Journal of Econometrics 177 (2013) 320–342

Contents lists available at ScienceDirect

Journal of Econometrics

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

A Markov-switching multifractal inter-trade duration model, with application to US equities

Fei Chen [a], Francis X. Diebold [b,c,*], Frank Schorfheide [b,c]

[a] *Huazhong University of Science and Technology, China*
[b] *University of Pennsylvania, United States*
[c] *NBER, United States*

ARTICLE INFO

Article history:
Available online 17 April 2013

JEL classification:
C41
C22
G1

Keywords:
High-frequency trading data
Point process
Long memory
Time deformation
Regime-switching model
Market microstructure
Liquidity

ABSTRACT

We propose and illustrate a Markov-switching multifractal duration (MSMD) model for analysis of inter-trade durations in financial markets. We establish several of its key properties with emphasis on high persistence and long memory. Empirical exploration suggests MSMD's superiority relative to leading competitors.

1. Introduction

It is a pleasure to help honor Hashem Pesaran with this paper, which fits well in the Pesaran tradition of dynamic econometric modeling in the presence of instabilities. In our case the dynamic econometric modeling focuses on inter-trade durations in financial markets (i.e., the times between trades), and the instabilities are fluctuations in conditionally expected trade arrival intensities, which are driven by regime-switching components. The general context, moreover, is financial econometrics, which has concerned Hashem more in recent years, and it involves extracting information from "Big Data", a hallmark of much of Hashem's work ranging from micro (e.g., Pesaran, 2006) to macro (e.g., Garratt et al., 2012).

Indeed the necessity of grappling with Big Data, and the desirability of unlocking the information hidden within it, is a key modern development – arguably *the* key modern development – in all the sciences.[1] Time-series econometrics, particularly time-series financial econometrics, is no exception. Big Data in financial econometrics has both cross-sectional and time-series aspects. In the cross-sectional dimension it arises from the literally hundreds of thousands of assets that trade in global financial markets. In the time-series dimension it arises from the similarly huge number of trades, often many per second, that are now routinely executed for financial assets in liquid markets. In this paper we are concerned in general with the time-series dimension, and in particular with the durations associated with inter-trade arrivals. Those trades are facilitated by modern hardware, software and algorithms, and they are continuously recorded electronically. The results are vast quantities of high-frequency (trade-by-trade) price data.

One might reasonably ask what, precisely, financial econometricians hope to *learn* from high-frequency trading data. Interestingly, such high-frequency data emerge as largely uninformative for some objects of interest (e.g., trend, or "drift", in log price), but highly informative for others (e.g., volatility), an insight traces at least to Merton (1980). In particular, although precise estimation of trend benefits greatly from a long calendar span but not from high-frequency sampling, precise estimation of volatility benefits immensely from high-frequency sampling.[2] Accurate volatility

* Corresponding author at: University of Pennsylvania, United States.
E-mail address: fdiebold@sas.upenn.edu (F.X. Diebold).

1 For background, see Diebold (2012).

2 Indeed the mathematical foundation of the modern financial econometrics "realized volatility" literature initiated by Andersen et al. (2001) and Barndorff-Nielsen and Shephard (2002) is precisely the convergence of empirical quadratic variation to population quadratic variation as sampling frequency increases.

0304-4076/$ – see front matter © 2013 Elsevier B.V. All rights reserved.
http://dx.doi.org/10.1016/j.jeconom.2013.04.016

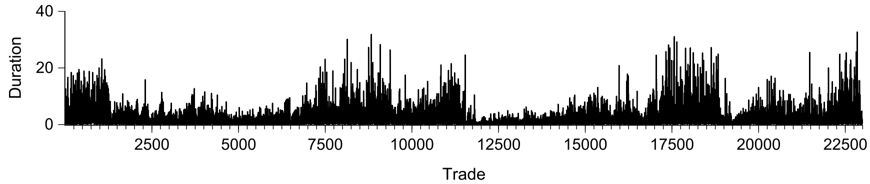

Fig. 1. *Citigroup duration time series*. We show a time-series plot of inter-trade durations between 10:00 am and 4:00 pm during February 1993, measured in seconds and adjusted for calendar effects. As indicated by the horizontal axis labeling, there were approximately 23,000 trades, and hence approximately 23,000 inter-trade durations.

estimation and forecasting, in turn, are crucial for financial risk management, asset pricing and portfolio allocation.

In this paper we are not directly interested in volatility; rather, as mentioned above, we are interested in inter-trade *durations*. However, high-frequency data is informative not only for the properties of volatility, but also for the properties of inter-trade durations. At one level that observation is trivial, as one obviously needs trade-by-trade data to infer properties of inter-trade durations. But at another level the observation is quite deep, linked to the insight that, via time-deformation arguments in the tradition of Clark (1973), properties of calendar-time volatility and transactions-time trade arrivals should be intimately related.

In particular, the time-deformation perspective suggests that serial correlation in calendar-time volatility is driven by serial correlation in calendar-time trade counts (i.e., the number of trades per unit of calendar time, such as an hour or a day). But serial correlation in calendar-time trade counts is driven by serial correlation in transactions-time trade-arrival intensity, which is ultimately driven by serial correlation in the information flow that drives trading. Hence the calendar-time "realized volatility" persistence revealed by high-frequency data should have parallels in calendar-time trade count persistence and transaction-time inter-trade duration persistence. Interestingly, the key and robust finding for realized volatility is not only high persistence, but *tremendously* high persistence, in the form of long memory.[3] Hence one suspects that long memory should be operative as well in trade-arrival intensity, yet the duration literature has not featured long memory prominently.

Against this background, we propose and evaluate a new model of inter-trade durations, closely-linked to the pioneering "multifractal" return volatility model of Mandelbrot et al. (1997), Calvet and Fisher (2001) and Calvet and Fisher (2004), as extended and surveyed by Calvet and Fisher (2008) and Calvet and Fisher (2012). As we will discuss subsequently in detail, our model's construction and implications are quite different from existing dynamic duration models, most notably the prominent autoregressive conditional duration (ACD) model of Engle and Russell (1998) and variants thereof. Our model, which we call the Markov-switching multifractal duration (MSMD) model, captures high persistence in duration clustering; indeed it intrinsically captures long memory while nevertheless maintaining covariance stationarity. It also captures additional important features of observed durations, such as over-dispersion. Finally, as we will also demonstrate, it is highly successful empirically.[4]

We proceed as follows. In Section 2, we review the important empirical regularities that routinely emerge in inter-trade durations in financial markets. In Section 3 we develop the MSMD model, and we characterize its properties with emphasis on its consistency with the empirical regularities. In Section 4 we contrast MSMD to various competing models. In Section 5 we apply MSMD to inter-trade duration data for a representative set of US equities, evaluating its performance in both absolute and relative terms. We conclude in Section 6.

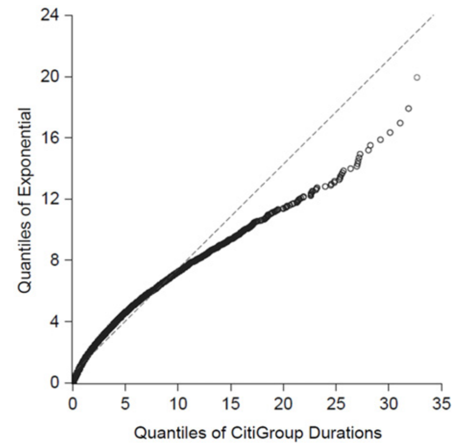

Fig. 2. *Citigroup duration distribution*. We show an exponential QQ plot for Citigroup inter-trade durations between 10:00 am and 4:00 pm during February 1993, measured in seconds and adjusted for calendar effects.

2. Empirical regularities in inter-trade duration data

Here we highlight three important properties of financial market inter-trade durations, illustrating them for a particular US equity (Citigroup). We shall be interested subsequently in econometric duration models that capture those properties.

The first property of inter-trade durations is serial correlation in duration dynamics. Durations are highly persistent, as is clear visually from the time-series plot of Citigroup durations in Fig. 1.[5,6] The second property of inter-trade durations, related to their high persistence, is called over-dispersion. Over-dispersion refers to the standard deviation exceeding the mean, in contrast to the equality that would obtain if durations were exponentially distributed. The exponential is an important benchmark for duration modeling, because i.i.d. exponential durations arise from time-homogeneous Poisson processes, as we discuss in greater detail subsequently in Section 3.1. In Fig. 2 we show an exponential duration Q–Q plot for the Citigroup durations, which clearly indicates a non-exponential duration distribution characterized by a longer right tail than

3 Findings of long memory in realized volatility run consistently from the early work of Andersen et al. (2001) through scores of subsequent studies, as surveyed for example in Andersen et al. (2013).

4 After the first version of this paper was written, we learned of contemporaneous and independent parallel work by Baruník et al. (2012), who use MSMD to study inter-trade durations in foreign exchange futures markets. Their interesting work complements ours and buttresses our claim that the MSMD model shows great empirical promise.

5 We have adjusted the durations for calendar effects and will provide details subsequently in Section 5.1.

6 The transactions-time duration clustering is matched by a corresponding calendar-time transactions-count clustering (i.e., clustering in the number of trades per unit of calendar time, such as an hour or a day), but to conserve space we do not pursue it here. There are, however, many interesting possibilities, including empirical assessment of moment-scaling laws of the form $E[(d_{i+1}+\cdots+d_{i+n})^q] = C_q n^{\tau(q)+1}$. See, for example, Calvet and Fisher (2012).