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

Yacine Aït-Sahalia
Princeton University

Jean Jacod
Université de Paris VI

1. Introduction

- Basic framework: X, often the log of an asset price, is assumed to follow an Itô semimartingale.

- 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.
  - Such a process will always generate a finite number of big jumps.
  - But it may give rise to either a finite or infinite number of small jumps.

- The model is

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

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

- μ is the jump measure of X, and its predictable compensator is the Lévy measure ν.

- The distinction between small and big jumps (ε) is arbitrary. What is important is that $\varepsilon > 0$ is fixed.

- In earlier work, we developed tests to determine on the basis of the observed sampled path on $[0, T]$:
  - whether a jump part was present
  - whether the jumps had finite or infinite activity
  - in the latter situation proposed a definition and an estimator of a degree of jump activity parameter
  - whether a Brownian continuous component was needed once infinite activity jumps are included
- In this talk, we show how these different results can be put in a common framework using a common methodology.

- We 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, 2006 returns on MSFT and INTC

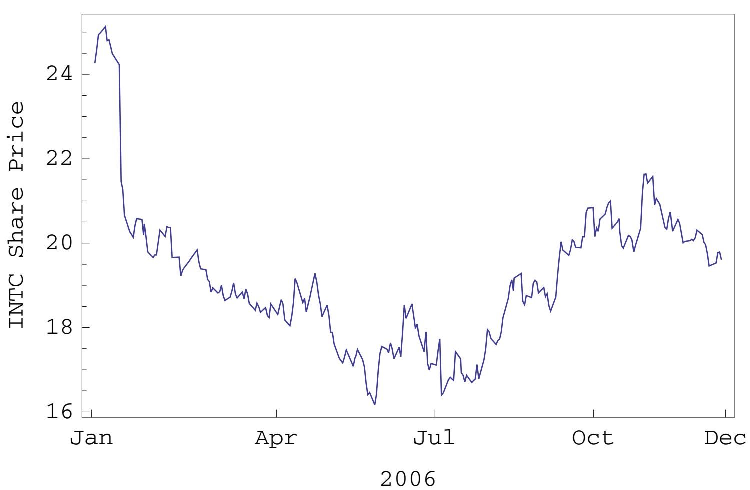

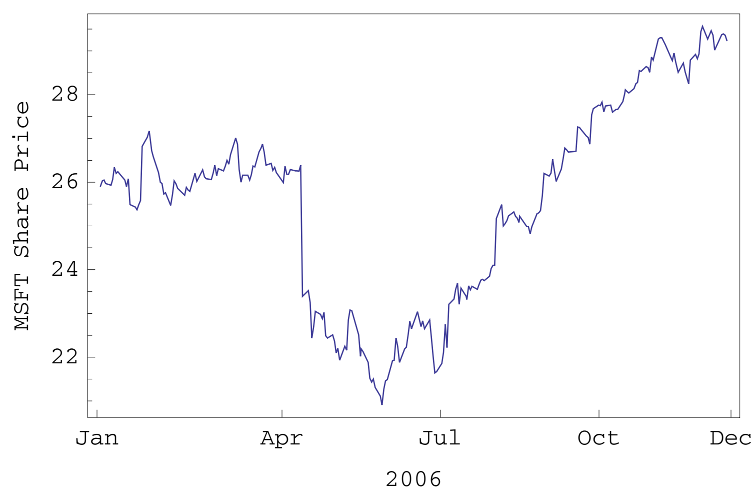

- And then design a set of statistical tools that can tell us something about specific components of the process that produced the observations

- 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's the light (visible or not)
  - here the data generating mechanism is assumed to be a semimartingale; in astrophysics it's whatever nuclear reactions inside the star are producing the light

- In astrophysics, one can look at a specific range of the light spectrum to learn something 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 something about the different components of the semimartingale that produced those returns

  - decide which component(s) need to be included in the model (jumps, finite or infinite activity, continuous component, etc.)

  - determine their relative magnitude

  - magnify specific components of the model if they are present, so we can analyze their finer characteristics (such as the degree of activity of jumps)

- From the time series of returns, we get the distribution of returns at time interval Δ_n

- 2006 returns on MSFT and INTC at 15 seconds

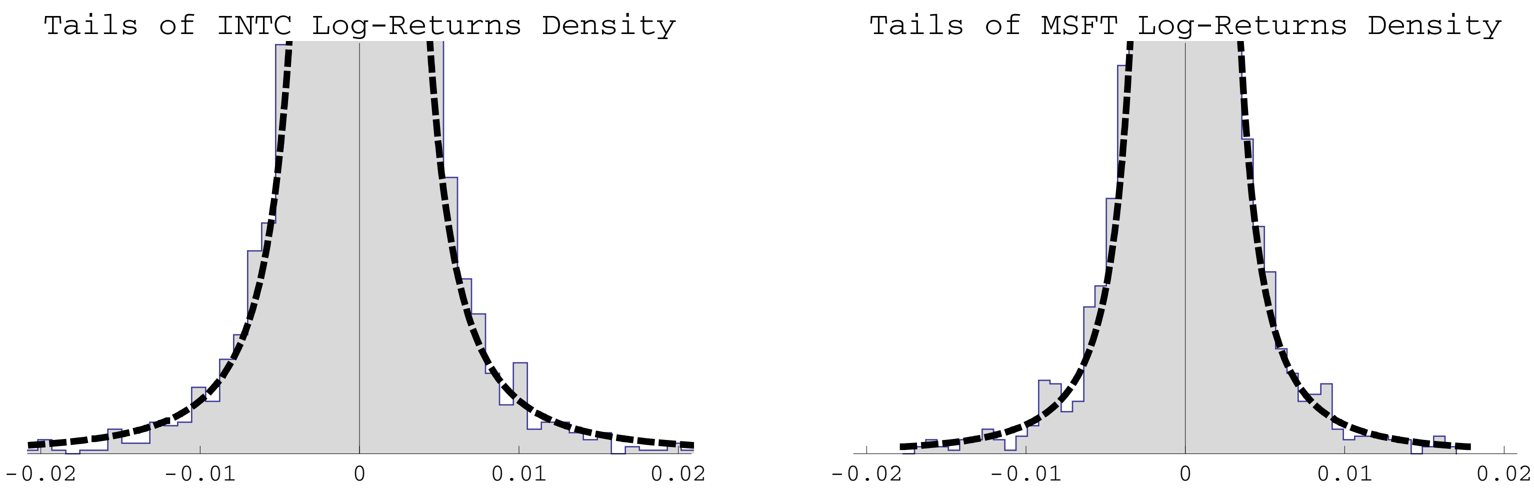

- From the previous plot, we would like to figure out which components should be included in the model

- And in what proportions

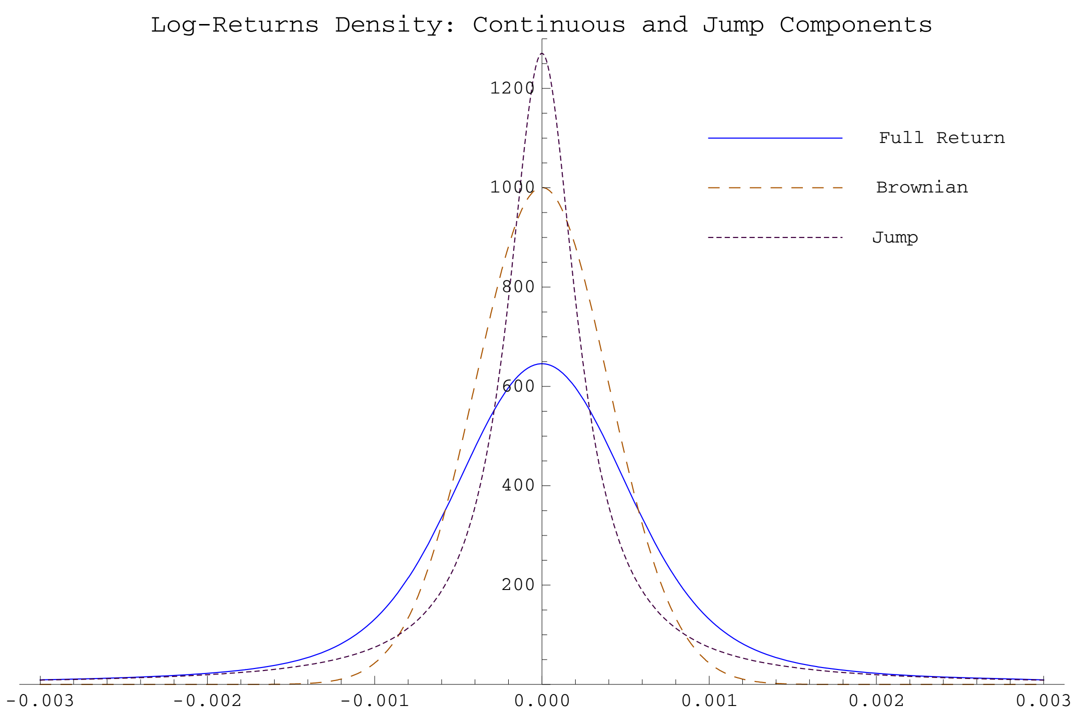

- Similarly to what is done in spectrographic analysis
  - we will emphasize visual tools
  - so we will only include the LLN here
  - and refer to the underlying papers for the formal derivations including regularity conditions and the CLT, as well as simulations.

2. The Measurement Device

- We construct power variations of the increments, suitably truncated and/or sampled at different frequencies.

- We exploit the different asymptotic behavior of the variations as we vary:
  - the power p
  - the truncation level u
  - the sampling frequency Δ

- 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 and Δ, is like having three knobs to adjust in the measurement device.

- Varying the power
  - Powers $p < 2$ will emphasize the continuous component of the underlying sampled process.
  - Powers $p > 2$ will conversely accentuate its jump component.
  - The power $p = 2$ puts them on an equal footing.

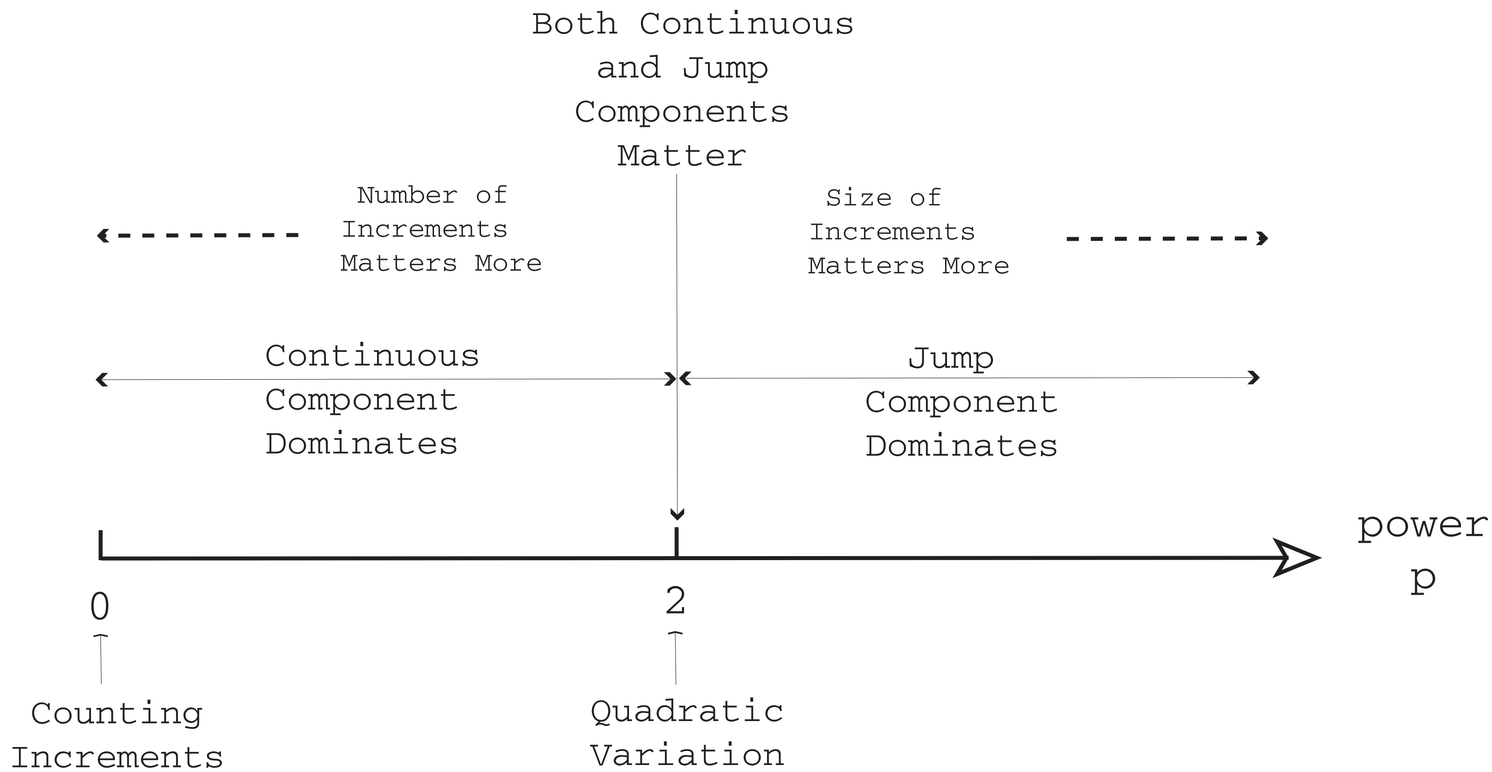
Both Continuous
and Jump
Components
Matter
Number of
Increments
Matters More
Size of
Increments
Matters More
Continuous
Component
Dominates
Jump
Component
Dominates
power
p
0
2
Counting
Increments
Quadratic
Variation

- Truncating the large increments at a suitably selected cutoff level can eliminate the big jumps when needed

- Early use of this device: Mancini (2001)

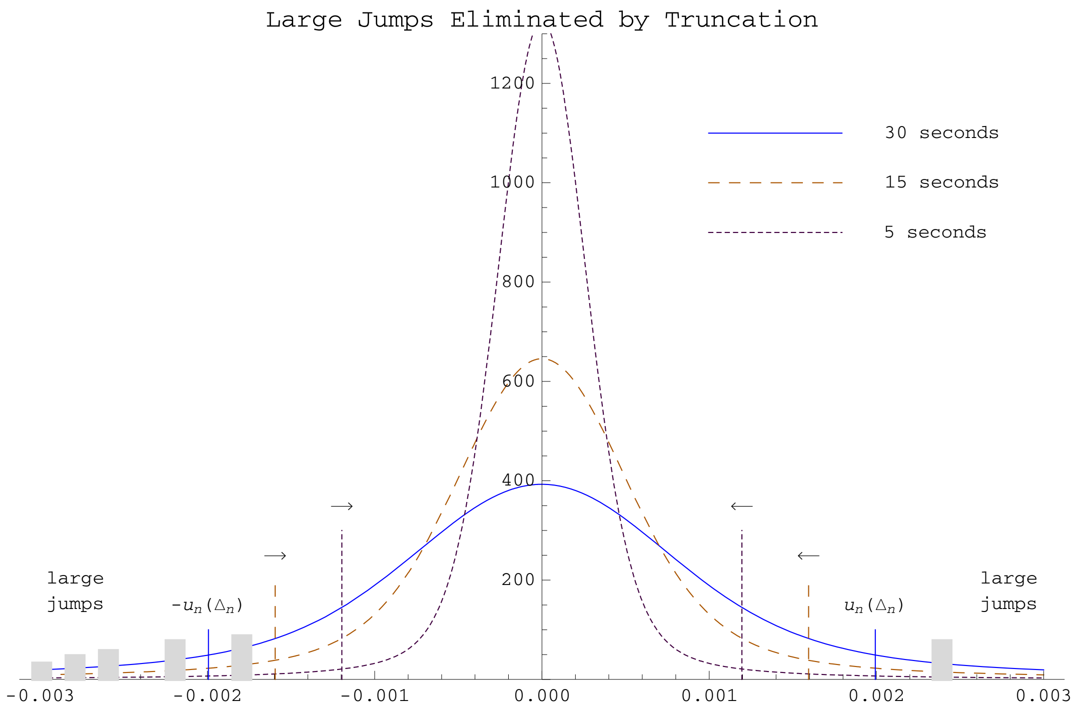

- Sampling at different frequencies can let us distinguish between situations where the variations:
  - converge to a finite limit;
  - converge to zero;
  - diverge to infinity.

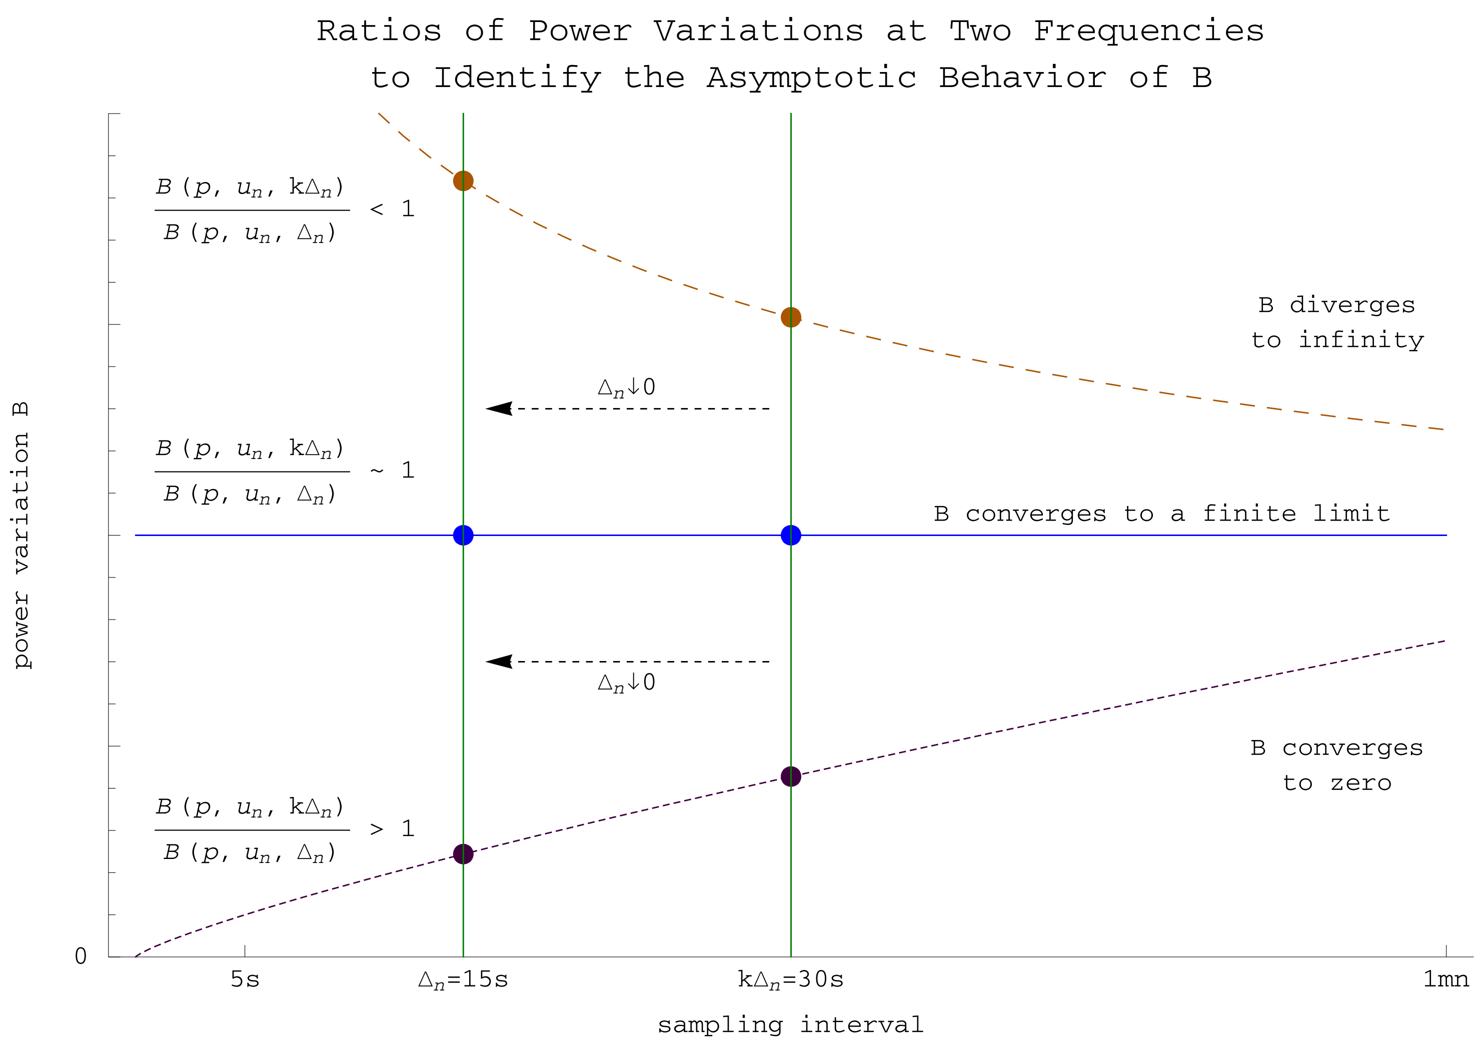

Ratios of Power Variations at Two Frequencies
to Identify the Asymptotic Behavior of B
$\frac{B(p,\, u_n,\, k\Delta_n)}{B(p,\, u_n,\, \Delta_n)} < 1$
B diverges
to infinity
$\Delta_n \downarrow 0$
$\frac{B(p,\, u_n,\, k\Delta_n)}{B(p,\, u_n,\, \Delta_n)} \sim 1$
B converges to a finite limit
$\Delta_n \downarrow 0$
B converges
to zero
$\frac{B(p,\, u_n,\, k\Delta_n)}{B(p,\, u_n,\, \Delta_n)} > 1$
power variation B
0
5s
Δ_n=15s
kΔ_n=30s
1mn
sampling interval

- These 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 effectively allow us to distinguish between all manners of null and alternative hypotheses.

- There are n observed increments of X on $[0, T]$, which are

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

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

$$\Delta X_s = X_s - X_{s-}$$

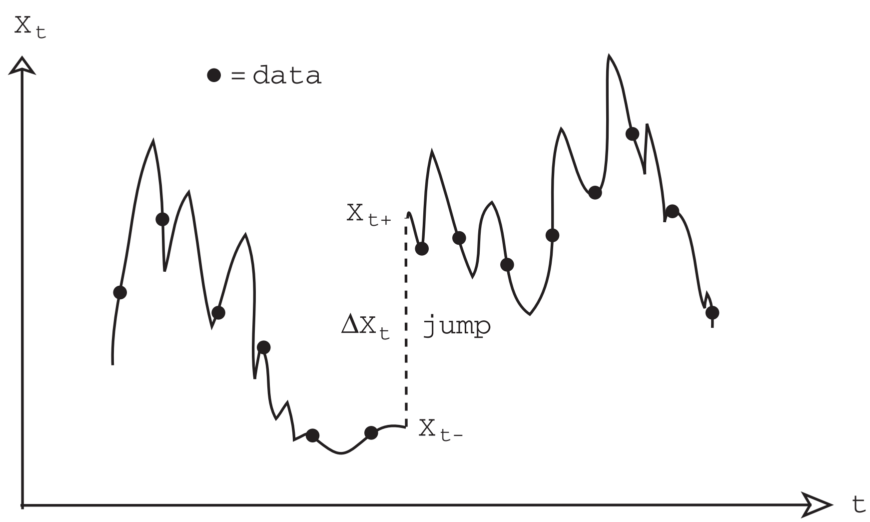

- For any real $p \geq 0$, the basic instruments are the sum of the p^{th}power of the increments of X, sampled at time interval Δ_n, and truncated at level u_n :

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

- The entire methodology relies only on the computation of B for various values of (p, u_n, Δ_n) :

```
B(p,u,del)=sum((abs(dX(del)).^p).*(abs(dX(del))<=u(del)))
```

- T is fixed, asymptotics are all with respect to $\Delta_n \to 0$.
- u_n is the cutoff level for truncating the increments
- $u_n \to 0$ when $n \to \infty$: in the form $u_n = \alpha \Delta_n^{\varpi}$ for some $\varpi \in (0, 1/2)$.
- $\varpi < 1/2$ to keep all the increments which contain a Brownian contribution.
- There will be further restrictions on the rate at which $u_n \to 0$, expressed in the form of restrictions on the choice of ϖ.
- If we don't want to truncate, we write $B(p, \infty, \Delta_n)$.

- Sometimes we will truncate in the other direction, that is retain only the increments larger than u :

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

- With $u_n = \alpha \Delta_n^{\varpi}$ and $\varpi < 1/2$, that can allow us to eliminate all the increments from the continuous part of he model.

- In terms of the power variations B :

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

- Sometimes, we will simply count the number of increments of X, that is, take the power $p = 0$

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

3. Which Component(s) Are Present

- Leaving aside the drift (effectively invisible at high frequency), the model has three components

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

$$\text{JUMPS} = \underbrace{\int_0^t \int_{\{|x| \leq 1\}} x(\mu - \nu)(ds, dx)}_{\text{small jumps}} + \underbrace{\int_0^t \int_{\{|x| > 1\}} x\mu(ds, dx)}_{\text{big jumps}}$$

- The analogy with spectrography would be that we are looking for three possible chemical elements (say, hydrogen, helium and everything else).

- Consider the sets

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

- Formally, $\Omega_T^W = \left\{\int_0^T \sigma_s^2 ds > 0\right\}$ and $\Omega_T^{\text{no}W} = \left\{\int_0^T \sigma_s^2 ds = 0\right\}$.

- We observe a time series and wish to determine in which set(s) the path was.

- There are theoretically many possible ways to do this, even if we restrict attention to power variations only.

- However, we wish to construct test statistics that are model-free in the sense that:

  – their implementation does not require that we estimate or calibrate the model, which can potentially be quite complicated (stochastic volatility, jumps, jumps in volatility, jumps in jump intensity, etc.)

  – so we want the distribution of the test statistics to be assessed using only power variations (of perhaps other powers, truncation levels and sampling frequencies)

3.1. Jumps: Present or Not

- Here are processes which measure some kind of variability of X and depend on the whole (unobserved) path of X:

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

where $p > 0$ and $\Delta X_s = X_s - X_{s-}$ are the jumps of X.

- $A(p)$ is finite for all $p > 0$. $B(p)$ is finite if $p \geq 2$ but often not when $p < 2$.

- The quadratic variation of the process is $[X,X]_T = A(2) + B(2)$.

- We have $\begin{cases} p>2, \text{ all } X & \Rightarrow \quad B(p,\infty,\Delta_n) \stackrel{\mathbb{P}}{\longrightarrow} B(p) \\ \text{all } p, \ X \text{ continuous} & \Rightarrow \quad \frac{\Delta_n^{1-p/2}}{m_p} B(p,\infty,\Delta_n) \stackrel{\mathbb{P}}{\longrightarrow} A(p). \end{cases}$

- 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:

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

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

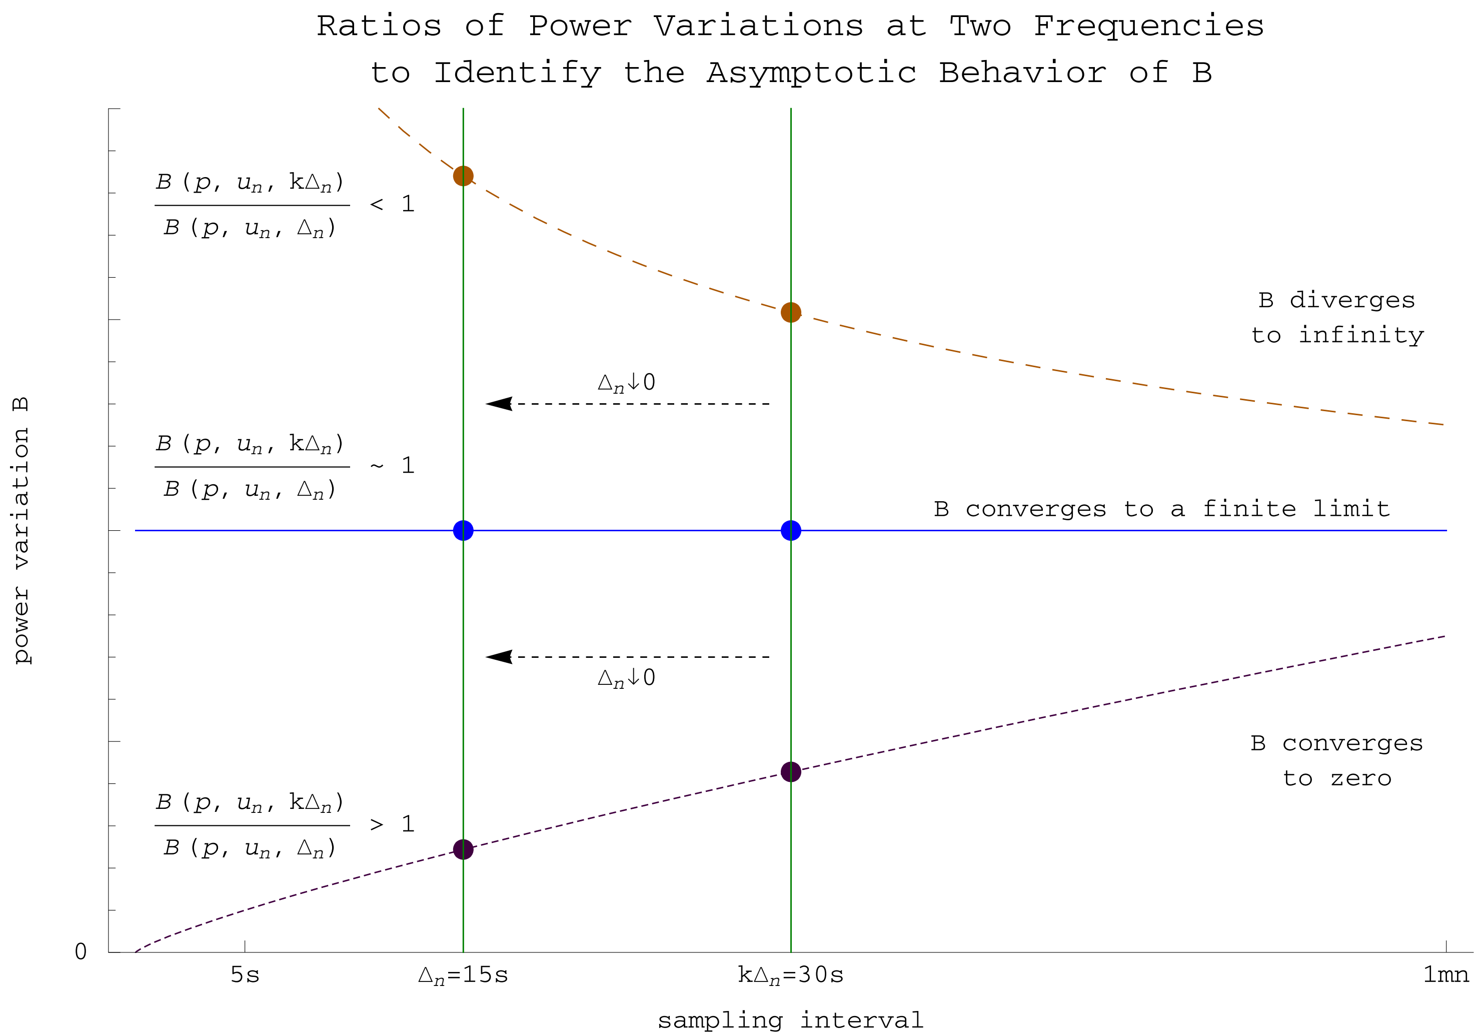
Ratios of Power Variations at Two Frequencies
to Identify the Asymptotic Behavior of B
$\frac{B(p, u_n, k\Delta_n)}{B(p, u_n, \Delta_n)} < 1$
B diverges to infinity
$\Delta_n \downarrow 0$
$\frac{B(p, u_n, k\Delta_n)}{B(p, u_n, \Delta_n)} \sim 1$
B converges to a finite limit
$\Delta_n \downarrow 0$
B converges to zero
$\frac{B(p, u_n, k\Delta_n)}{B(p, u_n, \Delta_n)} > 1$
power variation B
0
5s
Δ_n=15s
kΔ_n=30s
1mn
sampling interval

- Theorem

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

- This is valid on Ω_T^j whether the jump component include finite or infinite components, or both.

- We provide a CLT under Ω_T^c and one under Ω_T^j, so one can test either $H_0 : \Omega_T^c$ vs. $H_1 : \Omega_T^j$ or the reverse $H_0 : \Omega_T^j$ vs. $H_1 : \Omega_T^c$.

3.2. 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 oninfinite activity jumps: see for example Madan and Seneta (1990), Eberlein and Keller (1995), Barndorff-Nielsen (1998), Carr, Geman, Madan and Yor (2002), Carr and Wu (2003), etc.

3.2.1. 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 $H_1 : \Omega_T^i$.

- We choose an integer $k \geq 2$ and a real $p > 2$.

- The only difference is that we now truncate

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

- Without truncation, 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.

- 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 $\Delta_n \downarrow 0$, the truncation level $u_n = O(\Delta_n^{\rho})$ will have eliminated all the large jumps.

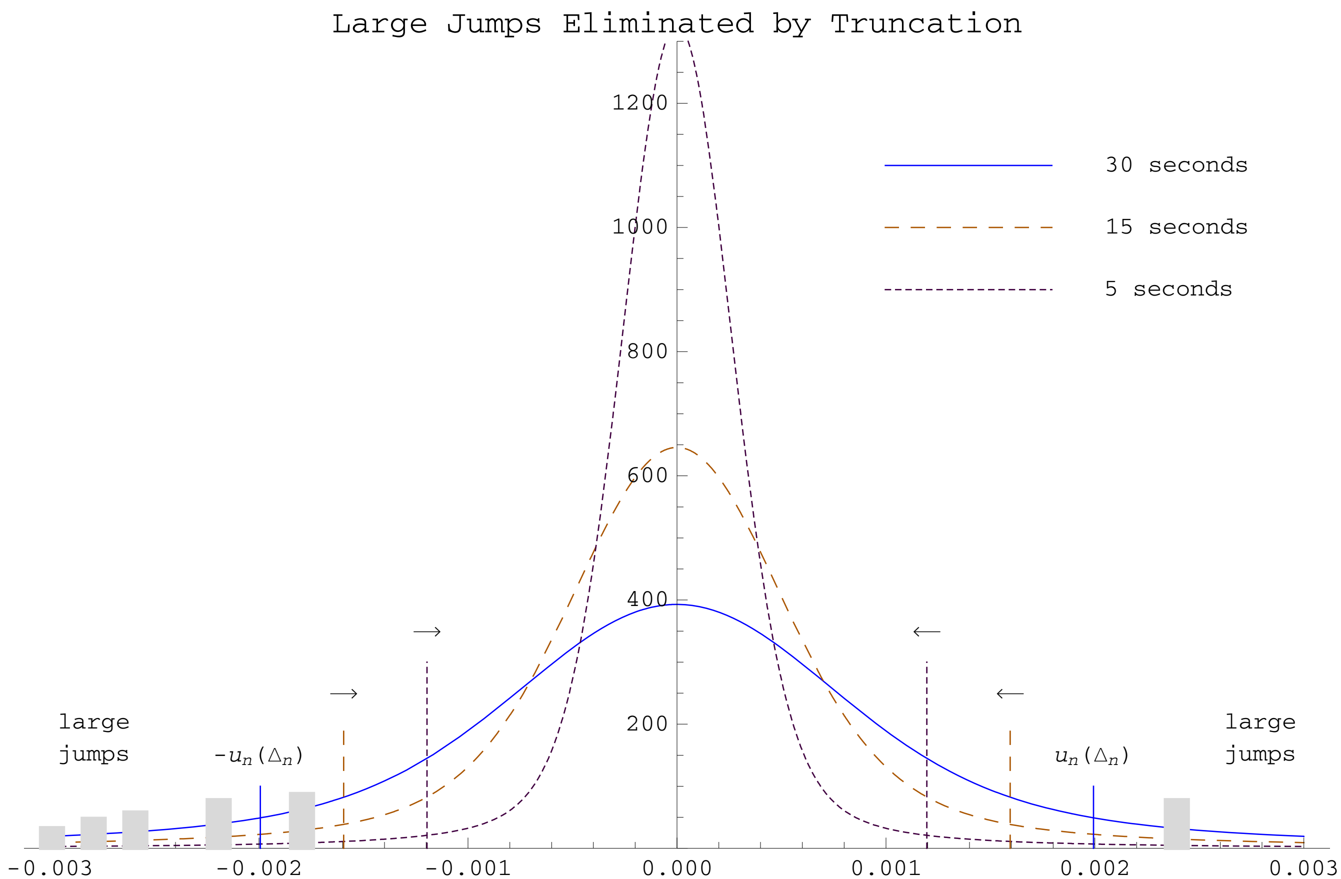
Large Jumps Eliminated by Truncation
1200
1000
800
600
400
200
30 seconds
15 seconds
5 seconds
large
jumps
$-u_n(\Delta_n)$
$u_n(\Delta_n)$
large
jumps
-0.003
-0.002
-0.001
0.000
0.001
0.002
0.003

- Then if there are only big jumps and the Brownian component, the two power variations $B(p, u_n, k\Delta_n)$ and $B(p, u_n, \Delta_n)$ will behave as if there were no jumps and the limit of the ratio will be 2 as in the test for jumps.

- But if there are small jumps, then the truncation cannot eliminate them because their size is Δ_n-dependent then each B truncated tends to the small of remaining jumps and the ratio tends to 1.

- Theorem: Under regularity conditions on u_n,

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

3.2.2. 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 $H_1 : \Omega_T^f \cap \Omega_T^W$.

- Why do we need different statistics? Because the distribution of S_{FA} is not model-free under Ω_T^i, and that of S_{IA} is not model-free under $\Omega_T^f \cap \Omega_T^W$.

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

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

- Theorem: Under regularity conditions on u_n,

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

3.3. Brownian Motion: Present or Not

- We 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: see for example Madan and Seneta (1990), Eberlein and Keller (1995), Carr, Geman, Madan and Yor (2002), Carr and Wu (2003), etc.

3.3.1. 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 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.

- Without the Brownian motion, however, and when $p > \beta$, 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.

- We choose an integer $k \geq 2$ and a real $p < 2$.

- We propose the test statistic

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

- Theorem: Under regularity conditions on u_n,

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

- An alternative: Tauchen and Todorov (2009)

- They employ the test statistic for jumps S_J, plot its logarithm for different values of the power argument and contrast the behavior of the plot above two and below two in order to identify the presence of a Brownian component.

- This works when there is a Brownian motion is present under the null.

3.3.2. Null Hypothesis: No Brownian Motion

- The null model is now pure jump (plus perhaps a drift) with jumps.
  - 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.

- Designing a test under this null is trickier
  - because we are aiming for a test that remains model-free even for this model.
  - that is, despite being driven by what is now a pure jump process, the behavior of the statistic should not depend on the characteristics of the pure jump process
  - such as for instance its degree of activity β
  - since those characteristics are a priori unknown.

- We choose a real $\gamma > 1$ to define two different truncation ratios

- And define a family of test statistics as follows:

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

- 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 $(H_1 : \Omega_T^W)$
  - 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 the Brownian motion is not present $(H_0 : \Omega_T^{\text{no}W})$
  - 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 $\frac{B(2,\gamma u_n,\Delta_n)}{B(2,u_n,\Delta_n)}$ in $S_{\text{no}W}$ will reflect it: it will now be $\gamma^{2-\beta}$.
  - 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 $\frac{U(0,u_n,\Delta_n)}{U(0,\gamma u_n,\Delta_n)}$ of the number of large increments, the U's.
  - 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.

- Theorem: Under regularity conditions on u_n,

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

4. The Relative Magnitude of the Components

- A typical "main sequence" star might be made of 90% hydrogen, 10% helium and 0.1% everything else.

- Here, what is the relative magnitude of the two jump and the continuous components?

- We can answer this question using the same device.

- It makes sense to consider $p = 2$ since this is the power where all the components are present together.

- We can then truncate to split the QV into its continuous and jump components

- And not truncate to estimate the full QV:

$$\frac{B(2,u_n,\Delta_n)}{B(2,\infty,\Delta_n)} = \text{\% of QV due to the continuous component}$$

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

- Alternative splitting of the QV based on bipower variation instead of truncating: Barndorff-Nielsen and Shephard (2004), Huang and Tauchen (2005), Andersen, Bollerslev and Diebold (2007).

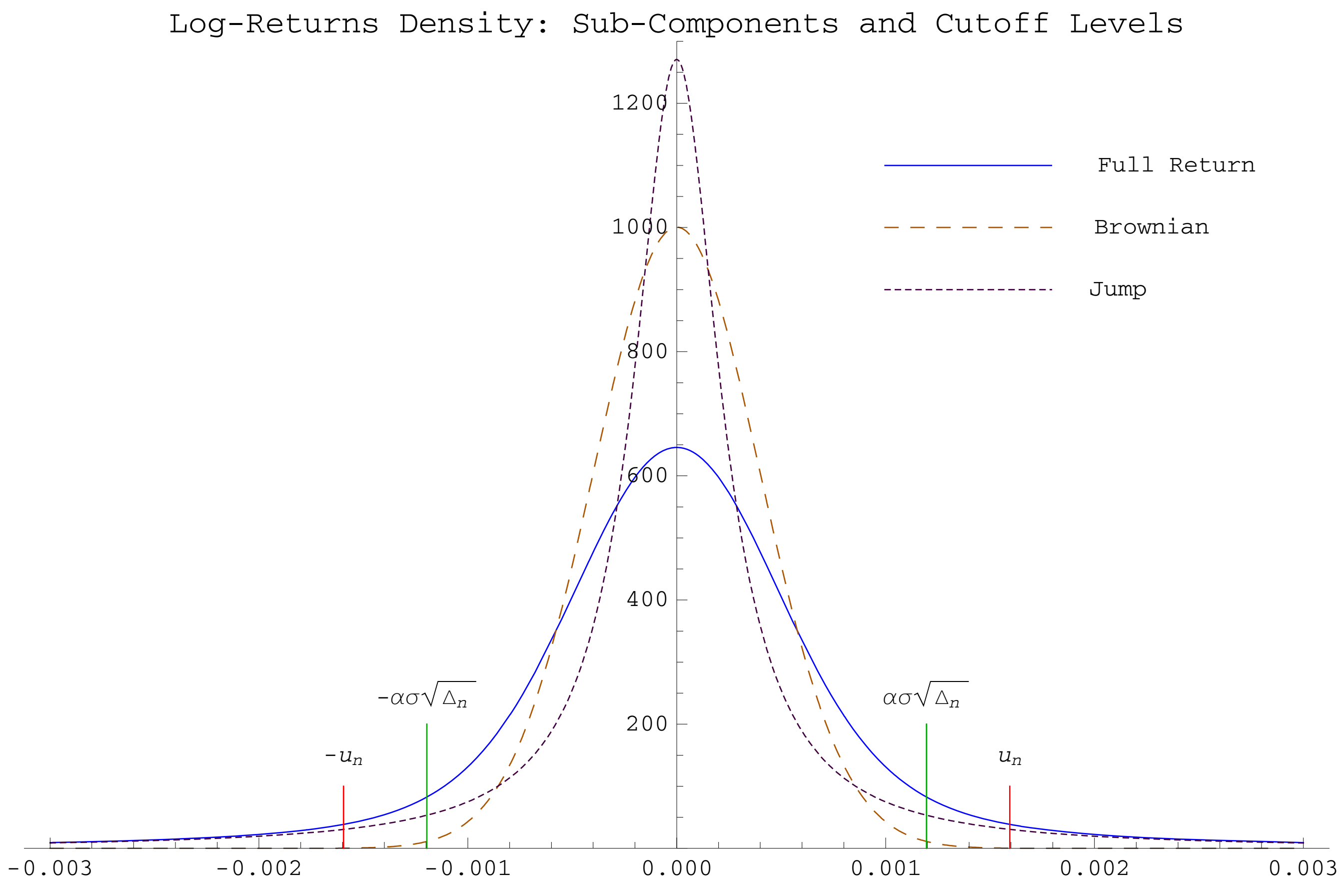
Log-Returns Density: Sub-Components and Cutoff Levels
1200
1000
800
600
400
200
Full Return
Brownian
Jump
$-\alpha\sigma\sqrt{\Delta_n}$
$\alpha\sigma\sqrt{\Delta_n}$
$-u_n$
u_n
-0.003
-0.002
-0.001
0.000
0.001
0.002
0.003

- We can then 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:

$$\frac{U(2,\varepsilon,\Delta_n)}{B(2,\infty,\Delta_n)} = \text{\% of QV due to big jumps}$$

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

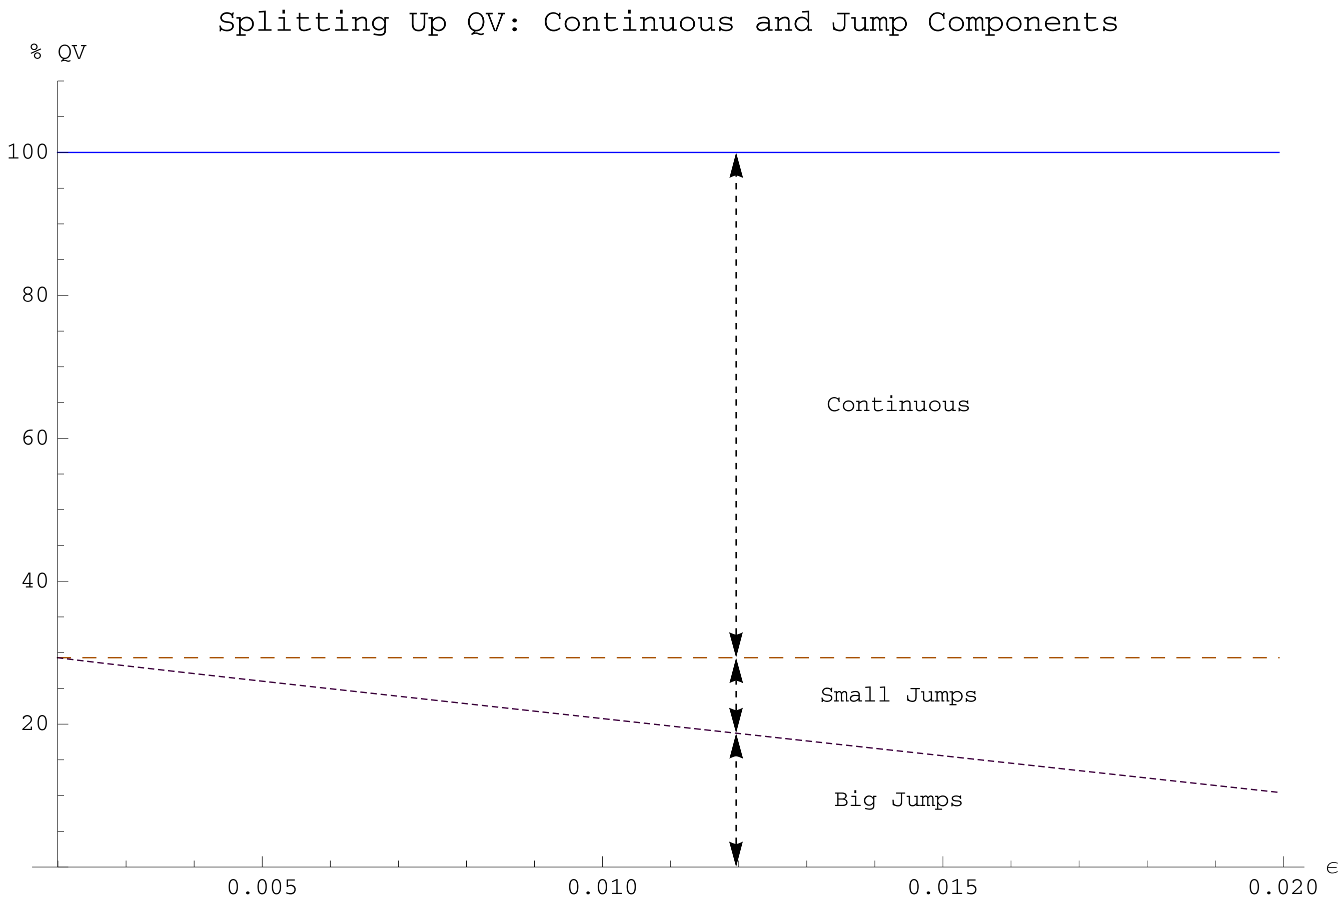
Splitting Up QV: Continuous and Jump Components
% QV
100
80
60
40
20
Continuous
Small Jumps
Big Jumps
0.005
0.010
0.015
0.020
ϵ

5. The Finer Characteristics of the Components

5.1. Defining an Index of Jump Activity

- Recall $B(p) = \sum_{s \leq T} |\Delta X_s|^p$.

- Define $I_T = \{p \geq 0 : B(p) < \infty\}$.

- Necessarily, the (random) set I_T is of the form $[\beta_T, \infty)$ or (β_T, ∞) for some $\beta_T(\omega) \leq 2$, and $2 \in I_T$ always.

- We call $\beta_T(\omega)$ the jump activity index for the path $t \mapsto X_t(\omega)$ at time T.

- We define this index in analogy with 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 $\int_{\{|x| \leq 1\}} |x|^r \nu(dx) < \infty$, where ν is the Lévy measure

  - This property shows that, for a Lévy process, the jump activity index coincides with the Blumenthal-Getoor index of the process.

  - In the further special case where X is a stable process, then β is also the stable index of the process.

- β captures an essential qualitative feature of ν, which is its level of activity: when β increases, the (small) jumps tend to become more and more frequent.
  - Processes with finite jump activity have $\beta = 0$.
  - Processes with infinite jump activity may also have $\beta = 0$ if the rate of divergence of the jump measure is sub-polynomial.
  - Processes with $\beta \in (0, 2)$ have infinite jump activity
  - And the higher β, the more active the jumps.

- Brownian motion has $\beta = 2$ in the limit.

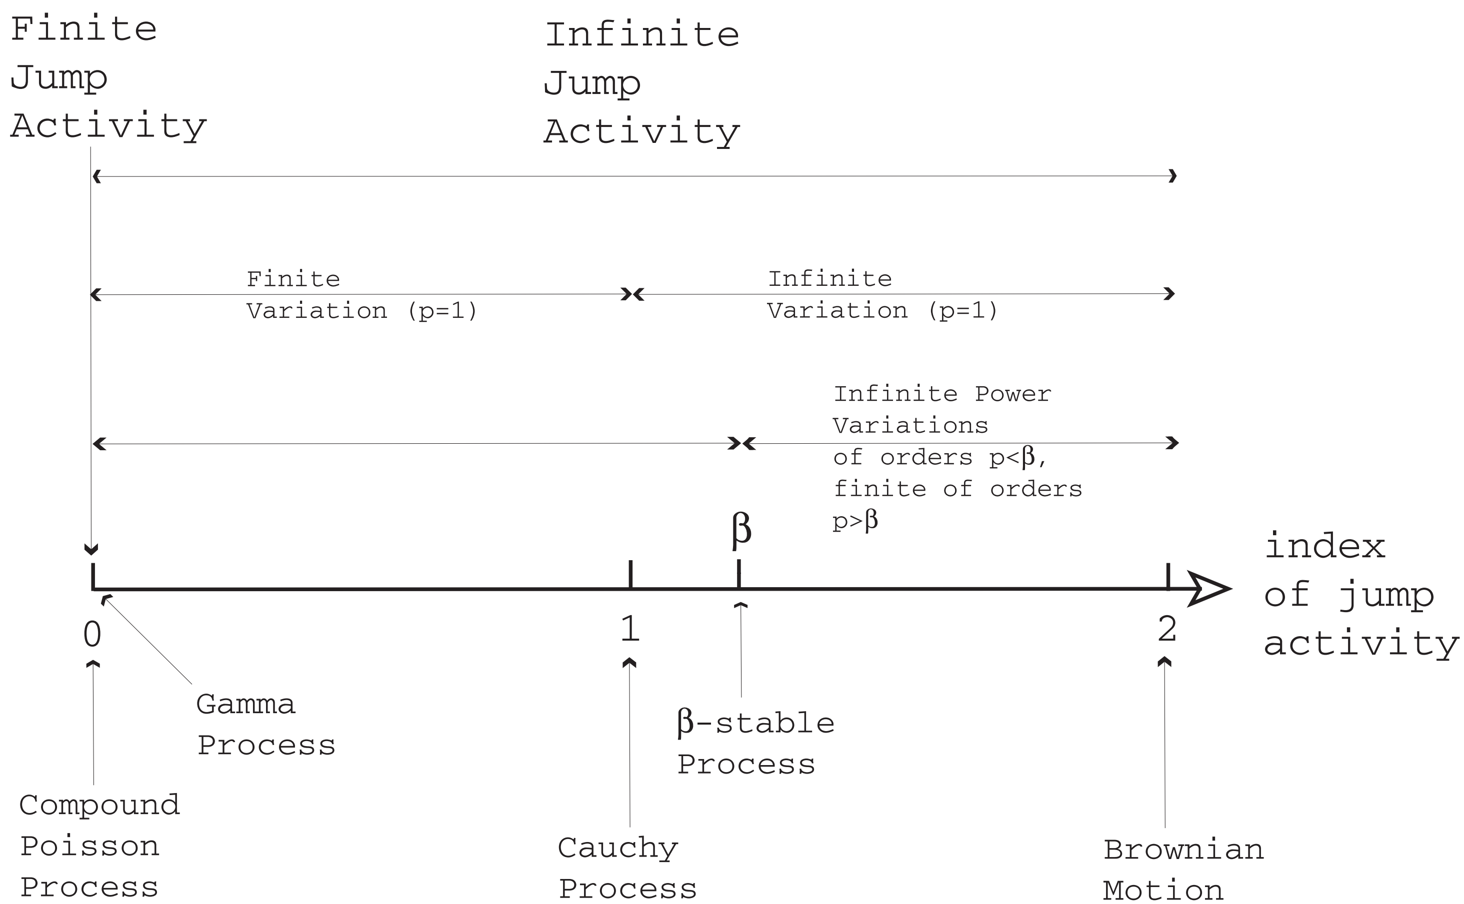
Finite
Jump
Activity
Infinite
Jump
Activity
Finite
Variation (p=1)
Infinite
Variation (p=1)
Infinite Power
Variations
of orders p<β,
finite of orders
p>β
β
index
of jump
activity
0
1
2
Gamma
Process
β-stable
Process
Compound
Poisson
Process
Cauchy
Process
Brownian
Motion

- The problem is made more challenging by the 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.
  - 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 adifferent range of the spectrum of returns

- Considering only returns that are larger than the cutoff $u_n = \alpha\Delta_n^{\varpi}$ for some $\varpi \in (0, 1/2)$.

- This allows us to eliminate the increments due to the continuous component.

- We can then use all values of p, not just those $p > 2$.

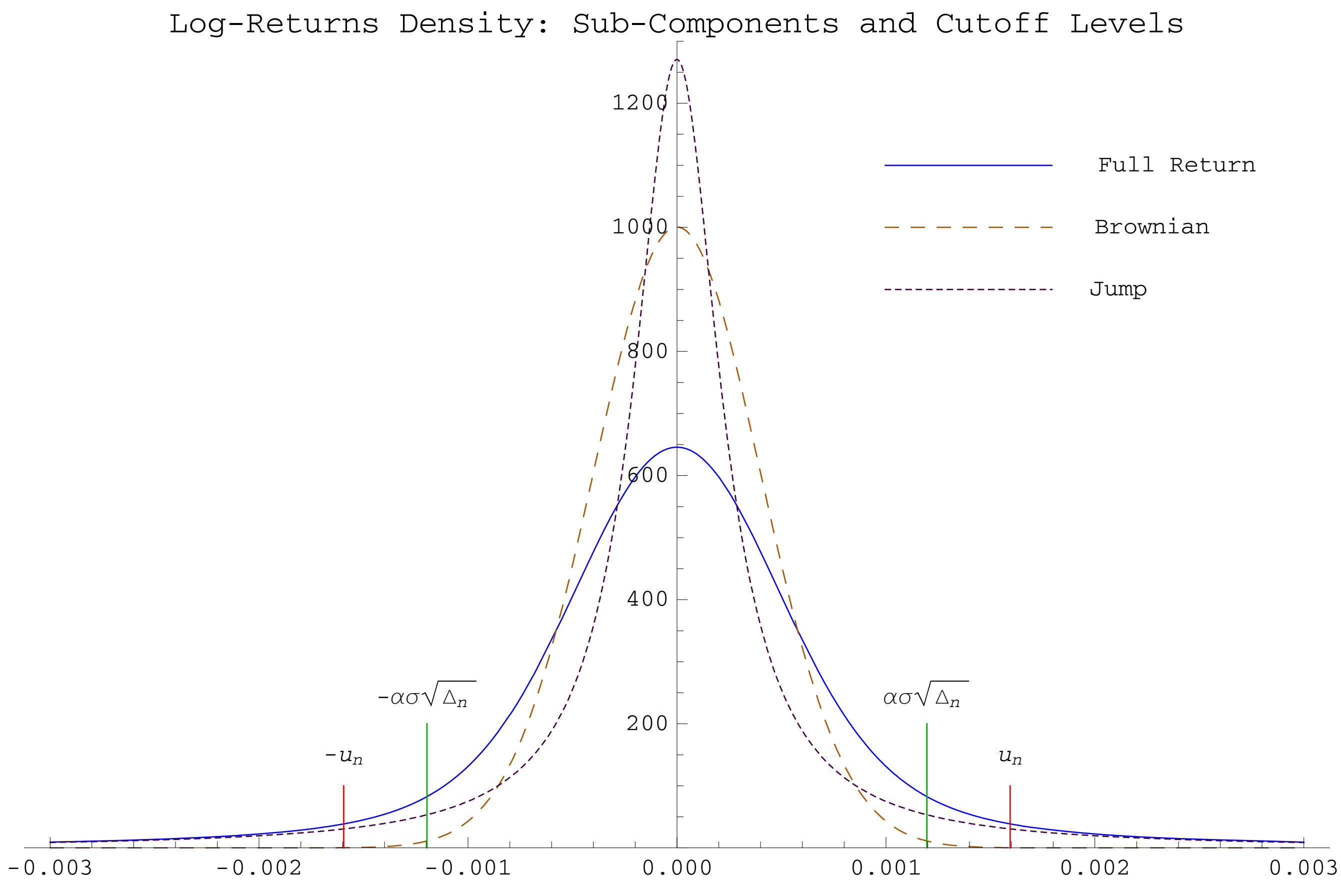
Log-Returns Density: Sub-Components and Cutoff Levels
1200
1000
800
600
400
200
Full Return
Brownian
Jump
$-\alpha\sigma\sqrt{\Delta_n}$
$\alpha\sigma\sqrt{\Delta_n}$
$-u_n$
u_n
-0.003
-0.002
-0.001
0.000
0.001
0.002
0.003

5.2. Estimating Jump Activity

- We propose two estimators of β based on counting the number of increments greater than the cutoff u_n.

- The first one: fix $0 < \alpha < \alpha'$ and consider two cutoffs $u_n = \alpha\Delta_n^{\varpi}$ and $u'_n = \alpha'\Delta_n^{\varpi}$ with $\gamma = \alpha'/\alpha$:

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

- The second one: sample on two time scales, Δ_n and $2\Delta_n$.

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

- Given consistent estimators and with a CLT

- We could test various hypotheses, for instance whether $\beta > 1$ or $\beta < 1$ which correspond to finite or infinite variation for X.

- Related methods: testing whether $\beta = 1$ (Cont and Mancini (2009)), testing whether $\beta = 2$ or $\beta < 2$ (Tauchen and Todorov (2009)).

6. Summary: (p, u, Δ)

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

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

H_1 \ H_0	Brownian Motion: Present or Not Ω_T^W	$\Omega_T^{\text{no}W}$
Ω_T^W	$\ddots$	$S_{\text{no}W}$: $\begin{pmatrix} p = 0,\ p' = 2 \\ u_n,\ \gamma u_n \\ \Delta_n \end{pmatrix}$
$\Omega_T^{\text{no}W}$	S_W : $\begin{pmatrix} p < 2 \\ u_n \\ \Delta_n,\ k\Delta_n \end{pmatrix}$	$\ddots$

Relative Magnitude of the Components
$\begin{pmatrix} p = 2 \\ u_n \\ \Delta_n \end{pmatrix}$

Estimating the Degree of Jump Activity β	
$\hat{\beta}$	$\begin{pmatrix} p = 0 \\ u_n,\ \gamma u_n \\ \Delta_n \end{pmatrix}$
$\hat{\beta}'$	$\begin{pmatrix} p = 0 \\ u_n \\ \Delta_n,\ k\Delta_n \end{pmatrix}$

7. Empirical Results: Intel & Microsoft 2006

7.1. The Data

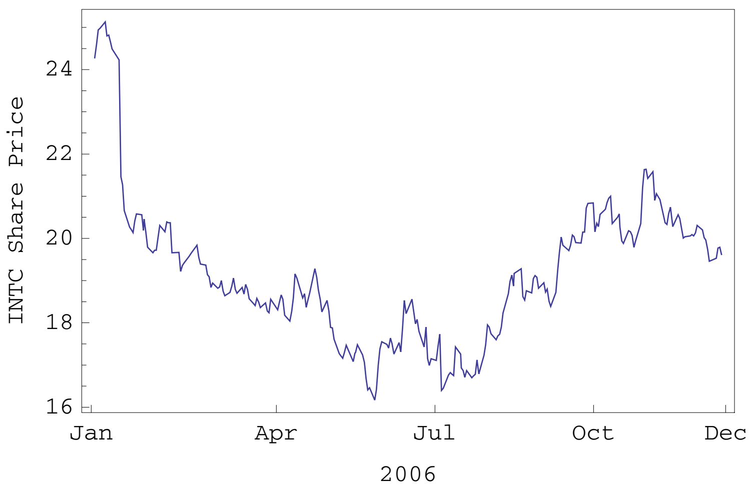

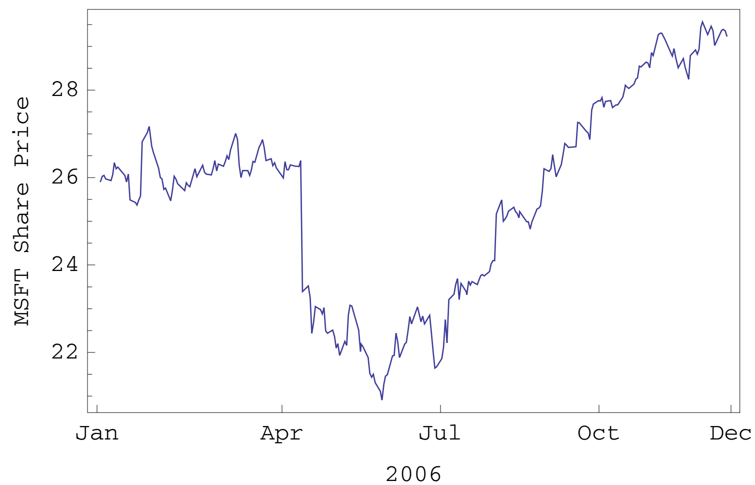

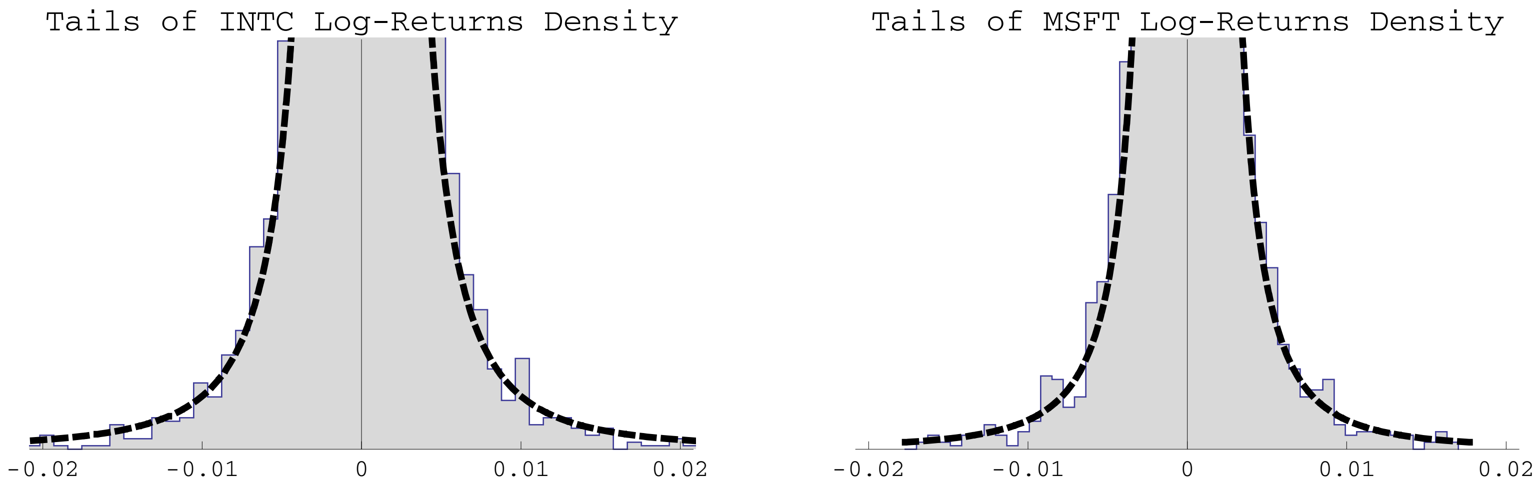
Tails of INTC Log-Returns Density
-0.02
-0.01
0
0.01
0.02
Tails of MSFT Log-Returns Density
-0.02
-0.01
0
0.01
0.02

- Whenever we need to truncate, we express the truncation cutoff level u_n in terms of a number of standard deviations of the continuous part of the semimartingale.

- We consider sampling frequencies up to 5 seconds.

- In real data, observations of the process X are blurred by market microstructure noise, which messes things up at very high frequency.

7.2. Jumps: Present or Not

- Two polar cases: observations are blurred with either an additive white noise or with noise due to rounding

  - Observations are affected by an additive noise, that is 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.

  - Or 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.

- We show that, in the presence of additive noise, $S_J(4, k, \Delta_n) \xrightarrow{\mathbb{P}} \frac{1}{k}$.

- In the presence of rounding error noise, the limit is $\frac{1}{k^{1/2}}$.

- So S_J has four possible limits: with $k = 2$ and $p = 4$,

$$\begin{array}{ccc} 1/2 & : & \text{additive noise dominates} \\ 1/2^{1/2} & : & \text{rounding error dominates} \\ 1 & : & \text{jumps present} \\ 2 & : & \text{no jumps present} \end{array}$$

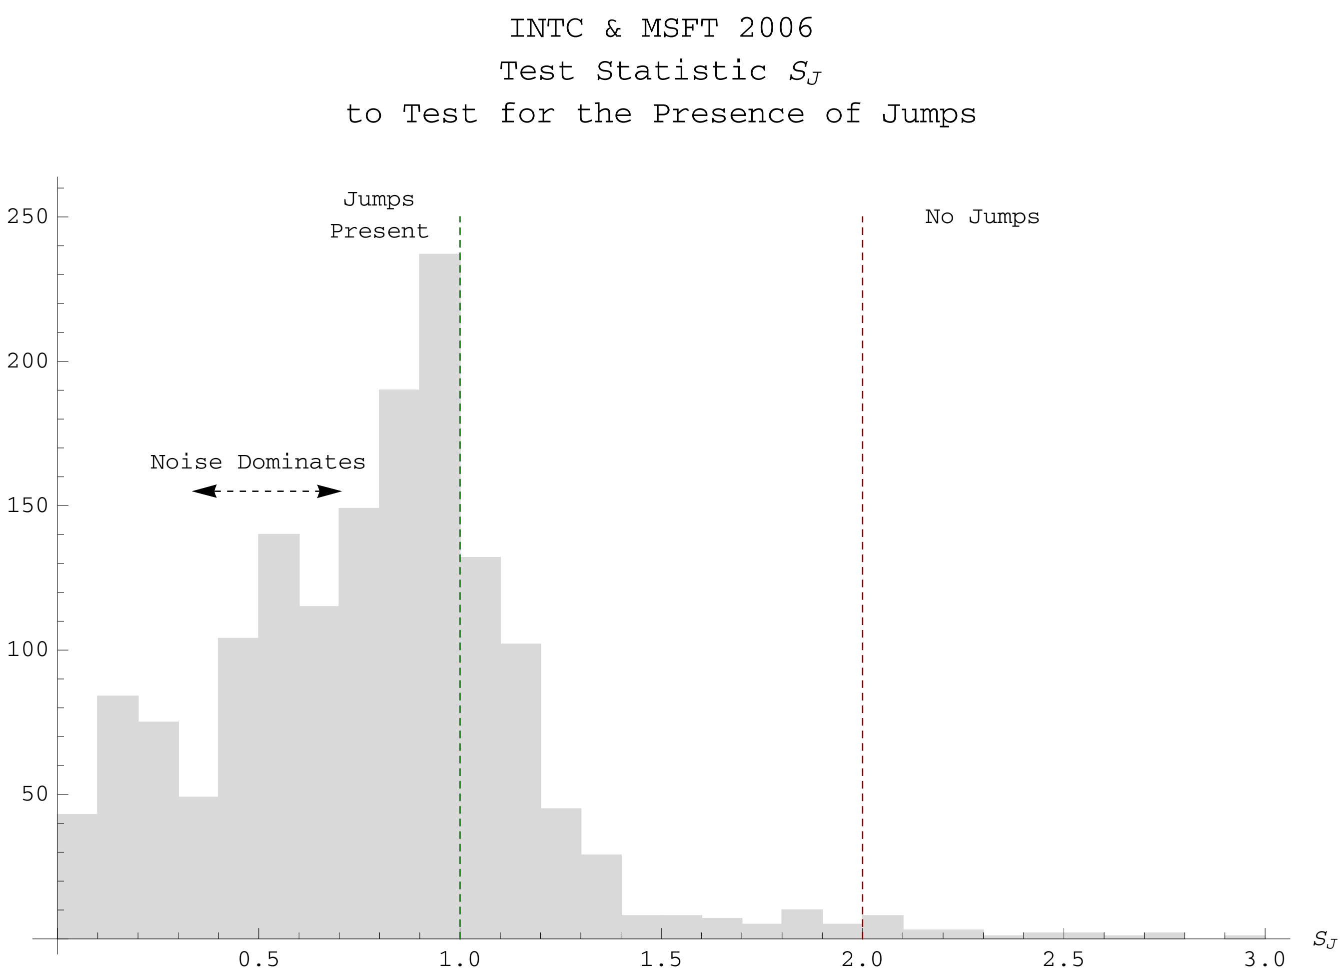
INTC & MSFT 2006
Test Statistic S_J
to Test for the Presence of Jumps
Jumps
Present
No Jumps
Noise Dominates
250
200
150
100
50
0.5
1.0
1.5
2.0
2.5
3.0
S_J

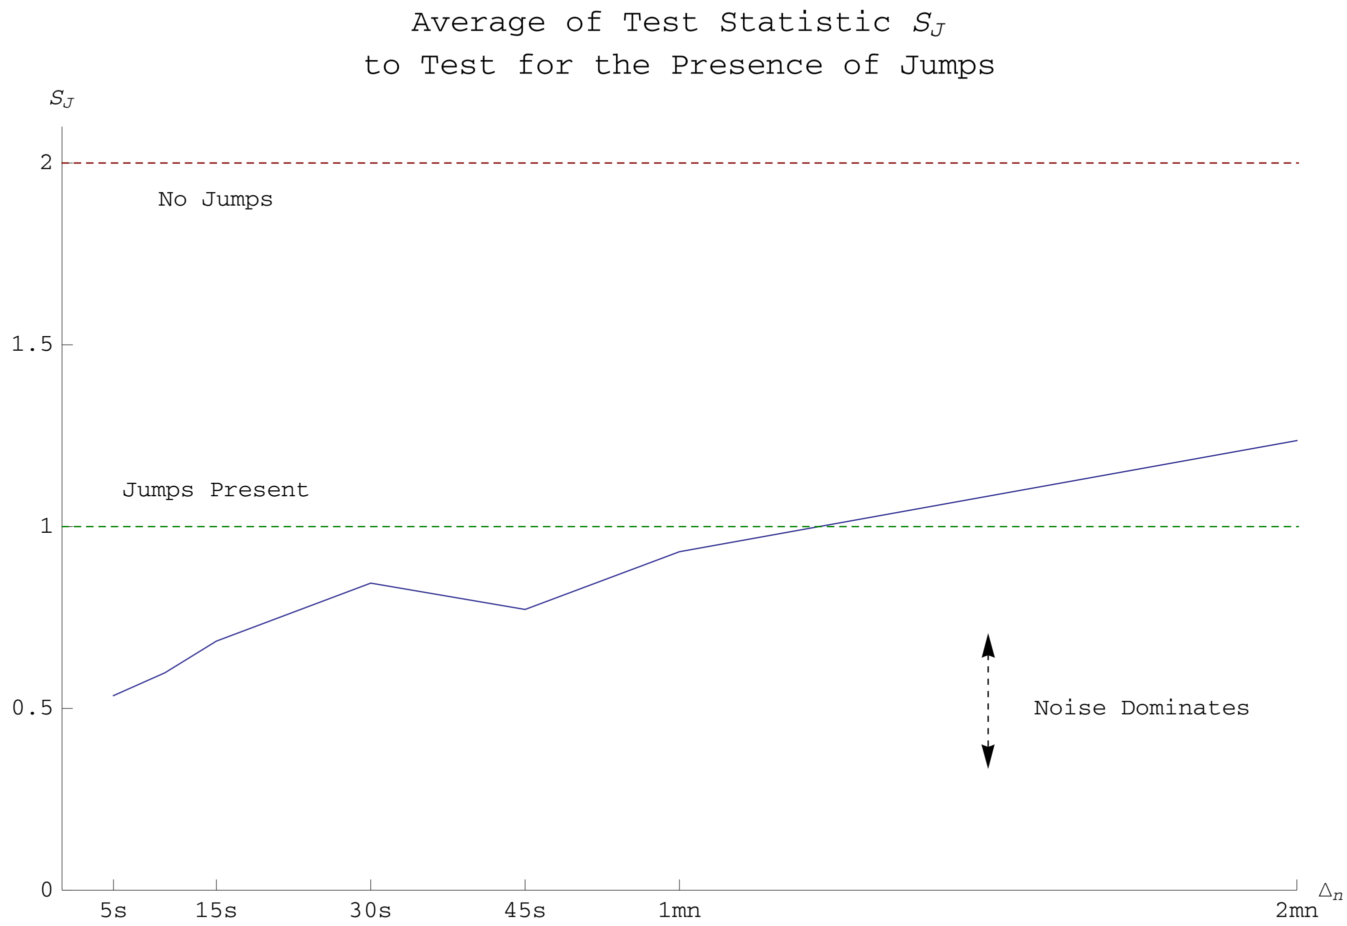
Average of Test Statistic S_J
to Test for the Presence of Jumps
S_J
2
1.5
1
0.5
0
No Jumps
Jumps Present
Noise Dominates
5s
15s
30s
45s
1mn
2mn
Δ_n

7.3. Jumps: Finite or Infinite Activity

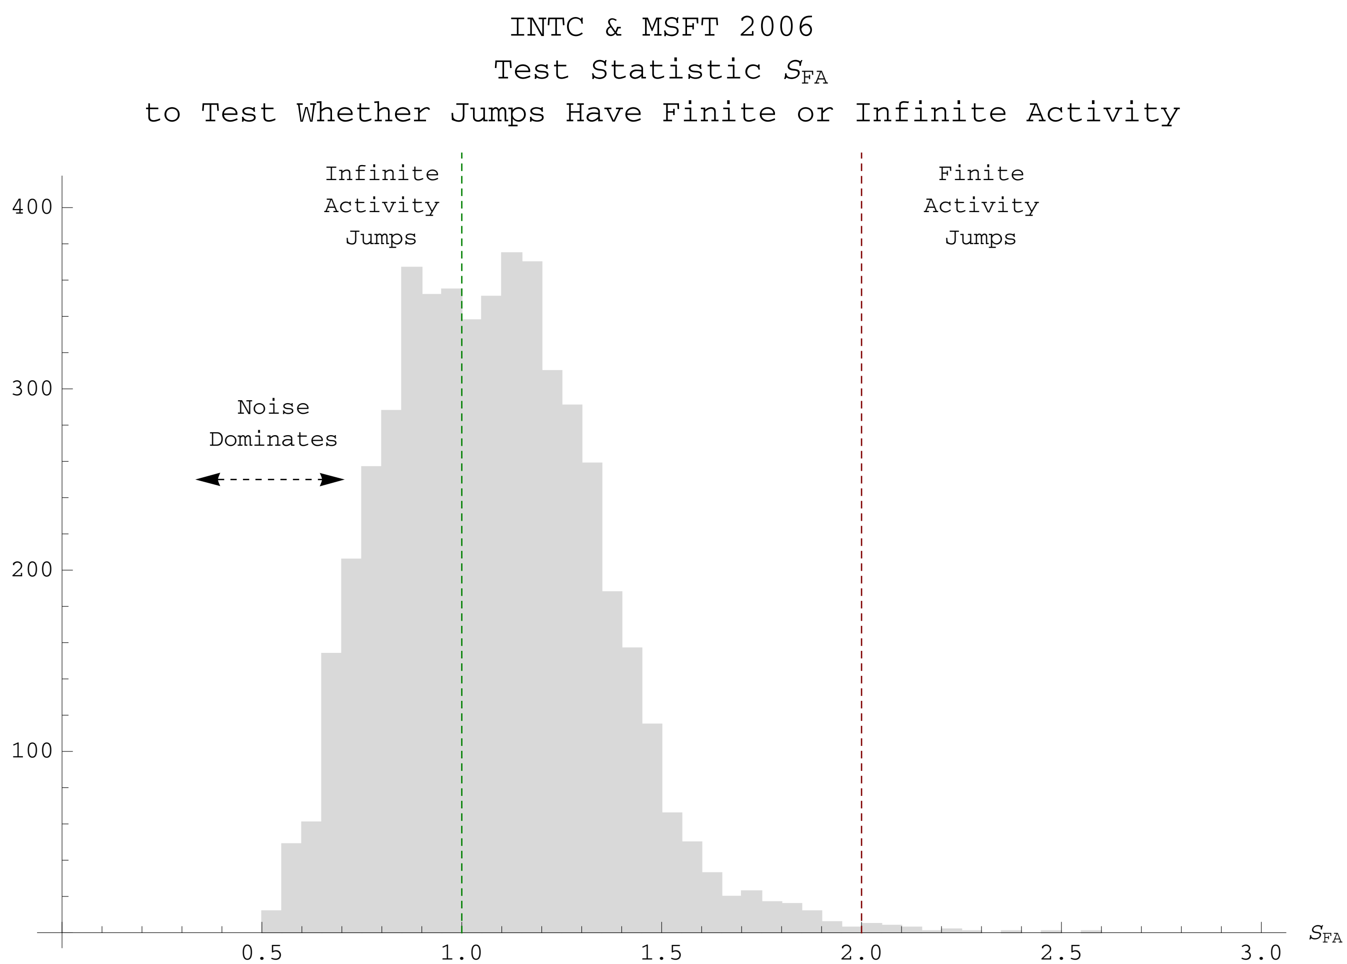

Average of Test Statistic S_{FA}
to Test Whether Jumps Have Finite or Infinite Activity

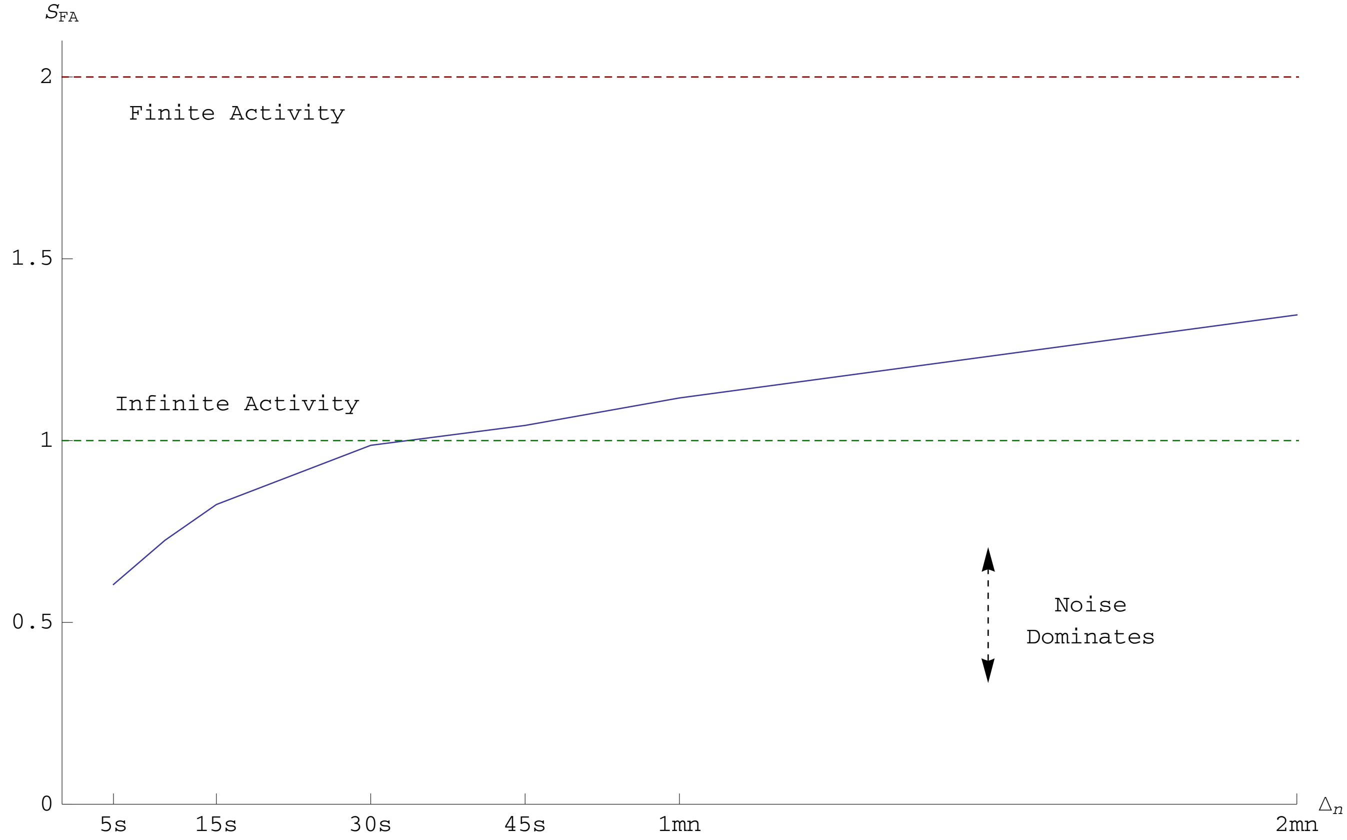

7.4. Brownian Motion: Present or Not

- Market microstructure noise with either an additive white noise or with noise due to rounding, the respective limits of S_W become 2 and $2^{1/2}$ with $k = 2$.

- S_W has four possible limits:

$$\begin{array}{ccc} 1 & : & \text{No Brownian motion} \\ k^{1-p/2} & : & \text{Brownian motion present} \\ k^{1/2} & : & \text{rounding error dominates} \\ k & : & \text{additive noise dominates} \end{array}$$

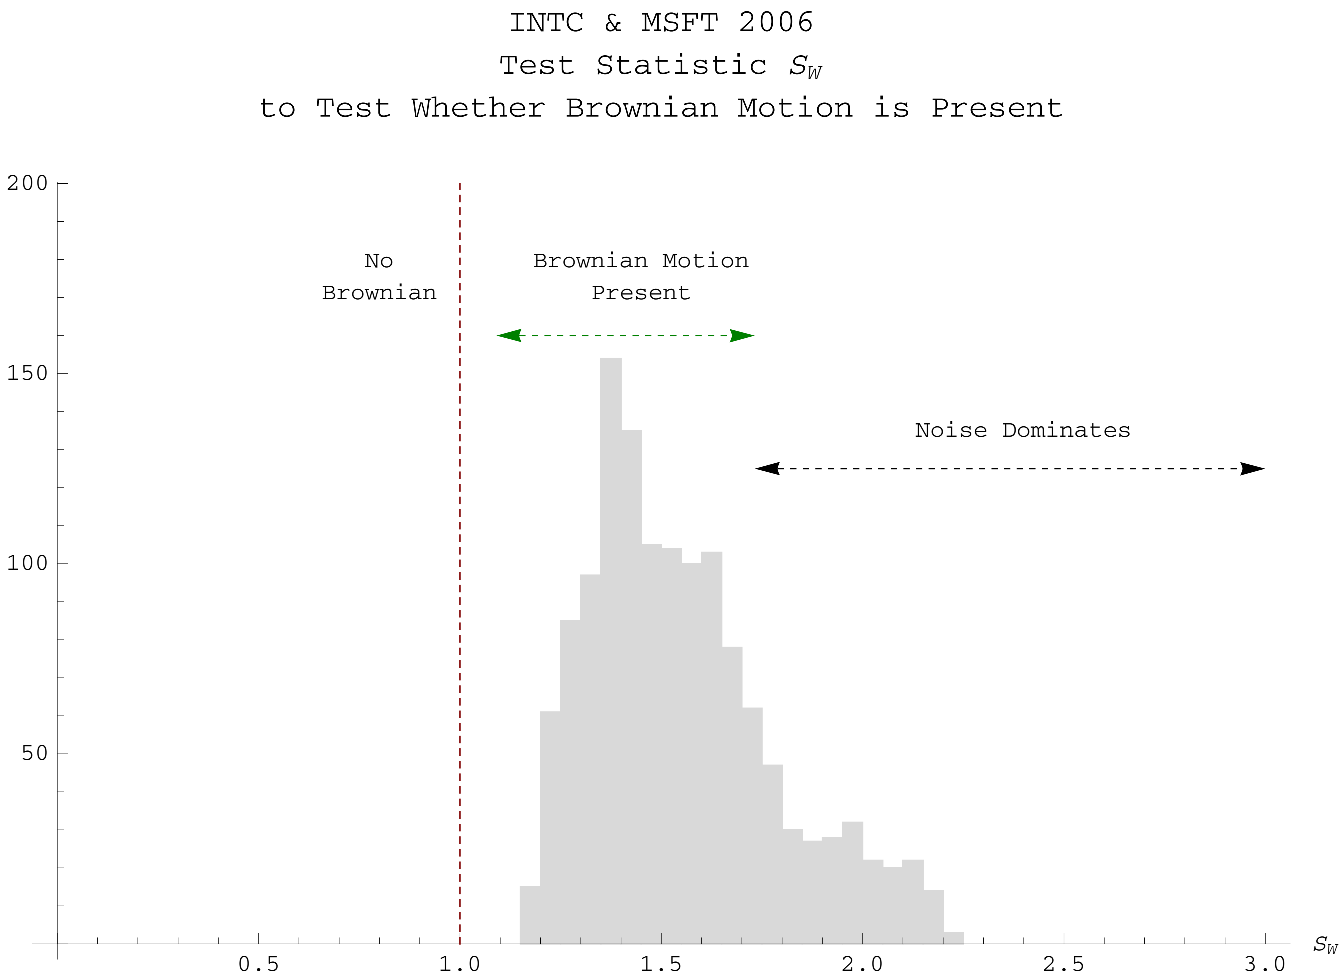
INTC & MSFT 2006
Test Statistic S_W
to Test Whether Brownian Motion is Present
No
Brownian
Brownian Motion
Present
Noise Dominates
200
150
100
50
0.5
1.0
1.5
2.0
2.5
3.0
S_W

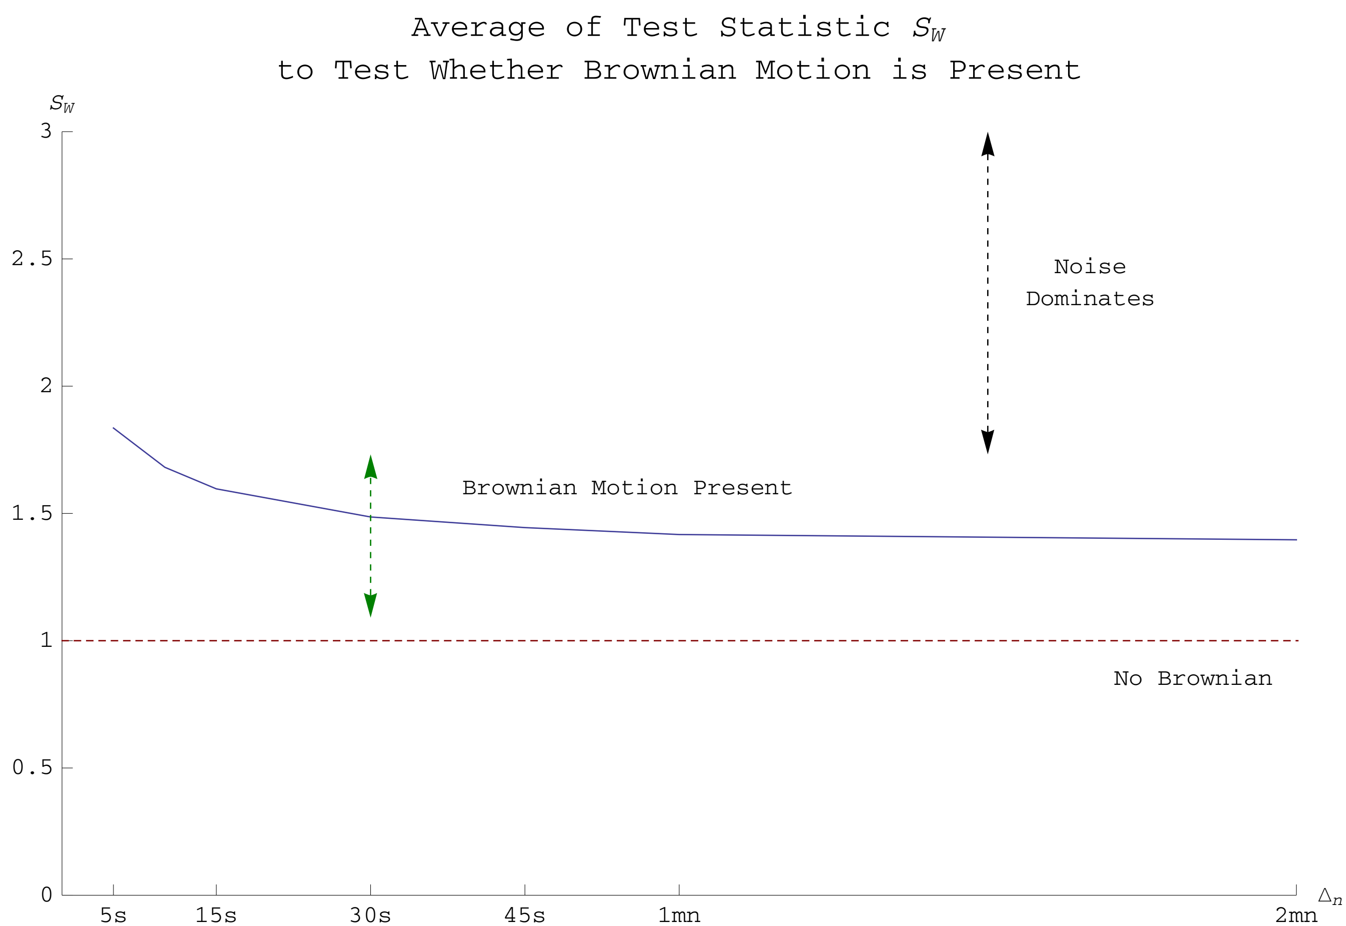
Average of Test Statistic S_W
to Test Whether Brownian Motion is Present
S_W
3
2.5
2
1.5
1
0.5
0
Noise
Dominates
Brownian Motion Present
No Brownian
5s
15s
30s
45s
1mn
2mn
Δ_n

7.5. QV Relative Magnitude

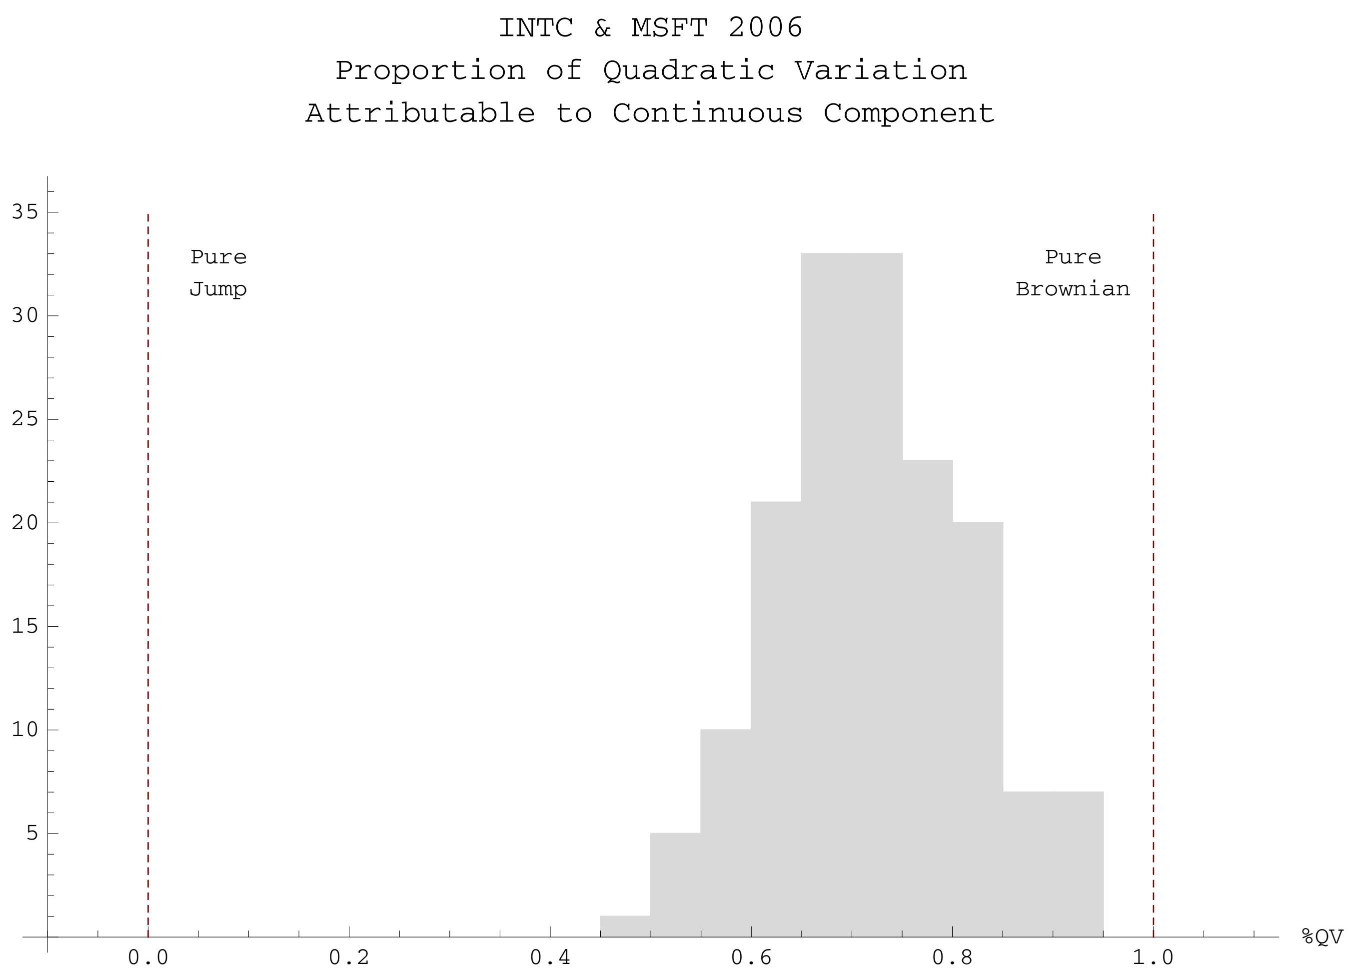

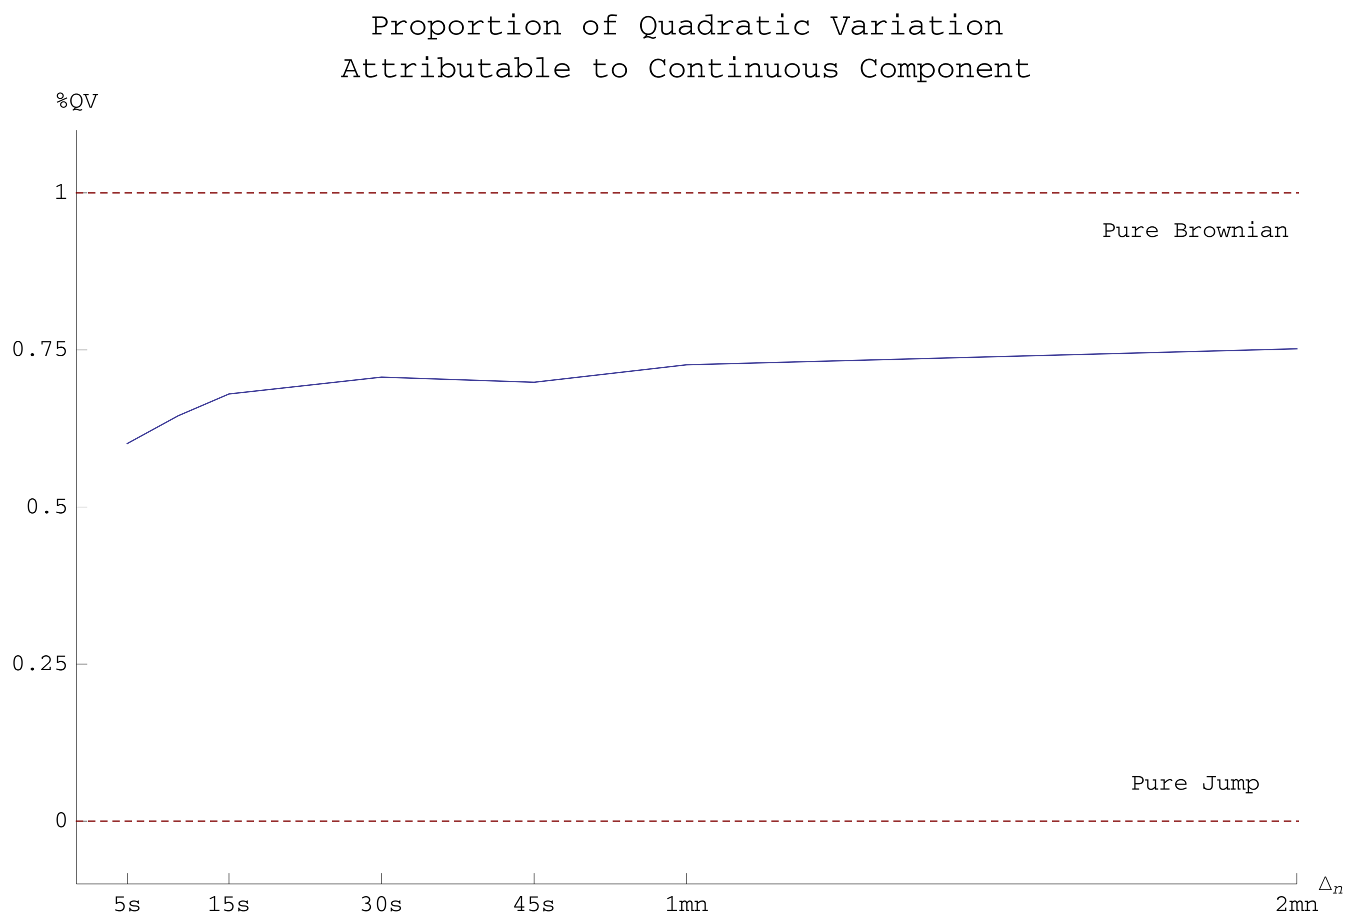
Proportion of Quadratic Variation
Attributable to Continuous Component
%QV
1
0.75
0.5
0.25
0
Pure Brownian
Pure Jump
5s
15s
30s
45s
1mn
2mn
Δ_n

7.6. Estimating Jump Activity

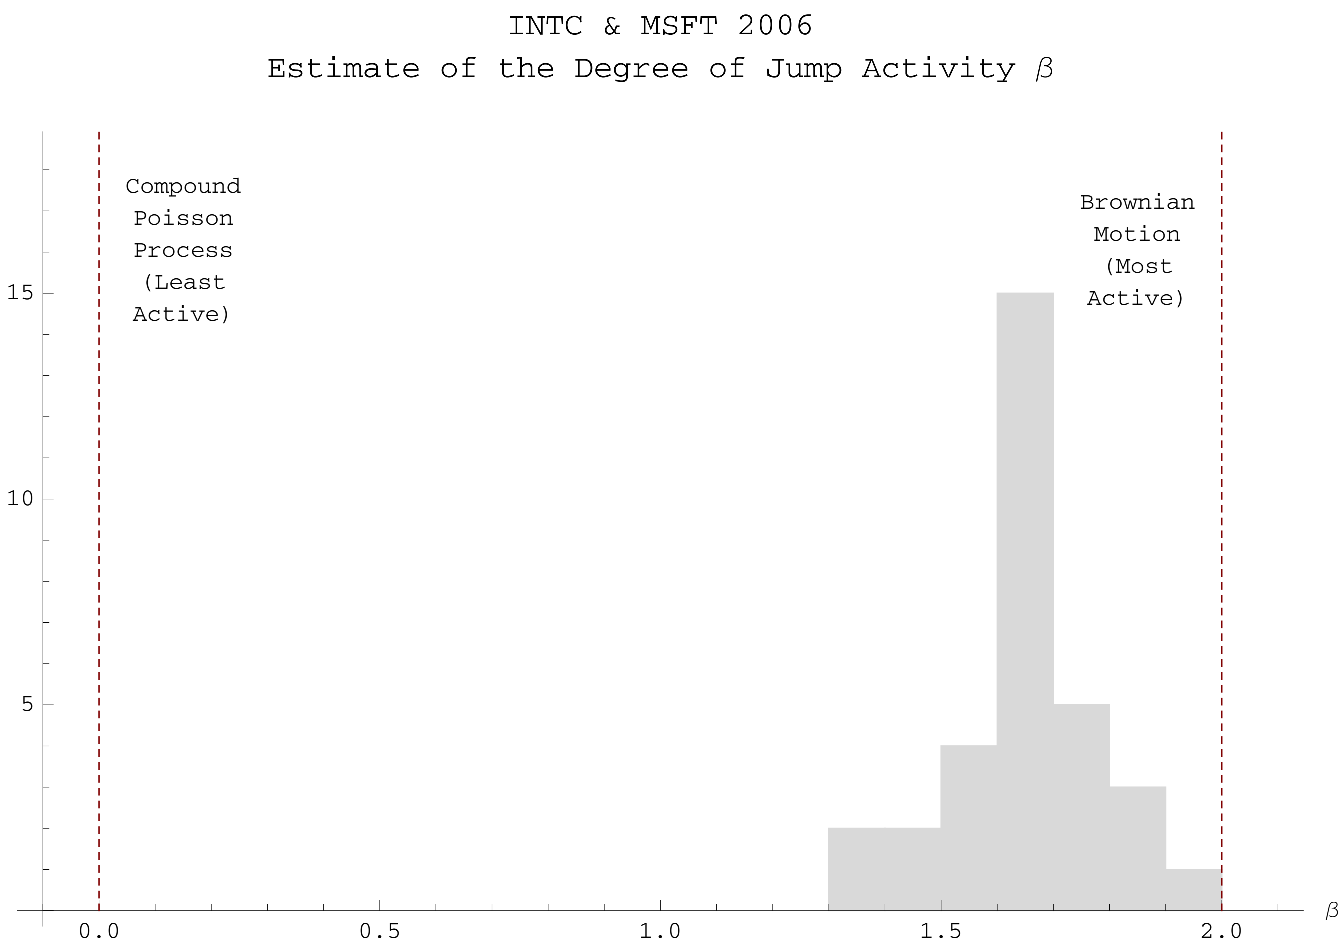

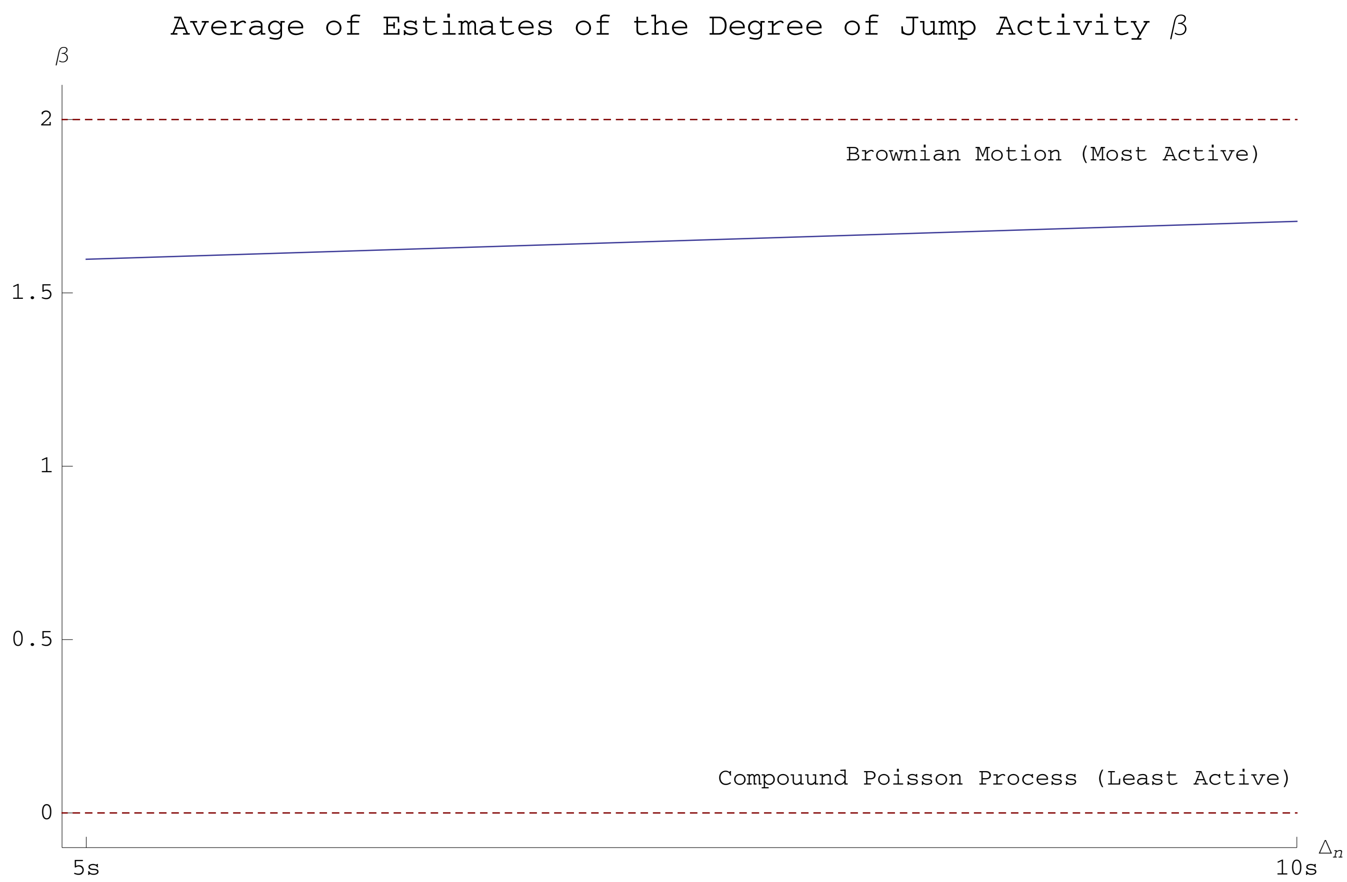
Average of Estimates of the Degree of Jump Activity β
β
2
1.5
1
0.5
0
Brownian Motion (Most Active)
Compouund Poisson Process (Least Active)
5s
10s
Δn

8. Conclusions

The empirical results for these data 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 somewhere around 1.5 or higher.
- Indicate that a continuous component is present.
- Representing about 3/4 of total QV.

- Pros
  - Unified methodology to address all these specification questions in a common framework
  - Symmetric treatment of null and alternative in each case, including distribution theory
  - Model-free
  - Extremely simple to implement
  - Impact of the noise on the statistics is characterized

- Cons
  - Not necessarily the optimal approach for each one of these questions taken individually.
  - Requires high frequency data (particularly the estimation of β)
  - Still to do: a full development of noise-robust statistics.