Journal of Financial Economics ∎ (∎∎∎∎) ∎∎∎–∎∎∎

Contents lists available at ScienceDirect

Journal of Financial Economics

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

Fact or friction: Jumps at ultra high frequency[☆]

Kim Christensen[a,*], Roel C.A. Oomen[b,c], Mark Podolskij[d,a]

[a] CREATES, Aarhus University, Department of Economics and Business, Fuglesangs Allé 4, 8210 Aarhus V, Denmark
[b] Deutsche Bank, 1 Great Winchester Street, London EC2N 2DB, United Kingdom
[c] University of Amsterdam, Department of Quantitative Economics, Valckenierstraat 65-67, 1018 XE Amsterdam, The Netherlands
[d] Aarhus University, Department of Mathematics, Ny Munkegade 118, 8000 Aarhus C, Denmark

ARTICLE INFO

Article history:
Received 14 May 2012
Received in revised form
6 January 2014
Accepted 16 January 2014

JEL classification:
C14
C80

Keywords:
Jump variation
High-frequency data
Microstructure noise
Pre-averaging
Realized variation

ABSTRACT

This paper shows that jumps in financial asset prices are often erroneously identified and are, in fact, rare events accounting for a very small proportion of the total price variation. We apply new econometric techniques to a comprehensive set of ultra high-frequency equity and foreign exchange tick data recorded at millisecond precision, allowing us to examine the price evolution at the individual order level. We show that in both theory and practice, traditional measures of jump variation based on lower-frequency data tend to spuriously assign a burst of volatility to the jump component. As a result, the true price variation coming from jumps is overstated. Our estimates based on tick data suggest that the jump variation is an order of magnitude smaller than typical estimates found in the existing literature.

1. Introduction

A long-standing consensus in the literature on asset pricing is that a realistic dynamic model should incorporate several, if not all, of the following stylized facts: random walk behavior (e.g., Fama, 1965) at a macroscopic level, market microstructure effects (e.g., Niederhoffer and Osborne, 1966) at a microscopic level, stochastic volatility (e.g., Mandelbrot, 1963), leverage (e.g., Black, 1976), and jumps (e.g., Press, 1967). Extensive support for these factors can be found both in the theory of finance and in the abundantly available financial market data. In this paper, we examine the role of the jump component by applying new econometric techniques to a comprehensive set of the finest resolution equity and foreign exchange tick data. The use of individual order-level tick data in the context of jump identification is novel to this paper and necessary to arrive at our main finding that the jump component is substantially smaller than currently thought.

☆ We would like to thank Bjørn Eraker (the referee), David Bates, Tim Bollerslev, Michael Brennan, Alvaro Cartea, Frank Diebold, Dobrislav Dobrev, David Leigh, Richard Louth, Andrew Patton, Kevin Sheppard, George Tauchen, Valeri Voev and the participants at the Fourth CSDA (Computational Statistics & Data Analysis) International Conference on Computational and Financial Econometrics, the Sixth Nordic Econometric Meeting, the Marie Curie High Frequency Research Training Workshop, the SITE workshop on Measuring and Modeling Risk with High Frequency Data, the 65th European Meeting of the Econometric Society, the 2012 North American Winter Meeting of the Econometric Society, and at seminars at CREATES, University of Aarhus, Cambridge University, Deutsche Bank, Queen Mary University of London, London School of Economics, Imperial College London, Duke University, Erasmus University Rotterdam, and Universidad Carlos III for helpful comments and insightful suggestions that greatly improved the paper. We are grateful to ICAP and the CME for giving permission to use the EBS foreign exchange and CME futures market data in this study. Kim Christensen and Mark Podolskij acknowledge financial support from CREATES, which is funded by the Danish National Research Foundation.

* Corresponding author. Tel.: +45 8716 5367.
E-mail addresses: kchristensen@creates.au.dk (K. Christensen), roel.ca.oomen@gmail.com (R.C.A. Oomen), mpodolskij@math.au.dk (M. Podolskij).

http://dx.doi.org/10.1016/j.jfineco.2014.07.007
0304-405X/© 2014 Elsevier B.V. All rights reserved.

Specifically, we show that jumps account for about 1% of total price variability (i.e., quadratic variation) in contrast to accepted estimates from lower-frequency data, which are an order of magnitude larger. Our microscopic view at the tick data provides the intuition for this result: A burst of volatility is often spuriously identified as a jump at the lower frequencies commonly used in the literature. No doubt exists that, in times of stress, asset prices do move sharply over short periods of time, and while occasionally genuine price jumps do occur, we find that more often than not price continuity is preserved even when accompanied with a severe deterioration of liquidity.

The foundations of most asset pricing models can be cast in the class of arbitrage-free Itô semimartingales. These processes are naturally decomposed into a continuous diffusive Brownian component and a discontinuous jump part. The importance of being able to distinguish between these two fundamentally different sources of risk is emphasized in Aït-Sahalia (2004). Specifically, jumps have a profoundly distinct impact on option pricing (e.g., Cox and Ross, 1976; Merton, 1976, Duffie, Pan, and Singleton, 2000), risk management (e.g., Duffie and Pan, 2001; Bakshi and Panayotov, 2010), and asset allocation (e.g., Jarrow and Rosenfeld, 1984; Liu, Longstaff, and Pan, 2003). Empirical work on identifying and modeling the jump component now spans nearly half a century. Table 1 provides a representative but necessarily incomplete overview. Starting with the influential paper of Press (1967), and continuing up to Jorion (1988), a number of papers estimate a (constant volatility) jump-diffusion model and report levels of jump variation (JV, expressed as a fraction of total return variation) in excess of 20%. An important shortcoming of the Press (1967)- or Merton (1976)-style jump-diffusion model is that the jump component is the only mechanism that can account for fat tails of the empirical return distribution so that—in the presence of stochastic volatility—the JV measurements are potentially inflated. From the 1990s onwards, a large body of work considers numerous generalizations of the jump-diffusion model to include one or several stochastic volatility factors as well as state-dependent jump components. Estimation methods for such models are often highly complex and numerically intensive (e.g., Eraker, Johannes, and Polson, 2003) but have the nice feature that they can exploit the information in the spot price (e.g., Andersen, Benzoni, and Lund, 2002), its associated derivative prices (e.g., Bates, 1996), or both (e.g., Pan, 2002). The majority of this literature concentrates on the US large-cap Standard & Poor's (S&P) 500 equity index and typically finds that the JV is around 10–20%. The corresponding figure for foreign exchange rates is comparable and that of Treasury bills and individual stocks is higher still.

The most recent work on jumps has seen a shift away from model-based inference on low-frequency data to model-free inference based on intraday data. In an influential series of papers, Barndorff-Nielsen and Shephard (2004, 2006) and Barndorff-Nielsen, Shephard, and Winkel (2006b) introduce the concept of (bi-) power variation—a simple but effective technique to identify and measure the variation of jumps from intraday data (see Aït-Sahalia and Jacod, 2009a,b and Mancini, 2004, 2009 for a related jump-robust threshold estimator). Using this, or variations thereof, a number of recent articles report model-free JV estimates of around 10% for the S&P 500 index and thus reinforce the earlier literature that the jump component is important (e.g., Huang and Tauchen, 2005; Andersen, Bollerslev, and Diebold, 2007).

By their nature, jumps are instantaneous and discrete moves in the price. Therefore, not surprisingly, identification is aided by the finest resolution price data. This is precisely what motivates the recent literature to use intraday data. However, the consensus five-minute frequency at which returns are typically sampled, instead of at the finest tick-by-tick resolution, reflects a compromise to ensure that the market microstructure effects are sufficiently benign for the theory to remain valid. Fig. 1 illustrates the cost of doing so: a diminished ability to distinguish jumps from bursts in volatility. From Panel A, in which returns are sampled at a five-minute frequency, one would probably conclude that the notorious flash-crash episode contains a number of very large jumps (see, e.g., Easley, de Prado, and O'Hara, 2011 for a discussion of the event). The widely used bi-power variation (BV) jump measure is highly significant indicating the presence of jumps on formal statistical grounds. Yet, when focusing on the relevant subperiod in Panel B, one can see that at tick frequency jumps are elusive. In fact, in a recent study using audit trail data, Kirilenko, Kyle, Samadi, and Tuzun (2011) characterize the flash-crash as a "brief period of extreme market volatility." The March 2011 earthquake in Japan led to similar scenarios in the US dollar–Japanese yen exchange rate, which experienced a flash-crash-type episode on March 16 and a rapid depreciation following a coordinated intervention by the Bank of Japan and other G7 central banks on March 18. Again, both events would be classified as exhibiting large jumps by conventional realized measures based on five-minute data, while the tick data reveal a period of heightened volatility instead of discrete price jumps (see Fig. 6).

This paper provides a comprehensive study into the magnitude of the jump component based on the highest resolution data available. To the best of our knowledge, we are the first to do so, and we find that the tick data have a fundamentally different story to tell. Our analysis employs the pre-averaging theory of Podolskij and Vetter (2009a) and Jacod, Li, Mykland, Podolskij, and Vetter (2009) to construct noise-robust jump measures from tick data. We apply these to a representative data set composed of US large-cap stocks (the 30 Dow Jones Industrial Average (DJIA) constituents), equity indexes (the S&P 500 and Nasdaq 100), and currency pairs (the Euro–US dollar (EURUSD), US dollar–Japanese yen (USDJPY), and US dollar–Swiss franc (USDCHF)). We start by confirming that when sampling at low frequencies the jump component appears substantial and in line with the extant literature at around 10%. In passing, we observe a near-perfect relation in which the jump component decreases in magnitude as the sampling frequency is increased from 15 minutes to five minutes. Next, we obtain our main result, in which most of the previously identified jumps vanish as we move to the tick frequency and we are left with highly volatile episodes instead. The overall JV measured across all instruments we consider is just over 1%.