Risk, jumps, and diversification

Tim Bollerslev^{a,1}, Tzuo Hann Law^b, George Tauchen^{a,*}

^aDepartment of Economics, Duke University, Durham, NC 27708-0097, USA

^b76, Lorong 7, Taman Suria, 34000 Taiping, Malaysia

Available online 13 February 2008

Abstract

We test for price discontinuities, or jumps, in a panel of high-frequency intraday stock returns and an equiweighted index constructed from the same stocks. Using a new test for common jumps that explicitly utilizes the cross-covariance structure in the returns to identify non-diversifiable jumps, we find strong evidence for many modest-sized, yet highly significant, *cojumps* that simply pass through standard jump detection statistics when applied on a stock-by-stock basis. Our results are further corroborated by a striking within-day pattern in the significant *cojumps*, with a sharp peak at the time of regularly scheduled macroeconomic news announcements.

© 2008 Published by Elsevier B.V.

JEL classification: C12; C22; C32; G12; G14

Keywords: Jump-diffusions; Stock returns; Diversification; Tests for jumps; Cojumps; High-frequency data; Bipower variation

1. Introduction

We examine the relationship between jumps in individual stocks and jumps in an aggregate market index. Several studies have recently presented strong non-parametric high-frequency data-based empirical evidence in favor of jumps in financial asset prices, thus discrediting the classical continuous time paradigm with continuous sample price paths in favor of one with jump discontinuities.² The tests that we implement here further corroborate this evidence for the presence of jumps at both the individual stock and aggregate market level. Of particular interest, however, are the contrasts between the outcome of tests for jumps conducted at the level of the individual stocks and the level of the index. Jumps in individual stocks can be generated by either stock-specific news or common market-level news. Meanwhile, from basic portfolio theory one might expect that jumps in a well-diversified index should only be generated by market-level news that induces *cojumps* across many stocks. Below, we formalize these ideas and develop a new test for *cojumps* motivated

*Corresponding author. Tel.: +1 919 660 1812.

E-mail addresses: bollev@econ.duke.edu (T. Bollerslev), tzuohann@gmail.com (T.H. Law), george.tauchen@duke.edu (G. Tauchen).

¹Tel.: +1 919 660 1846.

²Empirical studies and new statistical procedures include the work of Ait-Sahalia and Jacod (2008), Andersen et al. (2007b), Barndorff-Nielsen and Shephard (2004), Fan and Wang (2006), Huang and Tauchen (2005), Jiang and Oomen (2006), Lee and Mykland (2008), and Mancini (2004).

Fig. 1. Price of PG sampled every 30 s on August 3, 2004.

directly by these observations. As such, the test explicitly reflects all of the cross-covariances within a large panel of high-frequency returns.

Jumps are clearly of importance for asset allocation and risk management. A risk averse investor might be expected to shun investments with sharp unforeseeable movements. As an example of a jump, consider Fig. 1, which depicts the price of Proctor and Gamble (PG) on August 3, 2004 sampled once every 30 s. There is a sharp discontinuity in the evolution of the prices a little after 11 am, where PG gains about 30 cents over a period of just 2 min. Jumps like that are of great importance for standard arbitrage-based arguments and derivatives pricing in particular, as the effect cannot readily be hedged by a portfolio of the underlying asset, cash, and other derivatives.³

Of course, not all jumps are as sharply *ex post* identifiable as that shown in Fig. 1, so that a formal statistical methodology for identifying jumps is needed. In the results reported on below, we start with a univariate analysis of stock-by-stock jump tests based on the seminal work by Barndorff-Nielsen and Shephard (2004) (BN–S). The BN–S theory provides a convenient non-parametric framework for measuring the relative contribution of jumps to total return variation and for classifying days on which jumps have or have not occurred.

Our empirical investigation is based on high-frequency intraday returns for a sample of forty large-cap U.S. stocks and the corresponding equiweighted index of these same stocks over the 2001–2005 sample period. Consistent with previous empirical results, we find strong evidence for the presence of jumps in each of the individual stock price series as well as the aggregate index. Standard computations also indicate about 15–25 large jumps for each of the individual stocks scattered randomly across the five-year sample, with jumps accounting for 12% of the total variation on average. In contrast, for the aggregate index there are only seven highly significant large jumps across the whole sample, and jumps as a whole account for just about 9% of the total return variation. Although these specific numbers obviously rely on our use of the popular BN–S procedure for identifying jumps, the same basic findings of more frequent and larger sized jumps for the individual stocks compared to the index is entirely consistent with the limited empirical evidence based on the univariate threshold-type statistic recently reported in Lee and Mykland (2008).

³More formally, the effect of a jump on the price of a derivative is locally nonlinear and thereby cannot be neutralized by holding an *ex ante*-determined portfolio of other assets.

The fact that the data reveal a less important role for jumps in the index than in each of its components is not all surprising, since the idiosyncratic jumps should be diversified away in the aggregate portfolio. At the same time, however, there is at best a very weak positive association between the significant jumps in each of the individual stocks and the jumps in the index. The BN–S jump detection procedure relies on a daily z -statistic, with large positive values of the statistic discrediting the null hypothesis of no jumps on that day. We find the z -statistics for the individual stocks to be essentially uncorrelated with the z -statistics for the index, despite the fact that most of the individual stocks have β 's close to unity with respect to the index.

As one might conjecture, this low degree of association appears to be due to the large amount of idiosyncratic noise in individual returns. This noise masks the cojumps, making them nearly undetectable at the level of the individual return. To circumvent this problem, we develop a new cojump test applicable to situations involving large panels of high-frequency returns.⁴ Using this new test statistic we find strong evidence for many modest-sized common jumps that simply pass through the standard univariate jump detection statistic, while they appear highly significant in the cross-section based on the new cojump identification procedure.

As noted by Andersen et al. (2007b), Barndorff-Nielsen and Shephard (2006), Eraker et al. (2003), and Huang and Tauchen (2005) among others, many, although not all, of the statistically significant jumps in aggregate stock indexes coincide with macroeconomic news announcements and other *ex post* readily identifiable broad-based economic news which similarly impact financial markets in a systematic fashion.⁵ While macroeconomic events obviously also affect individual firms, individual stock prices are also affected by sudden unexpected firm-specific information that can force an abrupt revaluation of the firms' stock.⁶ Our results suggest that firm-specific news events are indeed the dominant effect of the two in terms of their immediate price impact at the individual stock level, and only by properly considering the cross-section of returns do the non-diversifiable cojumps become visible in a formal statistical sense. These results are further corroborated by our findings of strong intraday patterns in the importance of jumps across the day, with the peak in the pattern for the aggregate index closely aligned with the time of the release of regularly scheduled news announcements.

The rest of the paper proceeds as follows. Section 2 outlines the basic second-moment theory that underlies most of the recent univariate high-frequency data-based jump detection work, along with the relevant multivariate large-panel extensions. Section 3 briefly reviews the standard univariate (BN–S) jump test and develops our new test for cojumps based on the multivariate extensions from the preceding section. Section 4 discusses the high-frequency data and sampling schemes underlying our empirical analysis, with some of the details relegated to Data Appendix. Section 5 presents our main empirical findings, while Section 6 concludes with a few final remarks.

2. Generalized quadratic variation theory

2.1. Univariate second-moment variation measures

We start by considering the i th log-price process $p_i(t)$ from a collection of n processes $\{p_i(t)\}_{i=1}^n$ evolving in continuous time. We assume that $p_i(t)$ evolves as

$$dp_i(t) = \mu_i(t) dt + \sigma_i(t) dw_i(t) + d\mathcal{L}_{i,J}(t), \quad (1)$$

where $\mu_i(t)$ and $\sigma_i(t)$ refer to the drift and local volatility, respectively, $w_i(t)$ is a standard Brownian motion, and $\mathcal{L}_{i,J}(t)$ is a pure jump Lévy process.⁷ A common modeling assumption for the Lévy process is the compound Poisson process, or rare jump process, where the jump intensity is constant and the jump sizes are

⁴Alternative statistical procedures for identifying common jump arrivals in pairs of returns have recently been developed by Jacod and Todorov (2007) and Gobbi and Mancini (2007).

⁵Notable examples include the monthly employment report, FED interest rate changes, oil prices, legislative alterations, and security concerns.

⁶For instance, lawsuits against a cigarette company, announcements of war for a defense company, or legislation of privacy issues for an Internet search engine.

⁷This particular notation is adopted from Basawa and Brockwell (1982).

independent and identically distributed. Throughout the paper we adopt the timing convention that one time-unit corresponds to a trading day, so that $\{p_i(t-1+s)\}_{s \in [0,1]}$ represents the continuous log-price record over trading day t , where the integer values $t = 1, 2, 3, \dots$ coincide with the end of the day.

In practice the price process is only available at a finite number of points in time. Let $M+1$ denote the number of equidistant price observations each day; i.e., $p_i(t-1), p_i(t-1+1/M), \dots, p_i(t)$. The j th within-day return is then defined by

$$r_{i,t,j} = p_i\left(t-1 + \frac{j}{M}\right) - p_i\left(t-1 + \frac{j-1}{M}\right), \quad j = 1, 2, \dots, M, \quad (2)$$

for a total of M returns per day.

As discussed at length in, e.g., Andersen et al. (2008), the *realized variance*

$$RV_{i,t} = \sum_{j=1}^M r_{i,t,j}^2 \quad (3)$$

provides a natural measure of the daily *ex post* variation. In particular, it is well known that for increasingly finer sampling frequencies, or $M \rightarrow \infty$, $RV_{i,t}$ consistently estimates the total variation comprised of the integrated variance plus the sum of the squared jumps

$$\lim_{M \rightarrow \infty} RV_{i,t} = \int_{t-1}^t \sigma_i^2(s) ds + \sum_{k=1}^{N_{i,t}} \kappa_{i,t,k}^2, \quad (4)$$

where $N_{i,t}$ denotes the number of within-day jumps on day t , and $\kappa_{i,t,k}$ refer to the size of the k th such jump.

In order to separately measure the two components that make up the total variation in (4), Barndorff-Nielsen and Shephard (2004) and Barndorff-Nielsen et al. (2005) first proposed the so-called *bipower variation* measure

$$BV_{i,t} = \mu_1^{-2} \left(\frac{M}{M-1} \right) \sum_{j=2}^M |r_{i,t,j-1}| |r_{i,t,j}|, \quad (5)$$

where $\mu_1 = \sqrt{2/\pi} \approx 0.7979$. Under reasonable assumptions, it follows that

$$\lim_{M \rightarrow \infty} BV_{i,t} = \int_{t-1}^t \sigma_i^2(s) ds, \quad (6)$$

so that $BV_{i,t}$ consistently estimates the integrated variance for the i th price process, even in the presence of jumps. Thus, as such the contribution to the total variation coming from jumps may be estimated by $RV_{i,t} - BV_{i,t}$, or the *relative jump* measure advocated by Huang and Tauchen (2005),⁸

$$RJ_{i,t} = \frac{RV_{i,t} - BV_{i,t}}{RV_{i,t}}. \quad (7)$$

It follows that in the limit as $M \rightarrow \infty$, $RJ_t > 0$ only on days for which there are at least one jump, although for finite M sampling variation can occasionally result in $RJ_{i,t} < 0$.

2.2. A portfolio-theoretic multivariate extension

The preceding subsection lays out the by now fairly standard univariate second-moment variation theory. However, from the point of view of financial economics, and in particular portfolio theory, we are generally more interested in the behavior of a large ensemble of stock prices and returns. For a particular stock, say the i th, the realized bipower variation $BV_{i,t}$ in (5) measures the combined common and idiosyncratic components of the continuous part of the total daily variation, while the realized variance $RV_{i,t}$ in (3) measures this *plus* the contributions of both common jumps and idiosyncratic jumps. For a large portfolio that is well diversified in

⁸An equivalent ratio statistic, $-RJ_{i,t}$, was proposed and studied independently by Barndorff-Nielsen and Shephard (2006).

the sense of the *Arbitrage Pricing Theory*, one might expect that the idiosyncratic jumps are diversified away and only the common, or *cojumps* remain. This indeed turns out to be the case. For simplicity, we only consider an equiweighted portfolio, but the arguments below hold with equal force for any well-diversified portfolio; there would just be the notational nuisance of including the portfolio weights in the proper places in the formulas.

Consider the j th within-day return on an equiweighted portfolio of n stocks

$$r_{\text{EQW},t,j} = \frac{1}{n} \sum_{i=1}^n r_{i,t,j}.$$

Extending the arguments for the individual stocks with the same logic underlying equations (4) and (6), the realized variation for the EQW portfolio must satisfy

$$\begin{aligned} RV_{\text{EQW},t} &= \sum_{j=1}^M \left(\frac{1}{n} \sum_{i=1}^n r_{i,t,j} \right)^2 \xrightarrow{M \rightarrow \infty} \frac{1}{n^2} \sum_{i=1}^n \int_{t-1}^t \sigma_i^2(s) ds + \frac{1}{n^2} \sum_{i=1}^n \sum_{\ell=1, \ell \neq i}^n \int_{t-1}^t \sigma_i(s) \sigma_\ell(s) ds \\ &\quad + \frac{1}{n^2} \sum_{i=1}^n \sum_{k=1}^{N_{i,t}} \kappa_{i,t,k}^2 + \frac{1}{n^2} \sum_{i=1}^n \sum_{\ell=1, \ell \neq i}^n \sum_{k=1}^{N_t^*} \kappa_{i,t,k} \kappa_{\ell,t,k}, \end{aligned}$$

where the last sum only includes the N_t^* jumps (a random number) that occur simultaneously across all n stocks. Similarly,

$$\begin{aligned} BV_{\text{EQW},t} &= \mu_1^{-2} \left(\frac{M}{M-1} \right) \sum_{j=2}^M \left| \frac{1}{n} \sum_{i=1}^n r_{i,t,j-1} \right| \cdot \left| \frac{1}{n} \sum_{i=1}^n r_{i,t,j} \right| \\ &\xrightarrow{M \rightarrow \infty} \frac{1}{n^2} \sum_{i=1}^n \int_{t-1}^t \sigma_i^2(s) ds + \frac{1}{n^2} \sum_{i=1}^n \sum_{\ell=1, \ell \neq i}^n \int_{t-1}^t \sigma_i(s) \sigma_\ell(s) ds. \end{aligned}$$

Thus, it follows readily that

$$\begin{aligned} RV_{\text{EQW},t} - BV_{\text{EQW},t} &\xrightarrow{M \rightarrow \infty} \frac{1}{n^2} \sum_{i=1}^n \sum_{k=1}^{N_{i,t}} \kappa_{i,t,k}^2 + \frac{1}{n^2} \sum_{i=1}^n \sum_{\ell=1, \ell \neq i}^n \sum_{k=1}^{N_t^*} \kappa_{i,t,k} \kappa_{\ell,t,k} \\ &= \frac{1}{n} \overline{\kappa_t^2} + \frac{n-1}{n} \sum_{k=1}^{N_t^*} \overline{\kappa_{\cdot,t,k} \kappa_{\cdot,t,k}}, \end{aligned}$$

where the bars the $\cdot$ and $\cdot\cdot$ in the second term refer to the average over the relevant indices. For n large, the first term becomes negligible and $(n-1)/n \approx 1$ so that

$$RV_{\text{EQW},t} - BV_{\text{EQW},t} \approx \sum_{k=1}^{N_t^*} \overline{co_{t,k}}, \quad (8)$$

where $\overline{co_{t,k}}$ denotes the k th average cojump. In other words, in a large well-diversified portfolio only cojumps that occur simultaneously across many assets can cause the portfolio to jump; i.e., jumps that pervade the market.

3. Jump test statistics

3.1. The (univariate) BN–S approach

High-frequency data-based testing for jumps in a particular stock or financial instrument has received a lot of attention in the recent literature, and several different univariate test procedures exist for possible use. In this paper, however, we concentrate on the popular BN–S methodology developed in a series of influential papers starting with [Barndorff-Nielsen and Shephard \(2004\)](#). The BN–S methodology essentially kick-started

the entire line of inquiry, and it is by far the most developed and widely applied of the different methods. In addition to the extensive Monte Carlo validation provided in Huang and Tauchen (2005), empirical studies that have used the approach include the work of Andersen et al. (2007b) and Todorov (2006) among many others. As such, the BN–S methodology arguably represents *the* standard approach for non-parametric univariate jump detection on a day-by-day basis.

All jump test statistics in the BN–S methodology work by forming a measure of the discrepancy between $RV_{i,t}$ and $BV_{i,t}$ and then studentizing the difference. Of the various asymptotically equivalent possibilities,⁹ Barndorff-Nielsen and Shephard (2004) observe that a studentized version of the relative contribution measure $RJ_{i,t}$ in (7) can be expected to perform particularly well, since it largely mitigates the effects of level shifts in variance associated with time varying stochastic volatility. This conjecture is corroborated by the Monte Carlo evidence in Huang and Tauchen (2005), which indicates that the test statistic

$$z_{i,t} = \frac{RJ_{i,t}}{\sqrt{(v_{bb} - v_{qq}) \frac{1}{M} \max\left(1, \frac{TP_{i,t}}{BV_{i,t}^2}\right)}}, \quad (9)$$

where $v_{qq} = 2$, $v_{bb} = (\pi/2)^2 + \pi - 3 \approx 2.6090$, and the *tripower quarticity* is defined by

$$TP_{i,t} = \mu_{4/3}^{-3} M \left(\frac{M}{M-2} \right) \sum_{j=3}^M |r_{i,t,j-2}|^{4/3} |r_{i,t,j-1}|^{4/3} |r_{i,t,j}|^{4/3}, \quad (10)$$

and $\mu_{4/3} = 2^{2/3} \Gamma(\frac{7}{6}) / \Gamma(\frac{1}{2}) \approx 0.8309$, closely approximates a standard normal distribution under the null hypothesis of no jumps; i.e., $z_{i,t} \rightarrow N(0, 1)$.¹⁰ Moreover, the statistic also exhibits favorable power properties and, in general, provides an excellent basis for univariate jump detection.

3.2. The mcp cojump test

The discussion in Section 2.2 above suggests that the empirical results obtained by applying the BN–S z -statistic on a stock-by-stock basis, as we do farther below, need to be interpreted with caution in a portfolio context. Of course, there is nothing inherently wrong with the univariate z -statistic; the issues are more interpretative. Idiosyncratic jumps that cannot carry a risk premium can trigger the univariate test while, from a systematic risk perspective, we are primarily interested in the all-important cojumps in Eq. (8). Conversely, the large amount of noise in individual stock returns can potentially mask the contribution of a modest-sized cojump to a particular stock return, as appears to be the case in our empirical results below.

Testing for cojumps using high-frequency data is also an area of active interest. In early work, Barndorff-Nielsen and Shephard (2003) present a very clever way to adapt the univariate bipower approach to test for cojumps between a pair of returns, but the theory is difficult to implement due to the lack of a tractable way to estimate the divisor for the studentization in a jump-robust manner as in (9). More recently, Jacod and Todorov (2007) present a test for cojumps between a pair of returns based on higher order power variation, while Gobbi and Mancini (2007) propose a strategy to separate out the covariation between the diffusive and jump components of a pair of returns.

In this paper, however, we are not directly interested in cojumps between a particular pair of returns, but rather in the cojumps embodied in a large ensemble of returns. The expressions derived in Section 2.2 naturally suggest such a test statistic for cojumps. In particular, motivated by Eq. (8), we consider the *mean cross-product* statistic defined by the normalized sum of the individual high-frequency returns for each

⁹All asymptotic analysis is “fill-in” where the sampling interval goes to zero with the span of the observations held fixed.

¹⁰The denominator of the statistic in (9) also incorporates Jensen’s inequality-type adjustment for the (asymptotic) relationship between $TP_{i,t}$ and $BV_{i,t}^2$. Other jump-robust estimators for the integrated quarticity based on the summation of adjacent returns raised to powers less than two where the powers sum to four are also possible; see Barndorff-Nielsen and Shephard (2004). However, the Monte Carlo evidence in Huang and Tauchen (2005) indicates that the tripower quarticity in (10) performs quite well.

within-day period:

$$mcp_{t,j} = \frac{2}{n(n-1)} \sum_{i=1}^{n-1} \sum_{\ell=i+1}^n r_{i,t,j} r_{\ell,t,j}, \quad j = 1, 2, \dots, M. \quad (11)$$

The mcp -statistic provides a direct measure of how closely the stocks move together. It is entirely analogous to a U -statistic.

For each day in the sample we have a realization $\{mcp_{t,j}\}_{j=1}^M$ of length M . Summing the $mcp_{t,j}$ statistics across day t , it follows readily that

$$mcp_t = \sum_{j=1}^M mcp_{t,j} = \frac{1}{(n-1)} \left[nRV_{\text{EQW},t} - \frac{1}{n} \sum_{i=1}^n RV_{i,t} \right]. \quad (12)$$

Hence, from reasoning much like that underlying the BN-S bipower variation statistic in (5), we can expect the mcp -statistic to be reasonably insensitive to idiosyncratic jumps in the individual stocks. At the same time, the statistic is clearly very sensitive to cojumps and it will assume a large positive value whenever the ensemble of returns takes a large (positive or negative) move together. These characteristics become more apparent if we use $n/(n-1) \approx 1$, for large n , and re-write Eq. (12) as

$$mcp_t \approx BV_{\text{EQW},t} + (RV_{\text{EQW},t} - BV_{\text{EQW},t}) - \frac{1}{n} \overline{RV}_t, \quad (13)$$

where

$$\overline{RV}_t = \frac{1}{n} \sum_{i=1}^n RV_{i,t}.$$

Thus, mcp_t is essentially comprised of the overall continuous variance, $BV_{\text{EQW},t}$, and the overall jump contribution, $RV_{\text{EQW},t} - BV_{\text{EQW},t}$, as defined in (8); everything else is diversified away as $n \rightarrow \infty$. It is important to keep in mind that it is not possible using just observations on the equiweighted index alone to estimate directly the bipower and realized variations for EQW on a tick-by-tick basis, or for every j , as required by the mcp -statistic in (11). The large cross-section of returns is crucial for the implementation of the new mcp test.

Meanwhile, as is clear from Eq. (13), the cross-correlation among the diffusive components of the returns implies that the mcp -statistic in (11) does not have mean zero even in the absence of cojumps. Furthermore, the overall scale of the statistic fluctuates randomly with the general level of stochastic volatility in the market. We thus studentize the statistic as

$$z_{mcp,t,j} = \frac{mcp_{t,j} - \overline{mcp}_t}{s_{mcp,t}}, \quad j = 1, 2, \dots, M, \quad (14)$$

where

$$\overline{mcp}_t = \frac{1}{M} mcp_t = \frac{1}{M} \sum_{j=1}^M mcp_{t,j}$$

and

$$s_{mcp,t} = \sqrt{\frac{1}{M-1} \sum_{j=1}^M (mcp_{t,j} - \overline{mcp}_t)^2}.$$

In practice we find that the $mcp_{t,j}$ realizations are essentially serially uncorrelated for $j = 1, 2, \dots, M$. Hence, we simply standardize each of the within-day mcp -statistics by their corresponding daily sample standard deviation, $s_{mcp,t}$. The underlying presumption is that the location and scale are approximately constant within-days, but varies across days, which is consistent with the idea of a slowly varying stochastic return volatility.

As discussed further below, the presumption can be challenged by the known U-shaped pattern of within-day equity volatility, but it turns that this does not affect our main findings.

In order to implement the z_{mcp} -statistic in (14) as a test for cojumps, we need its distribution under the null of no jumps. Since the mcp -statistic in (11) is itself the average of $n(n-1)/2$ random variables, where $n = 40$ in our data set, one might expect the distribution of the z_{mcp} -statistic to be approximately standard $N(0, 1)$. In particular, Borovkova et al. (2001) give a central limit theorem for U -statistics for dependent data that would apply to the z_{mcp} -statistic for $n \rightarrow \infty$ with t and j fixed if the dependence (mixing conditions) among the individual stocks was sufficiently weak. However, the cross-correlation among stock returns is unlikely to satisfy these necessary conditions. For instance, it is fairly easy to show that under a simple one-factor representation for the returns, the mcp -statistic will, for large n , equal $\mu_\beta^2 \chi^2(1)$, where μ_β^2 is the square of the average (arithmetic) β . The average β is arguably unity, which gives an asymptotic $\chi^2(1)$ representation for the mcp -statistic. Still, the simple $\chi^2(1)$ representation does not hold in empirically more realistic multi-factor situations, and as such is of limited practical use. Lacking a viable approximate asymptotic distribution for the z_{mcp} -statistic, we thus adopt below a straightforward bootstrap procedure to get its distribution under the null of no jumps.

4. Data

4.1. Date source and sampling

Our original sample consists of all trades on 40 large capitalization stocks over the January 1, 2001 to December 31, 2005, five-year sample period.¹¹ While the BN–S jump detection scheme is based on the notion of $M \rightarrow \infty$, or ever finely sampled high-frequency returns, a host of practical market microstructure complications prevents us from sampling too frequently while maintaining the fundamental semimartingale assumption underlying equation (1). Ways in which to best deal with these complications and the practical choice of M are the subject of intensive ongoing research efforts; see, e.g., Ait-Sahalia et al. (2005), Bandi and Russell (2008), Barndorff-Nielsen et al. (2006), and Hansen and Lunde (2006). In the analysis reported on below, we simply follow most of the literature in the use of a coarse sampling frequency as a way to strike a reasonable balance between the desire for as finely sampled observations as possible on the one hand and the desire not to be overwhelmed by the market microstructure noise on the other hand. The volatility signature plots advocated by Andersen et al. (2000), as further detailed in the Data Appendix, suggest that a choice of $M = 22$, or 17.5 min sampling, strikes such a balance and largely mitigates the effect of the “noise” for all of the 40 stocks in the sample.¹²

In addition to high-frequency returns for each of the individual stocks, we also construct an equiweighted portfolio comprised of the same 40 stocks. We will refer to this index as EQW in the sequel. It is noteworthy that at the 17.5 min sampling frequency, the correlation between the return on EQW and the return on the exchange-traded SPY fund, which tracks the S&P 500, equals 0.93.¹³ As such, the return on the EQW index may be thought of as being representative of the return on the aggregate market, and we will sometimes refer to the EQW index as the market portfolio in the sequel.

4.2. An illustrative look at PG

Before presenting the results for all of the 40 stocks and the EQW index, it is instructive to look at the different variation measures and the BN–S test for a single stock and a few specific days in the sample. To this end, Fig. 2 shows the price (adjusted for stock splits) and returns for PG. While the price appears to be steadily

¹¹The criteria used in selecting the 40 stocks are discussed in more detail in the Data Appendix. The specific ticker symbols for each of the stocks are included in many of the figures.

¹²For simplicity we decided to maintain the identical sampling frequency across all stocks. However, $M = 22$ is clearly a conservative choice for many of the stocks, and we also calculated the same statistics based on finer sampled high-frequency returns. Subsampling techniques could also be used to further enhance the accuracy of the realized variation measures and test statistics. The results from some of these additional calculations are briefly discussed below.

¹³At the 5-min interval the correlation is 0.88.

Fig. 2. Prices and returns for PG from 2001 to 2005, $M = 22$.

Fig. 3. Return, realized variance, bipower variation, relative jump, and z_t -statistic for PG from 2001 to 2005, $M = 22$.

increasing over the 1,241 trading days in the 2001–2005 sample, the corresponding 27,302 high-frequency 17.5 min returns are all seemingly scattered around zero. At the same time, the return plot clearly indicates the presence of volatility clustering. This is further underscored by Fig. 3, which plots the RV_t and BV_t variation

Fig. 4. Prices and returns for PG on March 26–27, 2001, $M = 22$.

measures, the relative jump contribution RJ_t , along with the BN–S z -statistic in Eq. (9). Comparing the BN–S test statistic in the lower plot to the horizontal reference line for the 99.9% significance level included in the plot, indicates that PG jumped at least once during the active part of the trading day on 17 days in the sample.

To further illustrate the working of the jump detection scheme, Fig. 4 shows the intraday prices and returns for PG on March 26 and 27, 2001. At a first glance it appears as if the price evolves rather smoothly on the 26th, while it shows a rather sharp increase just before noon on the 27th. Nonetheless, the corresponding BN–S z -statistics equal to 3.63 and -0.20 for each of the two days, respectively, suggest at least one highly significant jump on the 26th and no jumps on the 27th. Importantly, the jump test statistic depends on both the magnitude of the largest price change(s) over the day and the overall level of the volatility for the day. As such, relatively small price changes may be classified as jumps on otherwise calm days, while apparent discontinuities may be entirely compatible with a continuous sample path process in a statistical sense on very volatility days. Indeed, looking at the plot of the returns in the lower panel, the overall level of the volatility was clearly much higher on the 27th than on the 26th, with the BV_t estimate for the continuous sample path variation for each of the two days equal to 5.11 and 1.04, respectively. Thus, while some of the statistically significant jumps identified by the BN–S test can rather easily be spotted by *ex post* visual inspection of the intraday prices that is not necessarily always the case.

5. Jumps and cojumps

5.1. Univariate analysis

We start with a univariate analysis using the standard BN–S approach applied to the index and each of the individual stocks on a stock-by-stock basis. Fig. 5 shows the number days over the 2001–2005 sample period on which the BN–S z -statistics indicate that the EQW index and each of the 40 stocks comprising the index jumped. A day is classified as containing at least one jump if the z -statistic exceeds the critical value of the Gaussian distribution at the 0.001 significance level. From the figure, the index is seen to jump on far fewer

Fig. 5. Number of flagged jump days for EQW and the 40 individual stocks from 2001 to 2005.

Table 1
Jump counts at different sampling frequencies

	Frequency		Jump count			Ratios	
	Minutes	M	EQW	SPY	STK	EQW/STK	SPY/STK
0.001 level	38.5	10	2	4	9.0	0.22	0.44
	35.0	11	5	7	11.7	0.43	0.60
	27.5	14	9	8	15.8	0.57	0.51
	17.5	22	7	11	22.1	0.32	0.50
	11.0	35	16	12	26.1	0.61	0.46
	7.0	55	33	24	31.5	1.05	0.76
	5.5	70	27	19	34.7	0.78	0.55
	5.0	77	32	38	36.7	0.87	1.04
	3.5	110	46	80	53.9	0.85	1.48
	2.5	154	46	100	87.2	0.53	1.15
	1.0	385	100	445	513.1	0.19	0.87
0.01 level	38.5	10	48	35	63.0	0.76	0.56
	35.0	11	47	47	66.2	0.71	0.71
	27.5	14	53	43	68.7	0.77	0.63
	17.5	22	38	45	71.8	0.53	0.63
	11.0	35	58	55	88.7	0.65	0.62
	7.0	55	77	78	90.6	0.85	0.86
	5.5	70	83	88	96.7	0.86	0.91
	5.0	77	112	123	101.1	1.11	1.22
	3.5	110	122	170	133.7	0.91	1.27
	2.5	154	136	219	192.0	0.71	1.14
	1.0	385	205	591	717.1	0.29	0.82
0.05 level	38.5	10	137	132	177.5	0.77	0.74
	35.0	11	156	158	180.6	0.86	0.87
	27.5	14	129	131	174.8	0.74	0.75
	17.5	22	110	121	175.5	0.63	0.69
	11.0	35	172	154	186.1	0.92	0.83
	7.0	55	202	208	204.0	0.99	1.02
	5.5	70	208	215	214.3	0.97	1.00
	5.0	77	255	250	220.2	1.16	1.14
	3.5	110	267	329	272.6	0.98	1.21
	2.5	154	290	398	355.8	0.82	1.12
	1.0	385	342	755	903.0	0.38	0.84

days than its components; the count for EQW is seven days while there are 22 jump days on average for the individual stocks. This finding is a simple consequence of diversification as discussed in Section 2.2.

As a further robustness check, and to guard against non-synchronous trading effects in the EQW portfolio returns, Table 1 shows the number of detected jumps in the EQW index, the SPY cash index, and the average over the 40 stocks for a wide range of different sampling frequencies (M ranging from 10 to 385) and three significance levels (0.001, 0.01, and 0.05). There is obviously a fairly close agreement between the number of jumps detected in the EQW index and the SPY down to the 5-min level, after which microstructure effects appear to cause a divergence. Our finding that the index jumps less than its components on average is also robust, as the ratio between the number of jumps in the index to the average number of jumps in the individual stocks systematically remains below one for most sampling frequencies and significance levels. Again, this holds true for the EQW index as well as the exchange-traded SPY fund.

The natural question is how jumps detected in the index relate to jumps in its components. The top panel of Fig. 6 shows the estimated β on EQW for each of the 40 stocks. The β 's are all fairly close to unity, which is

Fig. 6. Five-year high-frequency data-based β 's for each of the stocks, and regression slopes and correlations between $z_{Eqw,t}$ and $z_{t,i}$'s.

Fig. 7. Scatter plot of average z -statistics for the individual stocks and the corresponding z -statistic for EQW.

not surprising for large-cap stocks. The β 's suggest that the returns on each stock moves approximately one-for-one with the index return, apart from the idiosyncratic component. In particular, large movements, or jumps, in the index should show up as large movements or jumps in the stock prices. On the basis of the estimated β 's, one might expect a reasonably high correlation between a measure of the likelihood of a jump in the index and a jump in the individual stock prices. The BN-S z -statistic (9) is such a measure. The middle panel of Fig. 6 shows the 40 regression slope coefficients while the bottom panel shows the 40 estimated correlations between the stock z -statistics and the EQW z -statistic. Interestingly, the correlations are exceedingly low: rarely above 0.05 and frequently lower than 0.01, while the regression slope coefficients are similarly small and insignificant.

The low correlation between the jump-test z -statistics at the level of the individual stocks and the z -statistics for the EQW portfolio is further illustrated in Fig. 7. The figure shows a scatter plot of the average BN-S z -statistics for the 40 individual stocks against the z -statistic for EQW for each of the 1,241 days in the sample. The relationship is obviously very weak.¹⁴

5.2. Extracting evidence on cojumps from the cross-section

The outcomes of the standard univariate BN-S jump tests discussed immediately above should be interpreted with caution. In particular, as previously noted in Section 2.2, the tests for the individual stocks might not pick up systematic, or cojumps, due to the presence of much larger idiosyncratic jumps and potentially a low signal-to-noise ratio in the individual stock returns. We thus turn to the studentized *mean cross-product* statistic, z_{mcp} , in (14).

To implement the statistic, we bootstrap its null distribution under the assumption of no jumps. To do so, we computed values of the *mcp*-statistic in (11) along realizations of a 40×1 diffusion with zero drift and covariance matrix determined by the unconditional covariance matrix of the within-day returns. We simulated realizations of length equal to the sample size (22 steps per day for 1,241 days) replicated 1,000

¹⁴For those days on which the z -statistic for EQW is statistically significant at the 0.001 level, the average z -statistic across the 40 stocks is just about 1.0. Of course, as pointed out by Roel Oomen, under the null hypothesis of no jumps in *any* of the 40 stocks the distribution of the average of the individual z -statistics will not be standard normal, but rather normally distributed with a mean of zero and a variance of $\frac{1}{40}$. Thus, a value for the average in excess of 0.489 may reasonably be deemed significant at the 0.001 level when interpreted as a test for no jumps in *any* of the 40 stocks.

Fig. 8. The top panel shows the data-based bootstrap distribution of the z_{mcp} -statistic with $n = 40$ stocks and $M = 22$ within-day returns. The bottom panel shows the $Q-Q$ plot of the empirical distribution for the observed z_{mcp} -statistics relative to the bootstrap distribution.

times. This scheme generates just over 27.3 million simulated values under the null of no jumps. The top panel of Fig. 8 shows the resulting bootstrapped probability density of the z_{mcp} -statistic. The distribution is evidently highly non-Gaussian with a strong right skew.¹⁵ The 99.9% quantile, or critical value for a 0.001 level test for no cojumps, equals 4.145.

As a check on the sensitivity of this critical value to the covariance structure of the returns, the number of stocks, and the number of within-day returns, we recomputed the bootstrap distribution using equicorrelated returns based on different values of $\bar{\rho}$, n and M . The results displayed in Table 2 reveal that the critical value is quite insensitive to the level of correlation among the returns and to the number of stocks. It is, however, somewhat sensitive to the number of within-day returns, M . Intuitively, using the daily sample mean and standard deviation in studentizing the z_{mcp} -statistic in place of the population quantities, causes the distribution of the statistic to become more concentrated around zero.¹⁶ All together, however, it appears as if the bootstrap distribution shown in the top panel of Fig. 8 accurately captures the distribution of the z_{mcp} -statistic in the absence of jumps for the actual high-frequency panel analyzed here.

The $Q-Q$ plot in the bottom panel of Fig. 8, which depicts the quantiles of the empirical distribution versus the quantiles of the bootstrap distribution, reveals quite striking evidence for cojumps. The empirical distribution of the z_{mcp} -statistic is sharply right-shifted relative to the null distribution, and many of the sample z_{mcp} -statistics would be judged as statistically significant at most any commonly used significance level.

This strong empirical evidence for a significant number of cojumps is further corroborated by Fig. 9, which displays a scatter plot of the return on the EQW index against the z_{cp} -statistics. There are a total of

¹⁵ It has much the appearance of an affine transformation of a $\chi^2(1)$ which would be expected under a one factor model.

¹⁶ If the returns were perfectly correlated, the mcp -statistic would be distributed as $[\chi^2(1) - 1]/\sqrt{2}$ if population quantities were used to studentize; the 0.001 critical value would be 6.949. On the other hand, if $\rho = 0.995$, $M = 22$, and sample values are used to studentize, the bootstrap critical value equals just 4.266. Of course, for $\rho = 0.995$ and larger values of M the bootstrap critical value is much closer to the limiting value of 6.949.

Table 2
99.9% quantile of the bootstrapped z_{cp} distribution

	Average correlation $\bar{\rho}$	Quantile $Q_{0.999}$
<i>Data based</i>		
$n = 40, M = 22$	0.35	4.15
<i>Equicorrelated</i>		
$n = 40, M = 22$	0.00	4.12
	0.10	4.14
	0.50	4.15
	0.80	4.15
$n = 20, M = 22$	0.00	4.09
	0.10	4.14
	0.50	4.15
	0.80	4.15
$n = 40, M = 78$	0.00	5.83
	0.10	5.89
	0.50	5.89
	0.80	5.89
$n = 20, M = 78$	0.00	5.76
	0.10	5.87
	0.50	5.89
	0.80	5.89

1, $241 \times 22 = 27,302$ points in the figure. As noted above, the 99.9% bootstrap quantile of the distribution of the z_{cp} -statistic equals 4.145, and so we should expect only about 32 points in Fig. 9 to lie to the right of this cutoff. There are, however, far many more points than that to the right of the cutoff, although in many instances the associated return on the EQW index is only moderately large in magnitude. In other words, it appears that many modest-sized cojumps simply go undetected by the jump test statistic when applied to returns individually.

We believe the above conclusions concerning the presence of many modest-sized non-diversifiable cojumps to be quite persuasive. To buttress the empirical findings, we next discuss how the high-frequency data also reveal the existence of a strong intraday pattern in the occurrence of cojumps. We then present the results from additional Monte Carlo simulations designed to assess the performance of the *mcp*-test vis-a-vis a univariate analysis in situations with cojumps.

5.3. Intraday patterns

There is a long history dating back at least to Wood et al. (1985) and Harris (1986) documenting the existence of a distinct U-shaped pattern in equity return volatility over the trading day; i.e., volatility tend to be high at the open and close of trading and low in the middle of the day. This same general pattern, albeit more tilted towards the open, is also apparent from Fig. 10 and the plot of the unconditional variance and bipower variation for the EQW returns over the trading day.

Specifically, we compute the unconditional realized variance for tick j as

$$\frac{1}{T} \sum_{t=1}^T r_{EQW,tj}^2, \quad j = 1, \dots, M,$$

while the unconditional bipower variation centered at the j th tick is computed as

$$\frac{1}{T} \sum_{t=1}^T |r_{EQW,tj-1}|^{1/2} |r_{EQW,tj}| |r_{EQW,tj+1}|^{1/2}, \quad j = 2, \dots, M-1.$$

Fig. 9. Intraday EQW returns versus z_{mcp} -statistic.

Fig. 10. Average intraday realized and bipower variation for the EQW index.

Fig. 10 shows these unconditional variance measures at three increasingly finer sampling frequencies: $M = 22, 77$, and 385 , corresponding to 17.5, 5, and 1 min, respectively.¹⁷ The unconditional realized variance systematically lies above the unconditional bipower variation over the entire day, thus reflecting the existence of jumps across the day. Interestingly, however, the plot also reveals a sharp rise in the realized variation relative to the bipower variation at 10 am EST, corresponding to the time of the release of several regularly scheduled macroeconomic news announcements.¹⁸ This therefore indirectly suggests that some of the cojumps may be associated with these types of systematic news. This is also consistent with the work of Andersen et al. (2003, 2007a) among others, which document a significant response in high-frequency financial market prices to surprises in macroeconomic announcement immediately after the release of the news. Moreover, Andersen and Bollerslev (1998) among others have previously noted a sharp increase in the average total intraday volatility at the exact time of important macroeconomic news announcements. What is particularly noteworthy is the much less dramatic increase in the average within-day bipower variation measure, in turn attributing most of the variation at that specific time-of-day to jumps.

Instead of averaging all of the individual returns before calculating the intraday variation measures as in Fig. 10, Fig. 11 shows a similar plot in which the two variation measures are first computed on a stock-by-stock basis and then averaged across all of the 40 stocks in the panel. This explicitly excludes the effect of diversification, and as a result the vertical scale of Fig. 11 is much larger than that of Fig. 10. More importantly, however, comparing the general shape between the two sets of pictures, the increase in the within-day variation at 10 am is much less apparent for the individual stock averages. The relative importance of jumps also appears to be much more evenly distributed across the entire day, and as such lend further credence to our hypothesis of the cojumps drowning in the firm-specific variation inherent in the individual stock returns.

The z_{mcp} -statistic in (14) that we used in testing for cojumps was based on the assumption of constant scale within the day. The average within-day patterns in Fig. 11 clearly seem to violate this assumption. There is no obvious right way to adjust the individual returns for the intraday volatility patterns, as the relative importance of the jump variation to the total variation may be changing over the day. Fortunately, that does not seem to matter much for our main conclusions. The cojump statistics depicted in Fig. 9 were based on the raw unadjusted returns ignoring the intraday pattern. As a robustness check we redid the same calculations in which we scaled the return for the i th stock over the j th time-interval by the reciprocal of the square root of the corresponding unconditional bipower variation for that particular stock and time-interval. This adjustment is extreme, in that it deflates returns near the beginning and the end of the day while inflating returns in the middle of the day under the implicit presumption that the share of the jump variance remains constant over the day. Nonetheless, the resulting Fig. 12 depicting the returns on the EQW index against the adjusted z_{cp} -statistics, is so similar to the original Fig. 9 that we conjecture any reasonable adjustment for the intraday pattern would result in the same basic conclusions.

5.4. Power

Our main conclusions hinge on the argument that the z_{mcp} -statistic is more sensitive to modest-sized cojumps than the BN-S z -statistic applied stock-by-stock because it explicitly utilizes the cross-sectional information. To investigate this hypothesis further, we use another bootstrap-type procedure where we take the observed data set as given, sprinkle in additional simulated jumps, and then recompute the jump test statistics. Importantly, this circumvents the need to specify a complete data generating process for the full 40-dimensional vector return process, as would be required by a more traditional simulation-based procedure.

Specifically, for idiosyncratic jumps we simulated 40 independent Gaussian compound Poisson processes with intensity λ_i and magnitude $N(0, \sigma_{j,i}^2)$, while for the common jumps we simulated one Gaussian compound Poisson process with intensity λ and magnitude $N(0, \sigma_j^2)$. The idiosyncratic jumps are then added to the actually observed within-day 17.5 min returns for each of the individual stocks as are the common jumps

¹⁷For visual comparisons, we extended the bipower variation directly to the left for $j = 1$ and to the right for $j = M$.

¹⁸There is also a smaller less pronounced peak at 1:15 pm. However, other computations by Peter Van Tassel, not shown here, suggest that this early afternoon peak is more fragile. Also, the pattern in Fig. 10 for the EQW portfolio essentially mirrors that for the SPY.

Fig. 11. Average over the 40 stocks intraday realized and bipower variation.

Fig. 12. Intraday EQW returns versus z_{mcp} -statistic based on adjusted data.

multiplied by the stock's estimated β . This in turn yields 40 new returns series with the additional simulated jumps scattered throughout the sample. From these new series, we then recompute the EQW returns, the BN-S z -statistics, and the z_{mcp} -statistics. This whole process is replicated 1,000 times and the outcomes averaged across the replications. Since the baseline observed data already contains jumps, the simulations reveal the *incremental* strength, or power, of the tests to the different types of added jumps.

Fig. 13. Significant simulated common jumps; BNS_{EQW} : ·, $\overline{BNS}_i$: □, and z_{mcp} : △.

The findings for different jump sizes holding the intensities constant are readily described and intuitively plausible. In the absence of any common jumps, as the idiosyncratic jump sizes increase from 0% to 1.50%, the jump detection rate of the BN–S z -statistic applied to the individual stocks gradually increases. However, it remains unchanged when applied to the EQW return, as does the detection rate of the z_{mcp} -statistic. The idiosyncratic jumps are effectively diversified away. On the other hand, if the idiosyncratic jumps are left out while the size of the common jump increases from 0% to 1.50%, the detection rates of both the BN–S z -statistics and z_{mcp} -statistics increase, with the z_{mcp} -statistic being slightly more sensitive to larger jump sizes. Still, the contrasts between the two tests appear rather small, as they both fairly easily detect the occasional rare common jump.

The most revealing contrasts across the statistics are obtained by considering their ability to detect a relatively small common jump. Specifically, Fig. 13 shows the jump detection rates for a simulation experiment in which the idiosyncratic jumps are left out, the common jump size is set to 0.25%, and the intensity λ of the common jumps increases from zero to 200 per year. For the z_{mcp} -statistic, there are $M = 22$ within-day values, so any day on which the statistic is statistically significant is classified as a jump day; in the observed data set, there are no days on which more than one statistically significant z_{mcp} -statistic occurs. For the BN–S z -statistic, the figure shows the number of detected jumps using the BN–S z -statistic applied to the EQW returns as well as the average number of jumps detected from applying it one-by-one across the panel of 40 stocks. Note that for $\lambda = 0$ there are more jumps detected than the nominal size of the tests because the baseline data itself contains jumps. Nonetheless, as clearly seen from the figure, the detection rate of the z_{mcp} -statistic increases more sharply with λ than do the detection rates based on the BN–S statistic applied stock-by-stock. This directly confirms the intuition that utilizing the cross-covariance structure is important for reliably identifying the non-diversifiable cojumps.

6. Conclusion

Using popular high-frequency data-based jump detection procedures we document distinct differences in the number of significant jumps in a large panel of high-frequency individual stock returns and aggregate index returns. Jumps occur more than three times as often at the individual stock level. The fact that the index jumped less, on average, than the individual stocks is a simple reflection of diversification of idiosyncratic jumps. We also find, however, that the values of the jump test statistics for the individual stocks are largely

uncorrelated with the values of the test statistic for an index constructed from the very same stocks. This lack of correlation is interesting in view of the fact that all of the stocks have a β of about unity with respect to the index. Apparently, the high level of idiosyncratic noise in the individual stock returns masks many of the cojumps that cause the index to jump. To more effectively detect cojumps, we develop a new *cross-product* statistic, termed the *mcp*-statistic, that is motivated by portfolio theory and directly uses the cross-covariation structure of the high-frequency returns. Employing this statistic we successfully detect many modest-sized cojumps embodied in the panel of returns.

We also find new evidence for an important intraday pattern in the arrival of cojumps. The stocks in our panel show a strong tendency to move sharply together, i.e., cojump, around 10 am Eastern time, corresponding to the regularly scheduled release-time for many macroeconomic news announcements. This newly documented within-day shift in the relative importance of cojumps is overlaid on top of the usual distinct U-shaped pattern in equity return volatility over the trading day along with the secular longer-run day-to-day movements in the overall level of the volatility.

Documenting the presence of cojumps and understanding their economic determinants and dynamics are crucial from a risk measurement and management perspective. Basic portfolio theory implies that the only kind of jumps that can carry a risk premium are a non-diversifiable cojumps. Measuring the risk premium on cojumps is far beyond the scope of the present paper. However, using index-level data Todorov (2006) makes progress towards separating the aggregate jump risk premium from the continuous volatility risk premium and understanding its dynamics. The ideas and techniques developed here may prove especially useful in future work along these lines.

Acknowledgments

The work of Bollerslev was supported by a grant from the NSF to the NBER and CREATES funded by the Danish National Research Foundation. We would like to thank the editor, two anonymous referees, and Roel Oomen for their comments, which helped to improve the exposition and sharpen the focus of the paper. We would also like to thank members of the Duke University Econometrics and Finance Lunch Group and participants at the April 21–22, 2007 Stevanovich conference on “Volatility and High Frequency Data” and the May 18–19, 2007, “Financial Econometrics Conference” at Imperial College, London, for useful suggestions and comments. We are especially grateful to Viktor Todorov and Xin Huang for many discussions and encouragement along the way.

Appendix A. Data appendix

We initially selected the 50 most actively traded stocks on the New York Stock Exchange (NYSE) according to their 10-day trading volume (number of shares) at the beginning of June 2006. Of these 50 stocks, we were able to successfully download reliable high-frequency prices for 40. The ticker symbols for these 40 stocks are included in many of the figures.

A.1. Data source and cleaning

Data on all completed trades were obtained from the Trade and Quote Database (TAQ) available via the Wharton Research Data Services (WRDS). This includes trades from all North American exchanges as well as over-the-counter trades. Each exchange has its own distinct market structure which might affect the structure of observed prices. Hence, to homogenize the data, we decided to only consider trades on the NYSE. The NYSE also accounts for the majority of the trades for all of the stocks in the sample.

Our sample covers the period from January 1st 2001 to December 31st 2005. Trading frequency increased significantly in the late 1990s, and by the end of 2001 all NYSE listed stocks had moved from fractional to decimalized trading, in turn allowing for the extraction of highly reliable high-frequency prices. Illogical data values such as time stamp errors (e.g., hour #25, minute #78, month #43 and year #3001) and negative prices are removed from the data. All-in-all, these errors represent a relatively small number of data points. We also exclude trades that occur outside of 9.30 am and 4 pm, as well as days with only partial trading. Examples of

such days are September 11, 2001, and certain holidays when the NYSE is only open for part of the day. A listing of all of these dates is available on the NYSE web-site.

Because of the unusual activity associated with trading at the beginning of each day, we start our intraday sampling at 9.35 am, 5-min after the market officially opens. This ensures a more homogenous trading and information gathering mechanism for all of the prices. The price series are sampled every 30-s using a slightly modified version of the previous tick method from [Dacorogna et al. \(2001\)](#). The previous tick method simply fixes the time where prices are ideally sampled at regular intervals and selects a completed trade prior to the time should there be no trade at that particular time. For instance, a trade completed at 9:34:58 is used in place of 9:35:00 when there is no actual trade at 9:35:00. In this case, there is therefore a 2-s backtrack, as defined by the time difference between the ideal sampling time and the actual sampling time. The first trade of the day is used if there are no prior trades on that day. With 30-s sampling from 9.35 am to 4.00 pm this leaves us with 771 prices per day. Also, the sample period from January 1st 2001 through December 31st 2005 consists of 1,241 normal trading days, for a total of close to a million transaction prices for each of the 40 individual stocks.

The raw high-frequency prices invariably contains a number of mis-recordings and other data errors. In some cases these errors are obvious by visual inspection of time series plots of the data, sometimes they are not. Thus, in addition to manually inspecting and correcting the data, we also employed a threshold filter of 1.5%, which appears to work well for removing and cleaning the remaining data errors at the 30-s sampling interval.

A.2. Sampling frequency

The statistics used in the paper formally becomes more accurate as the sampling frequency increases. However, as noted in the main text of the paper, there is a limit to how finely we can sample the price process while maintaining the basic underlying semimartingale assumption as a host of market microstructure influences start to materially affect the observed price changes; most importantly features having to do with specifics of the trading mechanism, [Black \(1976\)](#) and [Amihud and Mendelson \(1987\)](#), and discreteness of the data ([Harris, 1990, 1991](#)). As discussed at length in [Hansen and Lunde \(2006\)](#), the design of new procedures and “optimal” ways in which to deal with these complications is currently the focus of extensive research efforts. Rather than employing any of these more advanced procedures, in the analysis reported on here, we simply rely on the volatility signature plots proposed by [Andersen et al. \(2000\)](#) as an easy-to-implement procedure for choosing the highest possible sampling frequency so that the realized variation measures remain unbiased for the unconditional daily variance.

The corresponding plots for each of the 40 stocks (available upon request) suggest that by sampling $M = 22$ times per day, or equivalently by using 17.5 min returns, the market microstructure influences have essentially ceased and the plots for the realized variation measures become flat. Of course, for many of the stocks we could safely sample more frequently, but for simplicity we decided to maintain the identical sampling frequency across all 40 stocks throughout the entire sample. Importantly, the choice of $M = 22$ also involves relatively little interpolation in the construction of the equidistant 17.5 min returns. The median backtrack from the previous tick method for each of the 40 stocks is just about 6.5 s, and for none of the stocks does the median backtrack exceed 12 s.

References

- Ait-Sahalia, Y., Jacod, J., 2008. Testing for jumps in a discretely observed process. *Annals of Statistics*, forthcoming.
- Ait-Sahalia, Y., Mykland, P., Zhang, L., 2005. How often to sample a continuous time process in the presence of market microstructure noise. *Review of Financial Studies* 18, 351–416.
- Amihud, Y., Mendelson, H., 1987. Trading mechanisms and stock returns: an empirical investigation. *Journal of Finance* 42, 533–553.
- Andersen, T., Bollerslev, T., 1998. Deutschemark-dollar volatility: intraday activity patterns, macroeconomic announcements, and longer run dependencies. *Journal of Finance* 53, 219–265.
- Andersen, T., Bollerslev, B., Diebold, F., Labys, P., 2000. Great realizations. *Risk* 13, 105–108.
- Andersen, T., Bollerslev, T., Diebold, F., 2008. Parametric and nonparametric volatility measurement. In: Ait-Sahalia, Y., Hansen, L.P. (Eds.), *Handbook of Financial Econometrics*. North-Holland, Amsterdam.

- Andersen, T., Bollerslev, B., Diebold, F., Vega, C., 2003. Micro effects of macro announcements: real-time price discovery in foreign exchange. *American Economic Review* 93, 38–62.
- Andersen, T., Bollerslev, B., Diebold, F., Vega, C., 2007a. Real-time price discovery in stock, bond and foreign exchange markets. *Journal of International Economics* 73, 251–277.
- Andersen, T., Bollerslev, T., Diebold, F., 2007b. Roughing it up: including jump components in the measurement, modeling and forecasting of return volatility. *Review of Economics and Statistics* 89, 701–720.
- Bandi, F., Russell, J., 2008. Microstructure noise, realised volatility, and optimal sampling. *Reviews of Economic Studies*, forthcoming.
- Barndorff-Nielsen, O., Shephard, N., 2003. Measuring the impact of jumps in multivariate price processes using bipower covariation. Working paper, Nuffield College, Oxford University.
- Barndorff-Nielsen, O., Shephard, N., 2004. Power and bipower variation with stochastic volatility and jumps. *Journal of Financial Econometrics* 2, 1–37.
- Barndorff-Nielsen, O., Shephard, N., 2006. Econometrics of testing for jumps in financial economics using bipower variation. *Journal of Financial Econometrics* 4, 1–30.
- Barndorff-Nielsen, O., Graversen, S., Jacod, J., Podolskij, M., Shephard, N., 2005. A central limit theorem for realised power and bipower variations of continuous semimartingales. Working paper, Nuffield College, Oxford University.
- Barndorff-Nielsen, O., Hansen, P., Lunde, A., Shephard, N., 2006. Designing realised kernels to measure the ex post variation of equity prizes in the presence of noise. Working paper, Nuffield College, Oxford University.
- Basawa, I., Brockwell, P., 1982. Non-parametric estimation for non-decreasing Lévy processes. *Journal of the Royal Statistical Society Series B* 44, 262–269.
- Black, F., 1976. Noise. *Journal of Finance* 41, 529–543.
- Borovkova, S., Burton, R., Dehling, H., 2001. Limit theorems for functionals of mixing processes with applications to *U*-statistics and dimension estimation. *Transactions of the American Mathematical Society* 353, 4261–4318.
- Dacorogna, M., Gencay, R., Muller, U., Olsen, R., Pictet, O., 2001. *An Introduction to High Frequency Finance*. Academic Press, San Diego.
- Eraker, B., Johannes, M., Polson, N., 2003. The impact of jumps in volatility and returns. *Journal of Finance* 58, 1269–1300.
- Fan, J., Wang, Y., 2006. Multi-scale jump and volatility analysis for high-frequency financial data. Working paper, Princeton University.
- Gobbi, F., Mancini, C., 2007. Identifying the covariation between the diffusion parts and the co-jumps given discrete observations. Working paper, Dipartimento di Matematica per le Decisioni, Università degli Studi di Firenze.
- Hansen, P., Lunde, A., 2006. Realized variance and market microstructure noise. *Journal of Business and Economic Statistics* 24, 127–161.
- Harris, L., 1986. A transaction data study of weekly and intradaily patterns in stock returns. *Journal of Financial Economics* 16, 99–117.
- Harris, L., 1990. Estimation of stock variance and serial covariance from discrete observations. *Journal of Financial and Quantitative Analysis* 25, 291–306.
- Harris, L., 1991. Stock price clustering and discreteness. *Review of Financial Studies* 4, 389–415.
- Huang, X., Tauchen, G., 2005. The relative contributions of jumps to total variance. *Journal of Financial Econometrics* 3, 456–499.
- Jacod, J., Todorov, V., 2007. Testing for common arrivals of jumps for discretely observed multidimensional processes. Working paper, Duke University.
- Jiang, G., Oomen, R., 2006. A new test for jumps in asset prices. Working paper, University of Warwick.
- Lee, S., Mykland, P., 2008. Jumps in financial markets: a new nonparametric test and jump dynamics. *Review of Financial Studies*, forthcoming.
- Mancini, C., 2004. Estimation of the characteristics of the jumps of a general Poisson-diffusion model. *Scandinavian Actuarial Journal*, 42–52.
- Todorov, V., 2006. Variance risk premium dynamics. Working paper, Duke University.
- Wood, R., McInish, T., Ord, J., 1985. An investigation of transaction data for NYSE stocks. *Journal of Finance* 25, 723–739.