



Frequency of observation and the estimation of integrated volatility in deep and liquid financial markets[☆]

Alain P. Chaboud^a, Benjamin Chiquoine^b, Erik Hjalmarsson^a, Mico Loretan^{c,*}

^a Division of International Finance, Federal Reserve Board, Washington DC 20551, USA

^b Investment Fund for Foundations, Cambridge MA 02138, USA

^c IMF Institute, International Monetary Fund, Washington DC 20431, USA

ARTICLE INFO

Available online 20 September 2009

JEL classification:

C22
F31
G12

Keywords:

Realized volatility
Integrated volatility
Critical sampling frequency
Market microstructure noise
Government bond markets
Foreign exchange markets
Liquidity
Kernel estimator
Robust estimator
Jumps

ABSTRACT

Using two newly available ultrahigh-frequency datasets, we investigate empirically how frequently one can sample certain foreign exchange and U.S. Treasury security returns without contaminating estimates of their integrated volatility with market microstructure noise. Using the standard realized volatility estimator, we find that one can sample dollar/euro returns as frequently as once every 15 to 20s without contaminating estimates of integrated volatility; 10-year Treasury note returns may be sampled as frequently as once every 2 to 3 min on days without U.S. macroeconomic announcements, and as frequently as once every 40s on announcement days. Using a simple realized kernel estimator, this sampling frequency can be increased to once every 2 to 5 s for dollar/euro returns and to about once every 30 to 40s for T-note returns. These sampling frequencies, especially in the case of dollar/euro returns, are much higher than those that are generally recommended in the empirical literature on realized volatility in equity markets. The higher sampling frequencies for dollar/euro and T-note returns likely reflect the superior depth and liquidity of these markets.

© 2009 Elsevier B.V. All rights reserved.

1. Introduction

Estimating the volatility of asset returns is important for many economic and financial applications, including risk management, derivative pricing, and analyzing investment choices and policy alternatives. As Mandelbrot (1963, p. 418) noted, volatility estimation is complicated by the fact that “large [price] changes tend to be followed by large changes—of either sign—and small changes tend to be followed by small changes,” i.e., that volatility tends to cluster. One approach to estimating volatility is to use a parametric framework, such as the class of ARCH, GARCH, and stochastic volatility models. If data on returns are available at sufficiently high frequencies, one can also estimate volatility nonparametrically by computing the realized volatility, which is the natural estimator of the ex post integrated volatility. This nonparametric method is appealing both because it is computationally simple and because it is a valid estimator under fairly mild statistical assumptions.

The higher the sampling frequency and thus the larger the sample size of intraday returns, the more precise the estimates of daily integrated volatility should become. In practice, however, the presence of so-called market microstructure features, which

[☆] The initial version of this paper was written while Chiquoine and Loretan were employed in the Division of International Finance of the Federal Reserve Board. The views expressed in this paper are solely the responsibility of the authors and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System, of any other person associated with the Federal Reserve System, or of any persons associated with the International Monetary Fund.

* Corresponding author. INSAN HQ2-04-007, International Monetary Fund, 700 19th Street NW, Washington DC 20431, USA. Tel.: +1 202 623 8784; fax: +1 202 589 8784. E-mail addresses: alain.p.chabound@frb.gov (A.P. Chaboud), bchiquoine@tiff.org (B. Chiquoine), erik.hjalmarsson@frb.gov (E. Hjalmarsson), mloretan@imf.org (M. Loretan).

arise especially if the data are sampled at very high frequencies, creates important complications. The finance literature has identified many such features. Among them are the facts that financial transactions—and hence price changes and non-zero returns—arrive discretely rather than continuously over time, that buyers and sellers usually face different prices (separated by the bid-ask spread), that returns to successive transactions tend to be negatively serially correlated (due to, for instance, the so-called bid-ask bounce), and that the initial impact of trades on prices is often at least partially reversed.¹

The first aim of our paper is to study, for two specific financial assets, how the standard estimator of integrated volatility is affected by the choice of sampling frequency and, as a result, by the bias caused by market microstructure features. The two asset price series we study are obtained from some of the deepest and most liquid financial markets in existence today. They are the spot exchange rate of the dollar/euro currency pair, provided by Electronic Broking Systems (EBS), and the price of the on-the-run 10-year U.S. Treasury note, which is traded on BrokerTec. Both of these markets are electronic order book systems, which quite likely represent the future of wholesale financial trading systems. Both markets are strictly inter-dealer. These markets are far larger in terms of total trading volume than markets for individual stocks, even the handful of most liquid stocks traded on the New York Stock Exchange, and bid-ask spreads in these markets are narrower than in typical stock markets. In 2005, the time period considered in this paper, bid-ask spreads averaged 1.04 basis points for dollar/euro spot transactions on EBS and 1.68 basis points for 10-year Treasury note transactions on BrokerTec. Prices for both time series are available at ultrahigh sampling frequencies—up to the second-by-second frequency.

Our main hypothesis is that in such deep and liquid markets, microstructure-induced noise should pose less of a concern for volatility estimation, in the sense that it should be possible to sample returns more frequently than, say, returns on individual stocks before estimates of integrated volatility encounter significant bias caused by the markets' microstructure features. We label this sampling frequency (provided, of course, that it exists) the *critical sampling frequency*. This thesis is indeed borne out by our empirical results. Using volatility signature plots, we find that the critical sampling interval lengths for dollar/euro returns are as short as 15 to 20 s. The corresponding critical sampling interval lengths for returns on 10-year Treasury notes are between 2 and 3 min. These intervals are considerably shorter than the sampling intervals of several minutes—usually five or more minutes—that have often been recommended in the empirical literature on estimating integrated volatility for a number of other financial markets. The shorter critical sampling intervals and the associated larger sample sizes afford a considerable gain in the precision with which the integrated volatility of returns may be estimated. We conclude that in very deep and liquid markets, microstructure-induced frictions may be much less of an issue for integrated volatility estimation than was previously thought.

We also analyze whether the presence or absence of scheduled U.S. macroeconomic news announcements influences the precision with which the integrated volatility of asset returns may be estimated. While confirming the results of several previous empirical studies that integrated volatility is systematically higher on announcement days than on non-announcement days, we find that the critical sampling frequency is also systematically higher on announcement days. We interpret this finding as an indication that the higher trading volumes that occur on announcement days, an especially prominent feature in the U.S. Treasury bills and notes markets, help reduce some of the frictions caused by market microstructure features, raising the critical sampling frequencies and hence allowing greater estimation precision.

Although the critical sampling frequencies are already very high for both time series we consider in this paper, we find that it is possible to further increase these critical sampling frequencies by using so-called kernel estimators, which are designed explicitly to control for the effects of market microstructure noise. We find that by using a very simple version of a kernel estimator, it is possible to sample dollar/euro returns at frequencies as high as once every 2 to 5 s, and that T-note returns can be sampled as frequently as once every 30 to 40 s without incurring noticeable bias generated by market microstructure noise. This kernel estimator, which is almost as easy to compute as the standard realized volatility estimator, therefore offers substantial additional gains in terms of both how frequently one can sample on an intraday basis and the accuracy with which integrated volatility may be estimated.

Finally, we also examine how certain robust estimators of integrated volatility perform for the two time series at hand. These alternative estimators are not based on functions of the standard quadratic variation process, but instead on functions of absolute variation and bipower variation processes. A reason for considering such methods is that they are, by construction, more robust than the standard estimator to outlier activity (heavy tails) in the data; such “outliers” are frequently generated by discontinuities or jumps in the time series of financial asset prices. In general, these estimators measure somewhat different (but highly relevant) aspects of daily variation than does the standard realized volatility estimator. We find empirically that these alternative methods are indeed more robust than the standard estimator to the presence of jumps. For instance, the volatility estimates show less dispersion across announcement and non-announcement days than do estimates that are based on squared variation. However, we find no evidence that these robust methods are also less sensitive than the standard estimator to bias imparted by market microstructure noise. To the contrary, our results indicate that one should typically sample *less frequently* when using the absolute variation base estimator, relative to the critical sampling frequency we found for the standard volatility estimator.

The remainder of our paper is organized as follows. Section 2 provides some motivation for the use of the standard estimator of integrated volatility, which is based on the quadratic variation of returns. The section also details how market microstructure noise may cause bias in the standard estimator, provides an introduction to kernel-based estimators designed to circumvent this problem, and sets out the use of estimators based on absolute and bipower variation processes. Section 3 provides an overview of the characteristics of the foreign exchange (FX) and bond market data used in our empirical work. Section 4 provides the empirical

¹ For an overview of many of these market microstructure issues and their importance for financial theory and practice, we refer the reader to Hasbrouck (2006), O'Hara (1995), Campbell et al. (1997, ch.3), as well as to Roll (1984), Harris (1990, 1991), and Hasbrouck (1991).

متن کامل مقاله

دریافت فوری ←

ISIArticles

مرجع مقالات تخصصی ایران

- ✓ امکان دانلود نسخه تمام متن مقالات انگلیسی
- ✓ امکان دانلود نسخه ترجمه شده مقالات
- ✓ پذیرش سفارش ترجمه تخصصی
- ✓ امکان جستجو در آرشیو جامعی از صدها موضوع و هزاران مقاله
- ✓ امکان دانلود رایگان ۲ صفحه اول هر مقاله
- ✓ امکان پرداخت اینترنتی با کلیه کارت های عضو شتاب
- ✓ دانلود فوری مقاله پس از پرداخت آنلاین
- ✓ پشتیبانی کامل خرید با بهره مندی از سیستم هوشمند رهگیری سفارشات