

NBER WORKING PAPER SERIES

HOW CRASHES DEVELOP:
INTRADAILY VOLATILITY AND CRASH EVOLUTION

David S. Bates

Working Paper 22028
<http://www.nber.org/papers/w22028>

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
February 2016

I am grateful for comments on earlier versions of the paper from seminar participants at Iowa, Northwestern, Houston, and Lugano, and from conference participants at the 2012 IFSID Conference on Structured Products and Derivatives and McGill University's 2014 Risk Management Conference. The University of Iowa provided research support for this project. There were no relevant nor material external financial relationships that could have influenced this research. The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2016 by David S. Bates. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

How Crashes Develop: Intradaily Volatility and Crash Evolution
David S. Bates
NBER Working Paper No. 22028
February 2016
JEL No. C22,G13

ABSTRACT

This paper explores whether affine models with volatility jumps estimated on intradaily S&P 500 futures data over 1983-2008 can capture major daily outliers such as the 1987 stock market crash. I find that intradaily jumps in futures prices are typically small, and that self-exciting but short-lived volatility spikes capture intradaily and daily returns better. Multifactor models of the evolution of diffusive variance and jump intensities improve fits substantially, including out-of-sample over 2009-13. The models capture reasonably well the conditional distributions of daily returns and of realized variance outliers, but underpredict realized variance inliers.

David S. Bates
Henry B. Tippie College of Business
Department of Finance
University of Iowa
Iowa City, IA 52242-1000
and NBER
david-bates@uiowa.edu

What is a crash? In the jump-diffusion model of Merton (1976), a crash is a rare event – a single adverse draw from a Poisson counter, with a vanishingly small probability of multiple adverse draws within a single day. While this model may be successful at capturing outliers in daily returns, it does not appear to capture the intradaily evolution of major market downturns. The 23% drop in the S&P 500 index on Monday, October 19, 1987 from the preceding Friday’s closing level did not occur within five minutes, for instance; it took all day to achieve the full decline. Indeed, papers such as Tauchen and Zhou (2011) that use bipower variation to decompose realized variance into diffusive and jump components suggest there were no jumps at all on October 19. Instead, it was a draw of two standard deviations from a day that happened to have an unusually high intradaily realized volatility of 12%.

While the increasing availability of high-frequency data generated some exploration of intradaily volatility evolution, including in stock markets, there has been little direct estimation of stochastic processes with stochastic volatility and jumps using intradaily data. Papers such as Andersen and Bollerslev (1997) focus on volatility dynamics; in particular, reconciling GARCH-based volatility evolution estimates from daily versus intradaily data. As described in Andersen (2004), the recognition that realized variance effectively summarizes intradaily volatility information *and* sidesteps the issues of fitting pronounced diurnal volatility patterns and announcement effects shifted the focus of most intradaily research to realized variance. Whether jumps are important has been assessed indirectly in this literature, using the bipower variation approach of Barndorff-Nielsen and Shephard (2004, 2006) to assess intradaily jump contributions to realized variance. That approach maintains the Merton (1976) presumption that jumps are rare.

This indirect evidence and more direct parametric estimates by Stroud and Johannes (2014) on intradaily data indicate a fundamental mismatch between jump magnitudes from intradaily versus from daily stock market data. Bates (2006, 2012), for instances, estimates daily jump standard deviations at about 3%, excluding the 1987 crash. Intradaily jump standard deviations inferred from bipower variation are an order of magnitude smaller: about 0.5% in Tauchen and Zhou. Stroud and Johannes estimate comparable intradaily magnitudes: 0.2% - 0.4% for the standard deviation of unexpected jumps in 5-minute returns, and similar magnitudes for predictable announcement effects. Such estimates require multiple intradaily jumps to

explain the occasional 4-10% daily return magnitudes, which is inconsistent with the rare-jump hypothesis.

This paper develops an affine multifactor multifactor stochastic volatility/jump model for the intradaily evolution of S&P 500 futures that potentially reconciles intradaily and daily results. The key feature of the model is “self-exciting” synchronous and correlated jumps in intradaily stock returns and volatility – essentially a stochastic-intensity and multifactor generalization of the Duffie, Pan and Singleton (2000) volatility-jump model. Every intradaily jump substantially increases the probability of more intradaily jumps in volatility and in returns; and these can accumulate into the major outliers in daily returns that we occasionally observe. The model is estimated on 15-minute S&P 500 futures returns over 1983-2008 via the Bates (2006, 2012) approximate maximum likelihood (AML) filtration methodology, taking into account special features of intradaily futures data. The methodology provides direct estimates of the frequency and distribution of intradaily jumps, as well as estimates of the intradaily evolution of volatility and jump intensities. The 2009-2013 period is used for out-of-sample tests of the model.

While assorted nonaffine continuous-time models have been proposed and estimated predominantly on daily stock market data,¹ affine models have a couple of advantages for exploring how crashes develop. First, affine models are closed under time aggregation. An affine model fitted to intradaily returns generates an affine model for daily returns, which can then be tested against observed daily data.

Second, affine models are well suited to the current research focus on realized variance. Realized variance asymptotically approaches quadratic variation at higher sampling frequencies; and quadratic variation is an affine random variable in affine models with stochastic volatility and jumps. This property holds for any affine model; standard assumptions in the bipower variation literature such as no leverage effects and rare large jumps are unnecessary. This paper derives the conditional characteristic function of daily quadratic variation for the proposed

¹Chernov et al (2003) estimate diffusive log variance models on daily stock market data. The models have substantial volatility feedback, which is similar to the self-exciting jumps feature here. Stroud and Johannes (2014) have jump-diffusive models for intradaily log variance, but with constant jump intensities rather than self-exciting jumps. Calvet and Fisher (2008) propose a tightly parameterized Markov Chain model for daily log variance that also lacks volatility feedback. Aït-Sahalia et al (2015) and Fulop et al (2015) have affine models with stochastic volatility and self-exciting volatility jumps, which they estimate on daily stock market data.

intradaily model, assesses how well it approximates the conditional characteristic function of daily realized variances from 15-minute returns, and uses it for various diagnostics of those realized variances.

I explore three issues. First is the specification issue of identifying the appropriate time series model, using an extensive 1983-2008 history of S&P 500 futures that includes extreme stock market movements in October 1987 and in the fall of 2008. I find that multifactor models with self-exciting but short-lived volatility spikes substantially improve model fits both in-sample and out-of-sample.

Second is the issue of time aggregation: how well do various proposed affine models estimated using 15-minute returns capture the statistical properties of daily returns and daily measures of intradaily realized variance? Any estimation methodology applied to 15-minute returns is perforce focused on fitting the conditional distributions of those data. Not all high-frequency phenomena remain important at longer horizons, however. Examining how well the models fit daily returns and realized variances are major goals of this paper, and serve as additional tests of model specification.

Third, I explore the informational content of realized variance, which is a noisy signal of intradaily variance when intradaily jumps are present. A priori, directly estimating (filtering) underlying volatility from the full set of observed intradaily returns should be more accurate than using any summary statistics such as realized variance or bipower variation based upon those same data. However, directly using intradaily returns involves significant costs, both in time and in requiring explicit modeling of phenomena such as minimum tick sizes and diurnal variance fluctuations. I therefore examine the magnitude of the informational loss when filtering latent volatility solely from observed daily realized variance, within the framework of explicit affine parametric models and a filtration methodology that identify how such filtration should be conducted.

Section 1 of the paper describes the intradaily and overnight data, the multifactor models and estimation methodology, and how well the models fit. Section 2 contains additional diagnostics using intradaily realized variance, while Section 3 concludes. Overall, the multifactor affine models with volatility spikes do a reasonably good job of matching the

properties of intradaily and daily S&P 500 futures returns, especially as more factors are added. Furthermore, the most general 3-factor model captures the occasionally extreme observations of realized variance reasonably well; and those are when extreme daily stock market returns occur. The models do underpredict the frequency of small realized variance observations, indicating that some specification error remains.

1. Data and models

1.1. Data

S&P 500 futures began trading at the Chicago Mercantile Exchange (CME) on April 21, 1982, using the open-outcry pit trading prevailing at the CME at that time for all futures contracts. Initial trading hours were 9 AM to 3:15 PM Central Standard Time, with CME pit trading typically extending 15 minutes beyond trading at the New York Stock Exchange (NYSE).² On September 30, 1985, the NYSE and CME shifted the opening time by a half-hour, to 8:30 AM CST. Starting in December 1990, both the NYSE and CME instituted fewer trading hours on trading days adjacent to Christmas, the Fourth of July, and Thanksgiving.

In 1992, the CME introduced after-hours electronic trading through its Globex trading platform. In 1997, the CME introduced “E-mini” S&P 500 futures contracts, which were 1/5th the size of regular S&P 500 futures contracts and traded exclusively on Globex, including during the day. Trading activity switched increasingly to electronic trading via Globex, which accounted for 84% of CME group volume by 2011.³

The CME provides data in two formats: “End-of-Day” daily summaries and “Time and Sales” data. The former contains open, high, low, close, and settlement prices, as well as volume and open interest, while the latter contains the time and price of every daily transaction in which the price changed from the previous transaction. Bid and ask prices are also recorded when the bid price is above or the ask price is below the price of the previous transaction. No information is provided regarding the volume of transactions at a particular price. I acquired both sets of data for the original full-sized S&P 500 futures contract for January 3, 1983 through December 31,

² The CME and NYSE closed at the same time over October 23 through November 6, 1987, in the aftermath of the 1987 stock market crash.

³ CME Group, “Twenty Years of CME Globex,” June 21, 2012. (<http://www.cmegroup.com/education/files/globex-retrospective-2012-06-12.pdf>).

2013. I discarded the bid and ask data, and also transactions that were subsequently cancelled. The data over 1983-2008 are used for parameter estimation, whereas those over 2009-2013 are used for out-of-sample testing.

S&P 500 futures contracts typically mature on the third Friday of March, June, September and December – except for March 2008 contracts, which matured a day earlier because of Good Friday. Of the available maturities, I selected the shortest maturity with nine or more days until the third Friday, that being the most actively traded contract according to the “End of Day” volume data. For instance, I used data for March 1983 futures maturing on March 18, 1983 up to the close of trading on Wednesday, March 9. Prices of June 1983 contracts were then used for computing overnight futures returns from Wednesday to Thursday, and for intradaily and overnight returns from Thursday, March 10 to the close of trading on Wednesday, June 8.

I constructed intradaily 15-minute log-differenced futures prices broadly along the lines of the 5-minute futures returns in Chan et al (1991) and Andersen and Bollerslev (1997), by taking the last observed future price in every 15-minute interval. 15-minute returns were used instead of 5-minute primarily to triple optimization speed on this large data set, but also to span some of the short-duration price limit constraints on futures returns over 1989-2002 that are discussed below. For after-hours trading, I extended the time window by one minute, typically up through 3:16 PM CST, because trades were often recorded shortly after trading had officially ended. Overnight futures returns were constructed as the difference between the log of the futures price at the end of the first 15-minute trading session (typically 8:30 – 8:45 AM CST) and the log of the futures price of comparable maturity at the end of the preceding day. This approach allows some time for the incorporation of overnight news into the futures price, during the especially volatile initial fifteen minutes of trading. Furthermore, skipping the first 15 minutes when computing overnight returns spans possible constraints on opening futures prices from the price limit system instituted by the CME in 1989. The final data set consists of 6,557 overnight returns and 168,297 intradaily returns, from December 31, 1982’s closing price through December 31, 2008.

The CME’s price limit system was created in response to the stock market crash of October 19, 1987, and paralleled the NYSE’s “circuit breaker” system. The CME’s price limits

typically involved four prespecified constraints: a relatively tight initial band relative to the previous day's short-maturity S&P 500 futures settlement price that temporarily constrained both upward and downward moves at the open, and three progressively lower price limits (Levels 1 through 3) that temporarily constrained downward moves during the day. Hitting a price limit triggers a specific time interval during which the price limit remains in effect, followed by a two-minute trading halt if the price limit is binding ("locked limit"), followed by the resumption of trading with a new and lower price limit in effect. For instance, price change constraints over 1989-1996 were 5, 12, 20, and 30 points, respectively, over which time the S&P 500 index rose from 300 to 700. The 5-point opening limit lasted only until 8:40, with a trading halt from 8:40 to 8:42 if binding. A 12-point downward move triggered an interval lasting 30 minutes or until 2:30 PM CST, during which time futures contracts could be traded at or above the limit but not below. If the price limit was binding at the end of the interval, a two-minute trading halt was declared, followed by the resumption of trading with the 20-point lower limit in place. Hitting the 20-point downward limit started another interval of 30 minutes or until 2:30, followed by a trading halt (if still binding) and the 30-point limit taking effect for the remainder of the day. Separate rules applied if the futures price was locked limit at the close of the preceding day.

Limits on the opening price change were removed on October 15, 1997, as part of a revision in the circuit breaker system. The levels of permissible downward price changes changed over time, partly because of the rise and fall in the level of the S&P 500. In addition, the price constraints were significantly widened on May 13, 2001 to roughly 5%, 10%, and 15% of the end-of-quarter S&P 500 level, with roughly 20% being the maximal permissible daily price movement. These limits were relaxed further in January 2008 to 10%, 20% and 30%, respectively, to be consistent with the Dow Jones-based percentages used on the NYSE. (A 5% price limit remained on after-hours trading on Globex.) These revisions in 2001 and 2008 considerably reduced the frequency of trading halts, with an apparent absence of such halts in intradaily data over 2003-2008.⁴ Even the 10.4% intradaily drop in the December 2008 futures price on October 15, 2008, to 898.5 from October 14's settlement price of 1002.3, did not exceed

⁴John Nyhoff at the CME kindly provided me with a list of dates and times when the S&P 500 futures price limits were hit. Stroud and Johannes (2014) note that the GLOBEX-traded E-mini contract hit its overnight limit on October 24, 2008.

the 120-point limit that the CME had set on September 30, 2008 as the 10% limit for the fourth quarter of 2008.

The corresponding time intervals triggered by hitting a price limit were also shortened to roughly 15 minutes in 1996, and to roughly 10 minutes in 1997. However, longer halts could ensue contingent upon trading halts at the NYSE. The major such incident was the tripping of all three levels of NYSE circuit breakers on October 27, 1997.

Trading at the CME also halted sometimes because of exogenous events in New York or Chicago. For instance, the explosion of a Con Edison transformer in New York delayed the open of the NYSE and CME on December 27, 1990. On April 13, 1992, the accidental flooding of utility tunnels in Chicago shut down the CME, but not the NYSE.⁵ Both the NYSE and CME closed following the attacks on September 11, 2001, and did not reopen until September 17.

Price limits can artificially constrain observed futures returns. Consequently, I extended the time interval whenever a price limit was hit until that limit had expired and was no longer potentially binding on the futures price. Similarly, I extended time intervals whenever trading was suspended until trading had resumed, and computed returns over the expanded interval. For instance, the CME suspended trading at 11:15 AM CST on October 20, 1987, the day after the 1987 stock market crash;⁶ and the December 1987 futures price rebounded when the trading resumed at 12:05 PM. I combined returns over 11:15 AM to 12:15 PM into a single one-hour interval, with an associated log-differenced futures price of 13.76%.⁷ There were 103 instances over 1983-2008 of expanded time intervals, out of 174,859 total observations.

The gravitation of trading activity to the smaller “E-mini” Globex-traded S&P futures contract did affect the quality of the data on the larger S&P 500 contract used for out-of-sample testing over 2009-13. In particular, 15-minute intervals without transactions occurred increasingly frequently: 5 in 2011, 37 in 2012, and 166 in 2013. These cases were treated as missing observations, with the return horizon expanded until a transaction was recorded.

⁵ A list of market closings at the NYSE is available at <http://www.nyse.com/pdfs/closings.pdf>. Additional CME trading suspensions were identified by searching the data for periods without reported trades, and matching them against news reports.

⁶ See Carlson (2006, p.11).

⁷ Andersen and Bollerslev (1997) follow Chan et al (1991) in omitting data over October 15 through November 13, 1987. They also linearly interpolate prices when data are missing, yielding roughly identical successive 5-minute returns.

1.2. Models

I model the continuous-time process for the log futures price $f_t = \ln F_t$ underlying observed intradaily and overnight S&P 500 futures returns as a potentially multifactor affine jump-diffusion of the form

$$\begin{aligned}
 df_t &= \mu_0 dt + \sum_{i=1}^I [(\mu_i - \frac{1}{2})V_{it} dt + \sqrt{V_{it}} dW_{it}] + \sum_{j=1-K}^J (\gamma_j dN_{jt} - \lambda_{jt} \bar{k}_j dt) \\
 dV_{1t} &= (\alpha - \beta_1 V_{1t}) dt + \sigma \sqrt{V_{1t}} dW_{Vt} + \gamma_{V1} dN_{1t} \\
 dV_{it} &= -\beta_i V_{it} dt + \gamma_{Vi} dN_{it} \text{ for } i > 1,
 \end{aligned} \tag{1}$$

where W_{it} and W_{Vt} are Wiener processes with correlation ρ , the N_{jt} 's are Poisson counters with stochastic intensities λ_{jt} that depend linearly upon the variance state variables, and $\bar{k}_j = E e^{\gamma_j} - 1$ is the expected percentage jump size.

A key empirical issue is how many state variables are necessary to adequately summarize the evolution of futures prices. Potential specifications are categorized as SVJ(I, J, K) models, where I is the number of underlying state variables $\mathbf{V}_t = (V_{1t}, \dots, V_{It})'$ that determine diffusive variance and jump intensities; J is the number of synchronous jump processes for futures prices and variance state variables; and $K \leq I$ is the number of jump processes for futures prices only.

Furthermore, the underlying state variables are divided into a core jump-diffusive variance state variable V_{1t} and additional pure-jump state variables V_{it} that capture transient variance shocks. Log futures jump distributions are modeled as in Duffie, Pan and Singleton (2000; henceforth **DPS**) as $\gamma_j = \rho_j \gamma_{Vj} + \gamma_{fj}$, where $\gamma_{Vj} \sim \text{Exp}(\bar{\gamma}_{Vj})$ is the exponentially distributed jump in spot variance conditional upon $dN_{jt} = 1$, $\gamma_{fj} \sim N(\bar{\gamma}_{fj}, \delta_j^2)$ is an independent normal shock, and ρ_j captures the degree to which synchronous jumps in futures prices and variance covary. The Poisson counter N_{0t} identifies futures price jumps (with a V_{1t} -sensitive jump intensity) that are unaccompanied by volatility jumps – an extension that substantially improves overall fits.

The multifactor specification for \mathbf{V}_t in equation (1) allows considerable flexibility in how S&P futures and total diffusive spot variance can evolve. Whereas univariate models such as DPS must have invariant jump distributions to be affine, multivariate models allow the distributions of jumps in stock market returns and in total diffusive spot variance to vary as the components of \mathbf{V}_t change. Those components can have different degrees of persistence, and different correlations with contemporaneous stock market returns.

Affine models such as (1) imply that the joint cumulant generating function of returns $y_{t+\tau} = f_{t+\tau} - f_t$ and future spot variances $\mathbf{V}_{t+\tau}$ conditional upon current \mathbf{V}_t is affine in \mathbf{V}_t ,

$$\begin{aligned} CGF(\Phi, \boldsymbol{\psi} | \mathbf{V}_t, \tau) &= \ln E[e^{\Phi y_{t+\tau} + \boldsymbol{\psi}' \mathbf{V}_{t+\tau}} | \mathbf{V}_t] \\ &= C(\tau; \Phi, \boldsymbol{\psi}) + \mathbf{D}(\tau; \Phi, \boldsymbol{\psi})' \mathbf{V}_t. \end{aligned} \quad (2)$$

$C(\tau; \cdot)$ and $\mathbf{D}(\tau; \cdot)$ solve a system of ordinary differential equations, with known analytical solutions for two broad cases: the stochastic-intensity models of Bates (2000) without variance jumps, and the constant-intensity model of Duffie, Pan, and Singleton (2000) with variance jumps. The more general specifications in which volatility jumps affect the intensity of future jumps – e.g., the “self-exciting” volatility-jump models of Carr and Wu (2008) and Andersen, Fusari and Todorov (2015) – lack known analytical solutions for $C(\tau; \cdot)$ and $\mathbf{D}(\tau; \cdot)$, which must be evaluated numerically.

Self-exciting jumps are plausibly a major explanation for the intradaily development of major daily outliers. I consequently use intermediate volatility-jump models with jump intensities

$$\lambda_{jt} = \begin{cases} \lambda_{j1} V_{1t^*} & \text{for } j \leq 1 \\ \lambda_{j0} + \lambda_{j1} V_{1t^*} + \lambda_{jj} V_{jt^*} & \text{for } j > 1. \end{cases} \quad (3)$$

The jump intensities λ_{jt} for each Poisson counter N_{jt} are assumed constant within each intradaily or overnight period, reset at the end of every period, and depend upon the spot variance levels \mathbf{V}_{t^*} prevailing at the start of the period. This allows jumps to be self-exciting across periods, while retaining the analytic tractability of the constant-intensity DPS model. The recursive structure in (3) for $j > 1$ allows the initiation of additional self-exciting volatility components V_{jt} to depend upon the level of core volatility V_{1t} .

By iterated expectations, the joint conditional cumulant generating function of future returns $y_{t+\tau} = f_{t+\tau} - f_t$ and future spot variances $\mathbf{V}_{t+\tau}$ conditional upon observing past data $\mathbf{Y}_t = \{y_1, \dots, y_t\}$ is of the form

$$\begin{aligned} CGF(\Phi, \boldsymbol{\psi} | \mathbf{Y}_t, \tau) &= \ln E\{\exp[\Phi y_{t+\tau} + \boldsymbol{\psi}' \mathbf{V}_{t+\tau}] | \mathbf{Y}_t\} \\ &= C(\Phi, \boldsymbol{\psi}; \tau) + g_{t|t}[\mathbf{D}(\Phi, \boldsymbol{\psi}; \tau)] \end{aligned} \quad (4)$$

where $g_{t|t}(\boldsymbol{\psi}) \equiv \ln E[e^{\boldsymbol{\psi}' \mathbf{V}_t} | \mathbf{Y}_t]$ is the cumulant generating function of \mathbf{V}_t conditional upon data \mathbf{Y}_t . The precise functional forms of $C(\cdot)$ and $\mathbf{D}(\cdot)$ for the SVJ(I, J, K) models are given in the Appendix.

1.3 Rounding models and filtration

Because S&P 500 futures had a price tick size of 0.05 through October 31, 1997, and 0.10 thereafter, observed intradaily futures returns are not drawn from a continuous distribution. Indeed, 5% of all intradaily 15-minute returns over 1983-2008 are exactly zero despite intervening intraperiod price changes,⁸ while 42% are over a range of ± 5 price ticks. Consequently, the above continuous-time models are assumed to represent the underlying conditional distribution of futures returns; and those returns are rounded to the futures returns actually observed.⁹ The scaled probability of an observed log futures return $y_{t+\tau} = \ln(F_{t+\tau}/F_t)$ over a horizon of length τ is computed via Fourier inversion as the integrated conditional density of all future realizations falling within $\pm 1/2$ price tick of the observed $F_{t+\tau}$:

$$\begin{aligned} P_{t+\tau} &= \frac{\text{Prob}\left[y_{t+\tau} \in \left(\underline{y}_{t+\tau}, \bar{y}_{t+\tau}\right) | \mathbf{Y}_t\right]}{\bar{y}_{t+\tau} - \underline{y}_{t+\tau}} \\ &= \frac{1}{2\pi} \int_{-\infty}^{\infty} e^{C(i\Phi, 0; \tau) + g_{t|t}[\mathbf{D}(i\Phi, 0; \tau)]} \frac{e^{-i\Phi \bar{y}_{t+\tau}} - e^{-i\Phi \underline{y}_{t+\tau}}}{\bar{y}_{t+\tau} - \underline{y}_{t+\tau}} d\Phi \end{aligned} \quad (5)$$

for $\underline{y}_{t+\tau} = \ln\left(\frac{F_{t+\tau} - 1/2\varepsilon}{F_t}\right)$, $\bar{y}_{t+\tau} = \ln\left(\frac{F_{t+\tau} + 1/2\varepsilon}{F_t}\right)$, and $\varepsilon =$ one price tick.

Parameters are then estimated by maximizing the log likelihood function $\ln L = \sum_t \ln P_{t+\tau}$, using the AML filtration methodology discussed below for recursive updating of $g_{t|t}(\cdot)$. The

⁸ There were on average 30 price changes per period for those intradaily periods with a futures return of zero.

⁹ See Campbell, Lo, and MacKinlay (1997, Section 3.3.2) for an overview of rounding models.

computed probabilities are scaled by $\bar{y}_{t+\tau} - \underline{y}_{t+\tau}$ to make the results comparable in magnitude to log likelihood values based upon conditional probability densities.

The conditional moment generating function $G_{t|t}(\boldsymbol{\psi}) = \exp[g_{t|t}(\boldsymbol{\psi})]$ that summarizes what is known about \mathbf{V}_t at time t can be updated recursively over time via a straightforward extension of the Approximate Maximum Likelihood (AML) methodology in Bates (2006, 2012):

$$G_{t+\tau|t+\tau}(\boldsymbol{\psi}) = \frac{1}{2\pi P_{t+\tau}} \int_{-\infty}^{\infty} e^{C(i\Phi, \boldsymbol{\psi}; \tau) + g_{t|t}[D(i\Phi, \boldsymbol{\psi}; \tau)]} \frac{e^{-i\Phi \bar{y}} - e^{-i\Phi \underline{y}}}{\bar{y} - \underline{y}} d\Phi. \quad (6)$$

Derivatives of Eq. (6) provide the noncentral posterior moments of $\mathbf{V}_{t+\tau}$ conditional upon adding the latest datum $y_{t+\tau}$ to the data set :

$$E[V_{i,t+\tau}^n | \mathbf{Y}_{t+\tau}] = \left. \frac{\partial^n G_{t+\tau|t+\tau}(\boldsymbol{\psi})}{\partial \psi_i^n} \right|_{\boldsymbol{\psi}=0}. \quad (7)$$

As in Bates (2006, 2012), the distribution of the latent V_{it} 's conditional upon information \mathbf{Y}_t are modeled as independent and conditionally gamma, with joint conditional cumulant generating function

$$g_{t|t}(\boldsymbol{\psi}) = - \sum_{i=1}^I v_{it} \ln(1 - \kappa_{it} \psi_i). \quad (8)$$

For an additional datum $y_{t+\tau}$, Equations (5) to (7) are used to update the posterior moments

$$\begin{aligned} E[V_{i,t+\tau} | \mathbf{Y}_{t+\tau}] &\equiv \hat{V}_{i,t+\tau|t+\tau} = \kappa_{i,t+\tau} v_{i,t+\tau} \\ \text{Var}[V_{i,t+\tau} | \mathbf{Y}_{t+\tau}] &\equiv P_{i,t+\tau|t+\tau} = \kappa_{i,t+\tau}^2 v_{i,t+\tau}. \end{aligned} \quad (9)$$

From these are computed the parameters $(\kappa_{i,t+\tau}, v_{i,t+\tau})$ for $i = 1, \dots, I$ that summarize the conditional CGF $g_{t+\tau|t+\tau}(\boldsymbol{\psi})$ for $\mathbf{V}_{t+\tau}$. The algorithm is initiated using moment-matching unconditional gamma distributions for the initial \mathbf{V}_0 .

1.4 Intradaily and overnight seasonals

The effective length of any time interval of course differs for overnight and intradaily 15-minute returns. In addition, there are intradaily variations in trading activity and volatility, as well as day-of-the-week effects for intradaily and overnight effects. Finally, the actual length of a given trading day occasionally varies because of late openings or early closings – especially the half-day trading that began in December 1990 on specific days adjacent to Christmas, the Fourth of July, and Thanksgiving.

Wednesday was arbitrarily selected as the benchmark day, with Tuesday close to Wednesday close representing one full business day. The effective division between overnight returns (including the first 15 minutes of intradaily market trading) and intradaily returns (for the remainder of the day, typically until 3:16 PM CST) was estimated as variance proportions $(1 - f_{daily}, f_{daily})$, respectively. Calendar time t increases on day n , time segment m by a daily amount $252\tau_{nm}$ constructed as follows:

$$252\tau_{nm} = \begin{cases} (1 - f_{daily})\exp(\boldsymbol{\beta}_{ON}'\mathbf{d}_n^{ON}) & \text{overnight } (m = 0) \\ f_{daily} \frac{\exp[f(m, M_n)]}{\sum_{m=1}^{M_n} \exp[f(m, M_n)]} \exp(\boldsymbol{\beta}_{ID}'\mathbf{d}_n^{ID}) & \text{intradaily } (m > 0) \\ 5 - f_{daily} \exp(\beta_{Monday}) & \text{for Sept. 10 -17, 2001 (close to open)} \end{cases} \quad (10)$$

where \mathbf{d}_n^{ON} are day-of-the-week, holiday, and weekend dummies for overnight returns; \mathbf{d}_n^{ID} are day-specific dummies for intradaily returns, including a half-day indicator for holiday-adjacent shortened trading days; and M_n is the number of 15-minute segments available on day n after the opening segment.

The function $f(m, M_n)$ follows Andersen and Bollerslev (1997) in using Gallant's (1981) flexible Fourier form approach to estimate the intradaily variance pattern:

$$f(m, M_n) = b_1 \frac{m}{Avg_n(m)} + b_2 \frac{m^2}{Avg_n(m^2)} + b_{ah} 1(m = ah) + \sum_{p=1}^2 \left[c_p \cos\left(2\pi p \frac{m}{M_n}\right) + d_p \sin\left(2\pi p \frac{m}{M_n}\right) \right] \quad (11)$$

where $Avg_n(m) = (M_n + 1)/2$, $Avg_n(m^2) = (M_n + 1)(M_n + 2)/6$, and $1(m = ah)$ is an indicator of the after-hours trading segment, typically at $m = M_n$.¹⁰

As discussed above, there were 102 occasions over 1989-2002 when intradaily futures returns were constrained by endogenous price limits or exogenous market closings. In such cases, I aggregated log futures returns over additional periods until the trading constraint or market closing was no longer in effect, and treated the yearly time interval associated with that aggregate log return as the sum of the spanned τ_{nm} 's. This approach is equivalent when filtering volatility to treating the subintervals of that full interval as missing observations.

Table 1 contains the Vjump1b estimates of the parameters affecting intradaily and overnight variance patterns, while Figure 1 shows the composite effect. Results from other models were almost identical. Overnight trading (including trading in the first 15 minutes of Wednesday morning) typically accounts for 18.7% of the return variance from Tuesday close to Wednesday close. Intradaily trading over the remainder of Wednesday accounts for the remaining 81.7%. These patterns are roughly comparable for other days of the week. As shown in Figure 1 and in the Table 1 estimates, there is a U-shaped weekly pattern to overnight return variance and an inverse U-shaped pattern to intradaily variance, with the sum of the two imparting a roughly flat but slightly increasing trend in business-day variance across the week. Holidays during the week or as part of a holiday weekend substantially increase the variance of "overnight" returns that span those holidays.

Figure 2 shows the intradaily variance pattern for separate 15-minute returns over 8:45 AM through 3:16 PM, excluding the initial 8:30 – 8:45 interval. The roughly U-shaped intradaily pattern replicates Andersen and Bollerslev's (1997) Figure 6b, which was estimated on 5-minute S&P 500 futures returns using a GARCH model. Variance peaks in the afternoon at 2:15 – 2:30 PM CST and falls off thereafter, especially in the after-hours 3 – 3:15 PM trading segment. Andersen and Bollerslev emphasize that it is critical to account for this periodic intradaily pattern, which if ignored would strongly affect estimates of variance mean reversion. Accounting for this intradaily periodicity is also important when estimating jump risk. Substantial distributional mixing is occurring, with opening and closing returns having roughly

¹⁰ There was no after-hours session on October 23 through November 6, 1987.

double the variance (40% higher volatility) of mid-day returns. Failing to account for this distributional mixing would exaggerate the conditional leptokurtosis of 15-minute returns, and boost the estimated magnitudes of intradaily jump risk.

1.4 Volatility and jump parameter estimates

Table 2 contains estimates of the other jump-diffusion parameters that describe how diffusive spot variances V_t evolve, and the distributions of jumps. Five models are considered, all sequentially nested:

Model	Variance processes	Jumps in log futures prices
SVJ1	One-factor diffusive process for V_{1t}	Normally distributed jumps, with V_{1t} -dependent jump intensity.
Vjump1a	One-factor jump-diffusion	Jumps are correlated with V_{1t} jumps, and have a V_{1t} -dependent jump intensity.
Vjump1b	One-factor jump-diffusion	Two jumps: one uncorrelated with V_{1t} jumps, one correlated. Both have V_{1t} -dependent intensities.
Vjump2	Two-factor additive variance process. V_{1t} follows a jump-diffusion, while V_{2t} is a pure-jump volatility spike process.	Three jumps, with the last two correlated with V_{1t} and V_{2t} jumps, respectively.
Vjump3	Three-factor additive variance process. V_{1t} follows a jump-diffusion, while V_{2t} and V_{3t} are pure-jump volatility spike processes.	Four jumps, with the last three correlated with V_{1t} , V_{2t} , and V_{3t} jumps, respectively.

Comparing the models in Table 2 and in Figure 3 indicates two important sources of overall improvement in fit. First is the issue of accurately modeling high-frequency price changes of 0-5 ticks, which account for 42% of the data. Of primary importance is the introduction in model Vjump1b of a second futures jump component N_{0t} that is uncorrelated with variance jumps. This futures jump component is estimated as a high-frequency small-amplitude jump component in Vjump1b, and is even higher in frequency and smaller in amplitude for the more general Vjump2 and Vjump3 models (Table 2, Panel C). While observationally equivalent for daily returns to just scaling up diffusive variance V_{1t} , the N_{0t} futures jump component helps model the 15-minute price changes of 0-5 ticks more precisely.

Second is the impact of the 1987 crash. Better modeling of intradaily developments on October 19 and 20 accounts for almost a quarter of the full-sample improvement in log likelihood of the most general Vjump3 model versus SVJ1. The more general models also fit the fall of 2008 progressively better, indicating the two largest stock market crises over 1983-2008 are substantially influencing parameter estimation. Turbulent markets at other times (the first Gulf War, fall 1997, fall 1998, post-9/11/2001) do not appear to have had as distinct an impact on log likelihood fits. However, log likelihoods are generally increasing throughout the full 1983-2008 sample.

The more general models also have better fits over the out-of-sample 2009-13 period. The progressively better fits of the Vjump1b, Vjump2 and Vjump3 models out-of-sample are entirely attributable to better modeling of 0-5 tick price movements.

The multifactor Vjump2 and Vjump3 models have progressively richer descriptions of variance dynamics relative to the nested Vjump1b model. In particular, these two models indicate that substantial and self-exciting volatility spikes are an important component of volatility evolution. In the Vjump2 estimates, V_{2t} is typically near zero and has a V_{1t} -dependent intensity that averages out to $\lambda_{20} + \lambda_{21}E(V_{1t}) = 48$ jumps per year. If a variance jump occurs of average size 0.07, V_{2t} increases to $(.265)^2$ and the jump intensity increases by $\lambda_{22}\bar{V}_{V2} = 1486$ jumps/year, or 6 jumps/day. This increase in jump intensity implies that $\lambda_{22}\bar{V}_{V2}/(\beta_2 - \lambda_{22}\bar{V}_{V2}) = 3.0$ additional variance jumps are expected from every average-sized jump in V_{2t} .¹¹ Furthermore, the positively skewed exponential distribution can have V_{2t} jumps that are two or three times the average jump magnitude \bar{V}_{V2} , with corresponding projections of six or nine additional variance jumps. The associated futures price jumps are roughly mean-zero, with a standard deviation of 0.36% and a strongly negative correlation (-.92) with the magnitude of the variance jumps. The V_{2t} spikes resulting from variance jumps are highly transient, with an estimated half-life of roughly 1/3 day.

The Vjump3 model decomposes variance jumps into small, medium, and large sizes, with different properties and different associations with the futures returns from which they are

¹¹ This calculation is based upon a continuous-time stochastic intensity $\lambda_{2t} = \lambda_{20} + \lambda_{21}V_{1t} + \lambda_{22}V_{2t}$, with reversion rate $\beta_2 - \lambda_{22}\bar{V}_{V2}$ for V_{2t} . The increase in expected future jumps (from an increase in expected future jump intensities) is approximately equal to but less than the number when jump intensities are constant within periods.

filtered. The moderately frequent V_{2t} jumps of average size 0.070 in the Vjump2 model are divided further in the Vjump3 model into frequent jumps in V_{2t} of average size 0.022, and more rare jumps in V_{3t} with average size 0.076. The V_{2t} jumps have little self-propagation, die off rapidly within the day (half-life of 0.08 days), and primarily add additional and transient noise to the combined spot variance process $V_{1t} + V_{2t}$. The large and relatively rare V_{3t} jumps have a -0.76 correlation with associated log futures jumps $\gamma_3 \sim [-0.2\%, (1.0\%)^2]$, indicating they are inferred from the largest 15-minute or overnight futures returns. V_{3t} reverts towards zero with a half-life of 1.6 days, implying some spillover across days. Sequences of V_{3t} jumps typically initiate from near-zero values of V_{3t} , with a V_{1t} -dependent intensity that averages out to $\lambda_{30} + \lambda_{31}E(V_{1t}) = 7.4$ jumps per year. Conditional upon initiation, each average-sized V_{3t} jump generates an expected additional 4.9 jumps, through its impact on expected future jump intensities.

Define V_{it}^{tot} as the total impact of the state variable V_{it} upon diffusive and jump variance, using the variance factor loadings of Table 2, Panel D. Figure 4 graphs the incremental contributions of those total variance components to end-of-day total spot volatility $(V_{1t}^{tot} + V_{2t}^{tot} + V_{3t}^{tot})^{1/2}$ over 1983-2013. The total estimated number of jumps each day (close to close) is also graphed, using a filtration procedure described in Bates (2006, Appendix A.4) for each intradaily and overnight period. The graph indicates that major but substantially transient volatility spikes from V_{3t} jumps tend to occur when core volatility $\sqrt{V_{1t}^{tot}}$ is relatively high. Furthermore, the volatility spikes are typically the outcome of *multiple* synchronous jumps in log futures prices and underlying volatility. For instance, the -32.7% change in the log futures price on October 19, 1987 from the previous Friday's closing value was the result of an estimated 34 jumps in V_{3t} . The +5.7% and +17.4% returns on October 20 and 21 were the outcome of 97 and 52 jumps in V_{3t} , respectively. The turbulent and predominantly falling market in the fall of 2008 was similarly the result of multiple intradaily and overnight jumps in V_{2t} and V_{3t} that accompanied predominantly negative stock market jumps.

Figure 5 shows the intradaily evolution of the three factors over October 12-30, 1987. A declining stock market over the week of October 12-16, including an accelerating drop on Friday, October 16, set the stage for the stock market crash of October 19. A volatility spike began developing late on Friday afternoon, and was stimulated further by the futures price

opening 5.8% down on the morning of October 19. Substantial intradaily drops over October 19 generated further substantial estimated increases in the volatility-spike factor $\sqrt{V_{3t}^{tot}}$, which ended the day at roughly 193% annualized (2.1% per quarter-hour). Major market turbulence on October 20 contributed to further estimated increases in $\sqrt{V_{3t}^{tot}}$, which peaked at 282% annualized (3.1% per quarter-hour) at midday before ultimately declining in the afternoon and on October 21. S&P 500 futures prices opened 13.7% down on October 22, prompting another short-lived volatility burst that decayed as the market rebounded and stabilized.

1.5 Return diagnostics

Whether the above models are broadly capturing the conditional distributions of returns can be examined by looking at the normalized returns $z_{t+\tau} = N^{-1}[CDF(y_{t+\tau} | \mathbf{Y}_t, \hat{\Theta})]$, where $CDF(\cdot)$ is the conditional cumulative distribution function derived from the specified models and computed by Fourier inversion of the associated conditional characteristic functions. Under correct model specification, these residuals are independent draws from a $N(0,1)$ distribution. Table 3 reports in-sample summary statistics for the z 's from various models, while Figure 6 presents related normal probability plots. Normalized returns are computed at two horizons: for intradaily and overnight returns conditional upon past data, and for future daily returns (close-to-close) conditional upon past intradaily and overnight data up through the close of the preceding day.

Table 3 and Figure 6 indicate that the multifactor models are doing an increasingly good job in-sample of capturing conditional distributions at the intradaily/overnight frequencies used when estimating the models. The first four moments in all cases are roughly those of a standard normal distribution, although there are statistically significant deviations in all cases. The graphs in the first column of Figure 6 show that the multifactor models Vjump2 and Vjump3 help capture intradaily outliers, which are often substantially sequential. Both models do well for $|z|$ -values less than 4, but have some difficulty in capturing the extreme tails ($|z| > 4$). The Vjump2 model underpredicts extreme negative returns while the Vjump3 model overpredicts those returns; and both models underpredict extreme positive returns ($z > 4$) in Figure 6.

The Table 3 statistics for daily data indicate that the one-factor variance models Vjump1a and Vjump1b have difficulty in capturing the stock market crash of October 19, 1987. Whereas

the Vjump1b model partially explains any single 15-minute intradaily outlier, the probability under that model of observing the sequence of outliers that culminated in the overall -32.7% change in log futures prices is low – the equivalent to observing a -7 standard deviation draw from a standard normal distribution. By contrast, the 2-factor model Vjump2 successfully captures the 1987 outlier as a volatility spike that is plausible under the (in-sample) parameter estimates. The Vjump3 model does even better at capturing that one day – at the expense of a generally worse fit for negative daily returns ($z < 0$) in Figure 6.

Autocorrelation estimates for z and $|z|$ are also given in Table 3, as tests of the independence implications of a correctly-fitted model. All models have a small but statistically significant negative autocorrelation in intradaily/overnight normalized returns – a result reflecting a small negative autocorrelation in the original log-differenced futures prices. Autocorrelations in absolute normalized returns are reduced (although still statistically significant) under the Vjump2 and Vjump3 models, indicating the multifactor models capture intradaily volatility dynamics better. At the daily horizon, the autocorrelation in normalized returns is smaller, and is not statistically significant. All models have a statistically significant -7% autocorrelation in daily normalized absolute returns $|z|$, indicating room for improvement in the models' implications for daily volatility dynamics.

2. Realized variance

2.1 Data description

Define

$$RV_n = \sum_{m=1}^{M_n} [\ln(F_{n,m}) - \ln(F_{n,m-1})]^2 \quad (12)$$

as the daily realized variance on date n from M_n sequential intradaily time intervals that begin 15 minutes after pit trading in the S&P 500 futures contract opens, and are typically 15 minutes in duration. Because the effective length of the intradaily trading period varies on a daily basis, it is desirable to rescale realized variance by the aggregate intradaily time interval

$$252\tau_n = \sum_{m=1}^{M_n} 252\tau_{nm} = f_{daily} \exp(\boldsymbol{\beta}_{ID}' \mathbf{d}_n^{ID}) \quad (13)$$

where $\exp(\boldsymbol{\beta}_{ID}' \mathbf{d}_n^{ID})$ are the intradaily day-of-the-week effects reported in Table 1, and 252 converts the time units from yearly to daily. In particular,

$$Rvol_n = \sqrt{\frac{RV_n}{252\tau_n}} \quad (14)$$

is the realized volatility on a daily (24-hour) basis that can be compared to conditional volatility assessments from daily close-to-close returns. The major impact of the daily $252\tau_n$ rescaling is to eliminate the impact of half-days adjacent to holidays, which would otherwise look like realized volatility inliers.

Table 4 describes characteristics of various transforms of realized volatility, while Figures 7A and 7B graph realized variance and volatility. Realized variance is an extremely skewed and leptokurtic random variable, with especially pronounced outliers during October 19-22, 1987. By contrast, a log transform for volatility (or variance) appears substantially better behaved: substantially less pronounced skewness and excess kurtosis in Table 4, and more discernable persistence in Figure 7B. Indeed, Figure 7B appears to offer significant support for the 2-factor log variance process estimated on daily data by Chernov, Gallant, Ghysels and Tauchen (2003), in contrast to the affine specifications explored here. However, as discussed in Appendix A.3, this paper's extended DPS volatility-jump model implies a heavy-tailed distribution for realized variance, which is potentially compatible with the occasional extreme outliers in Figure 7A.

A Box-Jenkins analysis indicates realized variance follows an ARMA(3,1) process, using a Bayes information criterion for model selection. The statistical reliability of the least-squares ARMA estimation methodology is questionable given the extreme nonnormality of realized variance data. Nevertheless, the result is supportive of the Vjump3 model, which implies an ARMA(3,2) process for realized variance.

2.2 The Fourier transform of realized variance

As illustrated in Andersen, Bollerslev and Diebold (2007), the realized variance literature relies heavily upon the asymptotic convergence of realized variance to quadratic variation as the sampling frequency increases. Over short intervals of length τ , squared log returns are approximately the intraperiod quadratic variation:

$$\begin{aligned} (f_{t+\tau} - f_t)^2 &\approx \int_{s=t}^{t+\tau} (df_s)^2 \\ &= \int_{s=t}^{t+\tau} \left(\sum_{i=1}^I V_{is} ds + \sum_{j=1-K}^J \gamma_j^2 dN_{js} \right). \end{aligned} \quad (15)$$

The mean and variance of the divergence are of order $O(\tau^2)$ for jump-diffusions with finite-activity jumps, such as those considered in this paper. Consequently, the divergence of daily realized variance computed using T/τ intradaily squared log returns from intradaily quadratic variation has mean and variance of order $O(\tau)$, yielding asymptotic convergence.

Quadratic variation is an affine random variable for affine processes. The joint Fourier transform of integrated future $(df_s)^2$ and of future $\mathbf{V}_{t+\tau}$ conditional upon current \mathbf{V}_t is exponentially affine in \mathbf{V}_t , with an analytical form derived in Appendix A.3 for the processes in this paper. This joint Fourier transform provides an approximate joint cumulant generating function for intradaily realized variance and of end-of-day spot variances conditional upon the data \mathbf{Y}_{n,M_n} available at the end of the previous day:

$$\begin{aligned} \ln E \left[e^{i\Phi R V_{n+1} + i\boldsymbol{\psi}' \mathbf{V}_{n+1, M_{n+1}}} \mid \mathbf{Y}_{n, M_n} \right] \\ \approx C^{QV}(i\Phi, i\boldsymbol{\psi}; \boldsymbol{\tau}_{n+1}) + g_{n, M_n}[D^{QV}(i\Phi, i\boldsymbol{\psi}; \boldsymbol{\tau}_{n+1})], \end{aligned} \quad (16)$$

where $\boldsymbol{\tau}_{n+1} = (\tau_{n+1,0}, \dots, \tau_{n+1, M_{n+1}})'$ is the vector of overnight and intradaily time intervals on date $n+1$, and $g_{n, M_n}(\boldsymbol{\psi}) \equiv \ln E[e^{\boldsymbol{\psi}' \mathbf{V}_{n, M_n}} \mid \mathbf{Y}_{n, M_n}]$ is the joint cumulant generating function given above in Equation (7) for the end-of-day spot variances \mathbf{V}_{n, M_n} conditional upon data \mathbf{Y}_{n, M_n} through the end of the day. The accuracy of this approximation is examined below, for the 15-minute sampling frequency used in this paper.

The joint cumulant generating function (16) provides several additional diagnostics of the merits of the proposed affine processes. First, it is possible to recursively generate filtered estimates $\widehat{V}_{n|n}^{tot}$ of total end-of-day latent variance from daily observations of intradaily realized variance, using the same AML methodology used above with intradaily returns. While daily realized variance perforce contains less information than the intradaily returns from which it is computed, the information loss may possibly be viewed as tolerable relative to the significant costs of developing a microstructure model consistent with those intradaily returns.

The magnitude of the information loss for total end-of-day spot variance is reported in Table 5, conditional upon the returns-based parameter estimates of Table 2. Returns-based filtration has an 18% lower error (average conditional variance of residuals) than RV-based filtration under Vjump2 model, and 31% lower under the Vjump3 model. However, the theoretical filtration R^2 's are close: 56% versus 48% for the Vjump2 model, and 85% versus 79% for the Vjump3 models. Realized variances can indeed subsume most of the information of intradaily returns.

Second, analytical or numerical derivatives of the joint cumulant generating function (16) provide conditional forecasts of future realized variances. Table 6 compares next-day realized variance forecasts with those from two other methodologies: ARMA models fitted to daily realized variance data, and the HAR-RV-J regressions of Andersen, Bollerslev, and Diebold (2003, Table 2B), who use bipower variation to decompose realized variance into diffusive and jump components:

$$BV_n = \frac{\pi}{2(252\tau_n)} \sum_{m=2}^{M_n} |\ln(F_{n,m}) - \ln(F_{n,m-1})| |\ln(F_{n,m+1}) - \ln(F_{n,m})| \quad (17)$$

$$J_n = \max(0, Rvol_n^2 - BV_n).$$

Over the full 1983-2008 interval, filtered realized variance forecasts outperformed the ARMA approach. The superior performance was especially pronounced for the Vjump3 model, whose R^2 of 0.271 substantially exceeded the ARMA's R^2 of 0.130. The ABD approach applied to 15-minute returns did even better in-sample, with an R^2 of 0.306.

However, the ABD regression approach lacks subsample stability. Out-of-sample performance over 2009-13 is abysmal, while an HAR-RV-J model fitted as in ABD (2003) to the 1990-2002 period does well over the quieter 2009-13 period but much worse over the full 1983-2008 period. The problem is the extreme outliers in 1987, which heavily affect regression fits and performances in-sample versus out-of-sample. The ARMA(3,1) and Vjump time series models categorize outliers better, and forecast substantially better than the ABD approach during the quieter out-of-sample 2009-13 period.

Finally, how well the proposed models match the overall conditional distributions of RV realizations can be evaluated via normal probability plots, in the same fashion used above in Figure 6 for intradaily and daily returns. Normalized residuals $z_{n+1}^{RV} = N^{-1}[CDF^{QV}(RV_{n+1}|\mathbf{Y}_{n,M_n}, \hat{\Theta})]$ were generated using quadratic variation CDFs evaluated by Fourier inversion of a transform of Eq. (15), evaluated at $\boldsymbol{\psi} = \mathbf{0}$. Their statistical properties are summarized in Table 7, while the corresponding normal probability plots are in Figure 8. As a further diagnostic, comparable statistics and normal probability plots were generated from 100 runs of simulated intradaily data over 1983-2008, using model-specific parameter estimates and estimated time gaps.

Three results stand out. First, quadratic variation's theoretical CDFs (green dotted line in Figure 8) match realized-variance CDFs from simulated data (solid red line) for median and above-median RV values, but not for below-median values. Quadratic variation is an imperfect approximation to 15-minute realized variance, especially with regard to the frequency of RV inliers. However, QV-based conditional CDF's accurately capture the theoretical frequency of RV outliers from simulated data; and those outliers are when extreme daily stock market returns are observed.

Second, the more general models match the conditional distributions better, when compared with simulation-based benchmarks that correct the biases of the QV-based CDF's. The Vjump2 model substantially reduces the frequency of RV inliers and outliers, relative to the Vjump1b model. The Vjump3 model captures RV outliers event better than the Vjump2 model, but does slightly worse with regard to RV inliers. In all cases, however, the RV realization of $(26.9\%)^2$ on October 20, 1987 remains a major positive outlier.

Finally, even the most general Vjump3 model has substantial and statistically significant deviations from simulation benchmarks, in Table 7 and Figure 8. While these affine models capture the conditional distributions of daily returns reasonably well, as indicated in Figure 6, there appears to be significant room for improvement in their predictions for the overall conditional distributions of realized variances – especially the frequency of RV inliers. The Vjump3 model does do a reasonably good job of capturing median, above-median, and extreme RV realizations, however.

3. Summary and conclusions

Finite-activity jump-diffusion models such as Merton (1976) posit that log-differenced asset prices are drawn from a mixture of distributions, with major daily outliers the outcome of a draw from a higher-volatility distribution. This model can be viewed as positing instantaneous but instantaneously transitory spikes in intradaily and daily volatility. The intradaily evidence does not support this description of volatility evolution. Instead, large daily market movements are the accumulation over the day of a series of rapid and self-exciting intradaily increases in conditional volatility, typically accompanied by a falling market. Volatility can accelerate rapidly from relatively low levels, and can subside quickly from high levels; but the shifts take time to develop within a given day, and can spill over into subsequent days. These short-lived volatility spikes affect both diffusive and jump volatility, with the former more important. The initiation of a volatility spike is more likely when core volatility is already high.

Furthermore, intradaily jumps tend to be small, with the largest ones in the Vjump3 model having a standard deviation of 1%. While these can look large at a 15-minute horizon, they are small relative to the 4-6% daily returns occasionally observed, let alone the 28% drop in December S&P 500 futures prices on October 19, 1987. It is the accumulation of higher diffusive volatility and more frequent modest intradaily jumps that generates the major daily outliers occasionally observed. These outliers are accompanied by sharp increases in observed realized variance.

The central intradaily dynamic proposed and estimated in this paper is self-exciting intradaily volatility spikes that accompany modest and predominantly negative stock market jumps. Such models have not to my knowledge been previously estimated on intradaily data; but

somewhat similar models that have been estimated on daily data. Probably the closest to this paper in spirit and results is Chernov et al (2003)'s two-factor diffusive log volatility process. They find that daily stock market outliers are generated by the second volatility factor, which possesses considerable volatility feedback and is negatively correlated with stock market returns. Aït-Sahalia et al (2015) and Fulop et al (2015) estimate stock market processes with self-exciting synchronous volatility and stock market jumps, using daily data. The generally lower number of state variables in those specifications precludes the volatility spikes estimated in this paper. Stroud and Johannes (2014) estimate multifactor models with synchronous volatility and stock market jumps on 24-hour S&P 500 E-mini futures data over 2007-9; but their models have constant-intensity rather than self-exciting jumps.

More research on identifying intradaily jumps has been conducted in a nonparametric framework, often building on the bipower variation approach of Barndorff-Nielsen and Shephard (2004, 2006). While this paper uses parametric rather than nonparametric estimation, it does have two major implications for that literature. First, the assumptions often underlying nonparametric jump detectors are not empirically supported. Lee and Mykland (2008) assume no volatility jumps; this paper finds pronounced intradaily volatility jumps. Barndorff-Nielsen and Shephard allow for volatility jumps but assume no leverage effects; this paper finds pronounced negative correlations between market returns and volatility increases. And although Huang and Tauchen (2005) find their jump detector is insensitive to diffusive leverage, whether such detectors are robust to negatively correlated and self-exciting jumps in volatility and asset prices remains to be established.

Second, applications of intradaily jump detectors to stock market data by Andersen et al (2007) and Tauchen and Zhou (2011) largely focus on whether an identifiable jump occurred within a given day. While of some use in rejecting the hypothesis of no jumps, a single intradaily jump typically contributes little to overall daily variance. Intradaily stock market jumps are modest in magnitude, as found in those papers and in this one. Their importance lies more in their dynamic implications – a signal that future overall intradaily volatility from both diffusive and jump sources is likely to rise substantially. It was such overall volatility increases that generated the market moves exceeding 10% in magnitude in 1997 and 2008.

References

- Abramowitz, M., Stegun, I.A., 1972. *Handbook of Mathematical Functions*. Dover, 1972.
- Aït-Sahalia, Y., Cacho-Diaz, J., Laeven, R.J.A., 2015. Modeling financial contagion using mutually exciting jump processes. *Journal of Financial Economics* 117, 585-606.
- Andersen, T.G., 2004. Discussion [of Barndorff-Nielsen and Shephard, 2004]. *Journal of Financial Econometrics* 2, 37-48.
- Andersen, T.G., Bollerslev, T., 1997. Intraday periodicity and volatility persistence in financial markets. *Journal of Empirical Finance* 4, 115-158.
- Andersen, T.G., Bollerslev, T., 1997, Diebold, F.X., 2007. Roughing it up: Including jump components in the measurement, modeling, and forecasting of return volatility. *The Review of Economics and Statistics*, 89, 701-720.
- Andersen, T.G, Fusari, N., Todorov, V., 2015. The risk premia embedded in index options. *Journal of Financial Economics*, 558-584.
- Barndorff-Nielsen, O.E., Shephard, N., 2004. Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics* 2, 1-37.
- Barndorff-Nielsen, O.E., Shephard, N., 2006. Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Econometrics* 4, 1-30.
- Bates, D.S., 2000. Post-'87 crash fears in the S&P 500 futures option market. *Journal of Econometrics* 94, 181-238.
- Bates, D.S., 2006. Maximum likelihood estimation of latent affine processes. *Review of Financial Studies* 19, 909-965.
- Bates, D.S., 2012. U.S. stock market crash risk, 1926-2012. *Journal of Financial Economics* 105, 229-259.
- Calvet, L., Fisher, A.J, 2008. *Multifractal volatility: theory, forecasting, and pricing*. Amsterdam, Academic Press.
- Carlson, M., 2006. A brief history of the 1987 stock market crash, with a discussion of the Federal Reserve response. Federal Reserve working paper, November.
- Carr, P., Wu, L., 2008. Leverage effect, volatility feedback, and self-exciting market disruptions: disentangling the multi-dimensional variations in S&P 500 index options. Baruch College working paper.
- Chan, K. Chan, K.C., Karolyi, G.A., 1991. Intraday volatility in the stock index and stock index futures market. *Review of Financial Studies* 4, 657-684.
- Chernov, M., Gallant, A.R., Ghysels, E., Tauchen, G., 2003. Alternative models for stock price dynamics. *Journal of Econometrics* 116, 225-257.

- Duffie, D., Pan, J., Singleton, K.J., 2000. Transform analysis and asset pricing for affine jump-diffusions. *Econometrica* 68, 1343-1376.
- Fulop, A., Li, J., Yu, J., 2015. Self-exciting jumps, learning, and asset pricing implications. *Review of Financial Studies* 28, 876-912.
- Gallant, A.R., 1981. On the bias in flexible functional forms and an essentially unbiased form: The Fourier flexible form. *Journal of Econometrics* 15, 211-245.
- Huang, X., Tauchen, G., 2005. The relative contribution of jumps to total price variance. *Journal of Financial Econometrics* 3, 456-499.
- Lee, S.S., Mykland, P.A., 2008. Jumps in financial markets: a new nonparametric test and jump dynamics. *Review of Financial Studies* 21, 2535-2563.
- Merton, R. C., 1976. Option pricing when underlying stock returns are discontinuous. *Journal of Financial Economics* 3, 125-144.
- Poppe, G.P.M., Wijers, C.M.J., 1990. More efficient computation of the complex error function. *ACM Transactions on Mathematical Software* 16, 38-46.
- Stroud, J.R., Johannes, M.S., 2014. Bayesian modeling and forecasting of 24-hour high-frequency volatility. *Journal of the American Statistical Association* 109, 1368-1384.
- Tauchen, G., Zhou, H., 2011. Realized jumps on financial markets and predicting credit spreads. *Journal of Econometrics* 160, 102-118.

Appendix

A.1 Joint characteristic functions for futures prices and spot variance

The fundamental building blocks for the multifactor and multi-jump models in this paper are based on the single-factor Duffie, Pan, and Singleton (2000) model, which has constant-intensity synchronous correlated jumps in log futures prices and in underlying spot variance:

$$\begin{aligned} df_t &= [\mu_0 + (\mu_1 - \frac{1}{2})V_t]dt + \sqrt{V_t}dW_t + (\gamma dN_t - \lambda \bar{k}dt) \\ dV_t &= (\alpha - \beta V_t)dt + \sigma \sqrt{V_t}dW_{Vt} + \gamma_V dN_t. \end{aligned} \quad (\text{A.1})$$

W_t and W_{Vt} are Wiener processes with correlation ρ , N_t is a Poisson counter with constant intensity λ , $\gamma_V \sim \text{Exp}(\bar{\gamma}_V)$ is the exponentially distributed jump in spot variance conditional upon a jump, and $\gamma = \gamma_f + \rho_J \gamma_V$ for $\gamma_f \sim N(\bar{\gamma}_f, \delta_f^2)$ is the correlated jump in log futures prices conditional upon a jump, with expected percentage jump size $\bar{k} = Ee^\gamma - 1 = e^{\bar{\gamma}_f + \frac{1}{2}\delta_f^2} / (1 - \rho_J \bar{\gamma}_V) - 1$. The key modification is to allow jump intensities to vary across periods, while maintaining the DPS assumption of a constant intraperiod jump intensity.

The generalized Fourier transform for future $\mathbf{z}_{t+\tau} = (f_{t+\tau}, V_{t+\tau}, N_{t+\tau})$ and complex (Φ, ψ, ξ) conditional upon current \mathbf{z}_t is

$$\begin{aligned} F(\tau; \Phi, \psi | \mathbf{z}_t) &\equiv E[\exp(\Phi f_{t+\tau} + \psi V_{t+\tau} + \xi N_{t+\tau}) | \mathbf{z}_t] \\ &= \exp[\Phi f_t + \xi N_t + C(\tau; \Phi, \psi, \xi) + D(\tau; \Phi, \psi)V_t]. \end{aligned} \quad (\text{A.2})$$

$F(\cdot)$ solves the backward Kolmogorov equation $E_t dF(\cdot) = 0$ associated with (A.1), implying that $C(\tau; \cdot)$ and $D(\tau; \cdot)$ solve the following ordinary differential equations subject to $C(0; \cdot) = 0$ and $D(0; \cdot) = \psi$:

$$\begin{aligned} C_\tau &= \mu_0 \Phi + \alpha D + \lambda E[\Phi, D(\cdot), \xi] \\ D_\tau &= N + MD + \frac{1}{2}\sigma^2 D^2 \end{aligned} \quad (\text{A.3})$$

where $N = \frac{1}{2}\Phi^2 + (\mu_1 - \frac{1}{2})\Phi$, $M = \rho\sigma\Phi - \beta$, and

$$\begin{aligned}
E[\Phi, D(\cdot), \xi] &\equiv \iint e^{\Phi(\gamma_f + \rho_J \gamma_V) + D(\cdot)\gamma_V + \xi} p(\gamma_f) p(\gamma_V) d\gamma_f d\gamma_V - 1 - \Phi \bar{k} \\
&= \frac{\exp(\Phi \bar{\gamma}_f + \frac{1}{2} \Phi^2 \delta_f^2 + \xi)}{1 - [\rho_J \Phi + D(\cdot)] \bar{\gamma}_V} - 1 - \Phi \bar{k}.
\end{aligned} \tag{A.4}$$

ξ is used when estimating the number of jumps (Bates 2006, Appendix A.4), and is otherwise set to zero. The solution for $D(\tau; \cdot)$ is¹²

$$D(\tau; \Phi, \psi) = \frac{\psi[M + R(\tau, \Phi)] + 2N}{R(\tau, \Phi) - M - \sigma^2 \psi} \tag{A.5}$$

where

$$\begin{aligned}
R(\tau, \Phi) &= \gamma(\Phi) \frac{e^{\gamma(\Phi)\tau} + 1}{e^{\gamma(\Phi)\tau} - 1} \text{ and} \\
\gamma(\Phi) &= \sqrt{M^2 - 2\sigma^2 N}.
\end{aligned} \tag{A.6}$$

The solution for $C(\tau; \cdot)$ is

$$C(\tau; \Phi, \psi, \xi) = \mu_0 \Phi \tau + C^{diff}(\tau; \Phi, \psi) + \lambda C^{jump}(\tau; \Phi, \psi, \xi) \tag{A.7}$$

where

$$\begin{aligned}
C^{diff}(\tau; \Phi, \psi) &= \alpha \int_{t=0}^{\tau} D(t; \cdot) dt = \\
&= \frac{\alpha \tau}{\sigma^2} [M - \gamma(\Phi)] - \frac{2\alpha}{\sigma^2} \ln \left\{ 1 + [M - \gamma(\Phi)] \frac{1 - e^{\gamma(\Phi)\tau}}{2\gamma(\Phi)} \right\} \\
&\quad - \frac{2\alpha}{\sigma^2} \ln \left[1 - \frac{\sigma^2 \psi}{R(\tau, \Phi) - M} \right]
\end{aligned} \tag{A.8}$$

$$\begin{aligned}
C^{jump}(\tau; \Phi, \psi, \xi) &= \int_{t=0}^{\tau} E[\Phi, D(t; \cdot), \xi] dt \\
&= e^{\Phi \bar{\gamma}_f + \frac{1}{2} \Phi^2 \delta_f^2 + \xi} \left\{ \frac{b}{d} \tau - e \ln \left[1 + \frac{c(e^{-\gamma(\Phi)\tau} - 1)}{c + d} \right] \right\} - (1 + \Phi \bar{k}) \tau
\end{aligned} \tag{A.9}$$

for

¹² If jumps have stochastic intensity $\lambda_t = \lambda + \lambda_1 V_t$, then $D(\tau; \cdot)$ solves $D_\tau = N + MD + \frac{1}{2} \sigma^2 D^2 + \lambda_1 E[\Phi, D(\cdot)]$. This is of cubic rather than quadratic order, and has an implicit but not explicit solution for $D(\tau; \cdot)$, precluding rapid evaluation. Self-exciting jump models typically have been solved numerically; e.g., Carr and Wu (2008).

$$\begin{aligned}
b &= [M - \gamma(\Phi)] + \sigma^2 \psi \\
c &= \{-s[M + \gamma(\Phi)] - 2\bar{\gamma}_V N\} + \psi\{-s\sigma^2 - \bar{\gamma}_V[M - \gamma(\Phi)]\} \\
d &= \{s[M - \gamma(\Phi)] + 2\bar{\gamma}_V N\} + \psi\{s\sigma^2 + \bar{\gamma}_V[M + \gamma(\Phi)]\} \\
e &= \frac{2\bar{\gamma}_V}{2\bar{\gamma}_V^2 N + 2\bar{\gamma}_V Ms + \sigma^2 s^2} \text{ and} \\
s &= 1 - \rho_j \bar{\gamma}_V \Phi.
\end{aligned} \tag{A.10}$$

For multifactor models, define $C_i^{diff}(\tau; \Phi, \psi_i)$ and $D_i(\tau; \Phi, \psi_i)$ as equations (A.8) and (A.5) evaluated at the diffusive parameters $\Theta_i^{diff} = (\mu_i, \alpha_i, \beta_i, \sigma_i, \rho_i)$ associated with variance process V_{it} . For $i > 1$, $(\alpha_i, \sigma_i, \rho_i) = \mathbf{0}$ and $C_i^{diff} = 0$. Let $E_j[\Phi, D_j(\tau; \Phi, \psi_j), \xi_j]$ for $j > 0$ be equation (A.4) evaluated at the jump parameters $\Theta_j^{jump} = (\bar{\gamma}_{Vj}, \rho_j, \bar{\gamma}_{fj}, \delta_j)$ associated with jump process N_{jt} . Its corresponding $C_j^{jump}(\tau; \Phi, \psi_j, \xi_j)$ for $j > 0$ is equations (A.9) – (A.10) evaluated at parameters $(\Theta_j^{diff}, \Theta_j^{jump}, \xi_j)$. Finally, let $E_0(\Phi, \xi_0) = \exp(\Phi \bar{\gamma}_{f0} + \frac{1}{2} \Phi^2 \delta_{f0}^2 + \xi_0) - 1 - \Phi \bar{k}_0$ for jump process N_{0t} (futures jumps unaccompanied by volatility jumps), with associated $C_0^{jump}(\tau; \Phi, \xi_0) = E_0(\Phi, \xi_0)\tau$. Because jump intensities are assumed constant intraperiod, the generalized Fourier transform for future $\mathbf{z}_{t+\tau} = (f_{t+\tau}, \mathbf{V}_{t+\tau}, \mathbf{N}_{t+\tau})$ conditional upon current \mathbf{z}_t is a multifactor extension of (A.2) and (A.7):

$$\begin{aligned}
F(\tau; \Phi, \psi | \mathbf{z}_t) &\equiv E[\exp(\Phi f_{t+\tau} + \psi' \mathbf{V}_{t+\tau} + \xi' \mathbf{N}_{t+\tau}) | \mathbf{z}_t] \\
&= \exp \left[\Phi f_t + \mu_0 \Phi \tau + \xi' \mathbf{N}_t + \sum_{i=1}^I [C_i^{diff}(\tau; \Phi, \psi_i) + D_i(\tau; \Phi, \psi_i) V_{it}] \right. \\
&\quad \left. + \lambda_{01} V_{1t} C_0^{jump}(\tau; \Phi, \xi_0) + \sum_{j=1}^J (\lambda_{j0} + \lambda'_j \mathbf{V}_t) C_j^{jump}(\tau; \Phi, \psi_j, \xi_j) \right].
\end{aligned} \tag{A.11}$$

A.2 Time aggregation

The assumption that jump intensities are constant intraperiod but vary across periods necessitates a concatenated approach for multiperiod returns. By iterating expectations over (A.11), the multiperiod conditional cumulant generating function

$$\ln E[e^{\Phi(f_T - f_t) + \psi V_T} | \mathbf{V}_t] = \mu_0 \Phi(T - t) + \sum_{i=1}^I C_{t,T}^i(\Phi, \psi) + D_{t,T}^i(\Phi, \psi) V_{it} \quad (\text{A.12})$$

satisfies the backwards recursion

$$\begin{aligned} D_{t,T}^1 &= D_1(\tau_t; \Phi, D_{t+1,T}^1) + \lambda_{01} C_0^{jump}(\tau_t; \Phi) + \sum_{j=1}^J \lambda_{j1} C_1^{jump}(\tau_t; \Phi, D_{t+1,T}^1) \\ D_{t,T}^i &= D_i(\tau_t; \Phi, D_{t+1,T}^i) + \lambda_{ii} C_i^{jump}(\tau_t; \Phi, D_{t+1,T}^i) \text{ for } i > 1 \\ C_{t,T}^i &= C_{t+1,T}^i + C_i^{diff}(\tau; \Phi, D_{t+1,T}^i) + \lambda_{i0} C_i^{jump}(\tau; \Phi, D_{t+1,T}^i) \text{ for } i \geq 1 \end{aligned} \quad (\text{A.13})$$

subject to the terminal condition $(C_{T,T}^i, D_{T,T}^i) = (0, \psi_i)$ for $i = 1, \dots, I$.

A.3 The joint Fourier transform of realized variance and future spot variances

A.3.1 Intraproduct quadratic variation

Define $Q_t \equiv \int_{s=0}^t (df_s)^2$ as the quadratic variation of the log futures price for the benchmark constant-intensity Duffie et al (2000) process (A.1) above. Its stochastic evolution is

$$\begin{aligned} dQ_t &= V_t dt + \gamma^2 dN_t \\ dV_t &= (\alpha - \beta V_t) dt + \sigma \sqrt{V_t} dW_{Vt} + \gamma_V dN_t, \end{aligned} \quad (\text{A.14})$$

which is an affine system. Consequently, the generalized joint Fourier transform of future quadratic variation and spot variance is an exponentially affine function of current V_t :

$$\begin{aligned} F^{QV}(V_t, \tau) &\equiv E[e^{\Phi Q_T + \psi V_T} | Q_t, V_t] \\ &= \exp[\Phi Q_t + C^{QV}(\tau; \Phi, \psi) + D^{QV}(\tau; \Phi, \psi) V_t] \end{aligned} \quad (\text{A.15})$$

for arbitrary complex Φ and ψ and $\tau = T - t$. $F^{QV}(\cdot)$ solves the backward Kolmogorov equation $E_t dF^{QV}(\cdot) = 0$ associated with (A.14), implying that $C^{QV}(\tau; \cdot)$ and $D^{QV}(\tau; \cdot)$ solve the following recursive system of ordinary differential equations:

$$D_\tau^{QV} = \Phi - \beta D^{QV} + \frac{1}{2} \sigma^2 (D^{QV})^2 \quad \text{subject to } D^{QV}(0; \cdot) = \psi \quad (\text{A.16})$$

$$C_\tau^{QV} = \alpha D^{QV} + \lambda E[e^{\Phi \gamma^2 + D^{QV} \gamma_V} - 1] \quad \text{subject to } C^{QV}(0, \cdot) = 0. \quad (\text{A.17})$$

Equation (A.16) is of the same form as in (A.3). $D^{QV}(\tau; \Phi, \Psi)$ consequently has an analytical expression given in equations (A.5) – (A.7) for $N = \Phi$ and $M = -\beta$.

The expectation in (A.17) requires a solution for the joint Fourier transform of the squared log futures jump γ^2 and the spot variance jump γ_V . Because γ_V has an exponential distribution, $\gamma^2 = (\gamma_f + \rho_J \gamma_V)^2$ has a heavy-tailed distribution, with infinite values for the moment generating function $E[e^{\Phi \gamma^2}]$ when Φ is real and positive. Quadratic variation Q_t is consequently also heavy-tailed for the DPS process, although the conditional characteristic function and moments of all orders exist. Integrating successively over the independent densities of $\gamma_f \sim N(\bar{\gamma}_f, \delta^2)$ and $\gamma_V \sim \text{Exp}(\bar{\gamma}_V)$ yields

$$\begin{aligned} E^{QV}(\Phi, D) &\equiv E \left[e^{\Phi(\gamma_f + \rho_J \gamma_V)^2 + D \gamma_V} - 1 \right] \\ &= \frac{1}{\sqrt{1 - 2\delta^2 \Phi}} E \left\{ \exp \left[\frac{\Phi(\bar{\gamma}_f + \rho_J \gamma_V)^2 + \gamma_V D}{1 - 2\delta^2 \Phi} \right] \right\} - 1 \\ &= \frac{\exp[a(\Phi)x^2]}{\rho_{sv} \bar{\gamma}_V} \sqrt{-\frac{\pi}{4\Phi}} w \left(\frac{2a(\Phi)x + \bar{\gamma}_V D - 1}{2\sqrt{a(\Phi)}} \right) - 1 \end{aligned} \quad (\text{A.18})$$

where $a(\Phi) = \frac{\Phi \rho_{sv}^2 \bar{\gamma}_V^2}{1 - 2\delta^2 \Phi}$ and $x = \frac{\bar{\gamma}_f}{\rho_{sv} \bar{\gamma}_V}$. The function $w(\mathbf{z})$ for complex-valued \mathbf{z} is the Faddeeva or plasma-dispersion function (equation 7.1.3 in Abramowitz and Stegun (1972, p.297)). It is a scaled version of the complex complementary error function ($w(\mathbf{z}) = e^{-z^2} \text{erfc}(-i\mathbf{z})$), and can be evaluated using algorithm 680 of Poppe and Wijers (1990).¹³ Equation (A.18) is well-defined for complex-valued Φ provided $\text{Re}[a(\Phi)] < 0$ – which is *not* true for small positive real values of Φ .

¹³Poppe and Wijers' algorithm is accurate to 14 significant digits, whereas the complex-valued Faddeeva function ZERFE in IMSL is accurate to only 10 digits. Furthermore, ZERFE does not appear to be fully reliable for all values of Φ , when compared with direct numerical integration over the exponential density.

$C^{QV}(\tau; \Phi, \Psi)$ can be evaluated by integrating (A.17):

$$\begin{aligned} C^{QV}(\tau; \Phi, \Psi) &= \alpha \int_0^\tau D^{QV}(t, \cdot) dt + \lambda \int_0^\tau E^{QV}[\Phi, D^{QV}(t, \cdot)] dt \\ &\approx C^{QVd}(\tau; \Phi, \Psi) + \frac{\lambda\tau}{2} \{E^{QV}[\Phi, \Psi] + E^{QV}[\Phi, D^{QV}(\tau; \Phi, \Psi)]\} \\ &\equiv C^{QVd}(\tau; \Phi, \Psi) + \lambda C^{QVj}(\tau; \Phi, \Psi) \end{aligned} \quad (\text{A.19})$$

where $C^{QVd}(\cdot)$ is the analytical solution $C^{diff}(\cdot)$ given above in equations (A.8) and (A.6), using $N = \Phi$ and $M = -\beta$. The approximation for $C^{QVj}(\cdot)$ uses trapezoidal integration, yielding a numerical solution that is accurate to order $O(\tau^{3/2})$.

A.3.2 Realized variance

The divergence

$$(f_{t+\tau} - f_t)^2 - \int_{s=t}^{t+\tau} (df_s)^2 \quad (\text{A.20})$$

of per-period squared returns from quadratic variation has mean and variance of order $O(\tau^2)$. Consequently, the divergence

$$\sum_{m=1}^{T/\tau} (\Delta f_m)^2 - \int_{s=t}^T (df_s)^2 \quad (\text{A.21})$$

of daily realized variance from intradaily quadratic variation has mean and variance of order $O(\tau)$. Moments and other characteristics of realized variance can therefore be approximated by the conditional characteristic function of daily quadratic variation. The use of trapezoidal integration in (A.19) above reduces the accuracy of the resulting approximate conditional characteristic function for realized variance to order $O(\sqrt{\tau})$, which is still asymptotically convergent as intradaily time intervals τ go to zero.

For the multi-state variable process (1) of this paper, per-period squared returns are approximately

$$(f_{t+\tau} - f_t)^2 \approx \int_{s=t}^{t+\tau} (df_s)^2 = Q_{t+\tau} - Q_t, \quad (\text{A.22})$$

with jump intensities $\lambda_{jt} = \lambda_{j0} + \lambda'_j \mathbf{V}_t$ that are linear functions of \mathbf{V}_t values at the beginning of the period. Using a multifactor and multi-jump extension of the joint Fourier transform for quadratic variation in equations (A.15) and (A.19) yields an approximate formula for the generalized Fourier transform of squared returns variation and $\mathbf{V}_{t+\tau}$ that is exponentially affine in \mathbf{V}_t :

$$\begin{aligned}
& E[\exp(\boldsymbol{\Phi}(f_{t+\tau} - f_t)^2 + \boldsymbol{\psi}'\mathbf{V}_{t+\tau}) | \mathbf{V}_t] \\
& \approx \exp \left[C_1^{QVd}(\tau; \boldsymbol{\Phi}, \boldsymbol{\psi}_1) + \sum_{i=1}^I D_i^{QV}(\tau; \boldsymbol{\Phi}, \boldsymbol{\psi}_i) V_{it} \right. \\
& \quad \left. + \lambda_{01} V_{1t} E_0^{QV}(\boldsymbol{\Phi}) \tau + \sum_{j=1}^J (\lambda_{j0} + \lambda'_j \mathbf{V}_t) C_j^{QVj}(\tau; \boldsymbol{\Phi}, \boldsymbol{\psi}_j) \right].
\end{aligned} \tag{A.23}$$

$[C_i^{QVd}, D_i^{QV}]$ are $[C^{QV}, D^{QV}]$ from (A.19) and (A.16) evaluated at the $\boldsymbol{\Theta}_i^{diff}$ diffusive parameters of the i th spot variance process V_{it} . $C_i^{QVd} = 0$ for $i > 1$ and is therefore omitted from the summation. $E_0^{QV}(\boldsymbol{\Phi})$ is $E^{QV}(\boldsymbol{\Phi}, 0)$ from (A.18) at jump parameters $\boldsymbol{\Theta}_0^{jump}$, including the absence of volatility jumps ($\gamma_{V0} \equiv 0$). $C_j^{QVj}(\cdot)$ is $C^{QVj}(\cdot)$ from (A.19) evaluated at the parameters $(\boldsymbol{\Theta}_j^{diff}, \boldsymbol{\Theta}_j^{jump})$ of the j th spot variance process V_{jt} .

Concatenating (A.23) backwards in time over a given day's subintervals in a similar fashion to (A.13) yields the generalized joint Fourier transform of future realized variance and spot variances conditional upon current \mathbf{V}_t . I suppress overnight squared returns in the recursion by setting $\boldsymbol{\Phi} = 0$ for the overnight period, when computing the joint transform of tomorrow's intradaily realized variance and end-of-day spot variances conditional upon information available at the end of today. This joint transform can be used for conditional diagnostics of realized variance, and for filtering end-of-day spot variance values conditional upon observed daily realized variance.

Table 1: Intradaily and overnight variance patterns of S&P futures returns; Vjump1b estimates

Parameter	Estimate	Std. error	t-statistic
f_{daily}	0.808	0.006	145.7
Intradaily variance pattern			
b_1	1.587	0.266	5.96
b_2	-0.622	0.095	-6.53
b_{ah}	-0.058	0.036	-1.61
c_1	0.750	0.050	14.90
c_2	0.091	0.012	7.29
d_1	-0.203	0.012	-16.90
d_2	-0.018	0.007	-2.49
Day-of-the-week influences on aggregate intradaily variances, relative to Wednesday. The t-statistics test the null hypothesis value of 1.			
Monday	0.874	0.013	-9.39
Tuesday	0.967	0.013	-2.48
Wednesday	1		
Thursday	0.987	0.014	-0.92
Friday	0.973	0.015	-1.84
Half-day	0.389	0.029	-21.20
Day-of-the-week influences on overnight variances, relative to Tuesday close -> Wednesday open. The t-statistics test the null hypothesis of 1.			
Monday -> Tuesday	1.107	0.051	2.11
Tuesday -> Wednesday	1		
Wednesday -> Thursday	1.078	0.050	1.56
Thursday -> Friday	1.490	0.072	6.78
2-day holiday	2.514	0.496	3.05
3-day weekend	1.347	0.065	5.38
4- or 5-day holiday weekend	3.191	0.301	7.29

Notes: f_{daily} is the fraction of 24-hour variance attributable to intradaily trading, excluding the first 15 minutes. Equation (11) in the paper gives the intradaily variance pattern for the above parameters. Day-of-the-week influences are relative to benchmark Wednesday values of 1. The $1-0.389 = 61.1\%$ lower intradaily variance on the 41 half-days adjacent to holidays includes the impact of fewer trading hours on those days.

Table 2: Estimates of the log futures price process

$$df_t = \mu_0 dt + \sum_{i=1}^I [(\mu_i - \frac{1}{2})V_{it}dt + \sqrt{V_{it}}dW_{it}] + \sum_{j=1-K}^J \gamma_j dN_{jt} - \lambda_{jt}\bar{k}_j dt$$

$$dV_{1t} = (\alpha - \beta_1 V_{1t})dt + \sigma\sqrt{V_{1t}}dW_{Vt} + \gamma_{V1}dN_{1t}$$

$$dV_{it} = -\beta_i V_{it}dt + \gamma_{Vi}dN_{it} \text{ for } i > 1$$
(1)

where I is the number of state variables in V_t , J is the number of synchronous jumps in spot variance and futures prices, and K is the number of jumps in futures prices only. Diffusive shocks dW_{1t} and dW_{Vt} have correlation ρ , while the intraperiod jump intensities λ_{jt} of Poisson counter N_{jt} are given below in Panel C.

Panel A: Log likelihoods of various models

Model	I	J	K	Number of parameters	$\ln L$	$\Delta \ln L$		
						0-5 ticks	>5 ticks	All
In-sample (1983-2008)								
SVJ1	1	0	1	28	891,425.32			
Vjump1a	1	1	0	30	892,063.47	740.48	-102.32	638.16
Vjump1b	1	1	1	33	892,575.37	1533.22	-1021.32	511.90
Vjump2	2	2	1	42	893,212.58	255.16	382.05	637.21
Vjump3	3	3	1	51	893,432.73	146.89	73.24	220.13
Out-of-sample (2009-13)								
SVJ1					172,284.08			
Vjump1a					172,313.58	-7.92	37.42	29.50
Vjump1b					172,414.90	352.49	-221.66	130.82
Vjump2					172,478.28	443.29	-249.10	194.20
Vjump3					172,535.48	505.86	-254.46	251.40

Panel B: Parameter estimates for specific models.

Parameter estimates are generally on an annualized basis, with Tuesday close to Wednesday close equal to 1/252 years. The half-lives $HL_i = 252 (\ln 2) / (\beta_i - \lambda_{ii} \bar{V}_{Vi})$ to variance shocks are reported in days.

Parameters\Model	Vjump1a	Vjump1b	Vjump2		Vjump3		
	V_{1t}	V_{1t}	V_{1t}	V_{2t}	V_{1t}	V_{2t}	V_{3t}
Conditional mean							
μ_0	-.10 (.02)	-.08 (.03)	-.08 (.04)		.01 (.04)		
μ_i	8.0 (1.5)	8.0 (2.5)	5.9 (4.3)	7.4 (2.95)	2.5 (9.0)	1.9 (11.7)	0.3 (3.7)
Variance processes							
$\sqrt{E(V_i)}$.155 (.003)	.131 (.003)	.108 (.003)	.082 (.004)	.092 (.004)	.070 (.004)	.071 (.009)
HL_i (days)	3.6 (0.2)	5.2 (0.3)	9.4 (0.7)	0.35 (0.03)	13.4 (1.1)	0.08 (0.01)	1.6 (0.4)
σ	.47 (.05)	.51 (.02)	.26 (.03)		.19 (.03)		
ρ	-.70 (.07)	-.63 (.03)	-.81 (.08)		-.99 (.13)		

Panel C: Jump parameters from specific models.

Annualized jump intensities for Poisson counter N_{jt} are

$$\lambda_{jt} = \begin{cases} \lambda_{j1}V_{1t}^* & \text{for } j \leq 1 \\ \lambda_{j0} + \lambda_{j1}V_{1t}^* + \lambda_{jj}V_{jt}^* & \text{for } j > 1 \end{cases}$$

where V_{jt}^* 's are annualized spot variances at the start of the period. Spot variance jumps are distributed $\gamma_{Vj} \sim \text{Exp}(\bar{\gamma}_{Vj})$, with no variance jumps ($\bar{\gamma}_{Vj} = 0$) for $j = 0$. Log futures price jumps $\gamma_j \sim (\bar{\gamma}_j, \delta_j^2)$ have correlation $\text{Corr}_j = \rho_j \bar{\gamma}_{Vj} / \delta_j$ with variance jumps. Poisson counter N_{0t} indicates high-frequency jumps in futures prices that are unaccompanied by volatility jumps. $E[\Delta N | J]$ is the expected number of additional jumps over a 1-day or infinite horizon conditional upon an average-sized variance jump $\bar{\gamma}_{Vj}$. The estimated average length of an intradaily 15-minute period was 1.20e-4 years. Standard errors are in parentheses.

	Vjump1	Vjump1b		Vjump2			Vjump3			
	N_{1t}	N_{0t}	N_{1t}	N_{0t}	N_{1t}	N_{2t}	N_{0t}	N_{1t}	N_{2t}	N_{3t}
λ_{j0}						5.4 (2.8)			9.8 (23.3)	0.0 (0.0)
λ_{j1} $\times 10^{-4}$		38.3 (3.7)		130 (30)		0.36 (0.06)	229 (82)	6.5 (2.2)	5.5 (1.3)	.09 (.01)
λ_{jj} $\times 10^{-4}$	2.0 (0.1)		0.31 (0.02)		5.6 (1.5)	2.1 (0.6)		6.5 (2.2)	2.0 (1.3)	0.7 (0.2)
$\bar{\gamma}_j \times 10^4$	-5.1 (0.7)	-1.0 (0.1)	-20.4 (4.0)	-0.8 (0.2)	-2.3 (1.4)	-2.5 (3.9)	-0.6 (0.2)	0.2 (1.7)	-3.2 (2.2)	-19.9 (8.6)
$\delta_j \times 10^4$	32.6 (0.6)	11.4 (0.3)	90.2 (2.7)	6.9 (0.5)	20.2 (1.6)	36.5 (2.5)	5.6 (0.5)	17.8 (1.6)	6.3 (6.4)	103.3 (7.1)
$\bar{\gamma}_{Vj}$.0063 (.0002)		.0115 (.0006)		.0016 (.0002)	.070 (.011)		.0006 (.0001)	.022 (.003)	.076 (.012)
ρ_j	-.27 (.01)		-.49 (.03)		-1.18 (0.22)	-.05 (.01)		-1.89 (0.40)	-.02 (.01)	-.10 (.02)
Corr_j	-.52 (.02)		-.62 (.03)		-.60 (.07)	-.92 (.05)		-.68 (.09)	-.83 (.96)	-.76 (.06)
$E[\Delta N^{\text{day}} J]$	0.46 (0.02)	16.4 (1.4)	0.13 (0.01)	5.1 (1.1)	0.22 (0.04)	2.6 (0.3)	5.6 (1.6)	0.16 (0.03)	0.20 (0.11)	1.7 (0.2)
$E[\Delta N^\infty J]$	2.63 (0.16)	131 (14)	1.05 (0.09)	72 (16)	3.1 (0.5)	3.0 (0.4)	112 (33)	3.2 (0.7)	0.20 (0.11)	4.9 (1.4)

Panel D: Variance factor loadings

Model and V_i 's	Constant variance $\lambda_i^0 E(\gamma_i^2)$	$E(V_{it})$	$SD(V_{it})$	Jump variance loading $\lambda_{ji} E(\gamma_j^2)$ on V_{it}				Total variance loading
				$j = 0$	1	2	3	
Vjump1a V_1		.0239 (.0008)	.0211 (.0008)		0.22 (0.01)			1.22 (0.01)
Vjump1b V_1		.0171 (.0008)	.0165 (.0010)	0.50 (0.03)	0.26 (0.02)			1.76 (0.04)
Vjump2 V_1		.0116 (.0007)	.0077 (.0005)	0.63 (0.07)	0.23 (0.03)	0.05 (0.01)		1.91 (0.09)
V_2	.0001 (.0000)	.0066 (.0007)	.0431 (.0041)			0.28 (0.07)		1.28 (0.07)
Vjump3 V_1		.0085 (.0008)	.0053 (.0005)	0.72 (0.13)	0.21 (0.04)	0.03 (0.05)	0.10 (0.02)	2.05 (0.17)
V_2	.0000 (.0000)	.0049 (.0006)	.0118 (.0009)			0.01 (0.02)		1.01 (0.02)
V_3	.0000 (.0000)	.0051 (.0013)	.0478 (.0115)				0.79 (0.14)	1.79 (0.14)

Table 3: Statistical properties of the normalized intradaily returns

$z_{t+1} = N^{-1}[CDF(y_{t+1}|Y_t, \hat{\Theta})]$ for parameter estimates $\hat{\Theta}$ from various models. Under correct model specification, the z_{t+1} 's should be independent draws from a Gaussian distribution with zero mean and unitary variance. Heteroskedasticity-consistent standard errors are in parentheses.

	Vjump1a	Vjump1b	Vjump2	Vjump3
Intradaily and overnight returns (174,854 observations)				
Maximum	6.49	5.46	5.54	4.78
Minimum	-7.38	-5.17	-5.01	-4.29
Mean	.064	.063	.053	0.050
(Std. error)	(.002)	(0.002)	(.002)	(0.002)
Std. deviation	.956	0.970	0.999	1.001
Skewness	.03	0.01	-0.01	-0.02
Excess kurtosis	-.11	-.20	0.02	0.00
Corr(z_{t+1}, z_t)	-.013	-.015	-.012	-.013
(Std. error)	(.003)	(.002)	(.002)	(.002)
Corr($ z_{t+1} , z_t $)	.016	.025	.006	-.005
(Std. error)	(.003)	(.003)	(.002)	(.002)
Daily returns (6,558 observations)				
Maximum	4.10	3.99	3.54	3.91
Minimum	-7.39	-7.11	-5.17	-4.84
Mean	0.004	0.013	0.019	0.007
Std. error	0.013	0.013	0.013	0.013
Std. deviation	1.051	1.037	1.028	1.045
Skewness	0.01	0.03	0.00	-0.04
Excess kurtosis	0.79	0.53	0.17	0.31
Corr(z_{t+1}, z_t)	-.004	-.006	-.004	-.009
(Std. error)	(.012)	(.012)	(.012)	(.012)
Corr($ z_{t+1} , z_t $)	-.070	-.068	-.070	-.070
(Std. error)	(.012)	(.012)	(.012)	(.012)

Table 4: Summary statistics for daily realized volatility measures

Realized daily volatility is $Rvol_n = \sqrt{\frac{\sum(\Delta \ln F)^2}{252\tau_n}}$, where $252\tau_n$ is the estimated length of the intradaily period in days, with an average value of 0.774 (Vjump3 estimates). ARMA models are estimated in RATS, using the Bayes information criterion for model selection.

	$Rvol_n^2$	$Rvol_n$	$\ln(Rvol_n)$
Maximum	0.027530	26.93%	-1.312
Minimum	0.000003	0.17%	-6.380
Mean	0.000130	0.91%	-4.825
(Std. error)	(0.000012)	(0.01%)	(0.006)
Median	0.000062	0.79%	-4.847
Std. deviation	0.000982	0.68%	0.477
Skewness	63.4	12.5	0.63
Excess kurtosis	4553	367	1.83
ARMA model	ARMA(3,1)	ARMA(3,1)	ARMA(3,2)

Major Rvol outliers: 26.9% (10/20/1987), 14.8% (10/19/1987), 11.4% (10/22/1987), 8.7% (10/10/2008)

Table 5: The informational content over 1983-2008 of intradaily/overnight returns or realized variances for total latent variances, including jump risk.

$$V_n^{tot} = \hat{V}_{n|n}^{tot} + \varepsilon_n, \quad \varepsilon_t \sim (0, P_{n|n})$$

$$R^2 \equiv 1 - \frac{E[E_n(\varepsilon_n^2)]}{Var[V_n^{tot}]} = 1 - \frac{Avg[P_{n|n}]}{Var[V_n^{tot}]}$$

Model & Data	$E[V_n^{tot}]$	$Avg[\hat{V}_{n n}^{tot}]$	$Var[V_n^{tot}]$	$Var[\hat{V}_{n n}^{tot}]$	$Avg[P_{n n}]$	R^2
Vjump2 Y_{n,M_n}	0.0307	0.0301	3.3e-3	5.7e-3	1.4e-3	0.56
Vjump2 RV_n	0.0307	0.0298	3.3e-3	11.3e-3	1.7e-3	0.48
Vjump3 Y_{n,M_n}	0.0315	0.0314	7.6e-3	11.6e-3	1.1e-3	0.85
Vjump3 RV_n	0.0315	0.0310	7.6e-5	15.6e-3	1.6e-3	0.79

Unconditional moments of latent end-of-day total variance V_n^{tot} are based upon model-specific parameter estimates. End-of-day filtered total variances $\hat{V}_{n|n}^{tot}$ use either all past intradaily and overnight returns (Y_{n,M_n}) or all past observed daily realized variances (RV_n), with the in-sample moments over 1983-2008 reported in the table.

Table 6: Forecasting regressions for realized variance over 1983-2013

Forecasting regressions are of the form

$$Rvol_{n+1}^2 = a + b E_n(Rvol_{n+1}^2) + \varepsilon_{n+1}$$

with R^2 's reported for $(a, b) = (0, 1)$. The table gives the Mincer-Zarnowitz (MZ) decomposition of the mean squared error into squared-bias ($a \neq 0$), inefficiency ($b \neq 1$) and residual-error ($Var(\varepsilon)$) components. Forecasts of next-day's realized variance are from six approaches: an ARMA(3,1) model fitted to daily realized variance over 1983-2008; RV forecasts from the Vjump1b, Vjump2, and Vjump3 models; the HAR-RV-J regression approach of Andersen, Bollerslev and Diebold (2003, Table 2B), for the full 1983-2008 period, and for their original 1990-2002 period.

The HAR-RV-J regression over 1983-2008 was

$$Rvol_{n+1}^2 = 0.05e-5 + 1.63Rvol_n^2 - 0.013Rvol_{n-5}^2 - 0.008Rvol_{n-22}^2 - 4.50J_n + \varepsilon_{n+1},$$

(5.03e-5) (0.88) (0.038) (0.015) (2.78)

where $J_n = \max(0, Rvol_n^2 - BV_n)$ is the daily jump component given daily bipower variation BV_n computed from 15-minute returns and divided by the intradaily horizon $252\tau_n$, measured in days. HAR-RV-J standard errors (in parentheses) use Newey-West heteroskedasticity-consistent estimation with 5 daily lags.

Model	MSE $\times 10^8$	MZ decomposition (in %)			R^2
		Bias ²	Ineff.	Resid.	
In-sample (1983-2008)					
ARMA(3,1)	84.17	0	0	100	0.130
Vjump1b	83.45	0	4	96	0.139
Vjump2	79.42	0	17	83	0.179
Vjump3	70.51	0	6	94	0.271
ABD (1983-2008)	67.14	0	0	100	0.306
ABD (1990-2002)	91.69	0	5	95	0.098
Out of sample (2009-13)					
ARMA(3,1)	2.06	1	6	94	0.392
Vjump1b	1.83	0	0	100	0.463
Vjump2	1.84	1	0	99	0.456
Vjump3	1.87	1	2	98	0.446
ABD (1983-2008)	4.67	0	51	49	-0.381
ABD (1990-2002)	2.10	0	1	99	0.378

Table 7: Summary statistics for the residuals $z_{n+1} = N^{-1}[CDF^{QV}(RV_{n+1}|Y_{n,M_n}, \hat{\Theta})]$ over 1983-2008, on actual and simulated data. Normalized residuals should be i.i.d. standard Gaussian conditional upon correct specification. Simulated data are from 100 runs of 6557 days each, using model-specific parameter estimates and time gaps. The t-statistics for testing compatibility between actual moments M^{act} and simulated moments M^{sim} are $[M^{act} - Avg(M^{sim})] / [SD(M^{sim})\sqrt{1.01}]$.

Statistic	Actual Data	Simulated data		t-statistic
		Average	Std. deviation	
Vjump1b				
Maximum	>7 ^a	3.80	0.34	9.41
Minimum	-7.55	-4.52	0.42	-7.17
Mean	-0.191	-0.106	0.012	-7.27
Median	-0.155	-0.093	0.017	-3.59
SD	1.265	1.158	0.011	9.99
Skewness	-.32	-.12	0.03	-8.17
Excess kurtosis	0.55	-.28	0.05	16.40
Vjump2				
Maximum	6.46	3.80	0.31	8.52
Minimum	-5.50	-5.13	0.52	-0.71
Mean	-0.225	-0.163	0.014	-4.35
Median	-0.121	-0.096	0.022	-1.16
SD	1.250	1.211	0.009	4.10
Skewness	-0.40	-0.28	0.03	-3.50
Excess kurtosis	0.16	-0.13	0.06	4.41
Vjump3				
Maximum	5.88	3.80	0.34	6.04
Minimum	-5.62	-5.01	0.40	-1.53
Mean	-0.190	-0.156	0.012	-2.77
Median	-0.119	-0.112	0.017	-0.39
SD	1.245	1.215	0.011	2.88
Skewness	-0.33	-0.20	0.03	-3.78
Excess kurtosis	0.20	-0.10	0.07	4.53

^aThe normalized residual observed on October 20, 1987 could not be computed for the Vjump1b model, but is in excess of 7. The summary statistics consequently use a value of 7 for that observation.

Figure 1: Overnight (ON) and intradaily (ID) return horizons, relative to Wednesday horizon (Tuesday close -> Wednesday close). Overnight returns for Monday (Friday close -> Monday open) and for Tuesday through Friday include the first 15 minutes of the opening day, as do overnight returns spanning 2-day holidays and 4- or 5-day holiday weekends.




Figure 2: Intradaily 15-minute horizons, as a percentage of the total intradaily horizon.





Figure 3: Cumulative log likelihood differences relative to the 1-factor SVJ1 model.

Figure 4: Components of annualized conditional volatility, and daily estimated number of jumps. Total end-of-day variance estimates V_{it}^{tot} from the 3-factor variance model Vjump3 include the contributions to both diffusion risk and jump risk.

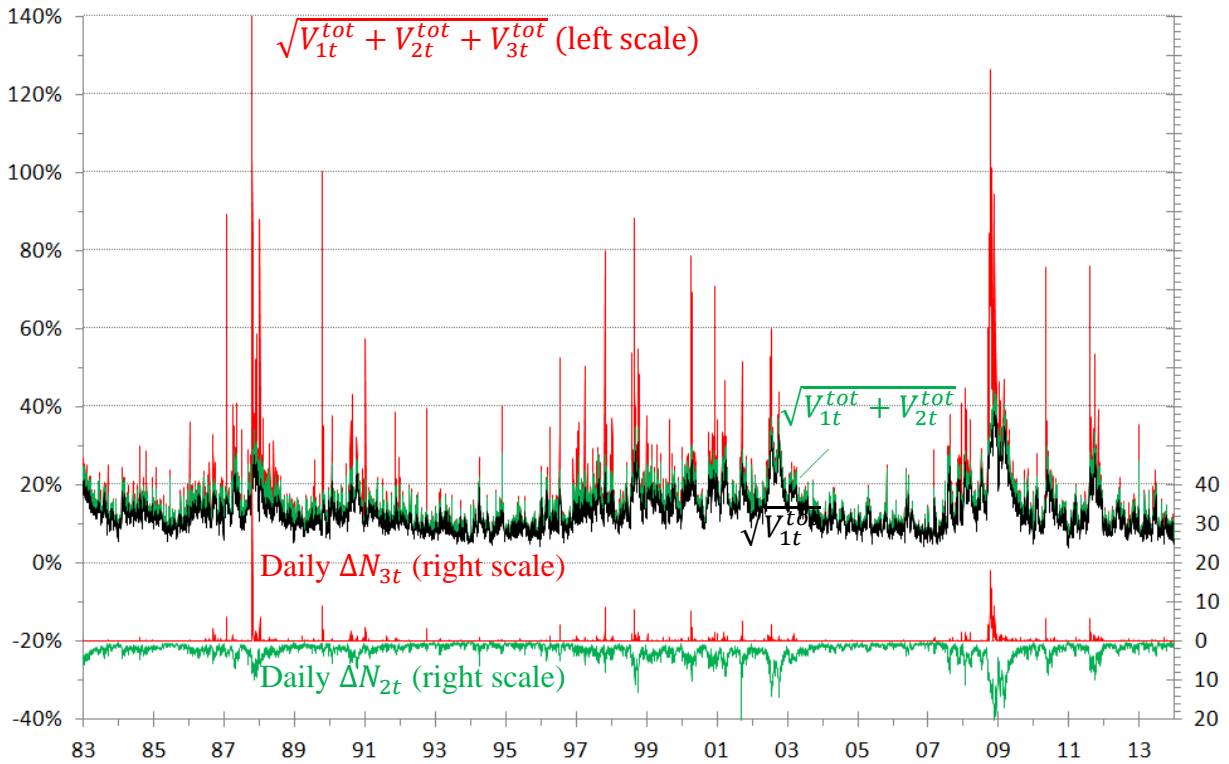


Figure 5: Annualized intradaily conditional volatility factors over October 12-30, 1987
 (100% annualized = 1.1% per ¼-hour). The December 1987 S&P 500 futures price (+) is on the right scale.

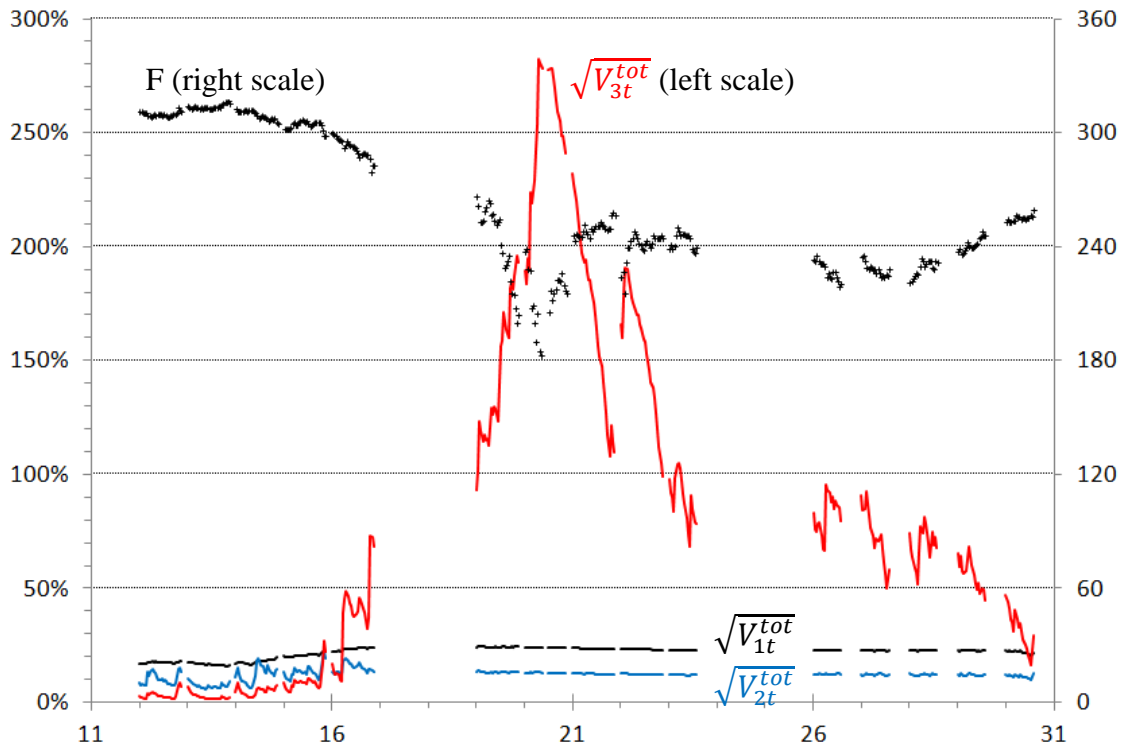




Figure 6: Normal probability plots for various models. These plot the ordered normalized residual values $z_{t+1} = N^{-1}[CDF(y_{t+1}|Y_t, \hat{\Theta})]$ (+ symbols on the vertical axis), against the models' predicted values (red diagonal line) over 1983-2008.

Intradaily and overnight returns


Vjump1b



Vjump2




Vjump3




Daily returns

Vjump1b



Vjump2



Vjump3




Figure 7A: Daily realized variance $RV_n / 252\tau_n$




Figure 7B: Daily realized volatility $\sqrt{RV_n / 252\tau_n}$, on a log scale






Figure 8: Normal probability plots for realized variance. These plot the ordered normalized residual values $z_{n+1} = N^{-1}[CDF(RV_{n+1}|Y_{n,M_n}, \hat{\theta})]$ (+ symbols on the vertical axis) against predicted values (green dotted lines) based upon quadratic variation's conditional CDFs, and against average values (red solid lines) from 100 runs of simulated data. The grey shaded areas are 95% confidence intervals for the deviation between observed data and average simulated values.

Vjump1b**Vjump2****Vjump3**