarding the volatility activity seem reasonably insensitive to the choice of c .

To summarize, the evidence suggests that the VIX index is a pure jump process without a continuous component and a relatively high activity index. The S&P 500 index itself, in contrast, is clearly a continuous plus jump process, which is consistent with findings in other studies regarding the characteristics of financial price indices (Todorov and Tauchen, 2010a, and the

Table 6: Robustness of estimated β_X for VIX index with respect to truncation level c .

p	Truncation $c = 0.5$		Truncation $c = 1.0$		Truncation $c = 1.5$	
	med($\hat{\beta}$)	MAD	med($\hat{\beta}$)	MAD	med($\hat{\beta}$)	MAD
0.50	1.76	0.111	1.75	0.115	1.73	0.112
0.70	1.81	0.107	1.79	0.113	1.77	0.105
0.95	1.84	0.099	1.84	0.111	1.83	0.107

Note: Notation as in Table 5. Truncation $c = 0.5$ corresponds to 3.57 standard deviations for a 5-minute intraday change in the VIX index in our sample.

references therein).

To the extent our evidence can be confirmed by future research, there would be important implications for modeling of the spot stochastic volatility process $\{\sigma_t^2\}$. First, the absence of a continuous component suggests that models such as the CGMY model are potentially plausible volatility models, and the pricing of volatility derivatives would be substantially model complicated as noted in (Cont and Tankov, 2004, Section III, pp. 245–494). Second, affine jump diffusions appear unlikely candidates for volatility, since the contrast between the top right panel of Figure 1 and the middle-left and bottom-left panels of Figure 2, together with the results in Table 5, suggest that this sort of model was unlikely to have generated the data. The same contrast appears for the other affine jump diffusion specifications of Table 1, whose QASF plots are not shown for reasons of space. Third, the pure jump models of Barndorff-Nielsen and Shephard (2001) would also be unlikely candidates. The driving Lévy process for these models must have an activity index less than unity, and the volatility series itself will have an activity index of at most unity due to the drift, which dominates, and we estimate activity levels well above unity. The most plausible class of models would seem to be the **EXP-OU-Lévy** discussed in Section 3, since these models can ensure positivity and accommodate a pure jump model with activity indices above unity, as we find in the data.

Finally, we should point out that our conclusions about the volatility modeling rely on an

estimate for the VIX index activity, which although less than 2, is nevertheless still very close to it. Therefore, our estimation results can potentially still be generated from a volatility process with a continuous martingale in it. However for this to happen, given our robustness checks of the estimation procedure, the continuous martingale should have a relatively small contribution in the power variation at the five-minute frequency (asymptotically, i.e., as we sample more frequently, the continuous martingale will eventually dominate the power variation). This is not the case for most parametric jump-diffusion volatility models used to date as we illustrated in our Monte Carlo. Thus, at the very least, our results indicate that jumps play a much more prominent role in volatility modeling.

6.2 Are Market Volatility and Price Jumps related?

Having detected the presence of jumps both in the S&P 500 index and the VIX index, a natural question arises about their dependence. We address this question in this section using the nonparametric tests developed in Jacod and Todorov (2009). Before presenting the tests and applying them to our data set, we briefly summarize previous findings based on parametric or semiparametric specifications. As mentioned in the introduction, the most commonly used model in finance which allows for jumps both in the price and the stochastic volatility is the double-jump model of Duffie et al. (2000). In their general specification, Duffie et al. (2000) allow for independent as well as dependent jumps in the index and its stochastic volatility. The studies that estimate double-jump models restrict them to arrive always together, see e.g. Chernov et al. (2003), Eraker et al. (2003). These papers, however find that the correlation between the jump sizes in the price and volatility is not statistically different from zero. On the other hand, using high-frequency data and in the context of a pure jump model for the volatility, Todorov (2009) finds strong semiparametric evidence for dependent price and volatility jumps although perfect dependence is rejected.

Determining whether the jumps in the price and volatility arrive together and if so whether

they are dependent is crucial from the perspective of successful risk management and consistent derivative pricing, see e.g. Cont and Kokholm (2009), as well as for determining the volatility and jump risk premia. Therefore, here we investigate this important question in a completely non-parametric framework. In doing so we rely on the VIX data and Theorem 1(b) linking the jump times of the VIX and the spot variance.

First, we investigate whether the jumps in the S&P 500 index and the VIX index arrive at the same time. For this, following Jacod and Todorov (2009), we use the following test statistic defined for two arbitrary processes X and Y observed over the time interval $(t-1, t)$ at frequency Δ_n

$$T_{cj}(t) = \frac{V_t(X, Y, 2, 2\Delta_n)}{V_t(X, Y, 2, \Delta_n)}, \quad (6.1)$$

where $V_t(X, Y, r, \Delta_n)$ is the following analogue of the realized power variation in a two-dimensional context

$$V_t(X, Y, r, \Delta_n) = \sum_{i=1}^{\lfloor 1/\Delta_n \rfloor} |X_{t-1+i\Delta_n} - X_{t-1+(i-1)\Delta_n}|^r |Y_{t-1+i\Delta_n} - Y_{t-1+(i-1)\Delta_n}|^r. \quad (6.2)$$

If there is common arrival of jumps in X and Y over the interval $(t-1, t]$, then this statistic converges to 1 (as $\Delta_n \rightarrow 0$), while if the jumps in the two series never arrive together the limiting value of $T_{cj}(t)$ is “around” 2. The intuition for that is that when common jumps are present then $V_t(X, Y, 2, \Delta_n)$ and $V_t(X, Y, 2, 2\Delta_n)$ converge to the same limit (which is $\sum_{s \in [t-1, t)} |\Delta X_s|^2 |\Delta Y_s|^2$). Under the alternative of no common jumps, as for the univariate results in (4.3), we will need rescaling of $V_t(X, Y, 2, \Delta_n)$ (which will depend on Δ_n) in order for it not to degenerate to zero (or infinity). For more details we refer to Jacod and Todorov (2009).

We calculated T_{cj} for each day in our sample. The median value of T_{cj} is 1.389, which is relatively close to the value of 1, corresponding to common arrival of jumps in the price and the stochastic volatility. More formally, we also conducted a formal test using T_{cj} and the testing procedure outlined in Jacod and Todorov (2009). For 5 percent significance we failed to reject the null of common arrival of jumps in 838 out of the 1212 days in the sample.

Another useful statistic that allows us to analyze cojumping in market volatility and market price level is the “realized” correlation between the squared jumps in those two series. For two arbitrary processes X and Y observed over the time interval $(t-1, t)$ at frequency Δ_n , the realized correlation is defined as

$$R_{cj}(t) = \frac{V_t(X, Y, 2, \Delta_n)}{\sqrt{V_t(X, 4, \Delta_n)V_t(Y, 4, \Delta_n)}}. \quad (6.3)$$

A value of zero of this statistic means disjoint arrival of jumps, while value close to 1 is evidence for a perfect dependence between the jumps in the two series over the given interval of time. This comes from the fact that when jumps are present we have $V_t(X, Y, 2, \Delta_n) \approx \sum_{s \in [t-1, t)} |\Delta X_s|^2 |\Delta Y_s|^2$ and $V_t(Z, 4, \Delta_n) \approx \sum_{s \in [t-1, t)} |\Delta Z_s|^4$ for $Z = X, Y$ (see Jacod and Todorov (2009) for more details).

The histogram of the (daily) realized correlation between the jumps in the S&P 500 index and the VIX index is plotted on Figure 4. As seen from the histogram, there is not only overwhelming evidence for common arrival of jumps, but also for a strong dependence between the realized jumps in the two series. This suggests that the jumps in volatility and market level should be modeled jointly. This result casts also doubt on the plausibility of empirical findings, based on affine jump diffusion models, for statistically insignificant dependence between the jump size of volatility and price jumps. Given the strong dependence between price and volatility jumps, we next explore whether the common jumps in the two series happen in the same direction. We do this by splitting $V_t(X, Y, r, \Delta_n)$ into cojump variation due to jumps in the same direction and one due to jumps in the opposite direction which we denote respectively as $V_t^+(X, Y, 2, \Delta_n)$ and $V_t^-(X, Y, 2, \Delta_n)$. The mean and the median of the ratio $\frac{V_t^-(X, Y, 2, \Delta_n)}{V_t(X, Y, 2, \Delta_n)}$ in our sample are respectively 0.921 and 0.997. Thus, almost all of the common jump variation in price and volatility is due to jumps in opposite directions. This is consistent with models generating dynamic leverage effect through jumps, e.g. Barndorff-Nielsen and Shephard (2001) and Todorov and Tauchen (2006), in which a negative price jump leads to an increase in the future volatility.




Figure 4: Histogram of daily realized correlation between price and volatility jumps.

7 Concluding Remarks

This paper shows in practical terms how to use high frequency options data (the VIX index) to make nonparametric inferences regarding the activity level of stock market volatility. The empirical implementation examines volatility dynamics using 5-min and 10-min level data on the VIX index and the S&P 500 index. The data are noisy and empirical conclusions are not unambiguously clear-cut, but nonetheless we present initial evidence suggesting a good stochastic volatility model could be one of the pure jump type whose driving jumps come from a very active Lévy process. Also, the volatility jumps and market price jumps occur in most cases at the same time and exhibit high negative dependence.

Our empirical findings, if further confirmed, can lead to several economically important conclusions. First, on an individual investor level, the pure jump dynamics of stochastic volatility would imply that hedging is quite complicated. This is in contrast with diffusive volatility dynamics in which a derivative instrument sensitive to the volatility suffices, see e.g. Liu and Pan (2003). A very active pure jump nature of volatility would mean that the volatility risk cannot be spanned with a handful of derivatives instruments. Also, the finding of strong dependence between the price and volatility jumps additionally complicates hedging. If volatility and price jumps were

independent, then the investor could use deep-out-of-the-money put options to hedge against the price jump risk and at-the-money options to hedge the volatility risk. Our findings suggest that volatility and jump risks share common origins and therefore such separate hedging cannot be expected to work well. Furthermore the two jump risks cannot be spanned with commonly traded derivative instruments, including variance swaps.

Second, on a macro level our empirical evidence has implications for the risk premia associated with price jumps and volatility risk. Typically these risk premia are modeled separately, e.g. price jump risk is modeled as a compensation for jump size risk only which is independent from the stochastic volatility. However, our results suggest that (negative) jumps on the market are associated with increase in the stochastic variance σ_t^2 and therefore at least part of the volatility risk either coincides or is highly correlated with the price jump risk. Thus, volatility and price jump risk premia share compensations for similar risks, and therefore should be modeled jointly.

8 Proof of Theorem 1

First, using e.g. Theorem V.32 in Protter (2004), we have that the vector \mathbf{f}_t is a strong Markov process. Therefore, the probability of \mathbf{f}_s under \mathbb{Q} conditioned on the filtration \mathcal{F}_t for $s > t$ is a function only of \mathbf{f}_t . Also, using the differentiability assumption on the functions $g_j^{(i)}(\cdot)$, we have that for $s \in [t, t + N]$, \mathbf{f}_s conditional on \mathcal{F}_t is a random function of \mathbf{f}_t which by Theorem V.40 in Protter (2004) is continuously differentiable. Therefore, $\mathbb{E}^{\mathbb{Q}}(\sigma_s^2 | \mathcal{F}_t)$ is a continuously differentiable function of \mathbf{f}_t for $s \geq t$ and from here we also have the continuous differentiability of $\mathbb{E}^{\mathbb{Q}}([S, S]_{t+N}^c - [S, S]_t^c | \mathcal{F}_t)$ in \mathbf{f}_t .

For the discontinuous part of the quadratic variation, using the definition of a jump compensator (see Jacod and Shiryaev (2003), Theorem II.1.8), we have that

$$\mathbb{E}^{\mathbb{Q}}([S, S]_{t+N}^d - [S, S]_t^d | \mathcal{F}_t) = \int_{\mathbb{R}} x^2 \eta(dx) \mathbb{E}^{\mathbb{Q}} \left(\int_t^{t+N} G^{(d)}(\mathbf{f}_s) \Big| \mathcal{F}_t \right),$$

and from here repeating the analysis for the continuous quadratic variation above, we have the

continuous differentiability of $\mathbb{E}^{\mathbb{Q}}([S, S]_{t+N}^d - [S, S]_t^d | \mathcal{F}_t)$ in \mathbf{f}_t as well. Hence ν_t is continuously differentiable in \mathbf{f}_t .

Part a. Given the continuous differentiability of ν_t (and the non-vanishing first derivatives of $F(\cdot)$) for an arbitrary ω in the probability space we have

$$k(\omega)V_t(\sigma^2, r, \Delta_n) \leq V_t(\nu, r, \Delta_n) \leq K(\omega)V_t(\sigma^2, r, \Delta_n), \quad t > 1, \quad r > 0, \quad (8.1)$$

for some finite constants $0 < k(\omega) \leq K(\omega)$, where we made use of the fact that the first derivatives of $G^{(c)}(\cdot)$, $G^{(d)}(\cdot)$ and $F(\cdot)$ are continuous functions of càdlàg processes and hence are locally bounded. From here, using the definition (4.1), we have the result in (4.8).

Part b. Given the monotonicity assumption on F and the fact that the sets of jump times of $f_t^{(i)}$ for $i = 1, \dots, k$ are almost surely disjoint (because of the independence of the driving Lévy processes $Z_{t_j}^{(i)}$), we have for every \mathbf{x} in the support of \mathbf{f}_t and $\mathbf{y} \in \mathbb{R}^k / \{\mathbf{0}\}$ that $F(\mathbf{x} + \mathbf{y}) \neq F(\mathbf{x})$.

□

Acknowledgements

We would like to thank the associate editor, two anonymous referees, and many conference and seminar participants for numerous suggestions and comments that significantly improved the paper.

References

- Ait-Sahalia, Y. and J. Jacod (2009a). Estimating the Degree of Activity of Jumps in High Frequency Data. *Annals of Statistics* 37, 2202–2244.
- Ait-Sahalia, Y. and J. Jacod (2009b). Testing for Jumps in a Discretely Observed Process. *Annals of Statistics* 37, 184–222.
- Andersen, T. and L. Benzoni (2007). Stochastic Volatility. Working paper, Northwestern University.
- Andersen, T., L. Benzoni, and J. Lund (2002). An Empirical Investigation of Continuous-Time Equity Return Models. *Journal of Finance* 57, 1239–1284.
- Andersen, T., T. Bollerslev, F. Diebold, and P. Labys (2008). Parametric and Nonparametric Volatility Measurement . In Y. Ait-Sahalia, L. P., and Hansen (Eds.), *Handbook of Financial Econometrics (Forthcoming)*. Amsterdam: North-Holland.

- Baillie, R., T. Bollerslev, and H. O. Mikkelsen (1996). Fractionally Integrated Generalized Autoregressive Conditional Heteroskedasticity. *Journal of Econometrics*, 3–30.
- Bandi, F. and R. Reno (2008). Nonparametric Stochastic Volatility. Working paper.
- Bandi, F. and R. Reno (2009). Nonparametric Leverage Effects. Working paper.
- Barndorff-Nielsen, O. and N. Shephard (2001). Non-Gaussian Ornstein–Uhlenbeck–Based Models and some of Their Uses in Financial Economics. *Journal of the Royal Statistical Society Series B*, 63, 167–241.
- Barndorff-Nielsen, O. and N. Shephard (2003). Realized Power Variation and Stochastic Volatility Models. *Bernoulli* 9, 243–265.
- Barndorff-Nielsen, O. and N. Shephard (2004). Power and Bipower Variation with Stochastic Volatility and Jumps. *Journal of Financial Econometrics* 2, 1–37.
- Barndorff-Nielsen, O. and N. Shephard (2007). Variation, jumps, Market Frictions and High Frequency Data in Financial Econometrics. In R. Blundell, T. Persson, and W. Newey (Eds.), *Advances in Economics and Econometrics. Theory and Applications, Ninth World Congress*. Cambridge University Press.
- Blumenthal, R. and R. Gettoor (1961). Sample Functions of Stochastic Processes with Independent Increments. *Journal of Math. Mech.* 10, 493–516.
- Bollerslev, T., T. Law, and G. Tauchen (2008). Risk, Jumps, and Diversification. *Journal of Econometrics* 144, 234–256.
- Britten-Jones, M. and A. Neuberger (2000). Option Prices, Implied Price Processes, and Stochastic Volatility. *Journal of Finance* 55, 839–866.
- Broadie, M., M. Chernov, and M. Johannes (2007). Specification and Risk Premiums: The Information in S&P 500 Futures Options. *Journal of Finance* 62.
- Carr, P. and L. Wu (2009). Variance Risk Premia. *Review of Financial Studies* 22, 1311–1341.
- Chernov, M., R. Gallant, E. Ghysels, and G. Tauchen (2003). Alternative Models for Stock Price Dynamics. *Journal of Econometrics* 116, 225–257.
- Comte, F. and E. Renault (1998). Long Memory in Continuous-Time Stochastic Volatility Models. *Mathematical Finance* 8, 291–323.
- Cont, R. and T. Kokholm (2009). A Consistent Pricing Model for Index Options and Volatility Derivatives. Technical report, Columbia University Center for Financial Engineering.
- Cont, R. and P. Tankov (2004). *Financial Modelling with Jump Processes*. Boca Raton, Florida, U.S.A.: Chapman and Hall.
- Corcuera, J., D. Nualart, and J. Woerner (2006). Power Variation of Some Integral Fractional Processes. *Bernoulli* 12, 713–735.
- Delbaen, F. and W. Schachermayer (1994). A General Version of the Fundamental Theorem of Asset Pricing. *Mathematische Annalen* 300, 463–520.
- Duffie, D., D. Filipović, and W. Schachermayer (2003). Affine Processes and Applications in Finance. *Annals of Applied Probability* 13(3), 984–1053.
- Duffie, D., J. Pan, and K. Singleton (2000). Transform Analysis and Asset Pricing for Affine Jump-Diffusions. *Econometrica* 68, 1343–1376.

- Eraker, B. (2004). Do Stock Prices and Volatility Jump?: Reconciling Evidence from Spot and Option Prices. *Journal of Finance* 59, 1367–1403.
- Eraker, B., M. Johannes, and N. Polson (2003). The Impact of Jumps in Volatility and Returns. *Journal of Finance* 58, 1269–1300.
- Haug, S. and C. Czado (2007). An Exponential Continuous Time GARCH Process. *Journal of Applied Probability* 44(4), 1–18.
- Heston, S. (1993). A Closed-Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options. *Review of Financial Studies* 6, 327–343.
- Jacod, J. and A. N. Shiryaev (2003). *Limit Theorems For Stochastic Processes* (2nd ed.). Berlin: Springer-Verlag.
- Jacod, J. and V. Todorov (2009). Testing for Common Arrivals of Jumps for Discretely Observed Multi-dimensional Processes. *Annals of Statistics* 37, 1792–1838.
- Jiang, G. and Y. Tian (2005). The Model-Free Implied Volatility and Its Information Content. *Review of Financial Studies* 18, 1305–1342.
- Liu, J. and J. Pan (2003). Dynamic Derivative Strategies. *Journal of Financial Economics* 69, 401–430.
- Merton, R. (1976). Option Pricing When the Underlying Stock Returns Are Discontinuous. *Journal of Financial Economics* 3, 125–144.
- Nelson, D. B. (1991). Conditional Heteroskedasticity in Asset Returns: A New Approach. *Econometrica* 59, 347–370.
- Protter, P. (2004). *Stochastic Integration and Differential Equations* (2nd ed.). Berlin: Springer-Verlag.
- Rogers, L. C. G. (1997). Arbitrage with Fractional Brownian Motion. *Mathematical Finance* 7, 95–105.
- Sato, K. (1999). *Lévy Processes and Infinitely Divisible Distributions*. Cambridge, UK: Cambridge University Press.
- Shephard, N. (2005a). General Introduction. In N. Shephard (Ed.), *Stochastic Volatility*. Oxford: Oxford University Press.
- Shephard, N. (2005b). *Stochastic Volatility*. Oxford: Oxford University Press.
- Todorov, V. (2009). Econometric Analysis of Jump-Driven Stochastic Volatility Models. *Journal of Econometrics*, forthcoming.
- Todorov, V. and G. Tauchen (2006). Simulation Methods for Lévy -Driven CARMA Stochastic Volatility Models. *Journal of Business and Economic Statistics* 24, 455–469.
- Todorov, V. and G. Tauchen (2010a). Activity Signature Functions for High-Frequency Data Analysis. *Journal of Econometrics* 154.
- Todorov, V. and G. Tauchen (2010b). Limit Theorems for Power Variations of Pure-Jump Processes with Application to Activity Estimation. *Annals of Applied Probability*, forthcoming.
- Wu, L. (2010). Variance Dynamics: Joint Evidence from Options and High-Frequency Returns. *Journal of Econometrics*, forthcoming.