

Analyzing the Spectrum of Asset Returns: Jump and Volatility Components in High Frequency Data*

Yacine Aït-Sahalia

Department of Economics
Bendheim Center for Finance
Princeton University and NBER

Jean Jacod

Institut de Mathématiques de Jussieu
CNRS UMR 7586
Université P. et M. Curie (Paris-6)

This Version: June 11, 2009

ABSTRACT

This paper describes a simple yet powerful methodology to decompose asset returns sampled at high frequency into their base components (continuous, small jumps, large jumps), determine the relative magnitude of the components, and analyze the finer characteristics of these components such as the degree of activity of the jumps.

Keywords: Continuous-time models; semimartingales; jumps; volatility; spectrum; high frequency financial returns.

JEL Code: G11.

*This research was partly funded by the NSF under grants DMS-0532370 and SES-0850533. MATLAB code to implement the methods described in this paper and sample data files can be downloaded from the authors' web pages.

I. Introduction

The basic model we analyze is the workhorse of modern asset pricing: X , typically the log of an asset price, is assumed to follow an Itô semimartingale. As is well known, for an asset pricing model to avoid arbitrage opportunities, asset prices must follow semimartingales (see Harrison and Pliska (1981), Delbaen and Schachermayer (1994)).

Semimartingales are very general models that nest most if not all continuous-time models used in asset pricing. A semimartingale can be decomposed into the sum of a drift, a continuous Brownian-driven part and a discontinuous, or jump, part. The jump part can in turn be decomposed into a sum of small jumps and big jumps. The continuous part can be scaled by a stochastic volatility process, which may be correlated with the asset price, may jump in conjunction or independently of the asset price, and in fact be a semimartingale itself.

This paper is devoted to analyzing the specification of semimartingales on the basis of high frequency financial returns. We wish to decide which component(s) need to be included in the model (jumps, finite or infinite activity, continuous component, etc.) and determine their relative magnitude. We may then magnify specific components of the model if they are present, so that we can analyze their finer characteristics (such as the degree of activity of jumps). While the underlying mathematical tools are heavily technical, and are developed elsewhere¹, the end result happens to be very simple from the point of view of applications. It requires little else than the recording of asset returns at high frequency, and the computation of a few quantities which we call truncated power variations. Our purpose here is to describe these tools, the intuition behind them, and show that they can be understood as part of a common framework.

To describe our methodology, it can be helpful to proceed by analogy with a spectrographic analysis. We observe a time series of high frequency returns (a single path) over a finite length of time $[0, T]$. For example, consider the time series of Microsoft (MSFT) and

¹In earlier work, we developed tests to determine on the basis of the observed log-returns whether a jump part was present (Aït-Sahalia and Jacod (2009b)), whether the jumps had finite or infinite activity (Aït-Sahalia and Jacod (2008b)), in the latter situation proposed a definition and an estimator of a degree of jump activity parameter (Aït-Sahalia and Jacod (2009a)), and finally whether a Brownian continuous component was needed when infinite activity jumps are included (Aït-Sahalia and Jacod (2008a)).

Intel (INTC) prices, two of the most actively traded equities in the US, shown in Figure 1 during 2006.

Using that time series as input, we will then design a set of statistical tools that can tell us something about specific components of the process that produced the observations. These tools play the role of the measurement devices used in astrophysics to analyze the light emanating from a star, for instance. Our observations are the high frequency returns; in astrophysics it would be the light, visible or not. Here, the data generating mechanism is assumed to be a semimartingale; in astrophysics it would be whatever nuclear reactions inside the star are producing the light that is collected.

Astrophysicists can look at a specific range of the light spectrum to learn about specific chemical elements present in the star. Here, we design statistics that focus on specific parts of the distribution of high frequency returns in order to learn about the different components of the semimartingale that produced those returns.

From the time series of returns, we can get the distribution of returns at time interval Δ_n : for example, Figure 2 show the tails of the distribution of the 2006 returns on MSFT and INTC sampled at 15 seconds. Using these returns as inputs, we would like to figure out which components should be included in the model (continuous? jumps? which types of jumps?) and in what proportions. That is, we would like to deconstruct the observed series of returns back into its original components, continuous and jumps, as described in Figure 3. Figure 3 cannot be produced by visual inspection alone of Figure 2. We need to run the raw data through some devices that will emphasize certain components to the exclusion of others, magnify certain aspects of the model, etc.

Similarly to what is done in spectrographic analysis, we will emphasize visual tools in this paper. In spectrography, one needs to be able to recognize the visual signature of certain chemical elements. Here, we need to know what to expect to see if a certain component of the model is present or not in the observed data. This means that we will need to have a law of large numbers, obtained by imagining that we had collected a large number of sample paths instead of a single one. This allows us to determine the visual signature of specific components of the model. We will not attempt here to measure the dispersion around the expected pattern, and instead refer to our papers in the reference list for the corresponding central limit theorems, the formal derivations of the results including

regularity conditions, as well as simulation evidence on the adequacy of the asymptotics. Those papers are technically demanding because of the very nature of semimartingales, but also because depending upon which component is included or not in the model – precisely the questions we wish to answer – the asymptotics are driven by components with very different characteristics. By contrast, the intuition is fairly clear and this is what this paper focuses on, with the objective of facilitating applications of the results rather than their derivation.

The methodology we propose will determine which components should be included in a given semimartingale model of asset returns. This knowledge has various implications for asset pricing. Many high frequency trading strategies rely on specific components of the model being present or absent. If jumps need to be included in the model, then the familiar consequences of market completeness for contingent claims valuation typically no longer hold. And changes of measure will vary depending upon the type of jumps that are included. Optimal portfolios will vary depending upon the nature of the underlying assets' dynamics. Risk management is also heavily dependent upon the underlying dynamics: clearly, a model with only a continuous component will yield very different risk measures than one with jump components also present.

One word on data considerations before we proceed: when implementing the method on returns data, we will rely on ultra-high frequencies, meaning that the sampling intervals we use are typically of the order of a few seconds to a few minutes. This has two consequences. First, obviously, it limits the analysis to data series for which such sampling frequencies are available. This is becoming less and less of a restriction as time passes, but it does limit our ability to use long historical series, or returns data from less liquid assets. Second, this means that even for liquid assets market microstructure noise is going to be at least potentially a concern. Continuing with the spectrography analogy, market microstructure noise plays the same role as the blurring of astronomical images due to the Earth's atmosphere or light pollution. And we do not have the equivalent of a space-based telescope enabling the direct observation of the true or fundamental asset price. We will in the course of our analysis examine the consequences of this noise on the various statistics we proposed.

The paper is organized as follows. Section II presents the common measurement device we designed to answer the various specification questions. In latter sections, we analyze

these questions one by one: which components are present (Section III), in what relative proportions (Section IV), and some of the finer characteristics of the jump component (Section V). Section VII reports the results of applying the analysis to the INTC and MSFT during 2006. Section VIII concludes.

II. The Measurement Device

The log-price X_t follows an Itô semimartingale, a hypothesis maintained throughout, and formally stated as

$$X_t = X_0 + \underbrace{\int_0^t b_s ds}_{\text{drift}} + \underbrace{\int_0^t \sigma_s dW_s}_{\text{continuous part}} + \text{JUMPS} \quad (1)$$

$$\text{JUMPS} = \underbrace{\int_0^t \int_{\{|x| \leq \varepsilon\}} x(\mu - \nu)(ds, dx)}_{\text{small jumps}} + \underbrace{\int_0^t \int_{\{|x| > \varepsilon\}} x\mu(ds, dx)}_{\text{big jumps}} \quad (2)$$

where as usual W denotes a standard Brownian motion, and μ is the jump measure of X , and its predictable compensator is the Lévy measure ν (both μ and ν are random positive measures on $\mathbb{R}_+ \times \mathbb{R}$). In the perhaps more familiar differential form,

$$dX_t = b_t dt + \sigma_t dW_t + dJ_t \quad (3)$$

where J_t is the jump term.

The distinction between small and big jumps is based on a cutoff level $\varepsilon > 0$ in (2) that is arbitrary. What is important is that $\varepsilon > 0$ is fixed. A semimartingale will always generate a finite number of big jumps on $[0, T]$. But it may give rise to either a finite or infinite number of small jumps. For any measurable subset A of \mathbb{R} at a positive distance of the origin, the increasing process $\nu([0, t] \times A)$ is increasing and “compensates” the number of jumps of X whose size is in A , in the sense that the difference of these two processes is a (local) martingale. Therefore, $\nu([0, t] \times (-\infty, -\varepsilon) \cup (\varepsilon, +\infty)) < \infty$, whereas $\nu([0, t] \times [-\varepsilon, \varepsilon])$ may be finite or infinite, although we must have $\int_{\{|x| \leq \varepsilon\}} x^2 \nu([0, t], dx) < \infty$. Note that we have compensated the small jumps part, but not the big jumps one. Compensating the big jumps part is not always possible because the moments may not exist, whereas summing small jumps without compensation may lead to a divergent sum.

We will assume that the model produces observations that are collected at a discrete sampling interval Δ_n : this means in particular that only "regular" sampling schemes are considered below, although the methodology can be extended to some non-regular sampling scheme, at the expense of – significantly more – mathematical sophistication. There are $[T/\Delta_n]$ (where $[x]$ denotes the integer part of the positive real x) observed increments of X on $[0, T]$, which are

$$\Delta_i^n X = X_{i\Delta_n} - X_{(i-1)\Delta_n}, \quad (4)$$

to be contrasted with the actual (unobservable) jumps of X

$$\Delta X_s = X_s - X_{s-}, \quad (5)$$

as described in Figure 4.

Our basic methodology consists in constructing realized power variations of these increments, suitably truncated and/or sampled at different frequencies. These realized power variations are defined as follows, where $p \geq 0$ is any nonnegative real and $u_n > 0$ is a sequence of truncation levels:

$$B(p, u_n, \Delta_n) = \sum_{i=1}^{[T/\Delta_n]} |\Delta_i^n X|^p 1_{\{|\Delta_i^n X| \leq u_n\}} \quad (6)$$

Throughout, T is fixed, and asymptotics are all with respect to $\Delta_n \rightarrow 0$. Typically the truncation levels u_n go to 0, and this is usually achieved by taking $u_n = \alpha \Delta_n^\varpi$ for some constants $\varpi \in (0, 1/2)$ and $\alpha > 0$. Setting $\varpi < 1/2$ allows us to keep all the increments which mainly contain a Brownian contribution. There will be further restrictions on the rate at which $u_n \rightarrow 0$, expressed in the form of restrictions on the choice of ϖ . In some instances, we do not want to truncate at all and we then write $B(p, \infty, \Delta_n)$. Sometimes we will truncate in the other direction, that is retain only the increments larger than u :

$$U(p, u_n, \Delta_n) = \sum_{i=1}^{[T/\Delta_n]} |\Delta_i^n X|^p 1_{\{|\Delta_i^n X| > u_n\}}. \quad (7)$$

With $u_n = \alpha \Delta_n^\varpi$ as above, that can allow us to eliminate all the increments from the continuous part of the model. Then obviously

$$U(p, u_n, \Delta_n) = B(p, \infty, \Delta_n) - B(p, u_n, \Delta_n). \quad (8)$$

Finally, we sometimes simply count the number of increments of X , that is, take the power $p = 0$

$$U(0, u_n, \Delta_n) = \sum_{i=1}^{\lceil T/\Delta_n \rceil} 1_{\{|\Delta_i^n X| > u_n\}}. \quad (9)$$

We exploit the different asymptotic behavior of the variations $B(p, u_n, \Delta_n)$ and/or $U(p, u_n, \Delta_n)$ as we vary: the power p , the truncation level u_n and the sampling frequency Δ_n . This gives us three degrees of freedom, or tuning parameters, with enough flexibility to isolate what we are looking for. Having these three parameters to play with, p , u_n and Δ_n , is akin to having three knobs to adjust in the measurement device.

A. *The First Knob: Varying the Power*

The role of the power variable is to isolate either the continuous or jump components, or to keep them both present. As illustrated in Figure 5, powers $p < 2$ will emphasize the continuous component of the underlying sampled process while powers $p > 2$ will conversely accentuate its jump component. The power $p = 2$ (which receives much attention in the form of measuring realized volatility) puts them on an equal footing, which turns out to be useful here only when we seek to measure the relative magnitude of the components.

B. *The Second Knob: Varying the Truncation Rate*

Truncating the large increments at a suitably selected cutoff level can eliminate the big jumps when needed. The key is that there is a finite number of large jumps. Asymptotically, as the sampling frequency increases, the cutoff level gets smaller. But the large jumps have a fixed size, so at some point along the asymptotics the cutoff level becomes smaller than the large jumps, which are thus no longer part of the realized power variation $B(p, u_n, \Delta_n)$, as illustrated in Figure 6.

Alternatively, we can truncate to eliminate the Brownian component if we use the upwards power variation $U(p, u_n, \Delta_n)$, since the continuous component is only capable of generating increments that are smaller than $u_n = \alpha \Delta_n^\varpi$ when $\varpi < 1/2$.

C. The Third Knob: Varying the Sampling Frequency

Sampling at different frequencies can let us distinguish between the three situations where the variations converge to a finite limit, converge to zero or diverge to infinity. We will achieve this by computing the ratio of two B 's evaluated at the biggest available frequency Δ_n and at the same time at some lower frequency $k\Delta_n$ where $k \geq 2$ is an integer. Sampling at frequency $k\Delta_n$ is obtained from the same data series, simply retaining one out of every k data points in Figure 4. As described in Figure 7, the limiting behavior of the ratio (1, less than 1 or greater than 1) will identify the underlying limiting behavior of B .

As we will see, the various limiting behaviors of the variations are indicative of which component of the model dominates at a particular power and in a certain range of returns (by truncation), just like certain chemical elements have a very specific spectrographic signature. So they will effectively allow us to distinguish between all manners of null and alternative hypotheses if we can identify which situation corresponds to which of the spectrographic signatures of B .

III. Which Component(s) Are Present

Leaving aside the drift, which is effectively invisible at high frequency, the model (1)-(2) has three components: a continuous part, a small jumps part and a big jumps part. The analogy with spectrography would be that we are looking for the signature of three possible chemical elements (say, hydrogen, helium and everything else) in the light being recorded. Here, based on the observed log-returns, what can we tell about which component(s) of the model are present?

Consider the following sets defined pathwise on $[0, T]$:

$$\begin{aligned}
 \Omega_T^c &= \{X \text{ is continuous in } [0, T]\} \\
 \Omega_T^j &= \{X \text{ has jumps in } [0, T]\} \\
 \Omega_T^f &= \{X \text{ has finitely many jumps in } [0, T]\} \\
 \Omega_T^i &= \{X \text{ has infinitely many jumps in } [0, T]\} \\
 \Omega_T^W &= \{X \text{ has a Wiener component in } [0, T]\} \\
 \Omega_T^{\text{no}W} &= \{X \text{ has no Wiener component in } [0, T]\}
 \end{aligned} \tag{10}$$

Formally, $\Omega_T^W = \left\{ \int_0^T \sigma_s^2 ds > 0 \right\}$ and $\Omega_T^{\text{no}W} = \left\{ \int_0^T \sigma_s^2 ds = 0 \right\}$, and the definition of the four other sets is clear.

We observe a time series originating in a given unobserved path, and wish to determine in which set(s) the path is. At any given – fixed – frequency this is a theoretically unanswerable question since for example any such time series can be obtained by discretization of a continuous path, and also of a discontinuous one. However we wish to construct test statistics that behave well asymptotically, as $\Delta_n \rightarrow 0$, and if possible under the only structural assumptions (1)–(2). That is, they should be model-free in the sense that their implementation *and their asymptotic properties* do not require that we specify or calibrate the model, which can potentially be quite complicated (stochastic volatility, jumps, jumps in volatility, jumps in jump intensity, etc.).

It turns out that this aim is achievable, using the power variations introduced above, for some of the problems. For others we need some additional structural assumptions, to be explained later when needed. Let us also mention that for all results one also needs some weak boundedness-like or smoothness-like assumptions on the coefficients, such as the process b_t should be (locally) bounded: as a rule, these assumptions are not explicitly stated here, and we refer to the original papers for the mathematically precise statements.

A. *Jumps: Present or Not*

The first question we address is whether the path of X contains jumps or not. There is by now a vast literature concerned with detecting jumps; see Aït-Sahalia and Jacod (2009b) for references. Using the methodology of power variations, we start with two processes which measure some kind of variability of X and depend on the whole (unobserved) path of X :

$$A(p) = \int_0^T |\sigma_s|^p ds, \quad B(p) = \sum_{s \leq T} |\Delta X_s|^p \quad (11)$$

where $p > 0$. The variable $A(p)$ is finite for all $p > 0$, and positive on the set Ω_T^W . The variable $B(p)$ is finite if $p \geq 2$ but often not when $p < 2$. The quadratic variation of X is $[X, X]_T = A(2) + B(2)$.

Of course, hoping to estimate $B(p)$ using $B(p, u_n, \Delta_n)$ is too naive in general, but it

works in specific cases. Namely, we have the following behavior of $B(p, \infty, \Delta_n)$

$$\begin{cases} p > 2, \text{ all } X & \Rightarrow B(p, \infty, \Delta_n) \xrightarrow{\mathbb{P}} B(p) \\ \text{all } p, \text{ on } \Omega_T^c & \Rightarrow \frac{\Delta_n^{1-p/2}}{m_p} B(p, \infty, \Delta_n) \xrightarrow{\mathbb{P}} A(p) \end{cases} \quad (12)$$

where m_p denotes the p th absolute moment of the standard normal variable.

So we see that, when $p > 2$, $B(p, \infty, \Delta_n)$ tends to $B(p)$: the jump component dominates. If there are jumps, the limit $B(p)_t > 0$ is finite. On the other hand when X is continuous, then the limit is $B(p) = 0$ and $B(p, \infty, \Delta_n)_t$ converges to 0 at rate $\Delta_n^{p/2-1}$.

These considerations lead us to pick a value of $p > 2$ and compare $B(p, \infty, \Delta_n)_t$ on two different sampling frequencies. Specifically, for an integer k , consider the test statistic S_J :

$$S_J(p, k, \Delta_n) = \frac{B(p, \infty, k\Delta_n)_T}{B(p, \infty, \Delta_n)_T}. \quad (13)$$

The ratio in S_J exhibits a markedly different behavior depending upon whether X has jumps or not:

$$S_J(p, k, \Delta_n)_t \rightarrow \begin{cases} 1 & \text{on } \Omega_T^j \\ k^{p/2-1} & \text{on } \Omega_T^c \cap \Omega_T^W \end{cases} \quad (14)$$

That is, in the context of Figure 7, under Ω_T^j the variation converges to a finite limit and so the ratio tends to 1 (the middle situation depicted in the figure) while under $\Omega_T^c \cap \Omega_T^W$ the variation converges to 0 and the ratio tends to a limit greater than 1, with value specifically depending upon the rate at which the variation tends to 0 (the lower situation depicted in the figure). The notion of a set $\Omega_T^c \cap \Omega_T^W$ may seem curious at first, but it is possible for a process to have continuous paths without a Brownian component if the process consists only of a pure drift. Because this would be an unrealistic model for financial data, we are excluding the set $\Omega_T^c \cap \Omega_T^{\text{no}W}$ from consideration.

If one desires a formal statistical test of $\Omega_T^c \cap \Omega_T^W$ vs. Ω_T^j , with a prescribed asymptotic level $\alpha \in (0, 1)$, one can use a CLT under $\Omega_T^c \cap \Omega_T^W$ and one under Ω_T^j , such CLT being available again in a model-free situation, apart from some additional smoothness assumptions: so one can in fact test either $H_0 : \Omega_T^c \cap \Omega_T^W$ vs. $H_1 : \Omega_T^j$ or the reverse $H_0 : \Omega_T^j$ vs. $H_1 : \Omega_T^c \cap \Omega_T^W$. Note that the first limit in (14) is valid on Ω_T^j whether the jump component includes finite or infinite components, or both. It is not designed to disentangle the two types of jumps. How to do this is the question we now turn to.

B. Jumps: Finite or Infinite Activity

Many models in mathematical finance do not include jumps. But among those that do, the framework most often adopted consists of a jump-diffusion: these models include a drift term, a Brownian-driven continuous part, and a finite activity jump part (compound Poisson process): early examples include Merton (1976), Ball and Torous (1983) and Bates (1991).

Other models are based on infinite activity jumps: see for example Madan and Seneta (1990), Madan and Milne (1991), Eberlein and Keller (1995), Barndorff-Nielsen (1997), Barndorff-Nielsen (1998), Carr, Geman, Madan, and Yor (2002), Carr and Wu (2003), Carr and Wu (2004) and Schoutens (2003), although with the exception of Carr, Geman, Madan, and Yor (2002) models of this type are justified primarily by their ability to produce interesting pricing formulae rather than necessarily an attempt at empirical realism.

So, which is it, based on the data? Our objective is now to discriminate between finite and infinite activity jumps.

B.1. Null Hypothesis: Finite Activity

We first set the null hypothesis to be finite activity, that is $H_0 : \Omega_T^f \cap \Omega_T^W$, whereas the alternative is $H_1 : \Omega_T^i$. As in the previous subsection, we rule out the set $\Omega_T^f \cap \Omega_T^{\text{no}W}$ which, for all models in use in finance, is empty. We choose an integer $k \geq 2$ and a real $p > 2$. The only difference with testing for jumps using S_J is that we now truncate

$$S_{FA}(p, u_n, k, \Delta_n) = \frac{B(p, u_n, k\Delta_n)}{B(p, u_n, \Delta_n)}. \quad (15)$$

Without truncation, as in S_J , we could discriminate between jumps and no jumps, but not among different types of jumps. Like before, we set $p > 2$ to magnify the jump component at the expense of the continuous component. But since we want to separate big and small jumps, we now truncate as a means of eliminating the large jumps. Since the large jumps are of finite size (independent of Δ_n), at some point in the asymptotics the truncation level $u_n = \alpha\Delta_n^\varpi$ will have eliminated all the large jumps: see Figure 6 earlier. Then if there are only big jumps and the Brownian component, the two truncated power

variations $B(p, u_n, k\Delta_n)$ and $B(p, u_n, \Delta_n)$ will behave as if there were no jumps, leaving only the Brownian component. The limit of the ratio will be $k^{p/2-1}$ as in the test for jumps when there are no jumps.

But if there are infinitely many jumps, which are necessarily small, then the truncation cannot eliminate them. This is because however small u_n is, there are still infinitely many jumps in each Δ_n -increment. The Brownian component is dominated in every increment by the small jumps because $p > 2$. Both $B(p, u_n, k\Delta_n)$ and $B(p, u_n, \Delta_n)$ behave like the sum of the p^{th} power of the jumps that are smaller than u_n , and although they both go to 0, their ratio tends to 1. In the context of Figure 7, we are in the limiting case where both B 's go to zero but at the same rate: hence the ratio is 1.

That is, we have:

$$S_{FA}(p, u_n, k, \Delta_n) \xrightarrow{\mathbb{P}} \begin{cases} k^{p/2-1} & \text{on } \Omega_T^f \cap \Omega_T^W. \\ 1 & \text{on } \Omega_T^i. \end{cases} \quad (16)$$

B.2. Null Hypothesis: Infinite Activity

We next set the null hypothesis to be infinite activity, that is $H_0 : \Omega_T^i$, whereas the alternative is $H_1 : \Omega_T^f \cap \Omega_T^W$. We need a different statistic, S_{IA} , because although S_{FA} goes to 1, the distribution of S_{FA} is not model-free under Ω_T^i . The problem comes from the fact that the behavior of the truncated power variations $B(p, u_n, \Delta_n)$ depend on the degree of activity of the jumps when there are infinitely many jumps. So we need to specify what we precisely mean by “degree of activity”.

To this end, recalling the definition of $B(p)$ given in (11), we consider now the set $I_T = \{p \geq 0 : B(p) < \infty\}$. This (random) set I_T is of the form $[\beta_T, \infty)$ or (β_T, ∞) for some $\beta_T(\omega) \in [0, 2]$, and $2 \in I_T$ always. It turns out that $\beta_T(\omega)$, that is the lower bound of the set I_T , is a sensible measure of jump activity for the path $t \mapsto X_t(\omega)$ at time T . In the special case where X is a Lévy process, then $\beta_T(\omega) = \beta$ does not depend on (ω, T) , and it is also the infimum of all $r \geq 0$ such that $\int_{\{|x| \leq 1\}} |x|^r F(dx) < \infty$, where F is the Lévy measure, and this number has been introduced by Blumenthal and Gettoor (1961) and by extension we call β_T the (generalized) Blumenthal-Gettoor index, or degree of jump activity, of the process.

A priori this degree of jump activity can be random and depend on time, but we assume for tractability that this index is in fact constant in time and non-random. More precisely, we assume that the Lévy measure ν in (2) is of the form

$$\nu(dt, dx) = \frac{1}{|x|^{1+\beta}} \left(a_t^+ 1_{(0, z_t^+]}(x) + a_t^- 1_{[-z_t^-, 0)}(x) \right) + \nu'(dt, dx), \quad (17)$$

where a_t^\pm are nonnegative and z_t^\pm are positive stochastic processes, and ν' is another Lévy measure whose Blumenthal-Gettoor index is smaller than β . Note that the assumption (17) is only about the local behavior of the jump measure ν near 0, that is, only about the behavior of the small jumps. The big jumps, controlled by ν' , are unrestricted. The processes a_t^\pm are intensity parameters: as they go up, there are more and more small jumps. The processes z_t^\pm control the range of returns over which the behavior of the overall jump measure is stable-like with index β .

Note that, necessarily, $\beta \in (0, 2)$ here. Then, if further $\int_0^T (a_s^+ + a_s^-) ds > 0$, the number β is the index of X on the full interval $[0, T]$. Note that when X is a (possibly asymmetric) stable process then it satisfies this assumption, β being the index of the stable process. In fact, this assumption amounts to say that the small jumps of X behave like the small jumps of a stable process, or more accurately as those of a process which is a stochastic integral with respect to a stable process, whereas the big jumps are governed by ν' . We call processes which satisfy (17) “proto-stable” processes.

Most models in finance which exhibit jumps of infinite activity are proto-stable. While we will propose estimators of β below, the true β is of course unknown, and our model-free requirement means here that we wish to construct a test which does not depend upon β , the processes a_t^\pm or z_t^\pm , nor the residual jump measure ν' .

Coming back to our problem, we consider the set

$$\Omega_T^{i\beta} = \left\{ \int_0^T (a_s^+ + a_s^-) ds > 0 \right\}$$

on which the jump activity index of X equals β . Note that $\Omega_T^{i\beta} \subset \Omega_T^i$, the inclusion possibly being strict. However testing the null being Ω_T^i is impossible without further restriction, and so we set the null to be $\Omega_T^{i\beta}$.

We choose three reals $\gamma > 1$ and $p' > p > 2$ and define a family of test statistics as

follows:

$$S_{IA}(p, u_n, \gamma, \Delta_n) = \frac{B(p', \gamma u_n, \Delta_n) B(p, u_n, \Delta_n)}{B(p', u_n, \Delta_n) B(p, \gamma u_n, \Delta_n)} \quad (18)$$

which has the following limits:

$$S_{IA}(p, u_n, \gamma, \Delta_n) \xrightarrow{\mathbb{P}} \begin{cases} \gamma^{p'-p} & \text{on } \Omega_T^{i\beta} \\ 1 & \text{on } \Omega_T^f \cap \Omega_T^W \end{cases} \quad (19)$$

Intuitively, under the alternative of finite jump activity, the behavior of each one of the four truncated power variations in (18) is driven by the continuous part of the semimartingale. The truncation level is such that essentially all the Brownian increments are kept. Then the truncated power variations all tend to zero at rates $\Delta_n^{p/2-1}$ and $\Delta_n^{p'/2-1}$ respectively and by construction the (random) constants of proportionality cancel out in the ratios, producing a limit 1 given under H_1 in (19).

If, on the other hand, jumps have infinite activity, then the small jumps are the ones that matter and the truncation level becomes material, producing four terms that all tend to zero but at the different orders in probability $u_n^{p-\beta}$, $u_n^{p'-\beta}$, $(\gamma u_n)^{p-\beta}$ and $(\gamma u_n)^{p'-\beta}$ respectively, resulting in the limit $\gamma^{p'-p}$ given under H_0 in (19). By design, that limit in S_{IA} is independent of β .

C. Brownian Motion: Present or Not

We now would like to construct procedures which allow to decide whether the Brownian motion is really there, or if it can be forgone with in favor of a pure jump process with infinite activity. When infinitely many jumps are included, there are a number of models in the literature which dispense with the Brownian motion altogether. The log-price process is then a purely discontinuous Lévy process with infinite activity jumps, or more generally is driven by such a process.

C.1. Null Hypothesis: Brownian Motion Present

In order to construct a test, we seek a statistic with markedly different behavior under the null and alternative. The idea is now to consider powers p less than 2, since in the presence of Brownian motion the power variation would be dominated by it while in its

absence it would behave quite differently. Specifically, the large number of small increments generated by a continuous component would cause a power variation of order less than 2 to diverge to infinity: recall Figure 7.

Without the Brownian motion, however, and when p is bigger than the Blumenthal-Gettoor index $\beta_T = \beta$, assuming the structural assumption (17), the power variation converges to 0 at exactly the same rate for the two sampling frequencies Δ_n and $k\Delta_n$, whereas with a Brownian motion the choice of sampling frequency will influence the magnitude of the divergence. Taking a ratio will eliminate all unnecessary aspects of the problem and focus on that key aspect.

So we choose an integer $k \geq 2$ and a real $p < 2$ and propose the test statistic

$$S_W(p, u_n, k, \Delta_n) = \frac{B(p, u_n, \Delta_n)}{B(p, u_n, k\Delta_n)} \quad (20)$$

which has the limits

$$S_W(p, u_n, k, \Delta_n) \xrightarrow{\mathbb{P}} \begin{cases} k^{1-p/2} & \text{on } \Omega_T^W \\ 1 & \text{on } \Omega_T^{\text{no}W} \cap \Omega_T^\beta, p > \beta \end{cases} \quad (21)$$

Note that the first convergence above, on Ω_T^W , do not require any specific assumptions on the jumps, only the second convergence requires (17).

C.2. Null Hypothesis: No Brownian Motion

When there are no jumps, or finitely many jumps, and no Brownian motion, X reduces to a pure drift plus occasional jumps, and such a model is fairly unrealistic in the context of most financial data series. But one can certainly consider models that consist only of a jump component, plus perhaps a drift, if that jump component is allowed to be infinitely active. If one wishes to set the null model to be a pure jump model (plus perhaps a drift), then the issue becomes to design a test statistic using power variations whose behavior is independent of the specific nature of the infinitely active pure jump process. In other words, we again assume (17), but we do not know β and wish to design a test that remains model-free in the sense that it does not depend on β , a_t^\pm or z_t^\pm in (17).

We choose a real $\gamma > 1$ to define two different truncation ratios and define a family of

test statistics as follows:

$$S_{\text{no}W}(p, u_n, \gamma, \Delta_n) = \frac{B(2, \gamma u_n, \Delta_n) U(0, u_n, \Delta_n)}{B(2, u_n, \Delta_n) U(0, \gamma u_n, \Delta_n)}. \quad (22)$$

To understand the construction of this test statistic, recall that in a power variation of order 2 the contributions from the Brownian and jump components are of the same order. If the Brownian motion is present ($H_1 : \Omega_T^W$) then once that power variation is properly truncated, the Brownian motion will dominate it if it is present. And the truncation can be chosen to be sufficiently loose that it retains essentially all the increments of the Brownian motion at cutoff level u_n and a fortiori γu_n , thereby making the ratio of the two truncated quadratic variations converge to 1 under the alternative hypothesis.

If on the other hand the Brownian motion is not present ($H_0 : \Omega_T^{\text{no}W} \cap \Omega_T^\beta$), then the nature of the tail of jump distributions is such that the difference in cutoff levels between u_n and γu_n remains material no matter how far we go in the tail and the limit of the ratio $B(2, \gamma u_n, \Delta_n)/B(2, u_n, \Delta_n)$ in (22) will reflect it: it will now be $\gamma^{2-\beta}$. But since absence of a Brownian motion is now the null hypothesis, the issue for constructing a test is that this limit depends on the unknown β .

Canceling out that dependence is the role devoted to the ratio $U(0, u_n, \Delta_n)/U(0, \gamma u_n, \Delta_n)$ of the number of large increments. The U 's are always dominated by the jump components of the model whether the Brownian motion is present or not. Their inclusion in the statistic is merely to ensure that the statistic is model-free, by effectively canceling out the dependence on the jump characteristics that emerges from the ratio of the truncated quadratic variations.

Indeed, the limit of the ratio of the U 's is γ^β under both the null and alternative hypotheses. As a result, the probability limit of $S_{\text{no}W}$ will be γ^2 under the null, independent of β :

$$S_{\text{no}W}(p, u_n, \gamma, \Delta_n) \xrightarrow{\mathbb{P}} \begin{cases} \gamma^2 & \text{on } \Omega_T^{\text{no}W} \cap \Omega_T^\beta \\ \gamma^\beta & \text{on } \Omega_T^W \end{cases} \quad (23)$$

Generally speaking, the statistic S_W is more robust than $S_{\text{no}W}$; similarly S_{FA} is more robust than S_{IA} . This is due to their simpler design, and the lesser reliance on subtle cancellations to achieve their respective objectives. As a result, we recommend using S_{FA} and S_W in practical applications.

IV. The Relative Magnitude of the Components

A typical “main sequence” star might be made of 90% hydrogen, 10% helium and 0.1% everything else. In astrophysics, a natural metric to compare different atoms and address the question of percentages of various components is atomic mass. Here, what is the relative magnitude of the two jump and continuous components? We can answer this question using the same device. The natural metric is to consider $p = 2$ since this is the power where all the components are present together and ask the question of percentages of total quadratic variation (QV) attributable to each component.

As illustrated in Figure 8, by using truncations at the right rate we can split the QV into its continuous and jump components, and not truncate to estimate the full QV:

$$\begin{cases} \frac{B(2, u_n, \Delta_n)}{B(2, \infty, \Delta_n)} = \% \text{ of QV due to the continuous component} \\ 1 - \frac{B(2, u_n, \Delta_n)}{B(2, \infty, \Delta_n)} = \% \text{ of QV due to the jump component} \end{cases} \quad (24)$$

The use of truncation to estimate the continuous part of the quadratic variation when there are jumps was proposed by Mancini (2001), who relied on the law of the iterated logarithm. Alternatively, one can split the QV based on bipower variations instead of truncating: see Andersen, Bollerslev, and Diebold (2003), Barndorff-Nielsen and Shephard (2004) and Huang and Tauchen (2005).

Note that (24) suggests that an alternative test for the presence of jumps can be constructed based on the ratio $B(2, u_n, \Delta_n)/B(2, \infty, \Delta_n)$. However, this would work only if the null hypothesis is that no jumps are present, and the null hypothesis is that the ratio is 1. With jumps under the null, one would have to specify exogenously as part of the null hypothesis how large the fraction of QV due to jumps is.

We can split the rest of the QV, which by construction is attributable to jumps, into a small jumps and a big jumps component. This depends on the cutoff level ε selected to distinguish big and small jumps:

$$\begin{cases} \frac{U(2, \varepsilon, \Delta_n)}{B(2, \infty, \Delta_n)} = \% \text{ of QV due to big jumps} \\ \frac{B(2, \infty, \Delta_n) - B(2, u_n, \Delta_n) - U(2, \varepsilon, \Delta_n)}{B(2, \infty, \Delta_n)} = \% \text{ of QV due to small jumps} \end{cases} \quad (25)$$

We can then obtain a plot that looks like Figure 9 and provides a split of the QV into the various components.

V. Estimating the Degree of Jump Activity

The method described in Section B is able to tell finite activity jumps from infinite activity ones. Among jump processes, however, finite activity are the exception rather than the norm. And “infinite activity” can mean quite different things depending upon “how infinite” that infinite jump activity is. In fact, the degree of activity is accurately measured by the Blumenthal-Gettoor index β_T introduced earlier: at one end of the spectrum, infinite activity jump processes such as the Gamma process can look like Poissonian jumps; at the other end, they can look almost like Brownian motion, which is to say extremely active. So it seems natural to try to estimate the index β_T .

The next issue is then to estimate β_T , or rather β under the somewhat restricted assumption (??). The problem is made more challenging by the potential presence in X of a continuous, or Brownian, martingale part. β characterizes the behavior of ν near 0. Hence it is natural to expect that the small increments of the process are going to be the ones that are most informative about β . But that is where the contribution from the continuous martingale part of the process is inexorably mixed with the contribution from the small jumps. In other words, we need to see through the continuous part of the semimartingale in order to say something about the number and concentration of small jumps.

So we are now looking in a different range of the spectrum of returns, namely by considering only returns that are larger than the cutoff $u_n = \alpha \Delta_n^\varpi$ for some $\varpi \in (0, 1/2)$, as opposed to those that are smaller than the cutoff. This allows us to eliminate the increments due to the continuous component. We can then use all values of p , not just those $p > 2$, despite the fact that we wish to concentrate on jumps see Figure 8. In fact, we will simply use the power $p = 0$.

We propose two estimators of β based on counting the number of increments greater than the cutoff u_n . The first one is based on varying the actual cutoff level: fix $0 < \alpha < \alpha'$ and consider two cutoffs $u_n = \alpha \Delta_n^\varpi$ and $u'_n = \alpha' \Delta_n^\varpi$ with $\gamma = \alpha'/\alpha$:

$$\hat{\beta}_n(\varpi, \alpha, \alpha') = \frac{\log(U(0, u_n, \Delta_n)/U(0, \gamma u_n, \Delta_n))}{\log(\gamma)}, \quad (26)$$

The second one is based on varying the sampling frequency: sample at two time scales, Δ_n

and $2\Delta_n$:

$$\widehat{\beta}'_n(\varpi, \alpha, k) = \frac{\log(U(0, u_n, \Delta_n)/U(0, u_n, k\Delta_n))}{\varpi \log k}. \quad (27)$$

These estimators are consistent for β , and we have derived CLTs for them. One can then test various hypotheses involving β .

Related approaches include Woerner (2006), who proposes an estimator of the jump activity index and of the Hurst exponent, Cont and Mancini (2008), who are testing whether $\beta > 1$ or $\beta < 1$, which correspond to finite or infinite variation for X , and Tauchen and Todorov (2008), who provide a graphical method to determine whether $\beta = 2$ or $\beta < 2$, and Belomestny (2008) who proposes a method based on low frequency historical and options data. None of these methods, however, provide an estimator of β in the presence of a continuous component of the model.

VI. Summary: Tuning Power, Truncation and Sampling Frequency

We have seen that settings the three knobs of power, truncation level and sampling frequency in various combinations allowed us to determine which component of the model was likely to be present, in what proportion, and estimate the degree of activity of the jumps. Tables I, II, III and IV summarize the choice of the three tuning parameters (p, u, Δ) for the corresponding tasks.

VII. Empirical Results: Intel & Microsoft 2006

A. *The Data*

We use a dataset consisting of all transactions on MSFT and INTC in 2006, already previewed in Figure 1 and 2. The data source is the TAQ database. Using the correction variables in the dataset, we retain only transactions that are labeled “good trades” by the exchanges: regular trades that were not corrected, changed, or signified as canceled or in error; and original trades which were later corrected, in which case the trade record contains

the corrected data for the trade. Beyond that, no further adjustment to the raw data are made.

We sample the price series every 5 seconds. The two stocks trade on average more than once every second, so we are not retaining every transaction by doing so, which avoids incorporating every bid-ask bounce. When no transaction is available at the exact time stamp, we use the closest one available. When more than one transaction is available at the same time stamp, we average them. We do not include the overnight returns.

Whenever we need to truncate, we will express the truncation cutoff level u_n in terms of a number α of standard deviations of the continuous part of the semimartingale. That initial standard deviation estimate is obtained by using $B(2, 4\sigma\Delta_n^{1/2}, \Delta_n)$ where σ is a fixed realistic value for the asset under consideration; we use $\sigma = 0.25$. The specific value of this number serves only to identify a reasonable range of values. We then use for the truncation level u_n different multiples of it. Finally, each one of the statistic below is computed separately for each quarter of 2006.

B. Market Microstructure Noise

We consider sampling frequencies up to 5 seconds. In real data, observations of the process X at ultra high frequencies are blurred by market microstructure noise, which has the potential to change the asymptotic behavior of many statistics at very high frequency, and can force us to downsample. We will consider the two polar cases where observations are blurred with either an additive white noise or with noise due to rounding, respectively.

When observations are affected by an additive noise, then instead of $X_{i\Delta_n}$ we really observe $Y_{i\Delta_n} = X_{i\Delta_n} + \varepsilon_i$, and the ε_i are i.i.d. with $E(\varepsilon_i^2)$ and $E(\varepsilon_i^4)$ finite. When rounding is introduced, we observe $Y_{i\Delta_n} = [X_{i\Delta_n}]_a$, that is X rounded to the nearest multiple of a , say 1 cent for a decimalized asset. Power variations are affected by either type of noise, in a rather drastic way, since it modifies the limit in probability of most of our statistics, not to speak about their second order behavior like CLTs: so we incorporate their effect when analyzing the empirical results.

C. Jumps: Present or Not

Analyzing the limit of $B(p, u_n, \Delta_n)$, we find that, in the presence of additive noise, $S_J(4, k, \Delta_n) \xrightarrow{\mathbb{P}} 1/k$ while in the presence of rounding error noise, the limit is $1/k^{1/2}$. So S_J has four possible limits:

$$\begin{aligned}
 1/k & : && \text{additive noise dominates} \\
 1/k^{1/2} & : && \text{rounding error dominates} \\
 1 & : && \text{jumps present and no significant noise} \\
 k^{p/2-1} & : && \text{no jumps present and no significant noise}
 \end{aligned} \tag{28}$$

The empirical (daily) values of S_J are shown in the form of a histogram in Figure 11. The data for the histogram are produced by computing S_J for the four quarters of the year, the two stocks, and for a range of values of p from 3 to 6, Δ_n from 5 seconds to 2 minutes, and $k = 2, 3$. As indicated in (28), values below 1 are indicative of noise of one form or another dominating. We find that this is the case for the highest sampling frequencies. Figure 12 displays the mean value of S_J (across values of p and k , and the four quarters) as a function of Δ_n . The average value of S_J starts below 1 before settling down around 1. The conclusion from S_J is that the noise is a concern at the ultra high frequencies, but once past this domain, the evidence points towards the presence of jumps.

This conclusion is not surprising per se. Simple visual inspection of the tails of the 15-second log-return distribution in Figure 2 combined with a back-of-the-envelope calculation suggests the presence of jumps. These stocks are trading around \$20 and \$25 respectively during that period. They were decimalized, so one tick or \$0.01 corresponds to a return of about 0.005%, or 0.00005 on the log-return scale indicated on the x -axis of the figure. The 0.5% mark or half of the 0.01 log-return indicated on Figure 2 corresponds roughly to a 10-tick move over 15 seconds. Clearly, a continuous component alone would be very unlikely to generate such returns in such numbers.

D. Jumps: Finite or Infinite Activity

S_J tells us that jumps are likely to be present, but it cannot distinguish between finite and infinite activity jumps. For this, we turn to the statistic S_{FA} which is like S_J with the

addition of truncation. The data for the histogram in Figure 13 are produced by computing for the four quarters of the year and the two stocks the value of S_{FA} for a range of values of p from 3 to 6, α from 5 to 10 standard deviations, Δ_n from 5 seconds to 2 minutes, and $k = 2, 3$. We find that the empirical values of S_{FA} are distributed around 1, which is indicative of infinite activity jumps. That is, even as we truncate, the statistic continues to behave as if jumps are present. If only a finite number of jumps had been present, then the statistic should have behave as if the process were continuous.

The impact of the noise on S_{FA} is given by

$$\begin{aligned}
1/k & : && \text{additive noise dominates} \\
1/k^{1/2} & : && \text{rounding error dominates} \\
1 & : && \text{infinite activity jumps and no significant noise} \\
k^{p/2-1} & : && \text{finite activity jumps and no significant noise}
\end{aligned} \tag{29}$$

Figure 14 displays the mean value of S_{FA} (across the four quarters, two stocks, and values of p, α and k) as a function of Δ_n . A pattern similar to that of Figure 12 emerges. For very small values of Δ_n , the noise dominates (limits below 1), then the limit is around 1 as Δ_n increases away from the noise-dominated frequencies. Unless we start downsampling more (reaching 5 to 10 minutes), the limit does not get close to $k^{p/2-1}$.

E. Brownian Motion: Present or Not

In light of the likely presence of infinite activity jumps identified by S_{FA} , it makes sense to ask empirically whether a Brownian component is needed at all. For this purpose, we turn to the statistic S_W which, taking market microstructure noise into account, has four possible limits:

$$\begin{aligned}
1 & : && \text{No Brownian motion and no significant noise} \\
k^{1-p/2} & : && \text{Brownian motion present and no significant noise} \\
k^{1/2} & : && \text{rounding error dominates} \\
k & : && \text{additive noise dominates}
\end{aligned} \tag{30}$$

Figure 15 displays a histogram of the distribution of S_W obtained by computing its value for the four quarters of the year for a range of values of p from 1 to 1.75, α from 5

to 10 standard deviations, Δ_n from 5 seconds to 2 minutes, and $k = 2, 3$. The empirical estimates are always on the side of the limit arising in the presence a continuous component. As the sampling frequency increases, the noise becomes more of a factor, as usual, and for very high sampling frequencies, the results are consistent with some mixture of the noise driving the asymptotics. This is confirmed by Figure 16 which displays the mean value of S_W (across the four quarters, two stocks, and values of p, α and k) as a function of Δ_n . As we downsample away from the noise-dominated frequencies, the average value of the statistic settles down towards the one indicating presence of a Brownian motion.

F. QV Relative Magnitude

The previous empirical results indicate that are likely in the presence of a jump as well as a continuous component. We can then ask what fraction of the QV is attributable to the continuous and jump components. Taking microstructure noise into account yields the following limits for our measure of the fraction of QV due to the continuous component:

$$\begin{aligned}
 0 & & : & \text{additive noise dominates} \\
 0 & & : & \text{rounding error dominates} \\
 \text{actual fraction of QV} & & : & \text{no significant noise}
 \end{aligned}
 \tag{31}$$

The histograms in Figure 17 are obtained from computing the fraction of QV from the Brownian component using the four quarters, two stocks, values of α ranging from 2 to 5 standard deviations, and Δ_n from 5 seconds to 2 minutes. We find values around 75%. In Figure 18 (similar but as a function of Δ_n), we see that the estimated fraction is fairly stable as we vary the sampling frequency.

G. Estimating the Degree of Jump Activity

Finally, we estimate the degree of jump activity β . We found above that infinite activity jumps were likely present in the data. We are now asking how active are those jumps among all infinitely active jumps.

Of all the empirical methods employed in the paper, estimating β is the one that requires the largest sample size due to its reliance on truncating from the right in the power variations

U . That is, the estimators of β discard by construction a large fraction of the original sample and to retain a sufficient number of observations to the right of a cutoff u_n given by 5 or more standard deviations of the continuous part, we need to have a large sample to begin with. So we will estimate β using only the first two sampling frequencies, 5 and 10 seconds on a quarterly basis. In the case of the previous statistics, we noted that these sampling frequencies were subject to market microstructure noise. Here, however, because we are only retaining the increments larger than the cutoff u_n instead of those smaller than the cutoff, this is likely to be less of a concern despite the ultra high sampling frequencies.

We find estimated β 's in the range from 1.5 to 1.8, indicating a very high degree of jump activity, in effect much closer to Brownian motion than to compound Poisson. Figure 19 reports the values of the estimator $\hat{\beta}$ computed for the four quarters of the year, the two stocks, a range of values of α from 5 to 10 standard deviations, and Δ_n from 5 to 10 seconds. Figure 20 reports the corresponding data against the limited range of Δ_n employed. Of all the methods employed so far, this is the one that requires the highest frequency of observation because we are relying on a small fraction of the sample due to the need to eliminate the contributions from the continuous part of the model. On the other hand, truncating upwards means that this method is likely to be less affected by the noise despite the high frequency of observation.

VIII. Conclusions

The empirical results appear to:

- Indicate that jumps are present in the data;
- Point towards the presence of infinite activity jumps;
- Of degree of jump activity that is around 1.5 or higher;
- Indicate that a continuous component is present;
- Which represents approximately 3/4 of the total QV.

Of course, we do not claim on the basis of this limited evidence that these empirical results are in any way “universal”: they are likely to depend upon the assets under consideration, the time period, what type of data are used (transactions, quotes, etc.), among other considerations.

In terms of methodology, our assessment at present is the following. The pros of the approach are: we have a unified methodology to address all these seemingly disparate specification questions in a common framework; the method allows for a symmetric treatment of null and alternative hypotheses in each case, including the full distribution theory; all the test statistics are model-free; they are extremely simple to implement in practical applications; and, finally, we are able to characterize the impact of two important types of market microstructure noise, additive and rounding, on the various statistics.

The cons of the approach, at present, are the flip side of the commonality of technique: if we take individually each one of the specification questions we have identified, then the statistic we have proposed for that problem, based on that common approach, is not necessarily the optimal approach for that individual specification question. Also, the practical implementation of this method does require high frequency data (particularly the estimation of the degree of jump activity β , which requires ultra-high frequency data). Finally, while we understand what happens to each one of the statistics when microstructure noise dominates, we do not yet have noise-robust statistics that would be fully immune to the presence of the noise.

To conclude, we certainly expect to see refinements of our approach (see e.g., Fan and Fan (2008)) or perhaps entirely different techniques proposed in the future for each one of these problems, some already existing. However, we think that the main appeal of our approach comes from the commonality of technique, as well as the resulting simplicity of implementation.

REFERENCES

- Aït-Sahalia, Yacine, and Jean Jacod, 2008a, Is Brownian motion necessary to model high frequency data?, Working paper, Princeton University and Université de Paris-6.
- Aït-Sahalia, Yacine, and Jean Jacod, 2008b, Testing whether jumps have finite or infinite activity, Working paper, Princeton University and Université de Paris-6.
- Aït-Sahalia, Yacine, and Jean Jacod, 2009a, Estimating the degree of activity of jumps in high frequency financial data, *Annals of Statistics*, forthcoming -, -.
- Aït-Sahalia, Yacine, and Jean Jacod, 2009b, Testing for jumps in a discretely observed process, *Annals of Statistics* 37, 184–222.
- Andersen, Torben G., Tim Bollerslev, and Francis X. Diebold, 2003, Some like it smooth, and some like it rough, Working paper, Northwestern University.
- Ball, Clifford A., and Walter N. Torous, 1983, A simplified jump process for common stock returns, *Journal of Financial and Quantitative Analysis* 18, 53–65.
- Barndorff-Nielsen, Ole E., 1997, Normal inverse Gaussian distributions and stochastic volatility modelling, *Scandinavian Journal of Statistics* 24, 1–13.
- Barndorff-Nielsen, Ole E., 1998, Processes of normal inverse Gaussian type, *Finance and Stochastics* 2, 41–68.
- Barndorff-Nielsen, Ole E., and Neil Shephard, 2004, Power and bipower variation with stochastic volatility and jumps (with discussion), *Journal of Financial Econometrics* 2, 1–48.
- Bates, David S., 1991, The crash of '87: Was it expected? The evidence from options markets, *Journal of Finance* 46, 1009–1044.
- Belomestny, Denis, 2008, Spectral estimation of the fractional order of a Lévy process, Working paper, Weierstrass Institute for Applied Analysis and Stochastics, Berlin, Germany.
- Blumenthal, R.M., and R.K. Gettoor, 1961, Sample functions of stochastic processes with stationary independent increments, *Journal of Mathematics and Mechanics* 10, 493–516.
- Carr, Peter, Hélyette Geman, Dilip B. Madan, and Marc Yor, 2002, The fine structure of asset returns: An empirical investigation, *Journal of Business* 75, 305–332.
- Carr, Peter, and Liuren Wu, 2003, The finite moment log stable process and option pricing, *Journal of Finance* 58, 753–777.
- Carr, Peter, and Liuren Wu, 2004, Time-changed Lévy processes and option pricing, *Journal of Financial Economics* 71, 113–141.

- Cont, Rama, and Cecilia Mancini, 2008, Nonparametric tests for probing the nature of asset price processes, Working paper, Universita di Firenze.
- Delbaen, Freddy, and Walter Schachermayer, 1994, A general version of the fundamental theorem of asset pricing, *Mathematische Annalen* 300, 463–520.
- Eberlein, Ernst, and Ulrich Keller, 1995, Hyperbolic distributions in finance, *Bernoulli* 1, 281–299.
- Fan, Yingying, and Jianqing Fan, 2008, Testing and detecting jumps based on a discretely observed process, Working paper, Princeton University.
- Harrison, Michael, and Stanley Pliska, 1981, Martingales and stochastic integrals in the theory of continuous trading, *Stochastic Processes and Their Applications* 11, 215–260.
- Huang, Xin, and George T. Tauchen, 2005, The relative contribution of jumps to total price variance, *Journal of Financial Econometrics* 4, 456–499.
- Madan, Dilip B., and Frank Milne, 1991, Option pricing with V.G. martingale components, *Mathematical Finance* 1, 39–56.
- Madan, Dilip B., and Eugene Seneta, 1990, The Variance Gamma (V.G.) model for share market returns, *Journal of Business* 63, 511–524.
- Mancini, Cecilia, 2001, Disentangling the jumps of the diffusion in a geometric jumping Brownian motion, *Giornale dell'Istituto Italiano degli Attuari* LXIV, 19–47.
- Merton, Robert C., 1976, Option pricing when underlying stock returns are discontinuous, *Journal of Financial Economics* 3, 125–144.
- Schoutens, Win, 2003, *Lévy Processes in Finance: Pricing Financial Derivatives*. (Wiley U.K.).
- Tauchen, George T., and Viktor Todorov, 2008, Activity signature functions for high-frequency data analysis, Working paper, Duke University.
- Woerner, Jeannette H.C., 2006, Analyzing the fine structure of continuous-time stochastic processes, Working paper, University of Göttingen.




Figure 1
Time series of INTC and MSFT prices, year 2006.




Figure 2
Unconditional distribution of log-returns from INTC and MSFT, year 2006, at the 15 second frequency.




Figure 3
Deconvoluting the log-returns distribution into continuous and jump components.




Figure 4
 Discretely sampled data at interval Δ_n vs. continuous-time sample path, and difference between increments and jumps.

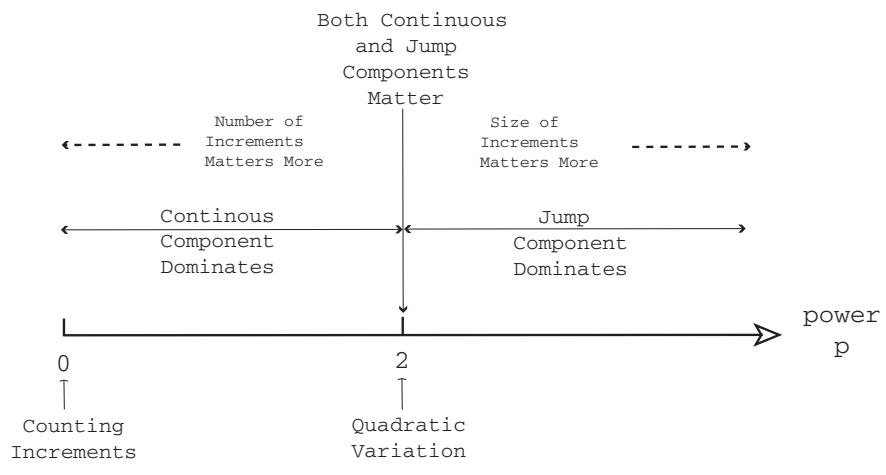


Figure 5
 Adjusting the power p and dominating component in the power variation.




Figure 6
Adjusting the truncation rate u_n and the asymptotic elimination of large jumps.

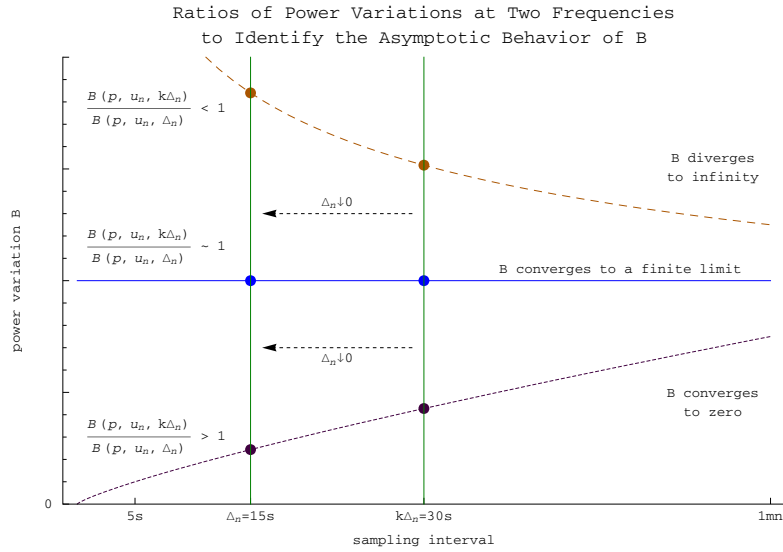


Figure 7
Three possible asymptotic behaviors of the power variation (diverge to infinity, converge to a finite limit, converge to zero) and means of identifying them.




Figure 8
Truncating to retain or avoid the contribution from the Brownian component of the model.

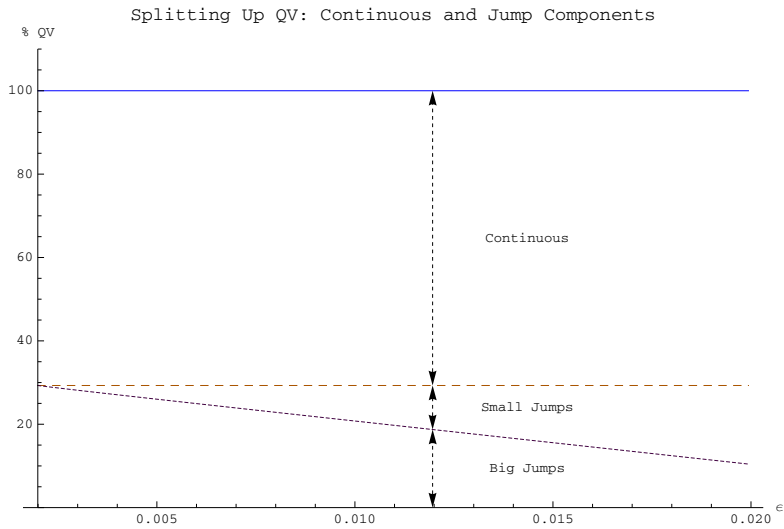


Figure 9
Splitting up the QV into continuous and jump components, and into small and big jumps as a function of the jump size cutoff ϵ .

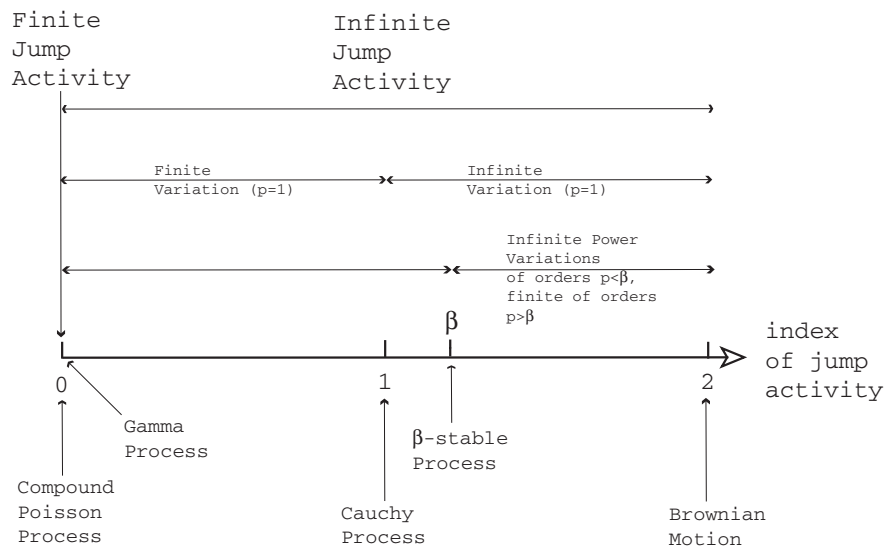


Figure 10
Index of jump activity β : Examples of processes.

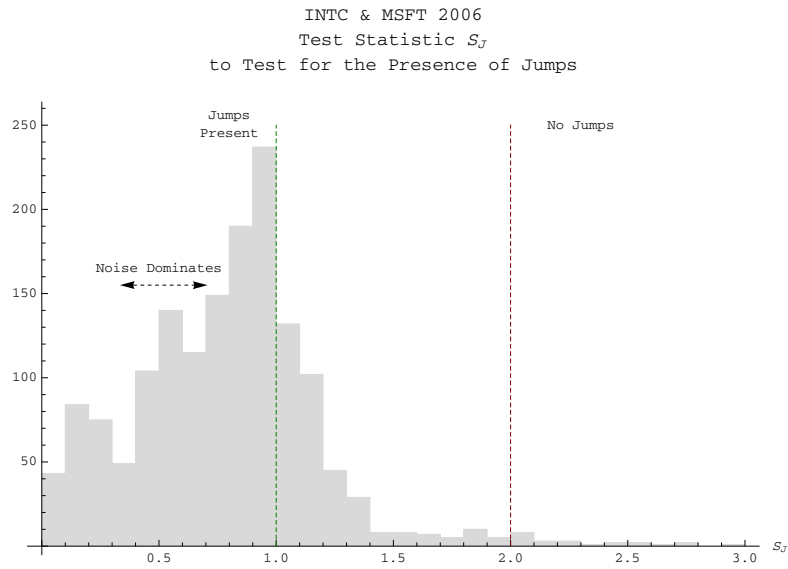


Figure 11
Empirical distribution of S_J for INTC and MSFT, 2006.




Figure 12
Average value of S_J as a function of the sampling interval for INTC and MSFT, 2006.




Figure 13
Empirical distribution of S_{FA} for INTC and MSFT, 2006.




Figure 14
Average value of S_{FA} as a function of the sampling interval for INTC and MSFT, 2006.

INTC & MSFT 2006
 Test Statistic S_W
 to Test Whether Brownian Motion is Present




Figure 15
 Empirical distribution of S_W for INTC and MSFT, 2006.




Figure 16
 Average value of S_W as a function of the sampling interval for INTC and MSFT, 2006.

INTC & MSFT 2006
 Proportion of Quadratic Variation
 Attributable to Continuous Component




Figure 17
 Empirical distribution of the proportion of QV attributable to the continuous component for INTC and MSFT, 2006.




Figure 18
 Average proportion of QV attributable to the continuous component as a function of the sampling interval for INTC and MSFT, 2006.

INTC & MSFT 2006
 Estimate of the Degree of Jump Activity β




Figure 19
 Empirical distribution of the index of jump activity β for INTC and MSFT, 2006.

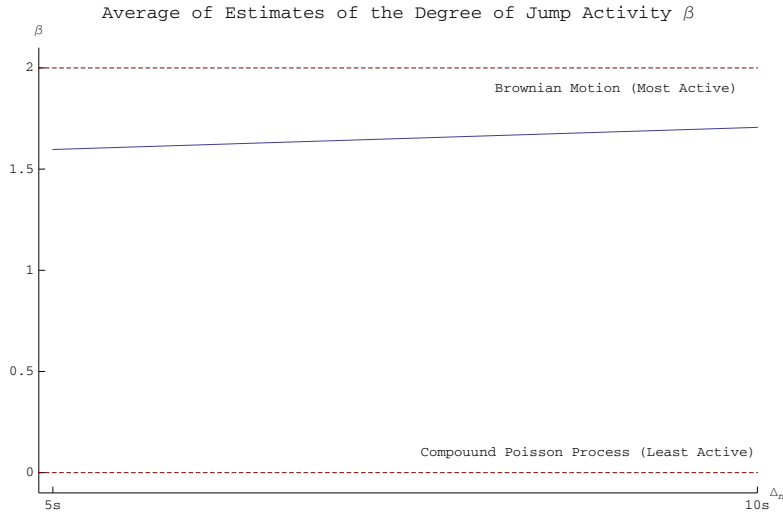


Figure 20
 Average value of the index of jump activity β as a function of the sampling interval for INTC and MSFT, 2006.

H_0		Jumps: Present or Not	
		Ω_T^c	Ω_T^j
H_1	Ω_T^c	\ddots	$S_J :$ $\begin{pmatrix} p > 2 \\ \infty \\ \Delta_n, k\Delta_n \end{pmatrix}$
	Ω_T^j	$S_J :$ $\begin{pmatrix} p > 2 \\ \infty \\ \Delta_n, k\Delta_n \end{pmatrix}$	\ddots

Table I
Combinations of (p, u, Δ) employed to test for the presence of jumps, and whether jumps have finite or infinite activity.

H_0		Jumps: Finite or Infinite Activity	
		Ω_T^f	Ω_T^i
H_1	Ω_T^f	\ddots	$S_{IA} :$ $\begin{pmatrix} p > 2, p' > 2 \\ u_n, \gamma u_n \\ \Delta_n \end{pmatrix}$
	Ω_T^i	$S_{FA} :$ $\begin{pmatrix} p > 2 \\ u_n \\ \Delta_n, k\Delta_n \end{pmatrix}$	\ddots

Table II
Combinations of (p, u, Δ) employed to test whether jumps have finite or infinite activity.

H_0		Brownian Motion: Present or Not	
		Ω_T^W	Ω_T^{noW}
H_1			
Ω_T^W	\dots	$S_{noW} :$ $\left(\begin{array}{c} p = 0, p' = 2 \\ u_n, \gamma u_n \\ \Delta_n \end{array} \right)$	
Ω_T^{noW}	$S_W :$ $\left(\begin{array}{c} p < 2 \\ u_n \\ \Delta_n, k\Delta_n \end{array} \right)$		\dots

Table III
Combinations of (p, u, Δ) employed to test whether a continuous component is present.

Relative Magnitude of the Components	Estimating the Degree of Jump Activity β
$\left(\begin{array}{c} p = 2 \\ u_n \\ \Delta_n \end{array} \right)$	$\hat{\beta} :$ $\left(\begin{array}{c} p = 0 \\ u_n, \gamma u_n \\ \Delta_n \end{array} \right)$
	$\hat{\beta}' :$ $\left(\begin{array}{c} p = 0 \\ u_n \\ \Delta_n, k\Delta_n \end{array} \right)$

Table IV
Combinations of (p, u, Δ) employed to estimate the relative magnitude of the components and the degree of jump activity.