VOLATILITY*

Federico M. Bandi and Jeffrey R. Russell
Graduate School of Business, The University of Chicago

First draft: January 2005
This draft: August 2006

Abstract

We provide a unified framework to understand current advances in two important fields in empirical finance: volatility estimation by virtue of microstructure noise-contaminated asset price data and transaction cost evaluation. In this framework, we review recently-proposed identification procedures relying on the unique possibilities furnished by asset price data sampled at high frequency. While discussing these procedures, we offer our perspective on the existing methods and findings, as well as on directions for future work.

JEL Classification: C22, C53, G11, G12, G13.

Keywords: High-frequency data, Realized volatility, Market microstructure noise, Transaction cost, Volatility and asset pricing, Liquidity and asset pricing.

*We survey new and fast-developing fields. We tried our best to keep this review updated until publication and apologize for unwanted omissions. We thank the William S. Fishman Faculty Research Fund at the Graduate School of Business of the University of Chicago (Bandi) and the Graduate School of Business at the University of Chicago (Russell) for financial support.

Table of contents

1 Introduction

Recorded asset prices deviate from their equilibrium values due to the presence of market microstructure frictions. Hence, the volatility of the observed prices depends on two distinct volatility components, i.e., the volatility of the *unobserved* equilibrium prices and the volatility of the equally *unobserved* market microstructure effects.

In keeping with this basic premise, this review starts from a model of price formation that allows for empirically relevant market microstructure effects to discuss current advances in the nonparametric estimation of both volatility notions using high-frequency asset price data.

Numerous insightful reviews have been written on volatility. The existing reviews concentrate on work that assumes observability of the equilibrium price and study its volatility properties in the absence of measurement error (see, e.g., Andersen, Bollerslev, and Diebold, 2002, and the references therein). Reviews have also been written on work that solely focuses on the measurement error and characterizes it in terms of frictions induced by the market's fine grain dynamics (see, e.g., Hasbrouck, 1996, and Stoll, 2000). Quantifying these frictions is of crucial importance to understand and measure the effective execution cost of trades. More recently, Barndorff-Nielsen and Shephard, 2007, have provided a discussion of current research on alternative nonparametric volatility estimators. While their review largely focuses on the frictionless case, it also offers interesting perspectives on the empirically more relevant case of equilibrium prices affected by market microstructure effects (see, also, McAleer and Medeiro, 2006).

The present review places emphasis on the volatilities of both unobserved components of a recorded price, i.e., equilibrium price and microstructure frictions. Specifically, our aim is to provide a unified framework to understand current advances in two important fields in empirical finance, namely equilibrium price volatility estimation and transaction cost evaluation. To this extent, we begin with a general price formation mechanism that expresses recorded (logarithmic) asset prices as the sum of (logarithmic) equilibrium prices and (logarithmic) market microstructure effects.

2 A model of price formation with microstructure effects

Write an observed logarithmic price as

$$p = p^* + \eta, \tag{1}$$

where p^* denotes the logarithmic equilibrium price, i.e., the price that would prevail in the absence of market microstructure frictions,[1] and η denotes a market microstructure contamination in the observed logarithmic price as induced by price discreteness and bid-ask bounce effects, for instance (see, e.g., Stoll, 2000). Fix a certain time period h (a trading day, for example) and assume availability of M equispaced high-frequency prices over h. Given Eq. (1) we can readily define continuously-compounded returns over any intra-period interval of length $\delta = \frac{h}{M}$ and write

[1] We start by being deliberately unspecific about the nature of the equilibrium price. We will add more economic structure to the model when discussing transaction cost evaluation (Section 7).

$$\underbrace{p_{j\delta} - p_{(j-1)\delta}}_{r_{j\delta}} = \underbrace{p^*_{j\delta} - p^*_{(j-1)\delta}}_{r^*_{j\delta}} + \underbrace{\eta_{j\delta} - \eta_{(j-1)\delta}}_{\varepsilon_{j\delta}}. \tag{2}$$

The following assumptions are imposed on the equilibrium price process and market microstructure effects.

Assumption 1. (The equilibrium price process.)

(1) *The logarithmic equilibrium price process p^*_t is a continuous stochastic volatility semimartingale. Specifically,*

$$p^*_t = \alpha_t + m_t, \tag{3}$$

where α_t (with $\alpha_0 = 0$) is a predictable drift process of finite variation and m_t is a continuous local martingale defined as $\int_0^t \sigma_s dW_s$, with $\{W_t : t \geq 0\}$ denoting standard Brownian motion.

(2) *The spot volatility process σ_t is càdlàg and bounded away from zero.*

(3) *The process $\int_0^t \sigma_s^4 ds$ is bounded almost surely for all $t < \infty$.*

Assumption 2. (The microstructure frictions.)

(1) *The microstructure frictions in the price process η have mean zero and are covariance-stationary with joint density $f_M(.)$.*

(2) *The variance of $\varepsilon_{j\delta} = \eta_{j\delta} - \eta_{(j-1)\delta}$ is $O(1)$ for all δ.*

(3) *The η's are independent of the $p^{*\prime}$s ($\eta \perp\!\!\!\perp p^*$).*

In agreement with classical asset-pricing theory in continuous time (see, e.g., Duffie, 1992), Assumption 1 implies that the equilibrium return process evolves in time as a stochastic volatility martingale difference plus an adapted process of finite variation. The stochastic spot volatility can display jumps, diurnal effects, high-persistence (possibly of the long-memory type), and nonstationarities. Leverage effects (i.e., dependence between σ and the Brownian motion W) are allowed.

Assumption 2 permits general dependence features for the microstructure friction components in the recorded prices. The correlation structure of the frictions can, for instance, capture first-order negative dependence in the recorded high-frequency returns as determined by bid-ask bounce effects (see, e.g., Roll, 1984) as well as higher order dependence as induced by clustering in order flow, for example. In general, the characteristics of the noise returns ε may be a function of the sampling frequency $\delta = \frac{h}{M}$. The joint density of the η's has a subscript M to make this dependence explicit. Similarly, the symbol $\mathbf{E}_M$ will be later used to denote expectations of the noise returns taken with respect to the measure $f_M(.)$.

While the equilibrium return process $r^*_{j\delta}$ is modelled as being $O_p\left(\sqrt{\delta}\right)$ over any intra-period time horizon of size $\delta = \frac{h}{M}$, the contaminations $\varepsilon_{j\delta}$ in the observed return process are $O_p(1)$. This result, which is a consequence of Assumptions 1(1) and 2(2), implies that longer period returns are less contaminated by noise than shorter period returns. Differently put, the magnitude of the frictions does not decrease with the distance between subsequent time stamps. Provided sampling does not occur between high-frequency price updates, the rounding of recorded prices to a grid (price discreteness) and the existence of different prices for buyers and sellers *alone* make this feature of the set-up presented here empirically compelling. As we discuss in what follows, the different stochastic orders of r^* and ε are important aspects of some recent approaches to equilibrium price variance estimation as well as to transaction cost evaluation.

2.1 The $MA(1)$ case

Sometimes the dependence structure of the microstructure friction process can be simplified. Specifically, one can modify Assumption 2 as follows:

Assumption 2b.

(1) *The microstructure frictions in the price process η are i.i.d. mean zero.*

(2) $\eta \perp\!\!\!\perp p^*$.

If the microstructure noise contaminations in the price process η are i.i.d., then the noise returns ε display an $MA(1)$ structure with a negative first-order autocorrelation. Importantly, the noise return moments do not depend on M, i.e., the number of observations over h or, equivalently, the sampling frequency $\frac{\delta}{M}$. This is an important feature of the $MA(1)$ model which, as we discuss below, has been exploited in recent work on volatility estimation.

The $MA(1)$ model, as typically justified by bid-ask bounce effects, is bound to be an approximation. However, it is known to be a realistic approximation in decentralized markets where traders arrive in a random fashion with idiosyncratic price setting behavior, the foreign exchange market being a valid example (see, e.g., Bai, Russell, and Tiao, 2005). It can also be a plausible approximation, capturing first-order effects in the data, in the case of equities when considering transaction prices and/or quotes posted on multiple exchanges. Bandi and Russell (2006b) provide additional discussions. Awartani, Corradi, and Distaso (2004) propose a formal test of the $MA(1)$ market microstructure model.

3 The variance of the equilibrium price

The recent availability of quality high-frequency financial data has motivated a growing literature devoted to the model-free measurement of the equilibrium price variance. We refer the interested reader to the review paper by Andersen, Bollerslev, and Diebold (2002) and the references therein. The main idea is to aggregate intra-period squared continuously-compounded returns and compute

$\widehat{V} = \sum_{j=1}^{M} r_{j\delta}^2$ over h. The quantity $\widehat{V}$, which has been termed "realized variance," is thought to approximate the increments of the quadratic variation of the semimartingale that drives the underlying logarithmic price process, i.e., $V = \int_0^h \sigma_s^2 ds$. The consistency result justifying this procedure is the convergence in probability of $\widehat{V}$ to V as returns are computed over intervals that are increasingly small asymptotically, i.e., as $\delta \to 0$ or, equivalently, as $M \to \infty$ for a fixed h. This result is a cornerstone in semimartingale process theory (see, e.g., Chung and Williams, Theorem 4.1, page 76, 1990).[2] More recently, the important work of Andersen, Bollerslev, Diebold, and Ebens (2001), Andersen, Bollerslev, Diebold, and Labys (2003) and Barndorff-Nielsen and Shephard (2002, 2004b) has championed empirical implementation of these ideas.

The theoretical validity of this identification procedure hinges on the observability of the equilibrium price process. However, it is widely accepted that the equilibrium price process and, as a consequence, the equilibrium return data are contaminated by market microstructure effects. Even though the early realized variance literature is aware of the potential importance of market microstructure effects, it has largely abstracted from them. The theoretical and empirical consequences of the presence of market microstructure frictions in the observed price process have been explored only recently.

3.1 Inconsistency of the realized variance estimator

Under the price formation mechanism in Section 2, the realized variance estimates are asymptotically dominated by noise as the number of squared return data increases over a fixed time period. Write

$$\widehat{V} = \sum_{j=1}^{M} r_{j\delta}^2 = \sum_{j=1}^{M} r_{j\delta}^{*2} + \sum_{j=1}^{M} \varepsilon_{j\delta}^2 + 2\sum_{j=1}^{M} r_{j\delta}\varepsilon_{j\delta}. \tag{4}$$

Since $r_{j\delta}^*$ is $O_p\left(\sqrt{\delta}\right)$ and $\varepsilon_{j\delta}$ is $O_p(1)$, the term $\sum_{j=1}^{M} \varepsilon_{j\delta}^2$ is the dominating term in the sum. Specifically, this term diverges to infinity almost surely as $M \to \infty$. The theoretical consequence of this divergence is a realized variance estimator that fails to converge to the increment of the quadratic variation (integrated variance) of the underlying logarithmic price process but, instead, increases without bound almost surely over any fixed period of time, however small: $\widehat{V} \overset{a.s.}{\to} \infty$ as $M \to \infty$ (or $\delta = \frac{h}{M} \to 0$ given h). This point has been made in independent and concurrent work by Bandi and Russell (2003, 2006a) and Zhang, Mykland, and Aït-Sahalia (2005).[3]

The divergence to infinity of the realized variance estimator over any fixed time period is an asymptotic approximations to a fairly pervasive empirical fact. When computing realized variance

[2] The corresponding weak convergence result is discussed in Jacod (1994), Jacod and Protter (1998), Barndorff-Nielsen and Shephard (2002), and Bandi and Russell (2003). Mykland and Zhang (2005) cover the case of irregularly-spaced data. Gonçalves and Meddahi (2004) discuss finite sample improvements through bootstrap methods (see, also, Gonçalves and Meddahi, 2006).

[3] This theoretical result is general and relies on the different stochastic orders of the equilibrium returns and noise returns. The result does not hinge on an $MA(1)$ structure for the noise return component ε, as implied by Assumption 2b(1). Also, importantly, the result does not hinge on the independence between the price process and the noise, as implied by Assumption 2(3) and Assumption 2b(2). Bandi and Russell (2003) clarify both statements.

estimates for a variety of sampling frequencies δ, the resulting estimates tend to increase as one moves to high frequencies (as $\delta \to 0$). In the terminology of Andersen, Bollerslev, Diebold, and Labys (1999, 2000), "the volatility signature plots," namely the plots of realized variance estimates versus different sampling frequencies,[4] are often upward sloping at high frequencies. Figure 1 shows volatility signature plots constructed using IBM mid-quotes obtained from (i) NYSE quotes and (ii) NYSE and Midwest exchange quotes. Figure 2 presents volatility signature plots for IBM using (i) NYSE and NASDAQ quotes and (ii) all quotes from the consolidated market. Figure 3 presents volatility signature plots using mid-quotes obtained from two NASDAQ stocks (Cisco Systems and Microsoft). The data are collected for the month of February 2002. In all cases the realized variance estimates increase as the sampling intervals decrease (see, also, the discussion in Bandi and Russell, 2006b).

Figure 1 about here
Figure 2 about here
Figure 3 about here

3.2 The mean-squared error of the realized variance estimator

The presence of market microstructure contaminations induces a bias-variance trade-off in integrated variance estimation through realized variance. When the equilibrium price process is observable, higher sampling frequencies over a fixed period of time result in more precise estimates of the integrated variance of the logarithmic equilibrium price (see, e.g., Andersen, Bollerslev, Diebold, and Labys, 2003, and Barndorff-Nielsen and Shephard, 2002). When the equilibrium price process is not observable, as is the case in the presence of microstructure frictions, frequency increases provide information about the underlying integrated variance but, inevitably, entail accumulation of noise that affects both the bias and the variance of the estimator (Bandi and Russell, 2003, 2006a, and Zhang, Mykland, and Aït-Sahalia, 2005).

Under Assumptions 1 and 2, absence of leverage effects ($\sigma \perp\!\!\!\perp W$), and unpredictability of the equilibrium returns ($\alpha_t = 0$),[5] Bandi and Russell (2003) provide an expression for the conditional (on the underlying volatility path) mean-squared error (MSE) of the realized variance estimator as

[4] See, also, Fang (1996).

[5] Both additional assumptions, namely absence of leverage effects and unpredictability of the equilibrium returns, can be justified. In the case of the latter, Bandi and Russell (2003) argue that the drift component α_t is rather negligible in practise at the sampling frequencies considered in the realized variance literature. They provide an example based on IBM. Assume a realistic annual constant drift of 0.08. The magnitude of the drift over a minute interval would be $0.08/(365*24*60) = 1.52 \times 10^{-7}$. Using IBM transaction price data from the TAQ data set for the month of February 2002, Bandi and Russell (2003) compute a standard deviation of IBM return data over the same horizon equal to 9.5×10^{-4}. Hence, at the one minute interval, the drift component is 1.6×10^{-4} or nearly $1/10,000$ the magnitude of the return standard deviation. Assuming absence of leverage effects is empirically reasonable in the case of exchange rate data. The same condition appears restrictive when examining high frequency stock returns. However, some recent work uses tractable parametric models to show that the effect of leverage on the unconditional MSE of the realized variance estimator in the absence of market microstructure noise is negligible (see Meddahi, 2002). This work provides some justification for the standard assumption of no-leverage in the realized variance literature. Andersen, Bollerslev, and Diebold (2002) discuss this issue.

a function of the sampling frequency δ (or, equivalently, as a function of the number of observations M), i.e.,

$$\mathbf{E}_M\left(\widehat{V}-V\right)^2 = 2\frac{h}{M}\left(Q+o(1)\right)+\Pi_M, \tag{5}$$

where

$$\Pi_M = M\mathbf{E}_M\left(\varepsilon^4\right)+2\sum_{j=1}^{M}(M-j)\mathbf{E}_M\left(\varepsilon^2\varepsilon^2_{-j}\right)+4\mathbf{E}_M(\varepsilon^2)V, \tag{6}$$

and $Q=\int_0^h\sigma_s^4ds$ is the so-called quarticity (see, e.g., Barndorff-Nielsen and Shephard, 2002). Notice that the bias of the estimator can be easily deduced by taking the expectation of $\widehat{V}$ in Eq. (4), i.e.,

$$\mathbf{E}_M\left(\widehat{V}-V\right) = M\mathbf{E}_M\left(\varepsilon^2\right). \tag{7}$$

As for the variance of $\widehat{V}$, we can write

$$\mathbf{E}_M\left(\widehat{V}-\mathbf{E}_M(\widehat{V})\right)^2 = 2\frac{h}{M}\left(Q+o(1)\right)+\Pi_M-M^2\left(\mathbf{E}_M\left(\varepsilon^2\right)\right)^2. \tag{8}$$

As we discuss below, the conditional MSE of $\widehat{V}$ can serve as the basis for an optimal sampling theory designed to choose M in order to balance bias and variance.

4 Solutions to the inconsistency problem

4.1 The early approaches: sparse sampling and filtering

Thorough theoretical and empirical treatments of the consequences of market microstructure contaminations in realized variance estimation are recent phenomena. However, while abstracting from in-depth analysis of the implications of frictions for variance estimation, the early realized variance literature is concerned about the presence of microstructure effects in recorded asset prices (see, e.g., the discussion on this topic in Andersen, Bollerslev, and Diebold, 2002).

In order to avoid substantial contaminations at high sampling frequencies, Andersen, Bollerslev, Diebold, and Ebens (2001), for example, suggest sampling at frequencies that are lower than the highest frequencies at which the data arrives. The 5-minute interval was recommended as a sensible approximate choice. Relying on the levelling off of the volatility signature plots at frequencies around 15 minutes, Andersen, Bollerslev, Diebold, and Labys (1999, 2000) suggest using 15- to 20-minute intervals in practise. If the equilibrium returns are unpredictable ($\alpha_t=0$), the correlation structure of the observed returns must be imputed to microstructure noise. Andersen, Bollerslev, Diebold, and Ebens (2001) and Andersen, Bollerslev, Diebold, and Labys (2003), among others, filter the data using an $MA(1)$ filter. An $AR(1)$ filter is employed in Bollen and Inden (2002).

4.2 MSE-based optimal sampling

More recently, an MSE-based optimal sampling theory has been suggested by Bandi and Russell (2003, 2006a). Specifically, in the case of the model laid out above, the optimal frequency $\delta^*=\frac{h}{M^*}$

at which to sample continuously-compounded returns for the purpose of realized variance estimation can be chosen as the minimizer of the MSE expansion in Subsection 3.2.

Bandi and Russell's theoretical framework clarifies outstanding issues in the extant empirical literature having to do with sparse sampling and filtering. We start with the former. The volatility signature plots provide useful insights about the bias of the realized variance estimates. The bias typically manifests itself in an upward sloping pattern as the sampling intervals become short, i.e., the bias increases with M (see Eq. (7)).[6] However, it is theoretically fairly arbitrary to choose a single optimal frequency solely based on bias considerations. While it is empirically sensible to focus on low frequencies for the purpose of bias reduction, the bias is only one of the components of the estimator's estimation error. At sufficiently low frequencies the bias can be negligible. However, at the same frequencies, the variability of the estimates might be substantial (see Eq. (8)). Figure 4 is a picture from simulations for parameter values consistent with IBM.

Figure 4 about here

The MSE-based sampling in Bandi and Russell (2003, 2006a) trades-off bias and variance optimally. As for filtering, while the dependence that the noise induces in the return data can be reduced by it, residual contaminations are bound to remain in the data. These contaminations continue to give rise to inconsistent realized variance estimates. Bandi and Russell (2003) make this point while studying the theoretical properties of both filtering at the highest frequencies at which observations arrive and filtering at all frequencies.

Bandi and Russell (2003) discuss evaluation of the MSE under Assumption 1 and 2 as well as in the $MA(1)$ case (i.e., under Assumptions 1 and 2b). In both cases, $\alpha_t = 0$ and $\sigma \perp\!\!\!\perp W$. When empirically justifiable, the $MA(1)$ case is very convenient in that the moments of the noise do not depend on the sampling frequency. Furthermore, the MSE simplifies substantially:

$$\mathbf{E}_M\left(\widehat{V} - V\right)^2 = 2\frac{h}{M}\left(Q + o(1)\right) + M\beta + M^2\alpha + \gamma, \tag{9}$$

where the parameters α, β, and γ are defined as

$$\alpha = \left(\mathbf{E}(\varepsilon^2)\right)^2, \tag{10}$$

$$\beta = 2\mathbf{E}\left(\varepsilon^4\right) - 3\left(\mathbf{E}(\varepsilon^2)\right)^2, \tag{11}$$

and

$$\gamma = 4\mathbf{E}(\varepsilon^2)V - \mathbf{E}(\varepsilon^4) + 2\left(\mathbf{E}(\varepsilon^2)\right)^2. \tag{12}$$

[6] The possible dependence between the equilibrium price p^* and the market microstructure frictions η complicates matters. Negative dependence, for instance, might drastically reduce the upward trend of the volatility signature plots at high sampling frequencies. Eq. (4) illustrates this point. The empirical relevance of negative dependence between equilibrium price and market microstructure frictions is discussed in Hansen and Lunde (2006). We focus on this issue in Subsection 5.2 below.

If M^* is large, the following approximation to the optimal number of observations applies:

$$M^* \approx \left(\frac{hQ}{(\mathbf{E}(\varepsilon^2))^2} \right)^{1/3}. \tag{13}$$

This approximation readily clarifies the nature of the microstructure-induced trade-off between the bias and variance of the realized variance estimator. If the signal coming from the underlying equilibrium price process (Q) is large relative to the noise determined by the frictions ($\left(\mathbf{E}(\varepsilon^2)\right)^2$), then sampling can be conducted at relatively high frequencies. Hence, M^* is effectively a signal-to-noise ratio.

In the $MA(1)$ case, evaluation of the MSE does not need to be implemented on a grid of frequencies and simply relies on the consistent estimation of the frequency-independent moments of the noise ($\mathbf{E}(\varepsilon^2)$ and $\mathbf{E}(\varepsilon^4)$) as well as on the estimation of the quarticity term Q.[7] In this case, Bandi and Russell (2003, 2006a) show that sample moments of the *observable* contaminated return data can be employed to identify the moments of the *unobservable* noise process at *all* frequencies. Thus, while realized variance is inconsistent in the presence of microstructure noise, appropriately defined arithmetic averages of the observed returns consistently estimate the moments of the noise. Under $\mathbf{E}\left(\eta^8\right) < \infty$, the following result holds

$$\frac{1}{M}\sum_{j=1}^{M} r_{j\delta}^{q} - \mathbf{E}\left(\varepsilon^{q}\right) \overset{p}{\to} 0 \qquad 1 \le q \le 4, \tag{14}$$

as $M \to \infty$.[8] We provide intuition for this finding in the case $q = 2$. The sum of the squared contaminated returns can be written as in Eq. (4) above, namely as the sum of the squared equilibrium returns plus the sum of the squared noise returns and a cross-product term. The price formation mechanism in Section 2 is such that the orders of magnitude of the three terms in Eq. (4) above differ since $r^*_{j\delta} = O_p\left(\sqrt{\delta}\right)$ and $\varepsilon_{j\delta} = O_p(1)$. Thus, the microstructure noise component dominates the equilibrium return process at very high frequencies, i.e., for small values of δ. This effect determines the diverging behavior of $\widehat{V}$, as discussed above. By the same logic,

[7] The quarticity term can be identified using the estimator proposed by Barndorff-Nielsen and Shephard (2002), namely

$$\widehat{Q} = \frac{M}{3h}\sum_{j=1}^{M} r_{j\delta}^{4}.$$

(See Barndorff-Nielsen and Shephard, 2004b, and Zhang, Mykland, and Aït-Sahalia, 2005, for alternative approaches.) However, $\widehat{Q}$ is not a consistent estimate of Q in the presence of noise. In practise, one could then sample the observed returns to be used in the definition of $\widehat{Q}$ at a lower frequency than the highest frequency at which observations arrive. Bandi and Russell (2006a) show by simulation that sampling returns in a reasonable (but possibly suboptimal) fashion for the purpose of quarticity estimation does not give rise to very imprecise sampling choices for realized variance. Using data, Bandi and Russell (2006a) find that sampling intervals for the quarticity between 10 and 20 minutes have a negligible effect on the resulting optimal frequency of the realized variance estimator. In light of the important role played by the quarticity in this and other identification procedures (see below), future research should study more efficient methods to estimate this term in the presence of realistic market microstructure noise effects.

[8] Importantly, this result is robust to the presence of a drift term ($\alpha_t \neq 0$), dependence between the frictions and the equilibrium price, and leverage effects (Bandi and Russell, 2003). Under assumptions, it is also robust to depedence in the frictions (Bandi and Russell, 2004). See Section 7 for additional discussions.

when we average the contaminated squared returns as in Eq. (14), the *average* of the squared noises constitutes the dominating term in the average. Naturally, then, while the remaining terms in the average vanish asymptotically due to the stochastic order of the equilibrium returns, i.e., $O_p\left(\sqrt{\delta}\right)$, the *average* of the squared noise returns converges to the second moment of the noise returns as implied by Eq. (14).

Using a sample of mid-quotes for the S&P 100 stocks over the month of February 2002, Bandi and Russell (2006a) report (average) daily optimal sampling frequencies that are between 0.5 minutes and about 14 minutes with a median value of about 3.5 minutes. The MSE improvements that the MSE-based optimal frequencies guarantee over the 5- or 15-minute frequency can be substantial. Not only do the optimal frequencies vary cross-sectionally, they also change over time. Using mid-quotes dating back to 1993 for three stocks with various liquidity features, namely EXXON Mobile Corporation (XOM), SBC Communications (SBC), and Merrill Lynch (MEL), Bandi and Russell (2006a) show that the daily optimal frequencies have substantially increased in recent times, generally due to decreases in the magnitude of the noise moments. This effect should in turn be attributed to an overall increase in liquidity.

In agreement with the analysis in Bandi and Russell (2003, 2006a), Oomen (2006) discusses an MSE-based approach to optimal sampling for the purpose of integrated variance estimation. However, some important novelties characterize Oomen's work. First, the underlying equilibrium price is not modelled as in Section 2 but as a compound Poisson process. Second, Oomen explores the relative benefits of transaction time sampling versus calendar time sampling.

Consider a Poisson process $N(t)$ with intensity $\lambda(t)$. In Oomen (2006) the observed logarithmic price process is expressed as

$$p_t = \underbrace{p_0 + \sum_{j=1}^{N(t)} \xi_j}_{p_t^*} + \sum_{j=1}^{N(t)} \eta_j, \tag{15}$$

where $\xi_j \sim$ i.i.d. $N(\mu_\xi, \sigma_\xi^2)$ and $\eta_j = \Delta\nu_j + \rho_2 \Delta\nu_{j-1} + ... + \rho_q \Delta\nu_{j-q+1}$ with $\Delta\nu_j = \nu_j - \nu_{j-1}$ and $\nu_j \sim$ i.i.d. $N(0, \sigma_\nu^2)$. The process $N(t)$ effectively counts the number of transactions up to time t. This process is assumed to be independent of both ξ and ν. Importantly, the equilibrium price p_t^* is equal to $p_0 + \sum_{j=1}^{N(t)} \xi_j$. Hence, p_t^* is a jump process of finite variation (in the tradition of Press, 1967) with integrated variance (i.e., the object of econometric interest) given now by $V = \sigma_\xi^2 \int_0^h \lambda(s) ds = \sigma_\xi^2 \Lambda(h)$. The microstructure noise contaminations η have an $MA(q)$ structure. Setting q equal to one yields a negative first-order autocorrelation of the calendar time continuously-compounded returns since

$$p_t - p_{t-\tau} = \sum_{j=N(t-\tau)}^{N(t)} \xi_j + \nu_{N(t)} - \nu_{N(t-\tau)-1}, \tag{16}$$

for any calendar time interval τ.

Oomen (2006) provides closed-form expressions for the MSE of the realized variance estimator under both calendar time sampling, as earlier, and transaction time sampling. Given M (the total number of observations), transaction time sampling leads to a sequence of prices $\{p_{t_i}\}_{i=0}^{M}$ with sampling times implicitly defined as $N(t_i) = i\left\lfloor \frac{N(h)}{M} \right\rfloor$, where $\lfloor x \rfloor$ is the integer part of x.[9] Oomen (2006) also discusses optimal choice of M in an MSE sense. Using IBM transaction prices from the consolidated market over the period between 2000 and 2004, he finds that transaction time sampling generally outperforms calendar time sampling. In his sample the average decrease in MSE that transaction time sampling induces is about 5%. Gains up to 40% can be achieved. As intuition suggests, the largest gains are obtained for days with irregular trading patterns.

4.3 Bias-correcting

The microstructure-induced bias of the realized variance estimator represents a large component of the estimator's MSE. This point is emphasized by Hansen and Lunde (2006). Hansen and Lunde (2006) propose a bias-adjustment to the conventional realized variance estimator. The bias-corrected estimator they suggest is in the tradition of robust covariance estimators such as those of Newey and West (1987) and Andrews and Monahan (1992). Its general form is

$$\widetilde{V} = \sum_{j=1}^{M} r_{j\delta}^2 + 2\sum_{h=1}^{q_M} \frac{M}{M-h} \sum_{j=1}^{M-h} r_{j\delta} r_{(j+h)\delta}, \tag{17}$$

where q_M is a frequency-dependent number of covariance terms. If the correlation structure of the noise returns has a finite order and $\alpha_t = 0$, under appropriate conditions on q_M the estimator in Eq. (17) is unbiased for the underlying integrated variance over the period, i.e., $\mathbf{E}_M\left(\widetilde{V}\right) = V$.

The intuition readily derives from the $MA(1)$ noise case. In this case the estimator takes the simpler expression

$$\widetilde{V}^{MA(1)} = \sum_{j=1}^{M} r_{j\delta}^2 + 2\frac{M}{M-1} \sum_{j=1}^{M-1} r_{j\delta} r_{(j+1)\delta}. \tag{18}$$

Under Assumption 1, Assumption 2b, and $\alpha_t = 0$, the covariance between $r_{j\delta}$ and $r_{(j+1)\delta}$, i.e., $\mathbf{E}_M(r_{j\delta}r_{(j+1)\delta})$, is the same at all frequencies and equal to $-\mathbf{E}(\eta^2)$. Hence, $\mathbf{E}_M\left(2\frac{M}{M-1}\sum_{j=1}^{M-1} r_{j\delta}r_{(j+1)\delta}\right) = -2M\mathbf{E}(\eta^2)$. The bias of the estimator $\widehat{V}$ is equal to $M\mathbf{E}(\varepsilon^2) = 2M\mathbf{E}(\eta^2)$ (see Eq. (7)). Therefore, the second term in Eq. (18) provides the required bias correction. Interestingly, the finite sample unbiasedness of Hansen and Lunde's estimator is robust to the presence of some dependence between the underlying local martingale price process (under $\alpha_t = 0$) and market microstructure noise, i.e., Assumption 2(3) or Assumption 2b(2) can be relaxed. In the $MA(1)$ case, again, it is easy to see that if $\mathbf{E}_M(r_{j\delta}^{*}\eta_{(j-s)\delta}) = 0$ for all $s \geq 1$ (implying that microstructure noise does not

[9] Equivalently, given M, business time sampling can be obtained by sampling prices at times t_i so that $\Lambda(t_i) = i\frac{\Lambda(h)}{M}$. Because $\lambda(.)$ is latent and since, conditionally on $\lambda(.)$, $\mathbf{E}(N(t)) = \Lambda(t)$, transaction time sampling can be interpreted as a feasible version of business time sampling (Oomen, 2006).

predict equilibrium returns) and $\mathbf{E}_M(r^*_{j\delta}\eta_{(j+1)\delta}) = 0$, then $\mathbf{E}_M\left(\widetilde{V}^{MA(1)}\right) = V$ (see the discussion in Bandi and Russell, 2006b). In other words, the contemporaneous covariances $\mathbf{E}_M(r^*_{j\delta}\eta_{j\delta})$ are not required to be zero. This is an important property.

Under an assumed $MA(1)$ noise structure, Zhou (1996) is the first to use the bias-corrected estimator in Eq. (18) in the context of variance estimation through high-frequency data. His original set-up assumes a constant return variance and Gaussian market microstructure noise. In this framework, Zhou (1996) characterizes the variance of the estimator and concludes that it can be minimized for a finite M. Under the more general assumptions in Section 2, but again in the presence of $MA(1)$ frictions, Hansen and Lunde (2006) have recently further studied the MSE properties of the estimator in Eq. (18). Using 5 years of DJIA price data from January 3, 2000, to December 31, 2004, they find that bias-correcting permits optimal sampling at higher frequencies than those obtained by Bandi and Russell (2006a) using the classical realized variance estimator. In addition, MSE improvements can be achieved. Consider Alcoa (AA), for example. They report an (average) daily optimal sampling frequency for their bias-corrected estimator equal to about 46 seconds. Their reported (average) optimal daily frequency for the realized variance estimator is about 9 minutes. Bias-correcting yields a reduction in the root MSE of about 33%.

Alternative bias-corrections can be provided in both the correlated noise case and in the $MA(1)$ case. The correlated noise case is studied in Bandi and Russell (2003) and Zhang (2006a) (see, also, Aït-Sahalia, Mykland, and Zhang, 2005b). The $MA(1)$ case is discussed in Bandi and Russell (2003) and Zhang, Mykland, and Aït-Sahalia (2005). For conciseness, here we only focus on the $MA(1)$ case. As we point out above, the bias of the realized variance estimator can be estimated consistently by computing an arithmetic average of the squared observed return data sampled at the highest frequencies (see Eq. (14)). The bias-corrected realized variance estimator is then equal to

$$\overleftrightarrow{V} = \widehat{V} - \overleftrightarrow{M}\frac{1}{M}\sum_{j=1}^{M} r^2_{j\delta}, \tag{19}$$

where M is the number of observations in the full sample and $\overleftrightarrow{M}$ is the number of observations used to compute $\widehat{V}$.[10] The approximate (MSE-based) optimal number of observations $\overleftrightarrow{M}^*$ of the estimator in Eq. (19) is now

[10] In a finite sample (when M is not large enough), the equilibrium return component in the estimated second moment $\frac{1}{M}\sum_{j=1}^{M} r^2_{j\delta}$ might be non-negligible. Specifically, conditional on the volatility path, the finite sample bias of $\frac{1}{M}\sum_{j=1}^{M} r^2_{j\delta}$, as an estimate of $\mathbf{E}(\varepsilon^2)$, is equal to $\frac{V}{M}$. Hence, this empirical moment can be purged of residual contaminations induced by the equilibrium price variance by substracting from it a quantity defined as $\frac{1}{M}\sum_{j=1}^{\widetilde{P}} r^2_{j\delta}$, where $\widetilde{P}$ is an appropriate number of low frequency returns calculated using 15- or 20-minute intervals, for instance. The quantity $\frac{1}{M}\sum_{j=1}^{\widetilde{P}} r^2_{j\delta}$ is roughly unbiased for $\frac{V}{M}$. The resulting estimator, i.e., $\frac{1}{M}\left(\sum_{j=1}^{M} r^2_{j\delta} - \sum_{j=1}^{\widetilde{P}} r^2_{j\delta}\right)$ has, of course, the same limiting properties as $\frac{\sum_{j=1}^{M} r^2_{j\delta}}{M}$ for any fixed $\widetilde{P}$. A similar correction is discussed in Bandi and Russell (2003, 2004) and Hansen and Lunde (2006). The presence of dependence between the frictions and the equilibrium price process complicates matters. Bandi and Russell (2004) discuss a bias-correction in this case.

$$\overleftrightarrow{M}^* \approx \left(\frac{hQ}{2\mathbf{E}\left(\varepsilon^4\right) - 3\left(\mathbf{E}(\varepsilon^2)\right)^2} \right)^{1/2}. \tag{20}$$

(Bandi and Russell, 2003).[11] In agreement with Hansen and Lunde's findings (Hansen and Lunde, 2006), this optimal frequency is generally higher then the optimal frequency of the realized variance estimator. Furthermore, it is associated with MSE improvements.

In the spirit of Zhou (1996) and Hansen and Lunde (2006), Oomen (2005) extends the framework in Oomen (2006) to the case of bias-corrected realized variance. Specifically, he studies the MSE properties of the estimator in Eq. (18) for the case of an underlying jump process of finite variation (as in Eq. (15)) and transaction time sampling. Using IBM transaction data from the consolidated market for the period January 2, 2003 - August 31, 2003, he confirms that (i) transaction time sampling can be beneficial in practise (as in Oomen, 2006) and (ii) bias-correcting can induce an increase in the optimal sampling frequency along with MSE gains. In the case of the bias-corrected estimator, he reports an (average) optimal daily frequency of about 12 seconds. The corresponding (average) optimal daily frequency of the classical realized variance estimator is around 2.5 minutes. While bias-correcting yields MSE gains of about 65% using his data, he reports that further gains (about 20%) can be obtained by employing transaction time sampling in place of calendar time sampling.

4.4 Sub-sampling

The bias-corrected estimators studied by Zhou (1996), Hansen and Lunde (2006), and Oomen (2005) are inconsistent. Biased in a finite sample, but consistent, is the estimator recently advocated by Zhang, Mykland, and Aït-Sahalia (2005) in the presence of $MA(1)$ market microstructure noise. (See, also, Aït-Sahalia, Mykland, and Zhang (2005a) for a study of consistent maximum likelihood estimation of the constant variance of a scalar diffusion process in parametric models with microstructure effects.) This promising approach relies on subsampling.[12] Assume availability of n, generally non-equispaced, observations. Define q non-overlapping sub-grids $G^{(i)}$ of the full grid of n arrival times with $i = 1, ..., q$. The first sub-grid starts at t_0 and takes every q-th arrival time, i.e., $G^{(1)} = (t_0, t_{0+q}, t_{0+2q}, ...,)$, the second sub-grid starts at t_1 and also takes every q-th arrival time, i.e., $G^{(2)} = (t_1, t_{1+q}, t_{1+2q}, ...,)$, and so on. Given the generic i-th sub-grid of arrival times, one can define the corresponding realized variance estimator as

$$\widehat{V}^{(i)} = \sum_{t_j, t_{j+} \in G^{(i)}} \left(p_{t_{j+}} - p_{t_j} \right)^2, \tag{21}$$

where t_j and t_{j+} denote consecutive elements in $G^{(i)}$. Zhang, Mykland, and Aït-Sahalia's sub-sampling approach entails averaging the realized variance estimates over sub-grids as well as bias-correcting them. To this extent, define

[11] The expression would be exact only if the estimator were defined as $\widehat{V} - \overleftrightarrow{M}\mathbf{E}(\varepsilon^2)$ which is, of course, infeasible in practise.

[12] Müller (1993), Zhou (1996), and the review of Politis et al. (1999) contain early discussions of similar ideas.

$$\widehat{V}^{sub} = \frac{\sum_{i=1}^{q} \widehat{V}^{(i)}}{q} - \overline{n}\widehat{\mathbf{E}}(\varepsilon^2), \tag{22}$$

where $\overline{n} = \frac{n-q+1}{q}$, $\widehat{\mathbf{E}}(\varepsilon^2) = \frac{\sum_{j=1}^{n}\left(p_{t_{j+}} - p_{t_j}\right)^2}{n}$ is a consistent estimate of the second moment of the noise return (as in Eq. (14)), and $\overline{n}\widehat{\mathbf{E}}(\varepsilon^2)$ is the required bias-correction (as in Eq. (7)). Under Assumption 1 with $\alpha_t = 0$ and Assumption 2b (i.e., the $MA(1)$ noise case), Zhang, Mykland, and Aït-Sahalia (2005) show that, as $q, n \to \infty$ with $\frac{q}{n} \to 0$ and $\frac{q^2}{n} \to \infty$, $\widehat{V}^{sub}$ is a consistent estimator of the integrated variance V over h. Provided $q = cn^{2/3}$, the rate of convergence of $\widehat{V}^{sub}$ to V is $n^{1/6}$ and the asymptotic distribution is mixed-normal with an estimable asymptotic variance. Specifically,

$$n^{1/6}\left(\widehat{V}^{sub} - V\right) \Rightarrow \left(\sqrt{8c^{-2}\left(\mathbf{E}(\eta^2)\right)^2 + c\frac{4}{3}Q}\right) N(0,1). \tag{23}$$

The proportionality factor c can be selected optimally in order to minimize the limiting variance in Eq. (23). This minimization leads to an asymptotically optimal number of subsamples given by

$$q^{asy} = c^{asy} n^{2/3} = \left(\frac{16\left(\mathbf{E}(\eta^2)\right)^2}{h\frac{4}{3}Q}\right)^{1/3} n^{2/3} \tag{24}$$

(Zhang, Mykland, and Aït-Sahalia, 2005). Both components of the factor c^{asy}, namely $\mathbf{E}(\eta^2)$ and Q, can be readily evaluated from the data. Specifically, the second moment of the noise η can be estimated by using a (standardized) sample average of squared continuously-compounded returns sampled at the highest frequencies as discussed in Subsection 4.2.[13] The quarticity term Q can be identified by employing the Barndorff-Nielsen and Shephard's quarticity estimator, namely $\widehat{Q} = \frac{M}{3h}\sum_{j=1}^{M} r_{j\delta}^4$ (Barndorff-Nielsen and Shephard, 2002), with continuously-compounded returns sampled at relatively low frequencies, among other methods. The 15- or 20-minute frequency have been shown to work reasonably well in practise.

The estimator of Zhang, Mykland, and Aït-Sahalia (2005) is effectively a “two-scale” estimator relying on very high-frequency return data to identify the bias component as well as on lower frequency return data to characterize the individual realized variances prior to averaging. In recent work, Zhang (2006a) has extended this approach to a "multi-scale" set-up. Her new estimator achieves the best attainable rate for this type of problems, $n^{1/4}$, and is robust to noise dependence in transaction time. See Aït-Sahalia, Mykland, and Zhang (2005b) for further discussions.

4.5 Kernels

The subsampling, or "two-scale" estimator, is a kernel-based estimator. Specifically, Barndorff-Nielsen, Hansen, Lunde, and Shephard (2005) have shown that it can be rewritten as a "modified" Bartlett kernel estimator, i.e.,

[13] Recall that, under the $MA(1)$ market microstructure model, $\mathbf{E}(\varepsilon^2) = 2\mathbf{E}(\eta^2)$. Hence, $\frac{1}{2M}\sum_{j=1}^{M} r_{j\delta}^2 \underset{M\to\infty}{\overset{p}{\to}} \mathbf{E}(\eta^2)$.

$$\widehat{V}^{sub} = \left(1 - \frac{n-q+1}{nq}\right)\widehat{\gamma}_0 + 2\sum_{s=1}^{q}\left(\frac{q-s}{q}\right)\widehat{\gamma}_s - \frac{1}{q}\vartheta_q, \tag{25}$$

where $\widehat{\gamma}_s = \sum_{j=1}^{n-s} r_j r_{j+s}$, $\vartheta_1 = 0$ and $\vartheta_q = \vartheta_{q-1} + (r_1 + ... + r_{q-1})^2 + (r_{n-q+2} + ... + r_n)^2$ for $q \geq 2$.

The modification $\frac{1}{q}\vartheta_q$ (called "edge-effect" in Barndorff-Nielsen, Hansen, Lunde, and Shephard, 2005), which mechanically derives from the construction in the previous subsection, is crucial for the estimator's consistency. Consider a more traditional, "unmodified" Bartlett kernel estimator defined as

$$\widehat{V}^{bartlett} = \left(\frac{n-1}{n}\frac{q-1}{q}\right)\widehat{\gamma}_0 + 2\sum_{s=1}^{q}\left(\frac{q-s}{q}\right)\widehat{\gamma}_s. \tag{26}$$

Under Assumption 1 (with $\alpha_t = 0$) and Assumption 2b (i.e., the $MA(1)$ noise case), Barndorff-Nielsen, Hansen, Lunde, and Shephard (2005) find that $\widehat{V}^{bartlett}$ is only "near-consistent" as $q, n \to \infty$ with $\frac{q}{n} \to 0$ and $\frac{q^2}{n} \to \infty$, namely under the same conditions yielding consistency of $\widehat{V}^{sub}$. The limiting variance of $\widehat{V}^{bartlett}$ is given by $4\left(\mathbf{E}(\eta^2)\right)^2$, which is small in practise when compared to V. [14]

Traditional kernel estimators can be rendered consistent. Barndorff-Nielsen, Hansen, Lunde, and Shephard (2006) have recently advocated unbiased symmetric kernels of the type

$$\widehat{V}^{BNHLS} = \widehat{\gamma}_0 + \sum_{s=1}^{q} w_s(\widehat{\gamma}_s + \widehat{\gamma}_{-s}), \tag{27}$$

where $\widehat{\gamma}_s = \sum_{j=1}^{n} r_j r_{j-s}$ with $s = -q, ..., q$, $w_s = k\left(\frac{s-1}{q}\right)$ and k is a function on $[0,1]$ satisfying $k(0) = 1$ and $k(1) = 0$ (see, also, Sun, 2006, for a class of unbiased, consistent estimators). If $q = cn^{2/3}$, this family of estimators is consistent (at rate $n^{1/6}$) and asymptotically mixed normal. Interestingly, when $k(x) = 1 - x$ (the Bartlett case), the limiting variance of $\widehat{V}^{BNHLS}$ is the same as that of the two-scale estimator. Hence, c can be chosen asymptotically as in the previous subsection. If, in addition, $k^{'}(0) = 0$ and $k^{'}(1) = 0$, then the number of autocovariances can be selected so that $q = cn^{1/2}$ and the estimator is consistent at rate $n^{1/4}$. When $k(x) = 1 - 3x^2 + 2x^3$, the limiting distribution of $\widehat{V}^{BNHLS}$ is the same as that of the multi-scale estimator of Zhang (2006a).

The limiting properties of the estimators in this subsection and in the previous subsection are derived under asymptotic conditions requiring the number of autocovariances (or subsamples) q and the number of observations n to diverge to infinity as $\frac{q}{n} \to 0$ and $\frac{q^2}{n} \to \infty$ (or $\frac{q^2}{n} \to c^2$ when $k^{'}(0) = 0$ and $k^{'}(1) = 0$). Whether these classical conditions in HAC estimation lead to valid asymptotic approximations to the finite sample distributions of these estimators is a question addressed in Bandi and Russell (2005b). As in Barndorff-Nielsen, Hansen, Lunde, and Shephard (2005, 2006), Zhang, Mykland, and Aït-Sahalia (2005), and Zhang (2006a), among others, Bandi and Russell (2005b) operate under Assumptions 1 (with $\alpha_t = 0$ and $\sigma \perp\!\!\!\perp W$) and Assumption 2b.

[14] The "unmodified" Bartlett kernel estimator and the "two-scale" estimator are quadratic estimators. A promising approach to consistent integrated variance estimation by virtue of unbiased quadratic estimators is contained in Sun (2006).

They recognize that, in practise, the number of autocovariances q is selected as a function of the number of observations n and set the ratio between q and n equal to a value ϕ such that $\phi \in (0,1]$. Subsequently, they derive the finite sample MSE properties of the Bartlett kernel estimator in Eq. (26), of the subsampling estimator in Eq. (22), and of the class of symmetric kernels in Eq. (27) as a function of ϕ. Finally, they optimize these properties by choosing ϕ as the minimizer of each estimator's finite sample MSE. Their main results can be summarized as follows:

1. The finite sample MSE properties of the consistent, two-scale estimator and of the inconsistent, "unmodified" Bartlett kernel estimator are similar.

2. A large component of their mean-squared error is bias-induced.

3. Asymptotic bandwidth selection methods (as in Eq. (24) above) can be very suboptimal in their case and, more generally, in the case of biased kernel estimators. Because their finite sample bias washes out asymptotically, asymptotic methods do not take the finite sample bias into account and have a tendency to select an excessively small number of bandwidths. A small number of bandwidths can lead to a large bias component in a finite sample.

4. This bias component can be reduced by choosing q in order to minimize the estimator's finite sample MSE. In the case of the modified (two-scale) and "unmodified" Bartlett kernel estimator, Bandi and Russell (2005b) propose a simple (MSE-based) rule-of-thumb which is given by:

   $$q^* \approx \left(\frac{3}{2} \frac{\frac{V^2}{n^2}}{Q} \right)^{1/3} n. \tag{28}$$

   Since the finite sample bias of these estimators does not depend on the moments of the noise (it only depends on the underlying variance process and the number of observations), this ratio should not be surprising when compared to Eq. (13). As earlier, the ratio trades-off bias and variance. If the bias component (in the numerator) is large relative to the variance component (in the denominator), then the number of autocovariances should be large. Preliminary (roughly unbiased) V and Q estimates can be obtained by using the classical realized variance estimator and the quarticity estimator with returns sampled at low 15- or 20-minute frequencies, for instance.

5. While the optimal finite sample MSE values of the two-scale estimator and of the "unmodified" Bartlett kernel estimator are generally smaller than the optimal finite sample MSE value of the classical realized variance estimator, the gains that these useful estimators can provide over the classical realized variance estimator might be either lost, or severely reduced, by suboptimally choosing the number of autocovariances.

6. The class of estimators proposed by Barndorff-Nielsen, Hansen, Lunde, and Shephard (2006) is unbiased. Asymptotic bandwidth selection criteria are expected to be less detrimental in

this case, i.e., suboptimal bandwidth choices will likely lead to smaller finite sample losses.

7. In general, even though all available consistent and "near-consistent" (in the terminology of Barndorff-Nielsen, Hansen, Lunde, and Shephard, 2005) estimators can yield accurate estimation of V when optimized, asymptotic approximations to their finite sample estimation error might be imprecise. A careful assessment of their accuracy requires paying close attention to their finite sample properties.

5 Equilibrium price variance estimation: directions for future work

5.1 Alternative integrated variance measures

The study of the implications of market microstructure noise for integrated variance estimation have largely focused on realized variance and its modifications. However, in the frictionless case promising alternative estimators have been studied in recent years. The (realized) range of Parkinson (1980) (Alizadeh, Brandt, and Diebold, 2002, Brandt and Diebold, 2006, Christensen and Podolskij, 2005, Martens and Van Dijk, 2005, and Rogers and Satchell, 1991, *inter alia*), the Fourier approach of Malliavin and Mancino (2002) (see, also, Barucci and Renò, 2002a, 2002b, and Kanatani, 2004a, 2004b), the realized power variation (Jacod, 1994, and Barndorff-Nielsen and Shephard, 2003) and the bypower variation of Barndorff-Nielsen and Shephard (2004a)[15] - more on the last two measures in what follows - are notable examples. It is now of interest to fully understand the properties of these estimators (and other estimators recently proposed) in the presence of realistic market microstructure frictions. Nielsen and Frederiksen (2005) and Huang and Tauchen (2005) contain interesting simulation work on the subject. Much is left for future research.

5.2 Relaxing the assumptions

Most of the current work on integrated variance estimation by virtue of noisy high-frequency data is conducted under an assumed diffusion process for the equilibrium price process and independent (of the equilibrium price process) $MA(1)$ market microstructure noise. While these assumptions provide a useful theoretical and empirical framework, important applications of interest might require a richer structure.

We start with the properties of the noise process and its relation with the underlying equilibrium price. Bandi and Russell (2003), Hansen and Lunde (2006), and Aït-Sahalia, Mykland, and Zhang (2005b) provide early, alternative approaches to noise persistence. A thorough discussion of the importance of allowing for dependent noise, mostly when sampling at very high frequencies, is contained in Hansen and Lunde (2006). Phillips and Yu (2006) emphasize that at high frequencies the noise process might even display dependence of the nonstationary type. The observations of

[15] Corradi and Distaso (2006) use this statistic in the context of specification tests for parametric volatility models. Barndorff-Nielsen, Graversen, Jacod, and Shephard (2006) contain a broad discussion of this and other measures as well as additional references.

Hansen and Lunde (2006) and Phillips and Yu (2006) can be understood in the context of the decomposition in Eq. (2):

$$r_{j\delta} = r^*_{j\delta} + \varepsilon_{j\delta}. \tag{29}$$

When gathering data at high frequencies, sampling between price updates occurring solely in correspondence with changes in the depth leads to observed returns that are equal to zero. In general, negligible observed returns $r_{j\delta}$ combined with unpredictable equilibrium returns $r^*_{j\delta}$ (as implied by our baseline model with $\alpha_t = 0$) induce highly persistent and potentially nonstationary microstructure noise components. Assume, for simplicity, $r_{j\delta} = 0$. Then,

$$0 = r^*_{j\delta} + \varepsilon_{j\delta} \Rightarrow \eta_{j\delta} = \eta_{(j-1)\delta} - r^*_{j\delta}. \tag{30}$$

A broader argument can be made: any factor inducing sluggishness in the adjustments to the observed prices *mechanically* determines persistence in the market microstructure noise components (Bandi and Russell, 2006b). Bandi and Russell (2006b) identify three main factors affecting the stickiness of the observed prices: the market structure (centralized versus decentralized markets), the type of price measurement (mid-quotes versus transaction prices), and the sampling method (calendar time sampling versus event time sampling). Mid-quotes that are posted on centralized markets and are sampled in calendar time are expected to have noise components that are substantially more dependent than transaction prices posted on decentralized markets and sampled in event time. In other words, the extent to which persistence is a first-order effect in the data depends on the economics of price formation as well as on the adopted sampling scheme.

In their articulate study of the properties of market microstructure noise, Hansen and Lunde (2006) also stress that attention should be paid to the dependence between the underlying equilibrium price process and market microstructure noise. Similarly to noise persistence, this dependence somewhat mechanically derives from the degree of stickiness in the observed prices (Bandi and Russell, 2006b). Eq. (29) shows that the more stable the observed prices are, the stronger is the negative dependence between equilibrium returns and noise returns. Hence, the factors inducing persistent noise components are also expected to be the factors inducing negative dependence between the noise returns and the equilibrium returns (Bandi and Russell, 2006b). Barndorff-Nielsen, Hansen, Lunde, and Shephard (2006) and Kalnina and Linton (2006) propose kernel-based approaches to the consistent estimation of integrated variance under some form of dependence between noise and equilibrium price. It is an important challenge for the literature on nonparametric variance estimation to study methods that provide satisfactory finite sample performance when noise persistence and dependence between noise and underlying equilibrium price are relevant effects in the data. In the case of kernel estimators, the issue of bandwidth selection is expected to be, as earlier, of first-order importance.

We now turn to models for the equilibrium price. The equilibrium price formation mechanism in Assumption 1 can be generalized to allow for a jump component in addition to the classical continuous semimartingale component. Barndorff-Nielsen and Shephard (2004a) have provided several stimulating theoretical results to show how to identify the integrated variance of the equilibrium price's continuous sample path component when finite activity jumps play a role (see, also, Mancini, 2003, 2004, for an alternative approach). Their main result is that realized power and realized by-power variation measures are, if properly constructed, "robust" to the presence of discontinuous components of this type. Assume the equilibrium price process is defined as in Assumption 1 and add a component to it expressed as $v_t = \sum_{j=1}^{N(t)} c_j$, where $N(t)$ is a finite activity, simple counting process and the $c_j^{'}$s are non-zero random variables.[16] Thus, $p_t^* = \alpha_t + m_t + v_t$. Now define the r, s - order bypower variation $BV_{(r,s)}$ as

$$BV_{(r,s)} = M^{-1+(r+s)/2} \sum_{j=1}^{M-1} \left| r_{j\delta}^* \right|^r \left| r_{(j+1)\delta}^* \right|^s . \tag{31}$$

In the absence of market microstructure frictions, Barndorff-Nielsen and Shephard (2004a) show that

$$BV_{(r,s)} \underset{M\to\infty}{\overset{p}{\to}} \mu_r \mu_s \int_0^h \sigma_s^{r+s} ds, \tag{32}$$

where $\mu_r = \mathbf{E}(|Z|^r) = 2^{r/2} \frac{\Gamma(\frac{1}{2}(r+1))}{\Gamma(\frac{1}{2})}$ with $Z \sim N(0,1)$, if $\max(r,s) < 2$.[17] This result readily implies that

$$\mu_r^{-1} \mu_{2-r}^{-1} BV_{(r,2-r)} \overset{p}{\to} V, \tag{33}$$

as $M \to \infty$. Since, when microstructure noise is assumed to be absent, realized variance converges to V plus the sum of the squared jumps over the period $\left(\sum_{j=1}^{N(h)} c_j^2 \right)$, subtracting $\mu_r^{-1} \mu_{2-r}^{-1} BV_{(r,2-r)}$ (with $r = 1$, for instance) from realized variance consistently estimates the sum of the squared jumps in the no noise case. This observation is employed by Andersen, Bollerslev, and Diebold (2005) and Huang and Tauchen (2005) in their analysis of the contribution of jumps to total price variance. Huang and Tauchen (2005) offer interesting simulation evidence about the robustness of this procedure to some form of market microstructure noise. More theoretical and empirical work ought to be done on the relative role played by jumps and continuous sample path price components in the presence of market frictions. Extensions to infinite activity jumps, and the impact of market frictions in this case, are also of interest. Barndorff-Nielsen, Shephard, and Winkel (2006) and Woerner (2006) are recent work on the subject in the frictionless case.

As discussed above, Oomen (2005, 2006) models the underlying equilibrium price as a pure jump process. In Large (2006), it is the observed price process which is modelled as a pure jump process with *constant* jumps whereas, coherently with Assumption 1, the underlying equilibrium

[16] If $N(t)$ is an homogeneous Poisson process and the $c_j^{'}$s are i.i.d., then v_t is a compound Poisson process.

[17] Realized r - order power variation is defined as $PV_{(r)} = M^{-1+r/2} \sum_{j=1}^{M} \left| r_{j\delta}^* \right|^r$. The limiting properties of $PV_{(r)}$ are studied in Jacod (1994) and Barndorff-Nielsen and Shephard (2003, 2004a). See Barndorff-Nielsen, Graversen, Jacod, and Shephard (2006) for discussions.

price process evolves as a stochastic volatility semimartingale. The difference between the observed price process and the underlying continuous semimartingale defines market microstructure noise. Write the observed price process as

$$p_t = p_0 + \int_0^t c_s dN_s, \tag{34}$$

where N_s is a simple counting process and c is an adapted process taking values k and $-k$, with $k > 0$. The quadratic variation of the observed price process $[p]_h$ can then be expressed as $k^2 N(h)$ since k represents the constant size of the jumps and $N(h)$ defines the number of jumps over the time interval h. Notice that the process $N(h)$ can be decomposed into the number of "continuations" $C(h)$, i.e., the number of jumps in the same direction as the previous jump, and the number of "alternations" $A(h)$, i.e., the number of jump reversals. Under assumptions, Large (2006) shows that the integrated variance of the underlying, unobservable semimartingale price process V can be consistently estimated using the quadratic variation of observed price process. Specifically, a consistent (in an asymptotic theory assuming increasingly frequent observations and small jumps) estimator can be defined by computing $[p]_h \frac{C(h)}{A(h)}$. While the quadratic variation of the observed price $[p]_h$ is generally a biased estimate of the quadratic variation of the underlying equilibrium price $[p^*]_h$, the bias can be corrected by using the factor $\frac{C(h)}{A(h)}$. The intuition goes as follows. The quadratic variation of the observed price process provides important information about the quadratic variation of the unobserved equilibrium price unless most of the jumps are alternations, for instance. In this case, $[p]_h$ will be an upward biased estimate of $[p^*]_h$. The correction factor $\frac{C(h)}{A(h)}$ will then act as a deflator.

In light of the local constancy of the observed price in the presence of an ever-evolving underlying equilibrium price, this approach captures the "mechanical effect" described in Bandi and Russell (2006b) yielding noise dependence and negative correlation between the noise and the underlying efficient price. The model's maintained assumption is that the observed prices change by fixed amounts or can be reduced, possibly by virtue of "rounding," to changes by fixed amounts. The practical applicability of this promising method will then depend on the nature of the data and hence on the price formation mechanism in specific markets. In general, an attentive analysis of the markets' fine grain dynamics has the potential to furnish important information about the process leading to market frictions. This information should be put to work to justify the use of different modelling and estimation approaches to integrated variance estimation.

5.3 Multivariate models

The provision of methods intended to identify integrated covariances and betas in the presence of market microstructure noise contaminations represents a necessary next step for the effective practise of portfolio choice and risk management through high-frequency asset price data. Barndorff-Nielsen and Shephard (2004b) study the asymptotic properties of realized covariance, i.e., the sum of the cross-products between two asset's calendar time returns over a period (a natural extension of the notion of "realized variance"), and realized beta in the frictionless case. To fix ideas, consider a second continuous stochastic volatility semimartingale price process $p^*_{(2)t}$ and re-define the original price process as $p^*_{(1)t}$. Assume, for simplicity, that the dynamics of the two price processes are driven by the same, scalar Brownian motion. The realized covariance (over h) between the original price 1 and price 2 is naturally defined as $\widehat{C}_{(1)(2)} = \sum_{j=1}^{M} r^*_{(1)j\delta} r^*_{(2)j\delta}$, where $r^*_{(u)j\delta} = p^*_{(u)j\delta} - p^*_{(u)(j-1)\delta}$ with $u = 1, 2$, and, as earlier, $\delta = h/M$. Similarly, the realized beta between asset 1 and asset 2 is defined as $\widehat{B}_{(1)(2)} = \widehat{C}_{(1)(2)} / \left(\sqrt{\widehat{V}_{(1)} \widehat{V}_{(2)}} \right)$. In the absence of frictions, Barndorff-Nielsen and Shephard (2004b) show that $\widehat{C}_{(1)(2)}$ is consistent for $\int_0^h \sigma_{(1)s}\sigma_{(2)s} ds$, i.e., (the increment of) the covariation process between price 1 and price 2 over h, and asymptotically mixed-normally distributed with an estimable limiting variance as $M \to \infty$. The corresponding results in the $\widehat{B}_{(1)(2)}$ case follow from the consistency of the realized covariance and variance estimates, as well as from their joint mixed normality, in the no noise case. Barndorff-Nielsen, Graversen, Jacod, and Shephard (2006) contains a comprehensive discussion of these (and other) findings.

New issues arise in practise when computing high-frequency estimates of integrated covariances and betas. Information arrives at different frequencies for different assets, thereby leading to an additional microstructure effect having to do with nonsynchronicity in the underlying price formation processes. Even abstracting from the presence of a noise component as in previous sections, nonsynchronous trading leads to downward biased realized covariance estimates when sampling continuously-compounded returns in calendar time at high frequencies. This is the so-called Epps effect (1979). The asset-pricing literature has long recognized the importance of this effect. Scholes and Williams (1977), Dimson (1979), and Cohen, Hawanini, Maier, Schwartz (1983), among others, use leads and lags in nonparametric covariance measures to adjust for nonsynchronous trading. Martens (2005) reviews the early work on the subject. In the realized covariance case, the adjusted estimator with U lags and L leads can be simply defined as $\widehat{C}^{UL}_{(1)(2)} = \sum_{j=1}^{M} \sum_{s=-L}^{U} r^*_{(1)j\delta} r^*_{(2)(j-s)\delta}$. The logic behind this adjustment is well-known (see, e.g., Cohen, Hawanini, Maier, Schwartz, 1983). Assume the equilibrium returns are martingale difference sequences ($\alpha_t = 0$). Then, $\widehat{C}^{UL}_{(1)(2)}$ is virtually unbiased for the true covariation over the period provided U and L are large enough. If U and L are small, then lack of price updates for either stock is bound to induce (downward) biases.

Initial work on realized covariance estimation in the presence of noisy high-frequency data is contained in Bandi and Russell (2005a) and Martens (2005). Bandi and Russell (2005a) study MSE-based optimal sampling for the purpose of realized covariance and beta estimation. Nonsynchronicity is accounted for by adding leads and lags to the optimized realized covariance estimator. Future research should study direct MSE-based optimization of the lead-lag estimator (for a certain number of leads and lags) as well as optimal choice of the number of leads and lags when noise is present. As is well-known, the inclusion of a large number of leads and lags improves the bias properties of the estimator but increases its variability. Martens (2005) study the MSE properties of a variety of covariance estimators (including realized covariance relying on equally-spaced returns and lead-lag estimators) through simulations based on Lo and MacKinlay's (1990) nonsynchronous trade model.

Recently, Hayashi and Yoshida (2005, 2006), Sheppard (2006), and Zhang (2006b), among others, have introduced promising, alternative approaches to high-frequency covariance estimation. The Hayashi and Yoshida's estimator, for instance, sums the products of all *overlapping* tick-by-tick returns rather than the products of the calendar time returns, as is the case for realized covariance. Specifically, the estimator is defined as

$$\sum_{j=1}^{\#}\sum_{s\in S_j} r^*_{(1)j} r^*_{(2)s}, \tag{35}$$

where $r^*_{(u)j} = p^*_{(u)j} - p^*_{(u)(j-1)}$ with $u = 1, 2$, $S_j = \left\{ s | \left(t^{(1)}_{j-1}, t^{(1)}_{j}\right) \cap \left(t^{(2)}_{s-1}, t^{(2)}_{s}\right) \neq 0 \right\}$, the $t_j^{\prime}$s are transaction times, and # denotes the number of transactions for asset 1 over the period. In the absence of classical microstructure noise contaminations, but in the presence of nonsynchronous trading, the Hayashi and Yoshida estimator is consistent and asymptotically normally distributed as the number of observations increases without bound over the trading day (Hayashi and Yoshida, 2005).

Voev and Lunde (2006) and Griffin and Oomen (2006) provide thorough finite sample studies of the MSE properties of several covariance estimators, including realized covariance, optimally-sampled realized covariance, and the Hayashi-Yoshida estimator, as well as recommendations for practical implementations.

Much remains to be done. While the first-order issues in high-frequency covariance estimation are likely to be fully understood, the methods and solutions are still in constant evolution. Arguably, the main goal of the literature is to provide reliable forecasts of *large* covariance matrices. We are far from this goal. On the one hand, the notion of reliability depends on the adopted metric (see below for discussions). On the other hand, the dimensionality of problems of practical interest continues to pose substantial issues when relying on high-frequency nonparametric estimates. Considerable effort is now being devoted to obtaining unbiased and efficient, in sample, high-frequency covariance estimates. We welcome this effort and emphasize that out-of-sample performance will ultimately be the judge.

5.4 Forecasting and economic metrics

Understandably, the initial work on integrated variance estimation by virtue of high-frequency data was largely motivated by volatility prediction (see, e.g., Andersen, Bollerslev, Diebold, and Labys, 2003, Andersen, Bollerslev, and Meddahi, 2004, 2005, and the references therein). In the no noise case, high-frequency volatility forecasting using alternative reduced-form models, as well as alternative integrated variance estimators, has been successfully conducted by Ghysels, Santa-Clara, and Valkanov (2006) and Forsberg and Ghysels (2004), among many others (Andersen, Bollerslev, Christoffersen, and Diebold, 2005, review this literature).

The noise case is now receiving substantial attention. Bandi and Russell (2006a) and Bandi, Russell, and Zhu (2006) employ reduced-form models to show that optimally-sampled realized variances (covariances) outperform realized variances (covariances) constructed using ad-hoc intervals in predicting variances (covariances) out-of-sample (see, also, Andersen, Bollerslev, and Meddahi, 2006). Ghysels and Sinko (2006) use the MIDAS approach of Ghysels, Santa-Clara, and Valkanov (2006) to evaluate the relative performance of realized variance based on fixed intervals, bias-corrected realized variance as in Eq. (17) above, and power variation. Confirming findings in Ghysels, Santa-Clara, and Valkanov (2006), their results point to the superior out-of-sample performance of power variation. Large (2006) employs the HAR-RV model of Corsi (2003), as in Andersen, Bollerslev, and Diebold (2005), to stress that his "alternation estimator" can have better forecasting properties than realized variance constructed using fixed, arbitrary intervals.

More work ought to be done. On the one hand, a comprehensive study using a variety of variance/covariance measures and reduced-form models appears to be needed. To this day, the literature appears to solely agree on the fact that realized variance constructed using ad-hoc fixed intervals is generally dominated, in terms of forecasting performance, by alternative measures. A complete comparison between these alternative measures, including optimally-sampled realized variance, optimally-sampled bias-corrected realized variance, and consistent kernel estimators, appears to be an important topic for future empirical work on the subject. On the other hand, as forcefully emphasized by Bandi and Russell (2006b), assessing the out-of-sample performance of alternative variance estimates using relevant economic metrics is arguably the single most important test in the literature. Thus far, two metrics have been proposed. Bandi and Russell (2006a) consider a *portfolio allocation* problem and the long-run utility that a mean-variance representative investor derives from alternative variance forecasts as the relevant performance criterion. The same portfolio-based approach has been recently implemented by Bandi, Russell, and Zhu (2006) and De Pooter, Martens, and Van Dijk (2006) in a multivariate context (see Fleming, Kirby, and Ostdiek, 2001, 2003, and West, Edison, and Cho, 1993, in the no noise case). Bandi and Russell (2005b, 2006c) study volatility forecasting for the purpose of *option pricing* in the context of a simulated derivative market (see Engle, Hong, and Kane, 1990, in the no noise case). In agreement with the forecasting results derived from reduced-form models, the use of economic metrics indicates that optimally-sampled realized variances (covariances) have the potential to substantially outperform realized variances (covariances) based on fixed intervals. In addition, optimally-sampled realized variances can yield more accurate forecasts than certain consistent kernel estimators (such as the two-scale estimator), when these estimators are implemented using asymptotic bandwidth selection methods. Consistent and "near-consistent" kernel estimators that are implemented at their full potential on the basis of finite sample criteria (as recommended by Bandi and Russell, 2005b) are likely to dominate optimally-sampled realized variance in the context of the above-mentioned metrics. Again, future work on the subject should provide a more comprehensive study focusing on a variety of suggested measures. In addition, alternative economic metrics should be investigated.

6 The variance of microstructure noise: a consistency result

Even though the classical realized variance estimator is not a consistent estimator of the integrated variance of the underlying equilibrium price, a re-scaled version of the standard realized variance estimator is, under assumptions, consistent for the variance of the noise return component (Bandi and Russell, 2003, and Zhang, Mykland, and Aït-Sahalia, 2005). More generally, sample moments of the observed return data can estimate moments of the underlying noise return process at high-frequencies (see Eq. (14) above). Bandi and Russell (2003) discuss this result and use it to characterize the MSE of the conventional realized variance estimator.

While the literature on integrated variance estimation focuses on the volatility features of the underlying equilibrium price, the empirical market microstructure literature places emphasis on the *other* component of the observed price process in Eq. (1), namely the price frictions η. When p is a transaction price, such frictions can be interpreted in terms of transaction costs in that they constitute the difference between the observed price p and the corresponding equilibrium price p^*.[18] Hasbrouck (1993) and Bandi and Russell (2004) provide related, but different, frameworks to use transaction price data in order to estimate the second moment of the transaction cost η (rather than moments of ε as generally needed in the integrated variance literature) under mild assumptions on the features of the price formation mechanism in Section 2. The implications of their results for measuring transaction costs are the subject of the next sections. We start with a discussion of traditional approaches to transaction cost evaluation.

7 The benefit of consistency: measuring market quality

7.1 Transaction cost estimates

Following Perold (1988), it is generally believed that an ideal measure of the execution cost of a trade should be based on the comparison between the trade price for an investor's order and the equilibrium price prevailing at the time of the trading decision. Although informed, individual investors can plausibly construct this measure, researchers and regulators do not have enough information to do so (see Bessembinder (2003) for a discussion).

Most available estimates of transaction costs relying on high-frequency data hinge on the basic logic behind Perold's original intuition. Specifically, there are three high-frequency measures of execution costs that have drawn attention in recent years, i.e., the so-called *bid-ask half spread*, the *effective half spread*, and the *realized half spread*. The bid-ask half spread is defined as half the difference between ask quote and bid quote. The effective half spread is the (signed[19]) difference between the price at which a trade is executed and the mid-point of reference bid-ask quotes. As for the realized half spread, this measure is defined as the (signed) difference between the transaction

[18] Measuring the execution costs of stock market transactions and understanding their determinants is of importance to a variety of market participants, such as individual investors and portfolio managers, as well as regulators. In November 2000, the Security and Exchange Commission issued Rule 11 Ac. 1-5 requesting market venues to widely distribute (in electronic format) execution quality statistics regarding their trades.

[19] Positive for buy orders and negative for sell orders.

price and the mid-point of quotes in effect some time after the trade.[20] In all cases, an appropriately chosen bid-ask mid-point is used as an approximation for the relevant equilibrium price.

The limitations of these measures of the cost of trade have been pointed out in the literature (the interested reader is referred to the special issue of the Journal of Financial Markets on transaction cost evaluation, JFM 6, 2003, for recent discussions). The bid-ask half spread, for example, is known to overestimate the true cost of trade in that trades are often executed at prices within the posted quotes. As for the effective and realized spreads, not only do they require the trades to be signed as buyer or seller-initiated, but they also require the relevant quotes and transaction prices to be matched.

The first issue (assigning the trade direction) arises due to the fact that commonly used high-frequency data sets (the TAQ database, for instance) do not contain information about whether a trade is buyer or seller-initiated. Some data sets do provide this information (the TORQ database being an example) but the length of their time series is often insufficient. Naturally, then, a considerable amount of work has been devoted to the construction of algorithms intended to classify trades as being buyer of seller-initiated simply on the basis of transaction prices and quotes (see, e.g., Lee and Ready, 1991, and Ellis, Michaely, and O'Hara, 2000). The existing algorithms can of course missclassify trades (the Lee and Ready method, for example, is known to categorize incorrectly about 15% of the trades), thereby inducing biases in the final estimates. See Bessembinder (2003) and Peterson and Sirri (2003) for discussions.

The second issue (matching quotes and transaction prices) requires potentially arbitrary judgment calls. Since the trade reports are often delayed, when computing the effective spreads, for example, it seems sensible to compare the trade prices to mid-quotes occurring before the trade report time. The usual allowance is 5 seconds (see, e.g., Lee and Ready, 1991) but longer lags can of course be entertained.

This said, there is a well-known measure which can be computed using low frequency data and does not require either the signing of the trades or the matching of quotes and transaction prices, i.e., *Roll's effective spread estimator* (Roll, 1984). Roll's estimator does not even rely on the assumption that the mid-points of the bid and ask quotes are valid proxies for the unobserved equilibrium prices. The idea behind Roll's measure can be easily laid out using the model in Section 2. Write the model in transaction time. Assume

$$\eta_i = sI_i, \tag{36}$$

where I_i equals 1 for a buyer-initiated trade and -1 for a seller-initiated trade with $p(I_i = 1) = p(I_i = -1) = \frac{1}{2}$. If Assumption 1 (with $\alpha_t = 0$) and Assumption 2b are satisfied, then

$$\mathbf{E}(r, r_{-1}) = -s^2. \tag{37}$$

Equivalently,

[20] The idea is that the traders possess private information about the security value and the trading costs should be assessed based on the trades' non-informational price impacts.

$$s = \sqrt{-\mathbf{E}(r, r_{-1})}. \tag{38}$$

Thus, the constant width of the spread can be estimated consistently based on the (negative) first-order autocovariance of the observed (low-frequency) stock returns.

Roll's estimator hinges on potentially restrictive assumptions. The equilibrium returns r^* are assumed to be serially uncorrelated. In addition, the microstructure frictions in the observed returns r follow an $MA(1)$ structure, as largely implied by bid-ask bounce effects, with a constant cost of trade s. Finally, the estimator relies on the microstructure noise components being uncorrelated with the equilibrium prices.

7.2 Hasbrouck's pricing errors

Hasbrouck (1993) assumes the price formation mechanism in Eq. (1). However, his set-up is in discrete time and time is measured in terms of transaction arrival times. Specifically, the equilibrium price p^* is modelled as a random walk while the η's, which may or may not be correlated with p^*, are mean-zero covariance stationary processes. Hence, he considerably relaxes the assumptions that are needed to derive the classical Roll effective spread estimator. Hasbrouck (1993) interprets the difference η between the transaction price p and the equilibrium price p^* as a *pricing error* impounding microstructure effects. The standard deviation of the pricing error σ_η is the object of interest. Because stocks whose transaction prices track the equilibrium price can be regarded as being stocks that are less affected by barriers to trade, σ_η is thought to represent a natural measure of market quality.

Using methods in the tradition of Beveridge-Nelson (1981) and Watson (1986) to study non-stationary time series (the observed price p in this case) expressed as the sum of a nonstationary component (the equilibrium price p^*) and a residual stationary component (the pricing error η), Hasbrouck (1993) provides estimates (and lower bounds) for σ_η. His empirical work focuses on NYSE stocks and utilizes transaction data collected from the Institute for the Study of Securities Markets (ISSM) tape for the first quarter of 1989. His (average) estimated σ_η value is equal to about 33 basis points. Under an assumption of normality, the average value for the expected transaction cost $\mathbf{E}\,|\eta|$ is equal to about 26 basis points $\left(\frac{2}{\sqrt{\pi}}\sigma_\eta \approx 0.8\sigma_\eta\right)$ in his data.

7.3 Full-information transaction costs

Bandi and Russell (2004) define an alternative notion of pricing error. Their approach imposes more economic structure on the model in Section 2. They begin by noting that in a rational expectation set-up with asymmetric information two equilibrium prices can be defined in general: an "efficient price," i.e., the price that would prevail in equilibrium given public information, and a "full-information price," the price that would prevail in equilibrium given private information. Both the efficient price and the full-information price are unobservable. The econometrician only observes transaction prices.

In this setting, two sources of "market inefficiency" arise. First, transaction prices deviate from efficient prices due to classical market microstructure frictions (see Stoll's AFA presidential address, Stoll, 2000, for discussions). Second, the presence of asymmetric information induces deviations between efficient prices and full-information prices. Classical approaches to transaction cost evaluation (in Subsection 7.1) and Hasbrouck's important approach to pricing error estimation (in Subsection 7.2) refer to the efficient price as the relevant equilibrium price. Hence, these methods are meant to only account for the first source of inefficiency.

A cornerstone of market microstructure theory is that uninformed agents learn about existing private information from observed order flow (see, e.g., the discussions in O'Hara, 1995). Since each trade carries information, meaningful revisions to the efficient price are made regardless of the time interval between trade arrivals. Hence, the efficient price is naturally thought of as a process changing discretely at transaction times. Contrary to the public-information set, the full-information set, by definition, contains all information used by the agents in their decisions to transact. Hence, the full-information price is unaffected by past order flow. Barring occasional news arrivals to the informed agents, the dynamic behavior of the full information price is expected to be relatively "smooth." As for the microstructure frictions, separate prices for buyers and sellers and discreteness of prices *alone* suggest that changes in the microstructure frictions from trade to trade are discrete in nature.

Bandi and Russell (2004) formalize these ideas by writing the model in Section 2 in transaction time. They add structure to the specification in Eq. (1) in order to account for the desirable properties of efficient price, full-information price, and microstructure frictions. Specifically, write

$$p_i = p^*_{t_i} + \eta_i \tag{39}$$

$$= p^*_{t_i} + \eta_i^{asy} + \eta_i^{fri}, \tag{40}$$

where $p^*_{t_i}$ is now the (logarithmic) full-information price, $p^*_{t_i} + \eta_i^{asy}$ is the discretely-evolving (logarithmic) efficient price, and η_i^{fri} denotes conventional (discrete) microstructure frictions. The deviation η_i includes a classical friction component η_i^{fri} and a pure asymmetric information component η_i^{asy}. The former is affected by both liquidity and asymmetric information,[21] the latter should only be affected by asymmetric information. As in Section 2, it is convenient to rewrite the model in terms of observed continuously-compounded returns, i.e.,

$$r_i = r^*_{t_i} + \varepsilon_i, \tag{41}$$

where $r_i = p_i - p_{i-1}$, $r^*_{t_i} = p^*_{t_i} - p^*_{t_{i-1}}$, and $\varepsilon_i = \eta_i - \eta_{i-1}$. At very high frequencies, the observed return data (the r_i's) are dominated by return components that are induced by the microstructure effects (the ε_i's) since the full-information returns evolve smoothly in time. Technically, $r^*_{t_i} = O_p\left(\sqrt{\max|t_i - t_{i-1}|}\right)$ and $\varepsilon_i = O_p(1)$. In this context, Bandi and Russell (2004) employ sample

[21] Market microstructure theory imputes classical frictions to operating (order-processing and inventory-keeping) costs and adverse selection. See, for example, the discussion in Stoll (2000).

moments of the *observed* high-frequency return data to identify the moments of the *unobserved* trading cost η_i. They do so by using the informational content of observed return data whose full-information return component $r^*_{t_i}$ is largely swamped by the component ε_i when sampling is conducted at the high frequencies at which transactions occur in practice.

Assume the covariance structure of the η's is such that $\mathbf{E}(\eta\eta_{-j}) = \theta_j \neq 0$ for $j = 1, ..., k < \infty$ and $\mathbf{E}(\eta\eta_{-j}) = 0$ for $j > k$. This structure accommodates temporal dependence possibly induced by clustering in order flow. It is then easy to show that

$$\sigma_\eta = \sqrt{\left(\frac{1+k}{2}\right)\mathbf{E}(\varepsilon^2) + \sum_{s=0}^{k-1}(s+1)\mathbf{E}(\varepsilon\varepsilon_{-k+s})}. \tag{42}$$

For every sample period (a trading day, for instance), an estimate of σ_η can thus be obtained by replacing the moments of the *unobserved* contaminations ε with the corresponding sample moments of the *observed* returns. At very high frequencies (represented here by a large number of observations for each period), the full-information return component of each sample moment is expected to be negligible. Formally,

$$\widehat{\sigma}_\eta = \sqrt{\left(\frac{k+1}{2}\right)\left(\frac{\sum_{i=1}^{\widetilde{M}} r_i^2}{\widetilde{M}}\right) + \sum_{s=0}^{k-1}(s+1)\left(\frac{\sum_{i=k-s+1}^{\widetilde{M}} r_i r_{i-k+s}}{\widetilde{M}-k+s}\right)} \underset{\widetilde{M}\to\infty}{\overset{p}{\to}} \sigma_\eta, \tag{43}$$

where $\widetilde{M}$ is now the total number of transactions over the period.[22] This consistency result only relies on the different stochastic orders of efficient price, full-information price, and classical frictions. These orders are simply meant to formalize the economics of price formation in markets with asymmetric information. The result is robust to predictability in the underlying full-information return process ($\alpha_t \neq 0$), presence of infrequent jumps in the full-information price, dependence between the full-information price and the frictions as well as time-dependence in the frictions.

When the number of observations for each time period is not large enough, the potential for (finite sample) contaminations in the estimates due to the presence of a non-negligible full-information price component is higher. Bandi and Russell (2004) suggest a finite sample adjustment.[23]

Bandi and Russell (2004) use the convention of calling the standard deviation $\widehat{\sigma}_\eta$, rather than the actual η, *full-information transaction cost,* or *FITC*. While the *FITCs* are standard deviations, one can either assume normality of the η's (as done in Hasbrouck, 1993, for similar purposes) or use the approach in Roll (1984) to derive expected costs. In the former case, a consistent estimate of $\mathbf{E}\,|\eta|$ can be obtained by computing $\frac{2}{\sqrt{\pi}}\widehat{\sigma}_\eta$. In the latter case, assume $\eta = sI$, where the random variable

[22] Under uncorrelatedness of the full-information returns, k can be estimated based on the dependence properties of the observed returns.

[23] In the absence of correlation between the frictions η and the full-information price p^* the bias-adjustment is relatively straightforward and can be implemented by using nonparametric estimates of the full-information price variance as described in Footnote 9, Subsection 4.3. In the presence of correlation between η and p^*, a complete bias-correction requires parametric assumptions on the underlying full-information price process. Bandi and Russell (2004) use the price formation mechanism proposed by Hasbrouck and Ho (1987) to quantify the estimates' finite sample bias.

I, defined in Subsection 7.1, represents now the direction (higher or lower) of the transaction price with respect to the full-information price and s is the constant full-information transaction cost. Then, $\widehat{\sigma}_\eta$ consistently estimates s.

Using a sample of S&P 100 stocks over the month of February 2002, Bandi and Russell (2004) report an average value for $\widehat{\sigma}_\eta$ equal to 14 basis points. Under normality, their estimated average $\mathbf{E}\,|\eta|$ is then equal to about 11 basis points. This value is larger that the corresponding average effective spread (about 6 basis points). Consistent with the economic interpretation underlying the construction of the *FITC*s, Bandi and Russell (2004) find that the *FITC*s are cross-sectionally more correlated with private information proxies, such as the *PIN* measure of Easley and O'Hara (see, e.g., Easley, Kiefer, O'Hara, and Paperman, 1996), the turn-over ratio (Stoll, 1989), and the number of analysts following the stock, than the average effective spreads and the average half bid-ask spreads. Furthermore, they find that the deviations of the efficient prices from the full-information prices, as determined by the existence of private information in the market place, can be as large as the departures of the transaction prices from the efficient prices.

Assume now σ_η is stochastic and latent. In keeping with the logic behind the vibrant and successful realized variance literature initiated by Andersen, Bollerslev, Diebold, and Labys (2003) and Barndorff-Nielsen and Shephard (2002), the high-frequency approach suggested in Bandi and Russell (2004) can be interpreted as providing a method to render the latent noise volatility observable (i.e., easily estimable without filtering) for each period of interest. While the realized variance literature has placed emphasis on the volatility of the underlying equilibrium price process, one can focus on the other volatility component of the observed returns, i.e., microstructure noise volatility. Treating the volatility of the noise component of the observed prices as being directly observable can allow one to address a broad array of fundamental issues. Some have a statistical flavor having to do with the distributional and dynamic properties of the noise variance and its relationship with the time-varying variance of the underlying equilibrium price process. Some have an economic importance having to do with the dynamic determinants of the cost of trade. Since the most salient feature of the quality of a market is how much agents have to pay in order of transact, much can be learned about the genuine market dynamics by exploiting the informational content of the estimated noise variances.

8 Volatility and asset pricing

Barring complications induced by the shorter observation span of asset price data sampled at high frequencies, the methods described in the previous sections can have important implications for the cross-sectional asset pricing literature.

A promising, recent strand of this literature has been devoted to assessing whether stock *market volatility* is priced in the cross-section of stock returns. Being innovations in volatility correlated with changes in investment opportunities, this is a relevant study to undertake. Ang, Hodrick, Xing, and Zhang (2005), Adrian and Rosenberg (2005), and Moise (2004), among others, find that the price of market volatility risk is negative. Volatility is high during recessions. Stocks whose returns covary with innovations in market volatility are stocks which pay off during bad times. Investors are willing to pay a premium to hold them. The results in Ang, Hodrick, Xing, and Zhang (2005), Adrian and Rosenberg (2005), and Moise (2004) are robust to the use of alternative, parametric and nonparametric, low-frequency volatility estimates. In virtue of the potential accuracy of the newly-developed high-frequency volatility measures, as described above, it is now of interest to re-evaluate the importance of market volatility as a systematic risk factor by studying the cross-sectional pricing implications of these measures. In this context, market microstructure issues ought to be accounted for.

Another strand of this literature has focused on the pricing implication of *market liquidity*. As is the case for market volatility, innovations in liquidity are correlated with the business cycle. Stocks yielding high returns when illiquidity is high provide a hedge. Not surprisingly, the price of market illiquidity risk is found to be negative (see, e.g., Acharya and Pedersen, 2005, and Pastor and Stambaugh, 2003). Liquidity is hard to measure. The recent advances in high-frequency volatility estimation provide a rich set of tools to separate liquidity-induced components (named "market microstructure frictions" earlier) from the estimated moments of the observed high-frequency asset returns. When aggregated across stocks (for each period of interest), these components have the potential to provide important information about the time-series properties of the overall level of market liquidity. These properties can be put to work to better understand the pricing of (il-)liquidity risk from a novel standpoint.

The pricing of idiosyncratic risk is also of interest. Since individuals are likely to take into account the cost of acquiring and rebalancing their portfolios, expected stock returns should somehow embed idiosyncratic transaction costs in equilibrium. This observation has given rise to a convergence between classical market microstructure work on price determination and asset pricing in recent years (the interested reader is referred to the recent survey of Easley and O'Hara, 2002). The studies on the cross-sectional relationship between expected stock returns and cost of trade largely hinge on liquidity-based theories of execution cost determination (Amihud and Mendelson, 1986, Brennan and Subrahmanyam, 1996, Datar, Naik, and Radcliffe, 1998, and Hasbrouck, 2003, among others). Alternatively, some recent studies rely on information-based approaches to the same issue (see, e.g., Easley, Hvidkjaer, and O'Hara, 2002). Much remains to be done. Full-information transaction costs, among other tools discussed earlier, may provide a promising bridge between liquidity-based and information-based arguments.

Generally speaking, the convergence between market microstructure theory and methods and asset pricing is still in its infancy. We are convinced that the recent interest in microstructure issues in the context of volatility estimation is providing, and will continue to provide, an important boost to this inevitable process of convergence.

References

[1] Acharya, V.V., and L.H. Pedersen (2005). Asset pricing with liquidity risk. *Journal of Financial Economics 77*, 375-410.

[2] Aït-Sahalia, Y., P. Mykland, and L. Zhang (2005a). How often to sample a continuous-time process in the presence of market microstructure noise. *Review of Financial Studies 18,* 351-416.

[3] Aït-Sahalia, Y., P. Mykland, and L. Zhang (2005b). Ultra high-frequency volatility estimation with dependent microstructure noise. *Working paper.*

[4] Alizadeh, S., M. Brandt, and F. Diebold (2002). Range-based estimation of stochastic volatility models. *Journal of Finance 57,* 1047-1091.

[5] Amihud, Y., and H. Mendelson (1986). Asset pricing and the bid-ask spread. *Journal of Financial Economics 17,* 223-249.

[6] Andersen, T.G., T. Bollerslev, and F.X. Diebold (2002). Parametric and nonparametric measurements of volatility. In Y. Aït-Sahalia and L.P. Hansen (Eds.) *Handbook of Financial Econometrics.* Elsevier, North-Holland. *Forthcoming.*

[7] Andersen, T.G., T. Bollerslev, and F.X. Diebold (2005). Some like it smooth, and some like it rough: disentagling continuous and jump components in measuring, modeling, and forecasting asset return volatility. *Working paper.*

[8] Andersen, T.G., T. Bollerslev, and N. Meddahi (2004). Analytic evaluation of volatility forecasts. *International Economic Review 45*, 1079-1110.

[9] Andersen, T.G., T. Bollerslev, and N. Meddahi (2005). Correcting the errors: a note on volatility forecast evaluation based on high-frequency data and realized volatilities. *Econometrica 73,* 279-296.

[10] Andersen, T.G., T. Bollerslev, and N. Meddahi (2006). Market microstructure noise and realized volatility forecasting. *Working paper.*

[11] Andersen, T.G., T. Bollerslev, P.F. Christoffersen, and F.X. Diebold (2005). Volatility and correlation forecasting. In G. Elliott, C.W.J. Granger, and Allan Timmermann (Eds.), *Handbook of Economic Forecasting.* Elsevier, North-Holland. *Forthcoming.*

[12] Andersen, T.G., T. Bollerslev, F.X. Diebold, and H. Ebens (2001). The distribution of realized stock return volatility. *Journal of Financial Economics 61,* 43-76.

[13] Andersen, T.G., T. Bollerslev, F.X. Diebold, and P. Labys (1999). (Understanding, optimizing, using, and forecasting) Realized volatility and correlation. *Working paper.*

[14] Andersen, T.G., T. Bollerslev, F.X. Diebold, and P. Labys (2000). Great Realizations. *Risk Magazine 13,* 105-108.

[15] Andersen, T.G., T. Bollerslev, F.X. Diebold, and P. Labys (2003). Modeling and forecasting realized volatility. *Econometrica 71*, 579-625.

[16] Andrews, D.W.K. and J.C. Monahan (1992). An improved heteroskedasticity and autocorrelation consistent covariance matrix estimator. *Econometrica 60*, 953-966.

[17] Ang, A., R. Hodrick, Y. Xing, and X. Zhang (2005). The cross-section of volatility and expected returns. *Journal of Finance 61,* 259-299.

[18] Awartani, B., V. Corradi, and W. Distaso (2004). Testing and modelling market microstructure effects with an application to the Dow Jones Industrial Average. *Working paper.*

[19] Bai, X., J.R. Russell, and G. Tiao (2005). Effects of non-normality and dependence on the precision of variance estimates using high-frequency data. *Working paper.*

[20] Bandi, F.M., and J.R. Russell (2003). Microstructure noise, realized volatility, and optimal sampling. *Working paper.*

[21] Bandi, F.M., and J.R. Russell (2004). Full-information transaction costs. *Working paper.*

[22] Bandi, F.M., and J.R. Russell (2005a). Realized covariation, realized beta, and microstructure noise. *Working paper.*

[23] Bandi, F.M., and J.R. Russell (2005b). Market microstructure noise, integrated variance estimators, and the accuracy of asymptotic approximations. *Working paper.*

[24] Bandi, F.M., and J.R. Russell (2006a). Separating microstructure noise from volatility. *Journal of Financial Economics 79,* 655-692.

[25] Bandi, F.M., and J.R. Russell (2006b). Comment on Hansen and Lunde. *Journal of Business and Economic Statistics 24,* 167-173.

[26] Bandi, F.M., and J.R. Russell (2006c). Microstructure noise, realized variance, and optimal sampling. *Working paper.*

[27] Bandi, F.M., J.R. Russell, and Yinghua Zhu (2006). Using high-frequency data in dynamic portfolio choice. *Econometric Reviews, forthcoming.*

[28] Barndorff-Nielsen, O.E., and N. Shephard (2002). Econometric analysis of realized volatility and its use in estimating stochastic volatility models. *Journal of the Royal Statistical Society, Series B 64*, 253-280.

[29] Barndorff-Nielsen, O.E., and N. Shephard (2003). Realized power variation and stochastic volatility. *Bernoulli 9*, 243-265.

[30] Barndorff-Nielsen, O.E., and N. Shephard (2004a). Power and bypower variation with stochastic volatility and jumps (with discussions). *Journal of Financial Econometrics 2,* 1-48.

[31] Barndorff-Nielsen, O.E., and N. Shephard (2004b). Econometric analysis of realized covariation: high frequency based covariance, regression, and correlation in financial economics. *Econometrica 72*, 885-925.

[32] Barndorff-Nielsen, O.E., and N. Shephard (2007). Variation, jumps, market frictions and high-frequency data in financial econometrics. In R. Blundell, P. Torsten and W. K. Newey (Eds.) *Advances in Economics and Econometrics. Theory and Applications.* Ninth World Congress. Econometric Society Monographs, Cambridge University Press.

[33] Barndorff-Nielsen, O.E., N. Shephard, and M. Winkel (2006). Limit theorems for multipower variation in the presence of jumps in financial econometrics. *Stochastic processes and their applications, forthcoming.*

[34] Barndorff-Nielsen, O.E., S.E. Graversen, J. Jacod, and N. Shephard (2006). Limit theorems for realized bypower variation in econometrics. *Econometric Theory 22*, 677-719.

[35] Barndorff-Nielsen, O.E., P. Hansen, A. Lunde and N. Shephard (2005). Regular and modified kernel-based estimators of integrated variance: the case with independent noise. *Working paper.*

[36] Barndorff-Nielsen, O.E., P. Hansen, A. Lunde and N. Shephard (2006). Designing realized kernels to measure ex-post variation of equity prices in the presence of noise. *Working paper.*

[37] Barucci, E. and R. Renò (2002a). On measuring volatility and the GARCH forecasting performance. *Journal of Internation Financial Markets, Institutions and Money 12,* 182-200.

[38] Barucci, E. and R. Renò (2002b). On measuring volatility of diffusion processes with high-frequency data. *Economic Letters 74,* 371-378.

[39] Bessembinder, H. (2003). Issues in assessing trade execution costs. *Journal of Financial Markets 6*, 233-257.

[40] Beveridge, S. and C. Nelson (1981). A new approach to the decomposition of economic time series into permanent and transitory components with particular attention to the measurement of the "Business Cycle." *Journal of Monetary Economics 7,* 151-174.

[41] Bollen, B., and B. Inder (2002). Estimating daily volatility in financial markets utilizing intraday data. *Journal of Empirical Finance 9*, 551-562.

[42] Brandt, M.W. and F.X. Diebold (2006). A no-arbitrage approach to range-based estimation of return covariances and correlations. *Journal of Business 79,* 61-74.

[43] Brennan, M.J., and A. Subrahmanyam (1996). Market microstructure and asset pricing: on the compensation for illiquidity in stock returns. *Journal of Financial Economics 41*, 441-464.

[44] Christensen, K. and M. Podolskij (2005). Asymptotic theory for range-based estimation of integrated volatility of a continuous semimartingale. *Working paper.*

[45] Chung, K.L., and R.J. Williams (1990). *Introduction to Stochastic Integration.* Second Edition. Birkhäuser.

[46] Cohen, K.J., G.A. Hawanini, S.F. Maier, R.A. Schwartz, and D.K. Whitcomb (1983). Friction in the trading process and the estimation of systematic risk. *Journal of Financial Economics 12*, 263-278.

[47] Corradi, V., and W. Distaso (2006). Semiparametric comparision of stochastic volatility models via realized measures, *Review of Economic Studies, forthcoming.*

[48] Corsi (2003). A simple long-memory model of realized volatility. *Working paper.*

[49] Datar, V.T., N.Y. Naik and R. Radcliffe (1998). Liquidity and stock returns: an alternative test. *Journal of Financial Markets 1*, 203-219.

[50] De Pooter, M., M. Martens, and D. Van Dijk (2006). Predicting the daily covariance matrix for S&P100 stocks using intraday data: but which frequency to use? *Econometric Reviews, forthcoming.*

[51] Dimson, E., 1979. Risk management when shares are subject to infrequent trading. *Journal of Financial Economics 7*, 197-226.

[52] Duffie, D. (1992). *Dynamic Asset Pricing Theory.* Princeton University Press. Princeton.

[53] Easley, D., and M. O'Hara (2002). Microstructure and asset pricing. In G. Constantinides, M. Harris, and R. Stulz (Eds.). *Handbook of Financial Economics*. Elsevier, North Holland.

[54] Easley, D., S. Hvidkjaer, and M. O'Hara (2002). Is information risk a determinant of asset returns? *Journal of Finance 57*, 2185-2221.

[55] Easley, D., N. Kiefer, M. O'Hara and J. Paperman (1996). Liquidity, information, and infrequently traded stocks. *Journal of Finance 1,* 1405-1436.

[56] Ellis, K., R. Michaely, and M. O'Hara (2000). The accuracy of trade classification rules: evidence from Nasdaq. *Journal of Financial and Quantitative Analysis 35*, 529-552.

[57] Engle, R.F., C.H. Hong, and A. Kane (1990). Valuation of variance forecasts with simulated option markets. NBER Working paper No: 3350.

[58] Epps, T.W, (1979). Comovements in stock prices in the very short run. *Journal of the American Statistical Association 74,* 291-298.

[59] Fang, Y. (1996). Volatility modeling and estimation of high-frequency data with Gaussian noise. *Unpublished Ph.D. thesis, MIT.*

[60] Fleming, J., C. Kirby, and B. Ostdiek (2001). The economic value of volatility timing. *Journal of Finance 56*, 329-352.

[61] Fleming, J., C. Kirby, and B. Ostdiek (2003). The economic value of volatility timing using "realized volatility." *Journal of Financial Economics 67*, 473-509.

[62] Forsberg, L. and E. Ghysels (2004). Why do absolute returns predict volatility so well? *Working paper.*

[63] Ghysels, E. and A. Sinko (2006). Comment on Hansen and Lunde. *Journal of Business and Economic Statistics 24,* 192-194.

[64] Ghysels, E., P. Santa-Clara, and R. Valkanov (2006). Predicting volatility: getting the most out of return data sampled at different frequencies. *Journal of Econometrics, forthcoming.*

[65] Goncalves, S. and N. Meddahi (2004). Bootstrapping realized volatility. *Working paper.*

[66] Goncalves, S. and N. Meddahi (2006). Box-Cox transforms for realized volatility. *Working paper.*

[67] Griffin, J. and R. Oomen (2006). Covariance measurement in the presence of nonsynchronous trading and market microstructure noise. *Working paper.*

[68] Hansen, P.R., and A. Lunde (2006). Realized variance and market microstructure noise (with discussions). *Journal of Business and Economic Statistics 24,* 127-161.

[69] Hasbrouck, J. (1993). Assessing the quality of a security market: a new approach to transaction cost measurement. *Review of Financial Studies 6*, 191-212 .

[70] Hasbrouck, J. (1996). Modelling market microstructure time series. In G.S. Maddala and C.R. Rao (Eds.) *Handbook of Statistics.* Elsevier, North-Holland.

[71] Hasbrouck, J. (2003). Trading costs and returns for U.S. securities: the evidence from daily data. *Working paper.*

[72] Hasbrouck, J. and T. Ho (1987). Order arrival, quote behavior, and the return generating process. *Journal of Finance 4,* 1035-1048.

[73] Hayashi, T. and N. Yoshida (2005). On covariance estimation of non-synchronously observed diffusion processes. *Bernoulli 11*, 359-379.

[74] Hayashi, T. and N. Yoshida (2006). Estimating correlations with nonsynchronous observations in continuous diffusion models. *Working paper.*

[75] Huang, X. and G. Tauchen (2005). The relative contribution of jumps to total price variation. *Journal of Financial Econometrics 3,* 456-499.

[76] Jacod, J. (1994). Limit of random measures associated with the increments of a Brownian semimartingale. *Working paper.*

[77] Jacod, J. and P. Protter (1998). Asymptotic error distributions for the Euler method for stochastic differential equations. *Annals of Probability 26,* 267-307.

[78] Kalnina, I. and O. Linton (2006). Estimating quadratic variation consistently in the presence of correlated measurement error. *Working paper.*

[79] Kanatani, T. (2004a). High-frequency data and realized volatility. *Unpublished Ph.D thesis, Kyoto University.*

[80] Kanatani, T. (2004b). Integrated volatility measuring from unevenly sampled observations. *Economics Bulletin 3,* 1-8.

[81] Large (2006). Estimating quadratic variation when quoted prices change by constant increments. *Working paper.*

[82] Lee, C. and M. Ready (1991). Inferring trade direction from intraday data. *Journal of Finance 46*, 733-746.

[83] Lo, A., and A.C. MacKinlay (1990). An econometric analysis of nonsynchronous trading. *Journal of Econometrics 45*, 181-212.

[84] Malliavin, P. and M.E. Mancino (2002). Fourier series method for measurement of multivariate volatilities. *Finance and Stochastics 6,* 49-61.

[85] Mancini, C. (2003). Statistics of a Poisson-Gaussian process. *Working paper.*

[86] Mancini, C. (2004). Estimation of the characteristics of jump of a general Poisson-diffusion process. *Scandinavian Actuarial Journal 1,* 42-52.

[87] Martens, M. (2005). Estimating unbiased and precise realized covariances. *Working paper.*

[88] Martens, M., and D. Van Dijk (2005). Measuring volatility with the realized range. *Working paper.*

[89] McAleer, M., and M. Medeiros (2006). Realized volatility: A review. *Working paper.*

[90] Meddahi, N. (2002). A theoretical comparison between integrated and realized volatility. *Journal of Applied Econometrics 17*, 475-508.

[91] Moise, C. (2004). Stochastic volatility risk and the size anomaly. *Working paper.*

[92] Müller, U.A. (1993). Statistics of variables observed over overlapping intervals. *Working paper.*

[93] Mykland, P.A., and L. Zhang (2005). ANOVA for diffusions. *Annals of Statistics, forthcoming.*

[94] Newey, W. and K. West (1987). A simple positive semi-definite, heteroskedasticity and autocorrelation consistent covariance matrix. *Econometrica 55*, 703-708.

[95] Nielsen, M.O., and P.H. Frederiksen (2005). Finite sample accuracy of integrated volatility estimators. *Working paper.*

[96] O'Hara, M. (1995). *Market Microstructure Theory.* Blackwell Publishers Ltd, Oxford.

[97] Oomen, R.C.A. (2005). Properties of bias-corrected realized variance under alternative sampling schemes. *Journal of Financial Econometrics 3,* 555–577.

[98] Oomen, R.C.A. (2006). Properties of realized variance under alternative sampling schemes. *Journal of Business and Economic Statistics 24,* 219-237.

[99] Parkinson, M. (1980). The extreme value method for estimating the variance of the rate of return. *Journal of Business 53*, 61-66.

[100] Pástor, L. and R. F. Stambaugh (2003). Liquidity risk and expected stock returns. *Journal of Political Economy 111*, 642-685.

[101] Perold, A (1988). The implementation shortfall. *Journal of Portfolio Management 14*, 4-9.

[102] Peterson, M. and E. Sirri (2002). Evaluation of the biases in execution cost estimates using trade and quote data. *Journal of Financial Markets 6*, 259-280.

[103] Phillips, P.C.B. and J. Yu (2005). Comment on Hansen and Lunde. *Journal of Business and Economic Statistics 24,* 202-208.

[104] Politis, D., J.P. Romano, and M. Wolf (1999). *Subsampling.* New York: Springer.

[105] Press, J.S. (1967). A compound events model for security prices. *Journal of Business 40*, 317-335.

[106] Rogers, L.C.G. and S.E. Satchell (1991). Estimating variance from high, low, open, and close prices. *Annals of Applied Probability 1*, 504-512.

[107] Roll, R. (1984). A simple measure of the effective bid-ask spread in an efficient market. *Journal of Finance 39*, 1127-1139.

[108] Scholes, M., and J. Williams (1977). Estimating betas from nonsynchronous data. *Journal of Financial Economics* 5, 309-327.

[109] Sheppard, K. (2006). Realized covariance and scrambling. *Working paper.*

[110] Stoll, H.R. (1989). Infering the components of the bid-ask spread: theory and empirical evidence. *Journal of Finance 44,* 115-134.

[111] Stoll, H.R. (2000). Friction. *Journal of Finance 55*, 1479-1514.

[112] Sun, Y. (2006). Best quadratic unbiased estimators of integrated variance. *Working paper.*

[113] Voev, V., and A. Lunde (2006). Integrated covariance estimation using high-frequency data in the presence of noise. *Working paper.*

[114] Watson, M. (1986). Univariate detrending methods with stochastic trends. *Journal of Monetary Economics 18,* 49-75.

[115] West, K.D., H.J. Edison, and D. Cho (1993). A utility-based comparison of some models of exchange rate volatility. *Journal of International Economics 35*, 23-45.

[116] Woerner, J. (2006). Power and multipower variation: inference for high-frequency data. In A.N. Shiryaev, M. do Rosario Grossinho, P. Oliviera, and M. Esquivel (Eds.) *Proceedings of the International conference on Stochastic Finance 2004.* Springer-Verlag.

[117] Zhang, L., P. Mykland, and Y. Aït-Sahalia (2005). A tale of two time scales: determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association 100,* 1394-1411.

[118] Zhang, L. (2006a). Efficient estimation of stochastic volatility using noisy observations: a multi-scale approach. *Bernoulli, forthcoming.*

[119] Zhang, L. (2006b). Estimating covariation: Epps effect, microstructure noise. *Working paper.*

[120] Zhou, B. (1996). High-frequency data and volatility in foreign-exchange rates. *Journal of Business and Economic Statistics 14*, 45-52.

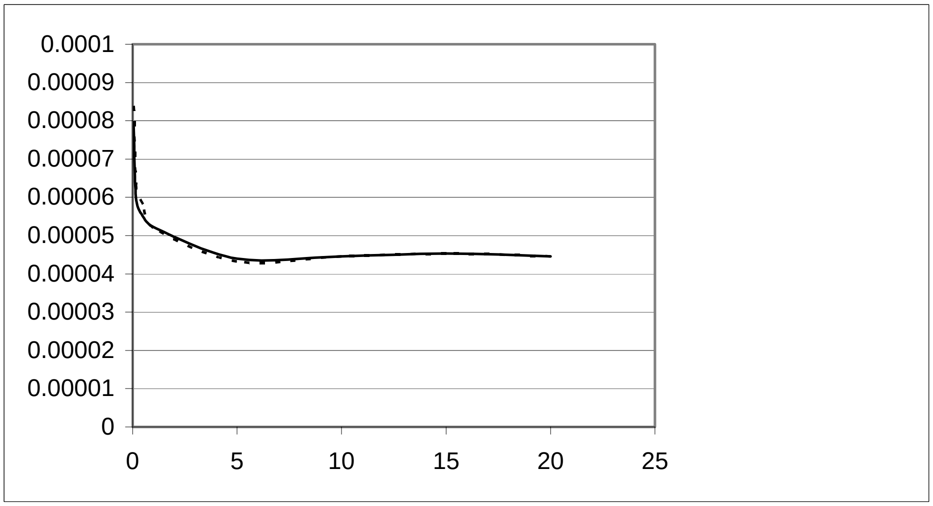

Figure 1. "Volatility signature plots" for IBM using mid-quotes from (i) the NYSE only (solid line) and (ii) the NYSE and the Midwest exchange (dotted line). We plot realized variance as a function of the sampling frequency (in minutes). The data are collected for the month of February 2002 using the filter discussed in Bandi and Russell (2006a).

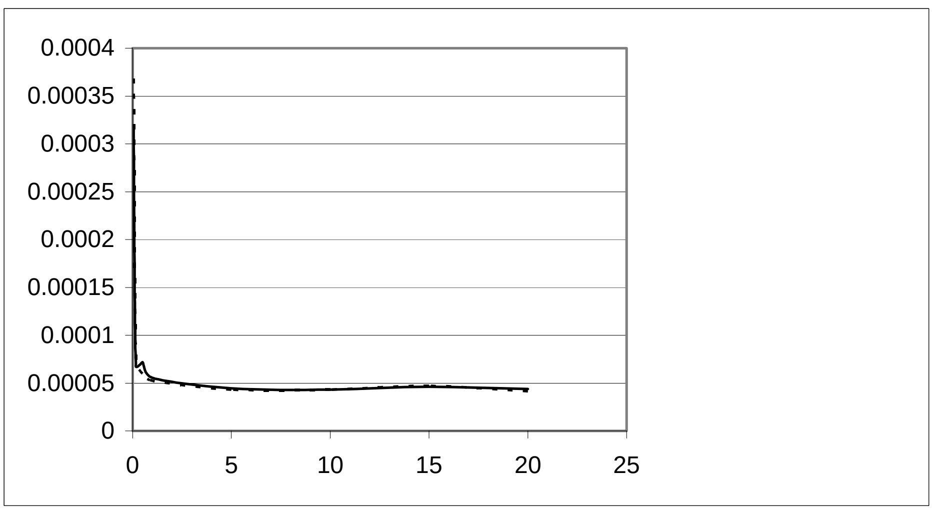

Figure 2. "Volatility signature plots" for IBM using mid-quotes from (i) the NYSE and NASDAQ (solid line) and (ii) the consolidated market (dotted line). We plot realized variance as a function of the sampling frequency (in minutes). The data are collected for the month of February 2002 using the filter discussed in Bandi and Russell (2006a).

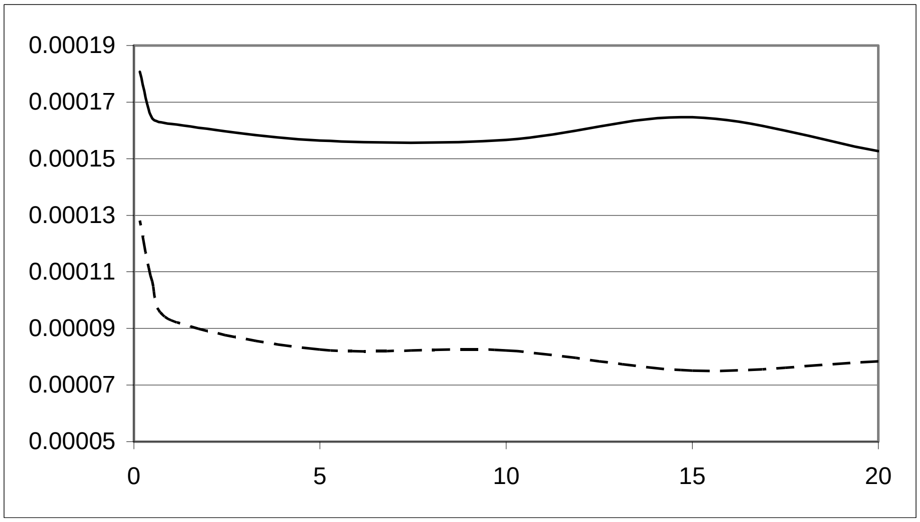

Figure 3. "Volatility signature plots" for Cisco Systems (dotted line) and Microsoft (dashed line). We plot realized variance as a function of the sampling frequency (in minutes). The data are mid-quotes collected for the month of February 2002. We use the filter discussed in Bandi and Russell (2006a).

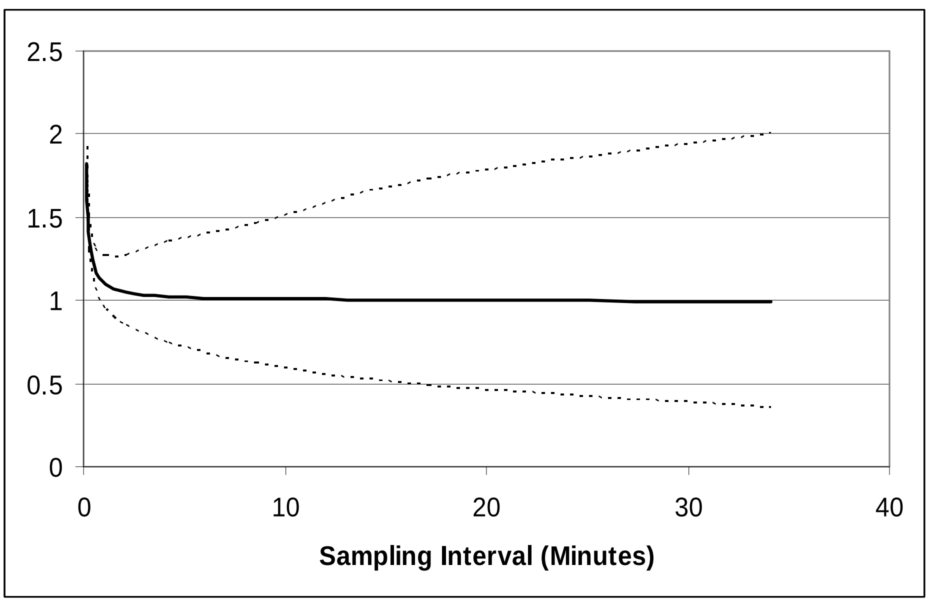

Figure 4. Simulated "volatility signature plot" from a stochastic volatility diffusion model with parameter values consistent with IBM (see Bandi and Russell, 2003, 2006a, for details). The solid line is the average (across simulations) of the realized variance estimates for each sampling interval (in minutes). The dotted lines are 95% empirical intervals from the simulations. The true integrated variance is standardized to 1.