From this, it follows that

$$\begin{aligned}\gamma_{1,2}(h) &= (1-e^{-\lambda_1\delta})^2(1-e^{-\lambda_2\delta})e^{-\lambda_1\delta(h-1)} \\ &\quad\times \Pr\{z_2<z_1|z_1,z_2<\delta\}E(z_1-z_2|z_2<z_1<\delta) \\ &\quad+(1-e^{-\lambda_1\delta})(1-e^{-\lambda_2\delta})e^{-\lambda_1\delta h} \\ &\quad\times \bigl(E(\min\{z_1,z_2\})+\delta-E(z_2|z_2<\delta)\bigr).\end{aligned}$$

The expectations on the second line are straightforward, that is,

$$E_\lambda(\min\{z_1,z_2\})=\frac{1}{\lambda_1+\lambda_2}$$

and

$$E_\lambda(z_2|z_2<\delta)=\frac{1-e^{-\lambda_2\delta}(1+\lambda_2\delta)}{(1-e^{-\lambda_2\delta})\lambda_2}.$$

To work out $E(z_1-z_2|z_2<z_1<\delta)$, we proceed as follows:

$$p(z_1,z_2|z_1,z_2<\delta)=\frac{\lambda_1\lambda_2 e^{-\lambda_1 z_1}e^{-\lambda_2 z_2}}{(1-e^{-\lambda_1\delta})(1-e^{-\lambda_2\delta})},\qquad z_1,z_2<\delta.$$

Integrating the foregoing density over z_2 from 0 to z_1, and then over z_1 from 0 to δ gives

$$\Pr(z_2<z_1|z_1,z_2<\delta)=\frac{1-e^{-\lambda_1\delta}-\frac{\lambda_1}{(\lambda_1+\lambda_2)}(1-e^{-(\lambda_1+\lambda_2)\delta})}{(1-e^{-\lambda_1\delta})(1-e^{-\lambda_2\delta})},$$

as required before. Next,

$$\begin{aligned}p(z_1,z_2|z_2<z_1<\delta) &= \frac{p(z_1,z_2|z_1,z_2<\delta)}{\Pr(z_2<z_1|z_1,z_2<\delta)} \\ &= \omega^{-1}\lambda_1\lambda_2 e^{-\lambda_1 z_1}e^{-\lambda_2 z_2}, \qquad (A.1)\end{aligned}$$

where $\omega=1-e^{-\lambda_1\delta}-\frac{\lambda_1}{(\lambda_1+\lambda_2)}(1-e^{-(\lambda_1+\lambda_2)\delta})$. From this, we then have

$$\begin{aligned}&E(z_1-z_2|z_2<z_1<\delta) \\ &\quad=\omega^{-1}\int_0^\delta\int_0^{z_1}(z_1-z_2)\lambda_1\lambda_2 e^{-\lambda_1 z_1}e^{-\lambda_2 z_2}\,dz_2\,dz_1 \\ &\quad=\omega^{-1}\bigl\{\lambda_2^2-\lambda_1^2 e^{-(\lambda_1+\lambda_2)\delta} \\ &\qquad\qquad -e^{-\lambda_1\delta}(\lambda_1+\lambda_2)(\lambda_1\lambda_2\delta-\lambda_1+\lambda_2)\bigr\} \\ &\qquad\times\bigl\{\lambda_1\lambda_2(\lambda_1+\lambda_2)\bigr\}^{-1}.\end{aligned}$$

Collecting the foregoing terms gives the required result.

ADDITIONAL REFERENCES

Bandi, F. M., and Russell, J. R. (2004), "Microstructure Noise, Realized Variance, and Optimal Sampling," unpublished manuscript, The University of Chicago, Graduate School of Business.

Epps, T. W. (1979), "Comovements in Stock Prices in the Very Short Run," *Journal of the American Statistical Association*, 74, 291–298.

Griffin, J. E., and Oomen, R. C. (2005a), "Realized Covariance Measurement in the Presence of Non-Synchronous Trading and Market Microstructure Noise," unpublished manuscript, University of Warwick, Warwick Business School.

——— (2005b), "Sampling Returns for Realized Variance Calculations: Tick Time or Transaction Time?" unpublished manuscript, University of Warwick, Warwick Business School.

Niederhoffer, V., and Osborne, M. F. M. (1966), "Market Making and Reversal on the Stock Exchange," *Journal of the American Statistical Association*, 61, 897–916.

Comment

Peter C. B. PHILLIPS
Cowles Foundation for Research in Economics, Yale University, New Haven, CT 06520,
University of Auckland, and University of York (*peter.phillips@yale.edu*)

Jun YU
School of Economics and Social Sciences, Singapore Management University, Singapore 178903
(*yujun@smu.edu.sg*)

1. INTRODUCTION

With the availability of ultra-high–frequency financial data, the task of finding an appropriate econometric model to describe the movement of financial variables at the tick-by-tick level has become an important goal in financial econometric research. The task has both theoretical and empirical dimensions. From an empirical perspective, the near-continuous recording of financial asset prices has opened up the intriguing possibility of fitting the quadratic variation process empirically, leading to what is possibly the most direct nonparametric measure of asset price volatility. The resulting quantity, known in the financial econometrics literature as realized variance (RV), measures the accumulated or integrated variance (IV) of the efficient price process from some given initialization. This quantity is now the focal point of much of the latest research on market volatility.

Compared with parametric methods of measuring volatility, this nonparametric approach basically trades off efficiency in exchange for robustness to specification bias. The theoretical justification of RV as a measure of volatility comes directly from standard stochastic process theory, according to which the empirical quadratic variation converges to IV as the infill-sampling frequency goes to zero. The empirical method inspired by this convergence has become popular only recently with the availability of ultra-high–frequency data, but the idea has been around for a long time, as indicated by Hansen and

Journal of Business & Economic Statistics
April 2006, Vol. 24, No. 2
DOI 10.1198/073500106000000116

Lunde (hereafter HL). In particular, it has been discussed by econometricians working with continuous-time models, an example being Maheswaran and Sims (1993).

Just as these exciting empirical possibilities have been recognized, a number of practical issues have arisen that challenge the suitability of conventional model specifications on which empirical quadratic variation measures depend. For example, although conventional wisdom may accept that efficient market prices can be well described by time-homogeneous continuous-time jump-diffusion models at daily or lower observation frequencies, it is well understood that at higher frequencies these models are usually too simplistic, and at ultra-high frequencies the presence of microstructure noise is a compelling complication that affects the dynamic properties of market prices and distorts empirical quadratic variation measures.

In practice, therefore, as some recently articulated arguments have emphasized, one should be careful in pushing the infill-sampling frequency to the limit, even though this is precisely what stochastic process theory would suggest in the ideal environment where the efficient price is observed. Indeed, the existence of market microstructure noise means that empirical quadratic variation measures are themselves contaminated with noise at high observation frequencies. Rather unsurprisingly, in the presence of noise, consistent estimation of IV inevitably depends on modifications to the empirical quadratic variation that take into account the dynamic structure of the market microstructure noise. In this regard, some recognition of the properties of microstructure noise in the data is desirable in designing the modifications. Because both the microstructure noise and the efficient price are latent variables, direct measurement of the noise is not possible, and thus the empirical and theoretical modeling issues involve some subtleties. Paramount among these is that careful attention to specification is required to achieve identification and the empirical separation of noise from IV measures. Moreover, it should be acknowledged that modifications to the empirical IV measures will usually stem directly from modeling assumptions made about the form of the latent microstructure noise. Accordingly, the resulting estimates may not be robust to relaxation of these assumptions.

HL's article provides useful theoretical analysis of the finite-sample and asymptotic properties of several IV measures and makes interesting empirical contributions to this emerging literature. It also documents some stylized facts about market microstructure noise and its relation to the hypothesized underlying efficient price process. These stylized facts are useful in the development of consistent and more efficient nonparametric and parametric estimators of IV. The authors' contribution to this literature is therefore most welcome.

In developing its theoretical and empirical results, HL impose a stationarity condition on market microstructure noise. This assumption has been used in other ongoing work in the field. Prima facie, this assumption seems quite reasonable, but we question its suitability. Some of HL's results are based on more specialized cases, such as pure microstructure noise and moving average noise. As in much earlier work, the signature plot is the main graphical tool used here to assess the validity of these microstructure noise assumptions, but conventional ACF and PACF plots are also used. HL focus on the bias in IV estimation induced by microstructure noise and provide useful methods for correcting for this bias. But because the corrections do not always retain positivity in the estimates, a clear practical recommendation does not emerge from their analysis.

In commenting on HL's article we begin by focusing on the modeling assumptions used to achieve microstructure noise separation and the identification of IV, giving particular attention to the advantages, limitations, and suitability of the commonly used pure microstructure noise assumption and its stationary extensions. Second, we discuss the use of the signature plots used by HL and other authors as a graphical diagnostic, and suggest an alternative graphical tool called the microstructure noise function that we are using in our own ongoing work (Phillips and Yu 2005a,b). The microstructure noise function has some advantages as a graphical device over the signature plot as a noise diagnostic and lends itself to nonparametric measurement. The final part of our comment outlines a new approach that enables us to study IV and microstructure noise in a panel regression framework. This approach provides a mechanism for analyzing and removing the effects of microstructure noise in a nonparametric way while still treating IV in a general way through the presence of a fixed effect.

2. STATIONARY AND PURE NOISE ASSUMPTIONS

Let $p^*(t)$, $p(t)$, and $u(t)$ denote the latent log-efficient price process, the observed log-price process, and the noise process. The time interval is standardized to [0, 1] and is partitioned into a grid of m subintervals as $\mathcal{G}_m = \{0 = t_{0,m}, t_{1,m}, \dots, t_{m,m} = 1\}$. The sampling interval on this grid is the mesh size $\Delta = \min_i |t_{i,m} - t_{i-1,m}|$, and for equispaced observations, we have $\Delta = t_{i,m} - t_{i-1,m} = \frac{1}{m}$ for all i.

Like other authors, HL assume that the data-generating mechanism in continuous time is given by the system

$$p(t) = p^*(t) + u(t), \tag{1}$$

$$dp^*(t) = \mu(t)\,dt + \sigma(t)\,dB(t), \tag{2}$$

where the efficient price $p^*(t)$ is a latent unobserved variable. A key assumption of HL (assumption 2) is that $u(t)$ is covariance stationary. Two subcases are given special attention: the pure noise assumption (assumption 3) and the moving average assumption (assumption 4).

The apparent advantage of the stationarity and pure noise assumptions is that they substantially simplify econometric analysis. To see this, first consider the pure noise assumption where $u(t)$ is taken to be iid $(0, \omega^2)$ over all grids such as $\mathcal{G}_m$ and also to be independent of $p^*(t)$. A direct consequence of this assumption is that the conditional expectation of the RV measure has the following very simple expression:

$$E_*\big(RV^{(m)}\big) = E\big(RV^{(m)}\big|\{p^*(t)\}_0^1\big) = IV + 2m\omega^2, \tag{3}$$

where $IV\ (= \int_0^1 \sigma^2(t)\,dt)$ is the integrated variance of the efficient price and $RV^{(m)} = \sum_{i=1}^m y_{i,m}^2$ with $y_{i,m} = p_{i,m} - p_{i-1,m}$ and $p_{i,m} = p(t_{i,m})$. Graphical plotting of the quantity $RV^{(m)}$ against the sampling interval Δ produces a volatility signature plot. For (3), it is apparent that the curve is simply a reciprocal function that has the simple form

$$E_*\big(RV^{(m)}\big) = IV + 2m\omega^2 = IV + \frac{2\omega^2}{\Delta} \tag{4}$$

for equispaced sampling, so that $E_*(RV^{(m)})$ asymptotes as $\Delta \to 0$. It is this asymptotic behavior in the neighborhood $\Delta \sim 0$ that is the characteristic feature of the signature plot, and empirical behavior of this type is often taken as confirmatory evidence that the pure noise model (1) and its central implication (3) conform to the observed data.

A further implication of (3) is that the second term dominates as $m \to \infty$ and $\Delta \to 0$. Thus, in this model, noise dominates as the infill sampling frequency increases and the quantity $RV^{(m)}/(2m)$ delivers a consistent estimate of ω^2, the noise variance. Moreover, because the pure noise assumption induces a unit root MA(1) structure on the return noise, the bias-corrected estimator,

$$RV_{\mathrm{AC}_1}^{(m)} = \sum_{i=1}^{m} y_{i,m}^2 + \sum_{i=1}^{m} y_{i,m} y_{i-1,m} + \sum_{i=1}^{m} y_{i,m} y_{i+1,m},$$

is naturally unbiased. On the other hand, when the noise follows a more general MA process, the return noise has a higher-order MA structure. This linkage explains the rationale for the use of a higher-order bias-corrected estimator, such as the one defined in HL's theorem 2.

Although the pure noise assumption clearly facilitates analysis, it suffers from the unhappy drawback that pure noise lacks physical realism in continuous time, involving a degree of instantaneous variability that is unimaginable in a physically realizable process. In fact, such a process in continuous time needs to be rigorously modeled as a generalized stochastic process, and physical realizations occur only in the form of linear functionals, such as temporal averages. Indeed, Gaussian pure noise may be interpreted as the derivative of Brownian motion, which does not exist as a conventional stochastic process, and realizable temporal averages of such processes take the form of stochastic integrals, such as $\int \psi(t-s)\, dW(s)$ for some Brownian motion W and weighting (or test) function ψ. Clearly, the mathematical form of (3) is a direct artifact of the pure noise assumption and corresponds with the fact that the quadratic variation of a pure noise process is infinite, so that the second term of (3) dominates as $m \to \infty$.

Notwithstanding the foregoing remarks, it is quite possible to proceed with a rigorous and physically realizable treatment under a pure noise assumption when attention is confined to a discrete time finite grid such as $\mathcal{G}_m$. In such a case, the realized microstructure noise process $u(t_{i,m})$ is a discrete-time iid process. Such an approach makes it possible to rigorously define the manner in which pure microstructure noise enters the system and indeed to study properties in the limit as $m \to \infty$. But although such an approach neatly finesses the mathematical difficulty of having to define pure noise in continuous time as a generalized process, the essential implication that the quadratic variation is infinite in the limit and that the second term of (3) dominates as $m \to \infty$ remains valid. The remaining issue of importance is whether such behavior is supported by empirical observation and the realities of practical trading in financial markets. We now turn to this issue.

It is our contention that pure microstructure noise and stationary microstructure noise both lack realism with the data. To illustrate, we make use of the dataset used by HL in their figure 2. The evidence from this data clearly contradicts the assumptions of stationary noise and pure noise. In our Figure 1

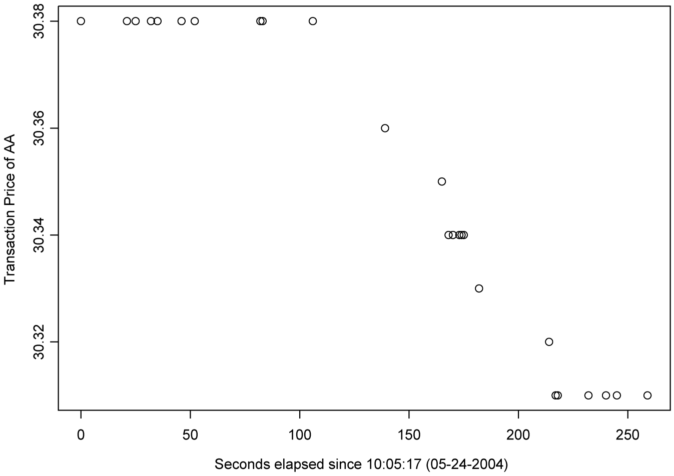

Figure 1. Time Series Plot of 33 Transactions That Occurred After 10:05:17 on May 24, 2004 for AA.

we plot 33 transactions that occurred after 10:05:17 on May 24, 2004 for Alcoa (AA) on the NYSE. This covers a sample period of just over 4 minutes of calendar time and is a typical segment of a trading day. It can be seen that the plot starts with a nearly 2-minute period of a flat transaction price, in which 11 transactions were recorded, and finishes with a 40-second period of another flat transaction price, in which 6 transactions were recorded. Such an observation record with periods of flat price trading is completely inconsistent with both stationary noise and iid noise at the tick-by-tick level. In contrast, this observation record is compatible with nonstationary noise, as we now explain.

When the efficient price follows a Brownian semimartingale as in (2) and the observed trading price is related to the efficient price according to (1), both conventional assumptions in this literature, then during the flat pricing period the microstructure noise must itself completely offset the efficient price process to produce a sustained flat transactions price. The noise process thus must inherit the same martingale-like behavior as the efficient price in continuous time over this subinterval. Consequently, the microstructure noise will be locally nonstationary and will have unit negative correlation with the efficient price process. These characteristics, it hardly needs to be said, are strongly at variance with pure noise or stationary noise assumptions.

As evident in column 4 of their table 1, HL also notice that observed prices normally display many spells of constancy within a trading day. So Figure 1 is quite typical. More particularly, for the DJIA stocks traded on the NYSE, the percentage of observations for which the transaction price was the same as the previous price ranges from 39% to 71%. The constancy is even more pronounced in mid-quotes. For example, for the DJIA stocks traded on the NYSE, the percentage of observations for which the quote price was the same as the previous price ranges from 45% to 87%.

These numbers are too substantial to be ignored in modeling the price process, yet it is now common practice to ignore this constancy in observed prices when working with ultra-high–frequency data and building models of microstructure noise. Ignoring constancy in trading prices inevitably induces specification error and can explain some of the anomalies that have arisen in applied work. In particular, when allowance is made for the local nonstationarity induced by constant price trading, signature plots no longer have the simple reciprocal function form of (4). In fact, local nonstationarity in microstructure noise can even induce nonmonotonicity in the signature plot, a feature that HL found in their own data analysis. In general, the larger the percentage of flat price trading, the further the departure from the reciprocal curve for the signature plot. This property may also help explain why signature plots calculated from the quote price are less like a reciprocal curve than those calculated from the transaction price.

If the locally nonstationary noise property is taken into account, then none of the bias-correction procedures examined by HL or those suggested elsewhere in the literature deliver an unbiased or consistent estimator of IV. To illustrate, consider the following simple Bernoulli model of trading:

$$p_{i,m} = \begin{cases} p^*_{i,m} + \epsilon_{i,m} & \text{with probability } \pi \\ p_{i-1,m} & \text{with probability } 1-\pi, \end{cases} \tag{5}$$

where $\epsilon_{i,m} \sim \text{iid}(0, \sigma_\epsilon^2)$, $p^*_{i,m}$ follows a random walk or a local-to-unity discrete time diffusion, and the process is initialized at $i = 0$ with $p_{0,m} = p^*_{0,m} = O_p(1)$. This model allows for flat trading with a constant probability of $1 - \pi$ and efficient price plus pure noise trading with probability π. Thus, when $\pi \in (0, 1)$, there is a positive probability of flat trading at each point on the temporal grid. Using conventional indicator notation, we may write the trading price in the form

$$p_{i,m} = (p^*_{i,m} + \epsilon_{i,m})\mathbb{1}_{(\zeta_i=1)} + p_{i-1,m}\mathbb{1}_{(\zeta_i=0)}, \tag{6}$$

where ζ_i is a Bernoulli variable that is unity with probability π and zero with probability $1 - \pi$. If $\pi = 1$, then $\mathbb{1}_{(\zeta_i=0)} = 0$ a.s., and the model reduces to the efficient price plus pure noise model. But for $\pi \in (0, 1)$, we have

$$p_{i,m} = p^*_{i,m}\mathbb{1}_{(\zeta_i=1)} + u_{i,m}, \qquad \text{with } u_{i,m} = \epsilon_{i,m}\mathbb{1}_{(\zeta_i=1)} + p_{i-1,m}\mathbb{1}_{(\zeta_i=0)}, \tag{7}$$

and it is clear that the "implied noise" process $u_{i,m}$ depends on $p_{i-1,m}$. In view of (6), we have the explicit expression

$$u_{i,m} = \sum_{j=0}^{i} p^*_{i-j,m}\{\mathbb{1}_{(\zeta_{i-j}=1)} - \delta_{0j}\}\prod_{k=0}^{j-1}\mathbb{1}_{(\zeta_{i-k}=0)} + \sum_{j=0}^{i}\epsilon_{i-j,m}\mathbb{1}_{(\zeta_{i-j}=1)}\prod_{k=0}^{j-1}\mathbb{1}_{(\zeta_{i-k}=0)},$$

where $\delta_{0j} = 1, 0$ as $j = 0, \neq 0$, and where we use the convention $\prod_{k=0}^{-1} = 1$. In this case the implied noise process $u_{i,m}$ clearly depends on the entire past history of shocks and efficient price realizations $\{\epsilon_{i-j,m}, p^*_{i-j-1,m}; j = 0, 1, \ldots, i\}$, and so assumption 4 of HL fails because $u_{i,m}$ has nonzero autocorrelations at all lags and inherits some persistence characteristics of the efficient price process. Consequently, the conventional kernel-based bias-correction procedures considered by HL will fail to correct for the induced serial correlation of $u_{i,m}$ and its first difference $u_{i,m} - u_{i-1,m}$, leading to biased estimates of IV.

Of course, discarding flat prices is one way to circumvent the problem of nonstationarity. However, by doing so, one must discard a very large amount of data. For example, from the data given in the HL's table 1, for the quote price of Philip Morris (MO) in 2000, one must throw away 87% of the data. Much of the present research on this topic, including that of Aït-Sahalia, Mykland, and Zhang (2005a,b) and Zhang, Mykland, and Aït-Sahalia (2005), is motivated by the desire to utilize all available high-frequency data rather than discard data to avoid difficulties stemming from microstructure noise. The problems outlined here fall very much within this category, and it seems highly desirable to seek better models of microstructure noise that accord with the observed transactions data and to use these models and all of the available data to provide better solutions to the problem of estimating the IV functional of the efficient price. In our work that is now underway (Phillips and Yu 2005a,b), we are attempting to move in that direction.

3. SIGNATURE PLOTS

Signature plots, which depict RV as a function of the infill sampling frequency (again, conveniently assuming an equidistant sampling scheme), have proved to be an effective and popular tool for assessing the degree of bias induced by microstructure noise. Without noise, one would expect a relatively flat curve corresponding to the convergence of empirical quadratic variation as $\Delta \to 0$. So any departure from a flat curve is taken as an indicator of the presence of microstructure noise. Under the pure noise assumption, the signature plots should diverge as $\Delta \to 0$ at the rate $O(\Delta^{-1})$ as in (4). Consequently, several authors have used the apparent explosive pattern (as $\Delta \to 0$) in observed signature plots as strong evidence in support of the pure noise assumption.

In constructing signature plots, it is conventional to aggregate data over a period, such as a month. Figure 2 shows the signature plot of AA from the consolidated market for May 2004, thereby giving an aggregate representation for the same month that includes the transactions data for May 24 that were plotted

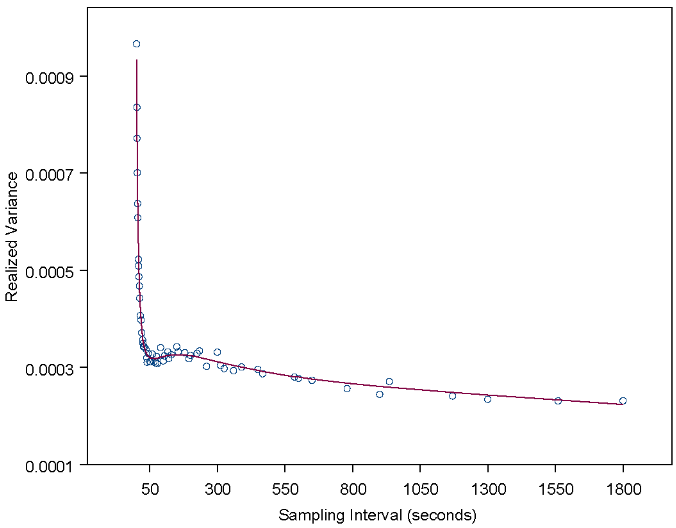

Figure 2. Signature Plot for AA From the Consolidated Market. The horizontal axis is the sampling interval ranging from 1 second to 1,800 seconds. The vertical axis is the averaged realized variance across all trading days in May 2004. Superimposed is the variable span smoother developed by Friedman (1984) and implemented by the S–PLUS command supsmu.

earlier in Figure 1. The explosive pattern in the signature plot of Figure 2 stands out in a dramatic way, although it is important to note that the vertical axis volatility measurements are very small numbers.

Although an explosive pattern (as $\Delta \to 0$) in a signature plot such as Figure 2 is sufficient to refute the noise-free assumption, it does not necessarily imply the presence of pure microstructure noise. This is because under the pure noise assumption, signature plots should behave specifically as reciprocal curves in Δ, as manifested in (4). Unfortunately, it is difficult to assess the rate of divergence through visual inspection of a signature plot such as our Figure 2. For this reason, we argue that signature plots are not effective graphical tools for checking the validity of specific microstructure noise assumptions like pure noise.

In ongoing work (Phillips and Yu 2005a), we suggest an alternative graphical method that is better suited for this purpose—a direct plot of the microstructure noise functions, where RV is treated as a function of the number of observations in a given sampling time frame. In such noise functions, shape characteristics are more evident as the number (m) of infill observations increases. Pure noise manifests as a linear relation in m as in (3), and departures from pure noise show up simply as nonlinearities or, more specifically, as concave shapes. Using DJIA transaction prices (from both a single market and the consolidated market), our estimates of microstructure noise functions have always turned out to be decisively concave, despite being monotonically increasing. This concavity indicates that the pure noise assumption is altogether too strong for most stock price data. Another advantage of plotting microstructure noise functions is that there is no need to do equidistant sampling.

Figure 3 shows the microstructure noise function for the same data as Figure 2. The departure from linearity is very evident in this plot and the (super smoother) fitted curve.

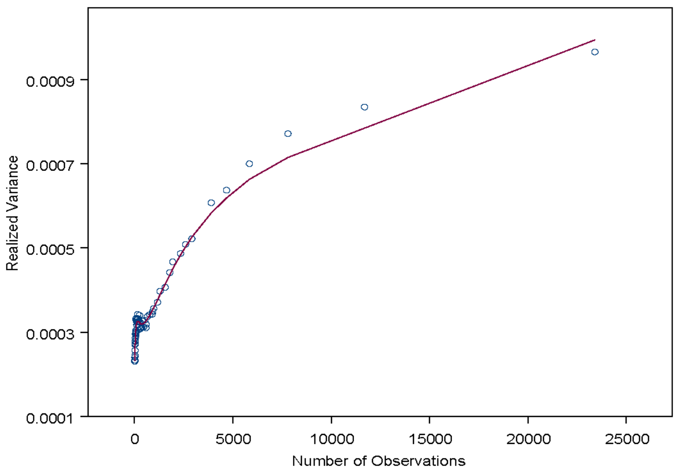

Figure 3. Market Microstructure Noise Function for AA From the Consolidated Market. The horizontal axis is the number of prices used to construct the realized variance. The vertical axis is the averaged realized variance across all trading days in May 2004. Superimposed is the variable span smoother developed by Friedman (1984) and implemented by the S–PLUS command supsmu.

4. PANEL APPROACH

Like other methods in this literature, the bias-correction procedures suggested by HL are all time series–based approaches. There are alternative approaches to dealing with noise and estimating IV. The microstructure noise function (3) has motivated us (Phillips and Yu 2005a) to use panel data methods to model the effects of microstructure noise nonparametrically and treat IV parametrically as a fixed effect in a panel regression. These panel data methods can be used to test more specialized specifications, such as that implied by the pure noise assumption. We briefly explain this idea here.

Suppose that $\mathcal{G}_n$ is a grid containing some or all the observation points, where $n+1 = \#(\mathcal{G}_n)$. Denote the empirical quadratic variation by $y_{nd} = [p,p]_d^{\mathcal{G}_n}$, using $d = 1, \dots, D$ as the date stamp for each day's observations with D days in total, where $[\cdot,\cdot]^{\mathcal{G}_n}$ represents the empirical quadratic variation on the grid $\mathcal{G}_n$. Similarly, let $\alpha_d = \int_0^1 \sigma_{dt}^2\, dt$ denote the IV for day d, and let σ_{dt}^2 be the diffusion function of the efficient price for day d. In general, IV is a random variable and varies from day to day. Often these variations are of interest in empirical work. On the other hand, it may be reasonable to assume the pure noise variance, ω^2, is constant across days over a relatively short time interval. Under these conditions, the microstructure noise function (3) may be formulated as

$$y_{nd} = \alpha_d + 2\sigma^2 n + \varepsilon_{nd} = \alpha_d + \beta n + \varepsilon_{nd}, \tag{8}$$

where $\beta = 2\sigma^2$, d is the daily date stamp, and ε_{nd} is an induced error process of the form

$$\varepsilon_{nd} = ([p^*,p^*]_d^{\mathcal{G}_n} - [p^*,p^*]_d) + 2[p^*,\epsilon]_d^{\mathcal{G}_n} + \left([\epsilon,\epsilon]_d^{\mathcal{G}_n} - E\{[\epsilon,\epsilon]_d^{\mathcal{G}_n}\}\right). \tag{9}$$

Because $[p^*,p^*]_d^{\mathcal{G}_n} - [p^*,p^*]_d \to_p 0$ as $n \to \infty$, and the empirical covariation $[p^*,\epsilon]_{Td}^{\mathcal{G}_n}$ has mean 0 and variance asymptotically proportional to $[p^*,p^*]_d$, the dominant term of ε_{nd} is the final component of (9). This component, $[\epsilon,\epsilon]_d^{\mathcal{G}_n} - E\{[\epsilon,\epsilon]_d^{\mathcal{G}_n}\}$, has mean 0 and variance $O(n)$.

By varying the number of observations used in the construction of RV, say $n = n_1, \dots, n_N$, the formulation (8) leads to a panel data model for RV. The model is complicated by the fact that the error process,

$$\varepsilon_{n_i d} = [\epsilon,\epsilon]_d^{\mathcal{G}_{n_i}} - E\left\{[\epsilon,\epsilon]_d^{\mathcal{G}_{n_i}}\right\} + O_p(1) = O_p\left(n_i^{1/2}\right)$$

is heterogeneous and autocorrelated over n_i. Nevertheless, even crude estimation methods, such as dummy variable least squares, can be used to estimate the model and deliver a consistent estimate of the slope coefficient $2\sigma^2$, because the regressor n_i is deterministic and has a stronger signal than $\varepsilon_{n_i d}$. To ensure consistency of estimates of the fixed effects and hence of IV, some sample splitting and jackknifing methods are required, as in the work of Zhang et al. (2005) and Aït-Sahalia et al. (2005b). One advantage of (8) is that it sets out a formal framework for studying such approaches.

As discussed earlier, it seems important in practical work to allow for the fact that the microstructure noise function may be nonlinear. We may generalize (8) by formulating a nonparametric noise function and attempting to fit this function empirically.

For example, suppose that f is continuous and asymptotically homogeneous of degree γ as $n \to \infty$. Then we can formulate the noise function in standardized form as

$$\frac{y_{n_i d} - \alpha_d}{n_N^\gamma} = f\left(\frac{n_i}{n_N}\right) + \frac{\varepsilon_{n_i d}}{n_N^\gamma}, \qquad i = 1, \ldots, N, \tag{10}$$

where f is taken to be common across days and n_1 is the smallest number of observations used in the calculation of RV. It is convenient for identification purposes to normalize the function f to pass through the origin so that

$$f\left(\frac{n_1}{n_N}\right) = 0. \tag{11}$$

For example, if $n_1 = 1{,}800$, which corresponds to a 30-minute sampling frequency, then the identification condition (11) is equivalent to assuming that RV calculated at the 30-minute sampling frequency yields an unbiased estimator. This is the benchmark that HL used. Averaging across days and using (11) leads to the estimable model

$$\frac{\bar{y}_{n_i \bullet} - \bar{y}_{n_1 \bullet}}{n_N^\gamma} = f\left(\frac{n_i}{n_N}\right) + \frac{\bar{\varepsilon}_{n_i \bullet} - \bar{\varepsilon}_{n_1 \bullet}}{n_N^\gamma}, \qquad i = 1, \ldots, N, \tag{12}$$

where we use the notation $\bar{y}_{n_i \bullet} = D^{-1} \sum_{d=1}^D y_{n_i d}$ and $\bar{\varepsilon}_{n_i \bullet} = D^{-1} \sum_{d=1}^D \varepsilon_{n_i d}$. When $\gamma > 1/2$, $n^{-\gamma}(\bar{\varepsilon}_{n_i \bullet} - \bar{\varepsilon}_{n_1 \bullet}) = o_p(1)$, and then (12) can be fitted consistently by kernel smoothing, making it possible to estimate the shape of the microstructure noise function and test the pure noise/linearity assumption. Because n_N^γ is simply a constant scaling factor in this regression, it is not necessary to standardize the data when performing the regression empirically, and the shape of the curve is invariant to the scaling factor. Hence γ does not need to be known a priori, and it is in fact implicitly determined within the nonparametric estimation of the function $f(\frac{n_i}{n_N})$. However, the standardization factor n_N^γ and the magnitude of γ do affect the limit theory.

There is an interesting link between the general model (10) and the specific approach taken by Zhang et al. (2005) to eliminate the effects of microstructure noise. When $\gamma = 1$ and the microstructure noise function is asymptotically linear as in (8), local level kernel estimation of (10) for a particular day leads to the estimate

$$\hat{f}(1) = \frac{\sum_{j=1}^N \frac{y_{n_j d} - \alpha_d}{n_N} K_h(\frac{n_j}{n_N} - 1)}{\sum_{j=1}^N K_h(\frac{n_j}{n_N} - 1)} \sim \frac{\sum_{j=1}^N \frac{y_{n_j d}}{n_N} K_h(\frac{n_j}{n_N} - 1)}{\sum_{j=1}^N K_h(\frac{n_j}{n_N} - 1)}, \tag{13}$$

where $K_h(\cdot) = h^{-1} K(\cdot/h)$ for some given kernel function $K(\cdot)$ and bandwidth parameter h. The estimate (13) is simply a locally smoothed version of the estimate

$$\tilde{f}(1) = \frac{y_{n_N d}}{n_N} = \frac{1}{n_N} [p, p]_d^{\mathcal{G}_{n_N}}$$

that Zhang et al. (2005) suggested for estimating the microstructure noise variance $2\sigma^2$ in the model with pure noise, that is, the slope coefficient in the linear regression (8). Nonparametric estimation of (10) and (12) may be considered a generalization of this approach that allows for a much wider class of microstructure noise.

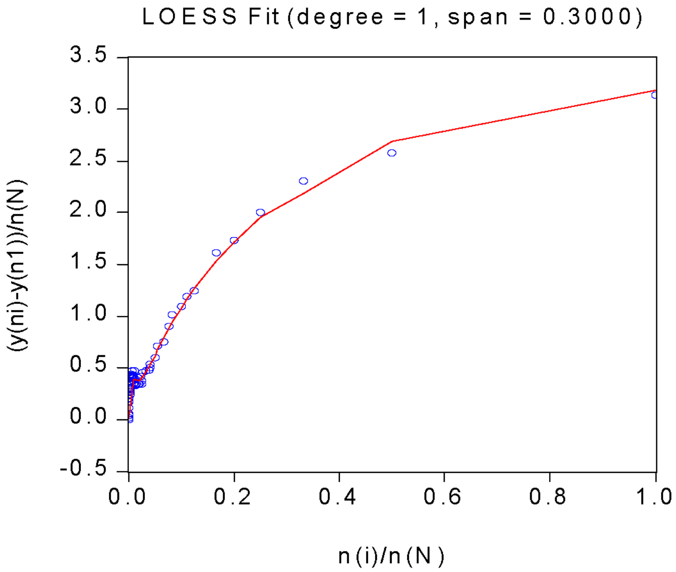

Figure 4. Locally Weighted Regression Estimate (with a tricube weight) of the Microstructure Noise Function f for AA for May 2004. The value $\gamma = 1$ is used, and the vertical axis is multiplied by 1×10^8.

Using the same data for AA as before, Figure 4 shows the nonparametric estimate of the function f obtained by locally weighted least squares regression of Cleveland and Devlin (1988) on (12). As before, the empirical evidence rather strongly rejects the pure noise assumption.

5. CONCLUSION

We find ourselves very much in agreement with the thrust of HL's message concerning the complexity induced by microstructure noise. In particular, we agree that noise is time dependent and correlated with the efficient price (features that in our view are a necessary consequence of the observed form of market transactions, as we have argued earlier) and that the properties of noise inevitably evolve over time, again just as the efficient price is itself evolutionary. We further agree that microstructure noise cannot be accommodated by simple specifications. Because microstructure noise at ultra-high infill sampling frequencies often offsets the actual transactions data to the latent efficient price, the complexity of microstructure noise includes local nonstationarity and perfect correlation with the efficient price. These properties are not permitted in the models and methods presently used in the literature. However, there are empirical procedures that are capable of addressing these additional complexities, as we have indicated in parts of our discussion.

We join the authors in saying there is still much to do in this exciting field, and look forward to further developments that build on the work that they and others have done.

ACKNOWLEDGMENT

Phillips thanks his hosts at SESS, Singapore Management University, for support and hospitality in March–April, 2005, when work on this article commenced, and acknowledges partial support from a Kelly Fellowship at the University of Auckland Business School and from the NSF under grant SES-04-142254. Yu gratefully acknowledge financial support from the Research Office at Singapore Management University.

ADDITIONAL REFERENCES

Cleveland, W. S., and Devlin, S. J. (1988), "Locally Weighted Regression: An Approach to Regression Analysis by Local Fitting," *Journal of the American Statistical Association*, 83, 596–610.

Friedman, J. H. (1984), "A Variable Span Smoother," Technical Report 5, Stanford University, Laboratory for Computational Statistics, Dept. of Statistics.

Maheswaran, S., and Sims, C. A. (1993), "Empirical Implications of Arbitrage-Free Asset Markets," in *Models, Methods and Applications in Econometrics*, ed. P. C. B. Phillips, London: Blackwell, pp. 301–316.

Phillips, P. C. B., and Yu, J. (2005a), "The Volatility Signature of Market Microstructure Noise," working paper, Singapore Management University.

——— (2005b), "Realized Volatility and Microstructure Noise Modeling," working paper, Singapore Management University.

Rejoinder

Peter R. HANSEN
Department of Economics, Stanford University, 579 Serra Mall, Stanford, CA 94305-6072

Asger LUNDE
Department of Marketing and Statistics, Aarhus School of Business, Fuglesangs Alle 4, 8210 Aarhus V, Denmark

1. INTRODUCTION

It is a privilege to receive this much feedback on our article from leading researchers whose work has inspired much of our research. We are grateful for the many insightful comments and interesting discussions. The comments not only elucidate key aspects of the problems that we have analyzed, but also provide us (and other readers alike) with a window onto many of the interesting topics that define the research frontier of this area.

The nine comments span a wide range of issues. In this rejoinder we first give a brief response to each comment, and then focus on key issues.

Aït-Sahalia, Mykland, and Zhang (AMZ) provide a very insightful discussion and challenge us to come up with better arguments and additional empirical evidence. Their discussion on identification is important and interesting. They argue that it may not be possible to identify parameters of interest and estimate the IV consistently for some types of dependent noise. This can be viewed as a shortcoming of a given model (as emphasized by AMZ) or as a limitation of the available information in the data. So an interesting question is: "What are the actual properties of the noise, and which limitations are implied for estimation of the quadratic variation?" Although this remains largely an open question, we believe that our results demonstrate that this problem is more complicated than that associated with independent noise. Of the nine comments, the one by AMZ is perhaps the most critical one. For example, AMZ express their doubts about our results on the dependence between efficient price and noise process and are critical of our procedure for cleaning the data for outliers. Many of the issues that we discuss in greater detail, such as in Sections 4 and 7, are motivated by their comment.

Much progress has been made since November 2003, when many works in this field were presented for the first time at a Montreal conference organized by Nour Meddahi. For example, Zhang, Mykland, and Aït-Sahalia (2005) have improved their original two-time scale estimator (TSRV) with a small-sample correction to deal with "small" noise and extended the applicability of the TSRV to a situation with dependent noise (see Aït-Sahalia, Mykland, and Zhang 2005). We have also benefitted from the "synergy effect," because the properties of the TSRV estimator challenged us to develop kernel-based estimators (in collaboration with Barndorff-Nielsen and Shephard) that are better estimators of IV. AMZ's comment is another addition to this active field of research, and, despite some differences of opinion, we found their comment immensely constructive.

Andersen, Bollerslev, Frederiksen, and Nielsen (ABFN) focus on the importance of jumps. They give a general discussion of realized power and bipower variation and suggest a way to robustify such quantities by skipping observations. They present several volatility signature plots for the robust bipower variation (BPV) that provide evidence of jumps in the efficient price (an issue that we ignored in our article). We give a brief discussion of jumps and BPV in relation to our results in Section 5.3 and in our discussion of Diebold's comment.

Bandi and Russell (BR) focus on properties of the noise process and include a very intuitive discussion of the results in our article. We very much agree with their call for the use of economic criteria for assessing the consequences of noise, and we enjoyed their perspective on issues that need to be addressed in future research. Their extensive empirical analysis adds clarity and additional evidence to some of our empirical findings. In our discussion of Figure 2 in Section 4, we add some insight to some of BR's volatility signature plots.

Barndorff-Nielsen and Shephard (BNS) stress that noise is very important in the multivariate context, an area that is currently undergoing much development (see, e.g., Renò 2003; Martens 2004; Barndorff-Nielsen and Shephard 2004; Hayashi and Yoshida 2005; Bandi and Russell 2005b; Griffin and Oomen 2005; Voev and Lunde 2006; Sheppard 2005). BNS also include a discussion of volatility in the context of non-Brownian stochastic volatility models and show that the bipower variation is another tool for dealing with the effects of market frictions.

Journal of Business & Economic Statistics
April 2006, Vol. 24, No. 2
DOI 10.1198/073500106000000170