Finance Research Letters 22 (2017) 42–48

Contents lists available at ScienceDirect

Finance Research Letters

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

Identifying events in financial time series – A new approach with bipower variation

György Andor, András Bohák*

Department of Finance, Budapest University of Technology and Economics, 1117 Budapest, Magyar tudosok korutja 2, Hungary

ARTICLE INFO

Article history:
Received 13 October 2016
Accepted 3 November 2016
Available online 4 November 2016

Keywords:
Jump detection
Event detection
Realized volatility
Bipower variation

ABSTRACT

We present a statistical test to identify significant events in financial price time series. In contrast to “jumps,” we define “events” as non-instantaneous, but nevertheless unusually fast and large, price changes. We show that non-parametric tests perform badly in detecting events so defined. We propose a new approach to explore the dependence of jump detection statistics on the sampling method used and find that our method improves the event detection rate of the standard test by a factor of three.

1. Introduction

The detection of discontinuities or jumps in the evolution of various financial variables (asset prices, interest rates, etc.) has attracted substantial interest in recent econometric literature. The reason behind the interest in ex-post identification of jumps is driven by multiple research fields. The first tries to extend the standard, purely continuous semi-martingale models with jump processes to better understand volatility and potentially improve the accuracy of option pricing (Aït-Sahalia, 2004) and realized volatility forecasting (Xingguo et al., 2016). The second area of research is using event studies as a tool for analyzing a broad range of financial phenomena like short-sale constraints (Yeh and Chen, 2014), informed trades around possible private information events (Ormos and Timotity, 2016) and intraday liquidity dynamics (Boudt and Petitjean, 2014). (De Moor and Sercu, 2015) show that extreme observations can bias the average return calculation and this bias affects small stocks more in CAPM alpha estimation.

All jump detection algorithms belong to one of two categories: parametric methods (e.g. Zawadowski et al. (2004)), which use user-defined parameters to define what qualifies as an “event” and non-parametric methods, where parameter definition is not necessary (e.g. the bipower variation (BPV) test proposed in Barndorff-Nielsen and Shephard (2006)). Parametric methods are not optimal for event studies, since the set of events identified and thus the results of the whole analysis are frequently very sensitive to the input parameters and it is hard to justify any particular parametrization.

In our paper, we extend the standard definition of “jumps” to include “events”: unusually large price changes that occur unusually quickly, but not in an instant. To illustrate this point, in Fig. 1, we plot the price (sampled every minute) of Wal-Mart Stores Inc. (WMT) on the three business days from the 21st to 23rd of October 2006.

Identifying price jumps can have many objectives, but whatever the purpose, the event seen in Fig. 1 should definitely be detected. This jump, however, is clearly not instantaneous. This observation leads to a very loose definition of the term event in contrast to the precise mathematical definition of jumps. We call an event any window of several minutes (even up

* Corresponding author. Fax: +36 1 463 2745.
E-mail addresses: andor@finance.bme.hu (G. Andor), bohak@finance.bme.hu (A. Bohák).

http://dx.doi.org/10.1016/j.frl.2016.11.003
1544-6123/

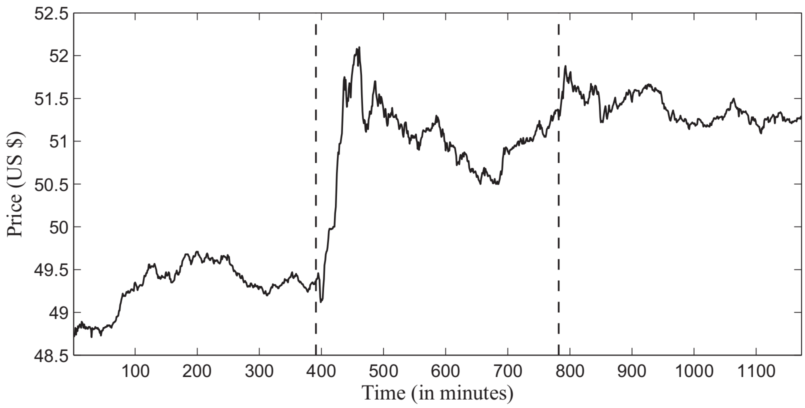

Fig. 1. Price of WMT stock on three business days between the 21st and 23rd of October 2006, clearly showing an event at the beginning of the second trading day. Black dashed lines signal day boundaries.

to an hour) where the change in price is "unusually high". We believe that this extension is crucial for event studies, since in the real world, even very fast price changes have some time duration. The main contribution of this paper is to modify non-parametric tests designed to identify jumps in a way which significantly improves their performance in identifying non-instantaneous events.

The main challenge we had to overcome was to prove that our proposed test is indeed superior to the ones already used in the literature. An obvious methodology to do this is a Monte Carlo analysis, where the time and size of the artificial price change are known and we can test if this jump is found by the algorithm or not (see Huang and Tauchen (2005) and the references therein). In one of the sections below, we show that our test outperforms the classic BPV test in terms of both size and power.

While Monte Carlo analysis is abundant in the literature, true empirical studies are limited at best and non-existent at worst. The problem with real data is that we do not know if an event occurred on a particular day, so to assess the power of an event detection algorithm, we have to find a benchmark method to compare to. However, if we accept that the benchmark method is correct, there is no point to test any algorithm against it: we could just use the benchmark method to identify events.

In our empirical analysis, we decided to start with an extensive dataset of almost a million trading days and use a parametric method with extreme parameters to define our benchmark event set. We do not believe that this benchmark method is ideal – or even acceptable: the parameters are chosen in a way that most of what we would like to call an event is probably not picked up by the algorithm. The events identified, however, are almost surely significant enough that any event detection algorithm should identify them (just like the one in Fig. 1). We show that our jump detection test significantly outperforms the original in this setup.

2. Theoretical framework and application problems

We use the framework of Barndorff-Nielsen and Shephard (2006) for our new event identification methodology. Their original test is based on the difference between the realized variance and the realized bipower variance of the returns within a trading day. The idea behind the test is that realized variance is increased by the jump component of the price movements, while bipower variation is not. The days with a significant jump component can thus be identified with a test statistic based on the appropriately scaled observed difference between the two. The family of test statistics proposed in the original article was extensively analyzed by Huang and Tauchen (2005). We use only the best performing member of the family, calculating a Z-score based on the realized variance (RV), the realized BPV and the realized tripower quarticity ($RTPQ$):

$$Z_t = \frac{(RV_t - BPV_t)/RV_t}{\sqrt{c\frac{1}{M}\max\left(1, \frac{RTPQ_t}{BPV_t^2}\right)}}$$

where c is a scaling constant.

The test statistic Z_t follows $N(0;1)$ in the case where there is no jump in the interval $[0,t]$. In the following empirical analysis, we will calculate the test statistic for each trading day of each stock separately. We refer to this method as the bipower variation test, or BPV test and we show that the choice of sampling frequency has a dramatic effect on its results.

The upper panel of Fig. 2 shows the price evolution of an event, constructed using a simple driftless geometric Brownian motion (GBM) with one-minute $\sigma = 0.11\%$ for one trading day (390 minutes) plus a non-instantaneous jump, i.e., a significant increase in the price from minute 189 to minute 199 at a rate of 0.44%/min. This could correspond to a surprising