Journal of Econometrics 168 (2012) 207–222

Contents lists available at SciVerse ScienceDirect

Journal of Econometrics

journal homepage: www.elsevier.com/locate/jeconom

Testing for jumps in noisy high frequency data☆

Yacine Aït-Sahalia[a,b,*], Jean Jacod[c], Jia Li[a]

[a] Department of Economics, Bendheim Center for Finance, Princeton University, Princeton, NJ 08544-1021, United States
[b] NBER, Princeton, NJ 08544-1021, United States
[c] Institut de Mathématiques, de Jussieu, CNRS – UMR 7586, Université Pierre et Marie Curie, Paris 6, 75 013 Paris, France

ARTICLE INFO

Article history:
Received 13 April 2010
Received in revised form
17 November 2011
Accepted 15 December 2011
Available online 22 December 2011

JEL classification:
C22

Keywords:
Semimartingale
Testing for jumps
High frequency data
Market microstructure noise
Pre-averaging

ABSTRACT

This paper proposes a robustification of the test statistic of Aït-Sahalia and Jacod (2009b) for the presence of market microstructure noise in high frequency data, based on the pre-averaging method of Jacod et al. (2010). We show that the robustified statistic restores the test's discriminating power between jumps and no jumps despite the presence of market microstructure noise in the data.

1. Introduction

The recent availability of observations on financial returns at increasingly higher frequencies has prompted the development of methodologies designed to test the specification of suitable models for these data. Motivated both by mathematical tractability and the need to avoid introducing arbitrage opportunities in the model, semimartingales are often employed.

We focus here on testing for the presence of jumps in a discretely observed semimartingale, which has been among the first issue to be considered in the literature. Existing tests for jumps include Aït-Sahalia (2002) (based on the transition function of the process), Carr and Wu (2003) (based on short dated options), Barndorff-Nielsen and Shephard (2004); Huang and Tauchen (2005) and Andersen et al. (2007) (based on bipower variations), Jiang and Oomen (2008) (based on a swap variance), Lee and Mykland (2008) and Lee and Hannig (2010) (based on detecting large increments) and Aït-Sahalia and Jacod (2009b) (based on power variations sampled at different frequencies).

When implemented on high frequency data, as most of them are designed to be, these tests are confronted by the presence of market microstructure noise. Furthermore, that measurement error tends to grow in proportion of the observed increments of the process as the sampling frequency increases, which distinguishes this problem from the classical measurement error problem in statistics.

This issue has received a fair amount of attention in the recent literature, but focused on the base case of quadratic variation estimation. Considering only methods that are robust to the simplest forms of market microstructure noise, there are currently four main approaches to quadratic variation estimation: maximum likelihood estimation (Aït-Sahalia et al., 2005; Xiu, 2010), linear combination of realized volatilities obtained by subsampling (Zhang et al., 2005; Zhang, 2006), linear combination of autocovariances (Barndorff-Nielsen et al., 2008) and pre-averaging (Jacod et al., 2009, 2010). The simplest forms of noise include additive errors and rounding, and combinations thereof. Robust estimators are available as long as the noise is sufficiently "smooth"; a pure rounding error is not. Attempting to generalize the type of noise allowed to an "unsmooth" setting raises a different set of issues that are beyond the scope of this paper (see Li and Mykland (2007) for a discussion).

All the tests developed so far for jumps assume away the presence of noise in high frequency data. In this paper, we examine

☆ Financial support from the NSF under grant SES-0850533 is gratefully acknowledged. We are also grateful to the Editor, Associate Editor and two referees for comments and suggestions.
* Corresponding author at: Department of Economics, Bendheim Center for Finance, Princeton University, Princeton, NJ 08544-1021, United States. Tel.: +1 609 258 4015; fax: +1 609 258 0719.
E-mail addresses: yacine@princeton.edu (Y. Aït-Sahalia), jean.jacod@upmc.fr (J. Jacod), jiali@princeton.edu (J. Li).

0304-4076/$ – see front matter © 2012 Elsevier B.V. All rights reserved.
doi:10.1016/j.jeconom.2011.12.004

the possibility of robustifying one of these tests for jumps, that of Aït-Sahalia and Jacod (2009b), using the pre-averaging method. The test, whose asymptotic properties were originally derived without allowing for the possibility of noise, is based on comparing variations of power greater than 2, at two different frequencies, and taking their ratio. If jumps are present, the two variations converge asymptotically as $\Delta_n \to 0$ to the same limit, which is simply the sum of the pth power of the jumps recorded between 0 and T; as a result their ratio converges to 1. On the other hand, if no jumps are present, the sum of the pth power of the jumps recorded between 0 and T is zero, and both variations then converge to 0. They do so at a rate that depends on the sampling interval Δ_n and so the ratio will pick up the difference between the two sampling frequencies: if the two sampling intervals are Δ_n and $k\Delta_n$, then the limit of the ratio will be $k^{p/2-1}$. Therefore, without noise, the test statistic has two sharply distinct limits depending upon whether jumps are present or not.

In the presence of noise, on the other hand, the theoretical limits of the statistic become respectively $1/k$ and $1/k^{1/2}$ in the two polar cases of additive noise and noise due to rounding error, and become so irrespectively of the presence or absence of jumps (see Aït-Sahalia and Jacod (2009b)). As a result, when either type of noise dominates, the test statistic loses its intended effectiveness at discriminating between presence and absence of jumps.

In this paper, we construct a robustified version of the statistic, using the pre-averaging approach, and show that this robustification restores the ability of the test to discriminate between jumps and no jumps, despite the presence of the noise. The results are nonparametric in nature, are valid for almost unrestricted semimartingales, and allow for a symmetric treatment of the two null hypotheses specifying either presence or absence of jumps.

The paper is organized as follows. Section 2 presents the model's setting and assumptions. Section 3 presents the test statistic, studies its properties when noise is taken into account, describes its robustification by pre-averaging and derives its asymptotic properties after robustification. Sections 4 and 5 report the results of simulations and of an empirical application to high frequency stock returns data. Section 6 concludes, while proofs are in the Appendix.

2. The model

2.1. The underlying process

We consider a one-dimensional underlying process $X = (X_t)_{t\geq 0}$, sampled at regularly spaced discrete times $i\Delta_n$ over a fixed time interval $[0, T]$, with a time lag which asymptotically goes to 0. In typical financial econometrics applications, X represents the logarithm of an asset price. The basic assumption is that X is an *Itô semimartingale* on a filtered space $(\Omega^{(0)}, \mathcal{F}^{(0)}, (\mathcal{F}_t^{(0)}), \mathbb{P}^{(0)})$, which means that it can be written as

$$X_t = X_0 + \int_0^t b_s ds + \int_0^t \sigma_s dW_s + (\delta 1_{\{|\delta|\leq 1\}}) \star (\underline{\mu} - \underline{\nu})_t + (\delta 1_{\{|\delta|>1\}}) \star \underline{\mu}_t, \tag{2.1}$$

where W is a Brownian motion, $\underline{\mu}$ is a Poisson random measure on $\mathbb{R}_+ \times E$ and its compensator is $\underline{\nu}(dt, dz) = dt \otimes \lambda(dz)$ where $(E, \mathcal{E})$ is an auxiliary space and λ is a σ-finite measure (all these are defined on the filtered space above and we refer for example to Jacod and Shiryaev (2003) for all unexplained terms). We further assume:

Assumption 1. (a) The process (b_t) is optional and locally bounded;

(b) The process (σ_t) is càdlàg (i.e., right-continuous with left limits) and adapted;

(c) The function δ is predictable, and there is a bounded function γ in $\mathbb{L}^2(E, \mathcal{E}, \lambda)$ such that the process $\sup_{z\in E}(|\delta(\omega^{(0)}, t, z)| \wedge 1)/\gamma(z)$ is locally bounded;

(d) We have almost surely $\int_0^t \sigma_s^2\, ds > 0$ for all $t > 0$.

In particular, when X is continuous, it has the form

$$X_t = X_0 + \int_0^t b_s ds + \int_0^t \sigma_s dW_s. \tag{2.2}$$

In this case, we will sometimes need a stronger assumption putting some further structure on the stochastic volatility process, namely:

Assumption 2. We have Assumption 1 and σ_t is also an Itô semimartingale which can be written as

$$\sigma_t = \sigma_0 + \int_0^t \tilde{b}_s ds + \int_0^t \tilde{\sigma}_s dW_s + M_t + \sum_{s\leq t} \Delta\sigma_s\, 1_{\{|\Delta\sigma_s|>v\}}, \tag{2.3}$$

where M is a local martingale orthogonal to W and with bounded jumps and $\langle M, M\rangle_t = \int_0^t a_s ds$, and the compensator of $\sum_{s\leq t} 1_{\{|\Delta\sigma_s|>v\}}$ is $\int_0^t a'_s ds$, and where $\tilde{b}_t$, a_t, and a'_t are optional locally bounded processes, whereas the adapted processes b_t and $\tilde{\sigma}_t$ are left-continuous with right limits.

Overall, these assumptions are standard and fairly unrestrictive. They do not significantly restrict the essential aspects of the process, allowing for stochastic volatility, jumps of finite or infinite activity, all manners of dependence between the characteristics of the process, etc. Of course, they do exclude some examples such as fractional Brownian motion or models without a continuous martingale part, given (d) in Assumption 1. We need the latter requirement to avoid degenerate limiting theorems under the null hypothesis where no jumps are present.

2.2. The noise

The main purpose of this paper is to test for the presence of jumps when the process X is observed with an error: instead of X_t we now observe

$$Z_t = X_t + \epsilon_t. \tag{2.4}$$

Of course, the observation error ϵ_t comes into the picture only at those observation times $t = i\Delta_n$, but it is convenient to have it defined for all t. We assume that the observation error is, conditionally on the process X, mean zero and mutually independent. Note however that the ϵ_t's are not necessarily unconditionally independent (the independence is only conditional on X). The assumption we will make on the noise term allows for an additive error of the white noise type, but also for noise involving rounding since the assumptions allow the noise ϵ_t to depend on X_t, or in fact even on the whole past of X up to time t.

Mathematically speaking, this can be formalized as follows: for each $t \geq 0$, we have a transition probability $Q_t(\omega^{(0)}, dz)$ from $(\Omega^{(0)}, \mathcal{F}_t^{(0)})$ into $\mathbb{R}$. The space $\Omega^{(1)} = \mathbb{R}^{[0,\infty)}$ is endowed with the product Borel σ-field $\mathcal{F}^{(1)}$ and the "canonical process" $(\epsilon_t : t \geq 0)$ and the probability $\mathbb{Q}(\omega^{(0)}, d\omega^{(1)})$ which is the product $\otimes_{t\geq 0} Q_t(\omega^{(0)}, \cdot)$. We introduce the filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ and the filtration $(\mathcal{G}_t)$ as follows:

$$\left.\begin{array}{l} \Omega = \Omega^{(0)} \times \Omega^{(1)}, \qquad \mathcal{F} = \mathcal{F}^{(0)} \otimes \mathcal{F}^{(1)}, \\ \mathcal{F}_t = \mathcal{F}_t^{(0)} \otimes \sigma(\epsilon_s : s \in [0, t)), \\ \qquad \mathcal{G}_t = \mathcal{F}^{(0)} \otimes \sigma(\epsilon_s : s \in [0, t)), \\ \mathbb{P}(d\omega^{(0)}, d\omega^{(1)}) = \mathbb{P}^{(0)}(d\omega^{(0)})\mathbb{Q}(\omega^{(0)}, d\omega^{(1)}). \end{array}\right\} \tag{2.5}$$