Comment on “Realized Variance and Market Microstructure Noise” by Peter Hansen and Asger Lunde

Federico M. Bandi and Jeffrey R. Russell
Graduate School of Business
University of Chicago
5807 South Woodlawn Avenue
Chicago, IL 60637

First draft: February 2005
This draft: August 2005

1 INTRODUCTION

If efficient asset prices follow continuous semi-martingales and are perfectly observed, their quadratic variation can be measured accurately from the sum of a large number of squared returns sampled over very finely spaced intervals, i.e., realized variance (Andersen et al., 2003, and Barndorff-Nielsen and Shephard, 2002). With the emergence of high-frequency data, it seems that we should be able to identify volatility rather easily. However, this identification hinges on being able to observe the true (or efficient) price process. Unfortunately, observed asset prices are affected by market microstructure effects, such as discreteness, different prices for buyers and sellers, “price impacts” of trades, and so forth. If we think of observed logarithmic prices as efficient logarithmic prices plus logarithmic market microstructure noise contaminations, then we face an interesting as well as complex econometric task in using high-frequency data to estimate quadratic variation from noisy observed asset price data.

Hansen and Lunde make several important contributions to this growing literature. First, they document empirical evidence regarding the dynamic features of microstructure noise. Second, they propose a clever procedure to remove microstructure-induced biases from realized variance estimates. We start by giving intuition for their proposed estimation method. We then turn to a point-by-point discussion of their empirical findings regarding the properties of the noise. We conclude by providing our views on the current status of the nonparametric literature on integrated variance estimation.

2 THE METHODOLOGY

Assume availability of $M+1$ equispaced price observations over a fixed time span $[0,1]$, a trading day, say, so that the distance between observations is $\delta = \frac{1}{M}$. Write an observed logarithmic price $\widetilde{p}_{j\delta}$ as the sum of an equilibrium logarithmic price $p^e_{j\delta}$ and a market microstructure-induced departure $\eta_{j\delta}$, namely

$$\widetilde{p}_{j\delta} = p^e_{j\delta} + \eta_{j\delta} \qquad j = 0, ..., M. \tag{1}$$

Both $p^e_{j\delta}$ and $\eta_{j\delta}$ are not observed by the econometrician. In the remainder of the text we will frequently refer to $\eta_{j\delta}$ as the "noise" component. Equivalently, in terms of continuously-compounded returns,

$$\widetilde{r}_{j\delta} = r^e_{j\delta} + \varepsilon_{j\delta} \qquad j = 1, ..., M, \tag{2}$$

where $\widetilde{r}_{j\delta} = \widetilde{p}_{j\delta} - \widetilde{p}_{(j-1)\delta}$, $r^e_{j\delta} = p^e_{j\delta} - p^e_{(j-1)\delta}$ and $\varepsilon_{j\delta} = \eta_{j\delta} - \eta_{(j-1)\delta}$.

2.1 INTUITION

In its general form, the estimator that Hansen and Lunde advocate is in the tradition of robust covariance estimators such as that of Newey and West (1987). Write

$$\widehat{V}_{q_M} = \underbrace{\sum_{j=1}^{M} \widetilde{r}^2_{j\delta}}_{RV} + \underbrace{2\sum_{h=1}^{q_M} \frac{M}{M-h} \sum_{j=1}^{M-h} \widetilde{r}_{j\delta}\widetilde{r}_{(j+h)\delta}}_{bias\ correction}, \tag{3}$$

where q_M is a frequency-dependent number of covariance terms. The first term in the right-hand side of Eq. (3) is the classical realized variance estimator, the second term is a correction intended to make the estimator unbiased for the integrated variance of the underlying price process in the presence of correlated noise. If the correlation structure of the noise return process is finite and the efficient price is a local martingale, then the estimator in Eq. (3) is unbiased for the integrated variance of the efficient price process over $[0,1]$, i.e., $V = \int_0^1 \sigma_s^2 ds$, under appropriate conditions on q_M. 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 and market microstructure noise.

Consider, for simplicity, the i.i.d. market microstructure noise case with noise independent of the underlying price process. Specifically, impose Assumption 1 and 3 in Hansen and Lunde's paper. Furthermore, set $q_M = 1$ to obtain

$$\widehat{V}_1 = \sum_{j=1}^{M} \widetilde{r}_{j\delta}^2 + 2\frac{M}{M-1} \sum_{j=1}^{M-1} \widetilde{r}_{j\delta}\widetilde{r}_{(j+1)\delta}. \tag{4}$$

The covariance between $\widetilde{r}_{j\delta}$ and $\widetilde{r}_{(j+1)\delta}$, i.e., $\mathbf{E}_M(\widetilde{r}_{j\delta}\widetilde{r}_{(j+1)\delta})$, is the same at all frequencies and equal to $-\mathbf{E}(\eta^2)$. Hence, $\mathbf{E}_M\left(2\frac{M}{M-1}\sum_{j=1}^{M-1}\widetilde{r}_{j\delta}\widetilde{r}_{(j+1)\delta}\right) = -2M\mathbf{E}(\eta^2)$. The bias of the realized variance estimator $\sum_{j=1}^{M}\widetilde{r}_{j\delta}^2$ is equal to $M\mathbf{E}(\varepsilon^2) = 2M\mathbf{E}(\eta^2)$. Therefore, the second term in Eq. (4) provides the required finite sample bias-correction.

Consider now the i.i.d. market microstructure noise case with noise dependent of the underlying price process. However, assume $\mathbf{E}_M(r_{j\delta}^e \eta_{(j-s)\delta}) = 0$ for all j and all $s \geq 1$, i.e., microstructure noise does not predict future efficient returns, and $\mathbf{E}_M(r_{j\delta}^e \eta_{(j+1)\delta}) = 0$ for all j. The bias of the realized variance estimator is now equal to

$$2\sum_{j=1}^{M} \mathbf{E}_M(r_{j\delta}^e \varepsilon_{j\delta}) + 2M\mathbf{E}(\eta^2) = 2M\mathbf{E}_M(r_{j\delta}^e \eta_{j\delta}) + 2M\mathbf{E}(\eta^2). \tag{5}$$

But

$$\begin{aligned}
&\mathbf{E}_M\left(\frac{M}{M-1}\sum_{j=1}^{M-1}\widetilde{r}_{j\delta}\widetilde{r}_{(j+1)\delta}\right) \\
=\; &\mathbf{E}_M\left(\frac{M}{M-1}\sum_{j=1}^{M-1}\left(r_{j\delta}^e + \varepsilon_{j\delta}\right)\left(r_{(j+1)\delta}^e + \varepsilon_{(j+1)\delta}\right)\right) && (6) \\
=\; &\frac{M}{M-1}\sum_{j=1}^{M-1}\mathbf{E}_M(r_{j\delta}^e\varepsilon_{(j+1)\delta}) + \frac{M}{M-1}\sum_{j=1}^{M-1}\mathbf{E}_M(\varepsilon_{j\delta}r_{(j+1)\delta}^e) + \frac{M}{M-1}\sum_{j=1}^{M-1}\mathbf{E}_M(\varepsilon_{j\delta}\varepsilon_{(j+1)\delta}) && (7) \\
=\; &-M\mathbf{E}_M(r_{j\delta}^e\eta_{j\delta}) - M\mathbf{E}(\eta^2) && (8)
\end{aligned}$$

which, again, provides the necessary finite sample bias-correction in the case of noise dependent on the price process.

The proposed estimator is theoretically interesting and, we believe, empirically useful. The cancelation that gives rise to bias-corrected integrated variance estimates, even in the presence of some correlation between the noise component and the underlying efficient price process, is clever.

The frequency δ at which to sample continuously-compounded returns for the purpose of integrated variance estimation through nonparametric estimates (like $\widehat{V}_{q_M}$ above) can be chosen optimally. This point has been made in recent research and estimable mean-squared error (MSE) expressions for a variety of nonparametric variance estimators have been provided (see Bandi and Russell, 2005a,b, and Zhang et al., 2005). Hansen and Lunde expand on this framework by considering an optimal (in an MSE-sense) frequency δ^* for their first-order bias-corrected estimator $\widehat{V}_1$ in the case of *(1)* i.i.d. noise with *(2)* noise independent of the price process. As shown above, the validity of their bias-correction does not rely on *(1)* and *(2)* in general.

3 THE PROPERTIES OF MICROSTRUCTURE NOISE

Hansen and Lunde's findings about the properties of the noise are summarized in the Introduction of the paper. They write:

(1) The noise is correlated with the efficient price.

(2) The noise is time-dependent.

(3) The noise is small.

(4) The noise has changed over time.

In this Section of our Comment we expand on these results and discuss them point-by-point.

3.1 THE NOISE IS CORRELATED WITH THE EFFICIENT PRICE

Hansen and Lunde study the relation between efficient price and noise from an interesting, new perspective (for a discussion of the early work on the subject we refer the reader to

Hasbrouck, 2006). As they point out, this relation has important theoretical and empirical implications for integrated variance estimation. This is a difficult relation to assess since the equilibrium prices and the noise components embedded in the observed asset prices are inherently unobserved (see Eq. (1) above). Hansen and Lunde largely use "signature plots" to study this relation. Signature plots are plots of realized variance versus the sampling frequency $\delta = \frac{1}{M}$ or, equivalently, versus the number of returns M (Andersen et al., 2000). The idea is simple. Assume $\mathbf{E}_M(r^e_{j\delta}\eta_{(j-s)\delta}) = 0$ for all j and all $s \geq 1$, as earlier. Consider the expected value of the realized variance estimator, namely

$$\mathbf{E}_M\left(\sum_{j=1}^{M}\widetilde{r}^2_{j\delta}\right) = \int_0^1 \sigma_s^2 ds + 2M\mathbf{E}_M(r^e_{j\delta}\eta_{j\delta}) + M\mathbf{E}_M(\varepsilon^2_{j\delta}). \tag{9}$$

If the noise is independent of the efficient price (i.e., if $\mathbf{E}_M(r^e_{j\delta}\eta_{j\delta}) = 0$), then realized variance becomes large as $\delta \to 0$ (or $M \to \infty$) since $M\mathbf{E}_M(\varepsilon^2_{j\delta}) \to \infty$. If the noise is negatively correlated with the efficient price, then realized variance might not increase at high-frequencies since $\mathbf{E}_M(r^e_{j\delta}\eta_{j\delta})$ and $\mathbf{E}_M(\varepsilon^2_{j\delta})$ operate in opposite directions. Hansen and Lunde report several signature plots that are not upward-sloping as the sampling frequency increases. These plots mainly refer to NYSE midquote price data for a sample of Dow Jones stocks.

We point out that, if the efficient price constantly evolves, factors inducing stickiness in the observed prices $\widetilde{p}$ *mechanically* determine a negative dependence between the unobserved components of the observed prices, p^e and η. This is simply due to the fact that any time the observed price remains fixed, any changes in the efficient price must be perfectly offset by movements in the opposite direction in the noise component. Hence, important aspects of the dependence between the efficient price and the noise, as well as aspects of the conditional and unconditional properties of the noise (see next subsection), can be understood by considering the factors inducing stickiness in the observed prices, primarily *i*) the market structure, *ii*) how prices are defined, and *iii*) the method used to sample prices.

MARKET STRUCTURE The market structure, or rules of trade, will impact the features of the noise. Assume $\widetilde{p}$ refers to quotes. If agents arrive in a random fashion with differing perceptions of the value of an asset and different motives, then continual random adjustments to the posted prices are expected to be made. These random price adjustments should induce noise components that are roughly uncorrelated with the underlying efficient price. This structure is consistent with foreign exchange markets, for instance, where banks around the world post quotes. Alternatively, specialist-driven markets, such as the NYSE, are more centralized with a single agent posting quotes. Without prices being set by multiple agents with dissenting views, the observed prices can remain fixed over longer durations. This results in stickiness in the posted quotes (the $\widetilde{p}$'s in our framework), thereby contributing to explain the negative correlation between efficient returns and noise returns as documented in Hansen and Lunde's study. Interestingly, if one combines quotes from multiple exchanges, then the random price setting behavior of a more decentralized market can be replicated.

Here we consider the foreign exchange market and NASDAQ as examples of fairly decentralized markets. The NYSE is used as an example of a relatively centralized market.

Fig. 1 contains calendar-time midquote variance plots for the Deutchmark/Dollar and Yen/Dollar exchange rates. The data is 5 minute data from all Tuesdays, Wednesdays, and Thursdays from the year 1996.

Fig. 2 contains calendar-time midquote variance plots for Cisco and Microsoft.

Fig. 3 contains calendar-time midquote variance plots for IBM and GE using *i*) NYSE and *ii*) NYSE and MIDWEST.

Fig. 4 contains calendar-time midquote variance plots for IBM and GE using *i*) NYSE and NASDAQ and *ii*) the consolidated market.

Fig. 3 confirms Hansen and Lunde's observation. Midquotes posted on the NYSE have a residual noise component that is negatively related to the unobserved efficient price component. This negative correlation can be mechanically induced by relatively stable quotes in a specialist-driven market. In the other cases, the signature plots are upward-sloping at high-frequencies. The negative correlation between noise and efficient price is dominated, and sometimes strongly so, by the second moment of the noise component. When dealing with quotes, we find support for this result in the foreign exchange market (our highest frequency here is only 5 minutes due to data limitations), in the case of NASDAQ, as well as in the case of combined quotes from multiple exchanges such as the NYSE and the MIDWEST, for instance. This result is particularly strong when considering the consolidated US market.

PRICE MEASUREMENT The noise associated with bid-ask mid-quotes is different from the noise associated with transaction prices. Even within centralized markets, the price formation mechanism leading to observed trade prices differs from the mechanism leading to posted quotes. The former change more randomly than the latter in that they are more affected by the random arrival of agents in the marketplace. Hansen and Lunde recognize that the negative correlation between r^e and η in the case of midquotes appears not to be a first-order effect when dealing with transaction prices. This result is easily understandable in the context of our assumed price formation mechanism. The observed price $\widetilde{p}$ is relatively stickier when measured using midquotes rather than transaction prices *regardless* of *i*) the market structure and *ii*) the method used to sample prices.

Fig. 5 presents calendar-time midquote and transaction variance plots for Cisco and Microsoft.

Fig. 6 presents calendar-time transaction variance plots for IBM using *i*) NYSE, *ii*) NYSE and MIDWEST, *iii*) NYSE and NASDAQ, and *iv*) the consolidated market. This figure should be compared to Figs. 3 and 4.

Fig. 7 presents calendar-time transaction variance plots for GE using *i*) NYSE, *ii*) NYSE and MIDWEST, *iii*) NYSE and NASDAQ, and *iv*) the consolidated market. This figure should be compared to Figs. 3 and 4.

SAMPLING METHOD Sampling returns in calendar time induces properties of the noise that are different from the properties of the noise obtained by sampling returns in trade time. (See Oomen, 2005a, for an interesting treatment of business-time sampling in the context of realized variance estimation.) At very high frequencies, calendar time sampling will inevitably result in sampling between quote updates, thereby artificially inducing stickiness in the $\widetilde{p}$'s. This stickiness, again, leads to an artificially negative correlation between noise returns and efficient returns. Here we compare the variance signature plots obtained from event time sampling to those previously obtained from calendar time sampling.

Fig. 8 contains event-time midquote variance plots for Cisco and Microsoft. This figure should be compared to Fig. 2.

Fig. 9 presents event-time midquote variance plots for IBM and GE using *i*) NYSE and *ii*) NYSE and MIDWEST. This figure should be compared to Fig. 3.

Fig. 10 presents event-time midquote variance plots for IBM and GE using *i*) NYSE and NASDAQ, and *ii*) the consolidated market. This figure should be compared to Fig. 4.

3.2 THE NOISE IS TIME-DEPENDENT

It is not a coincidence that, as reported by Hansen and Lunde, the noise appears to be more persistent when the dependence between noise and efficient price is more negative. We argue that price stickiness can also *mechanically* induce persistence in the microstructure noise. In fact, stable observed returns determine persistent noise contaminations in the presence of unpredictable efficient returns. This again can be easily seen. Consider the model in Eq. (1). For the sake of argument, assume $\widetilde{r}_{j\delta} := 0$ for all j. Then,

$$0 = r^e_{j\delta} + \varepsilon_{j\delta}, \tag{10}$$

thereby implying

$$\eta_{j\delta} = \eta_{(j-1)\delta} - r^e_{j\delta}. \tag{11}$$

As earlier, the extent of the persistence in the price contaminations will depend on the market structure, on the price measurement, and on the sampling scheme. The convenient i.i.d. market microstructure model can be a poor approximation when using midquotes from a highly centralized market like the NYSE as in Hansen and Lunde's study. However, it can a be a satisfactory approximation in other circumstances. For example, adding additional quote information from other exchanges to NYSE midquotes increases substantially the (negative) first-order autocorrelation of the recorded returns.

3.3 THE NOISE IS SMALL

We believe the size of the noise depends on the metric of interest. If the relevant metric is the MSE of an integrated variance estimator, for instance, then it is hard to claim that the size of the noise is small. In fact, integrated variance estimators that are more robust to market microstructure noise than the classical realized variance estimator have conditional (on the volatility path) MSEs that are substantially smaller than the MSE of the realized variance estimator (Bandi and Russell, 2005a,d; Hansen and Lunde, 2006).

More generally, we think it is important to ask what are the economic implications of alternative variance estimates (we will come back to this observation in Section 4). If sensible economic criteria suggest that nonparametric variance estimates that are more robust to market microstructure noise than the classical realized variance estimator provide economic gains, then it would certainly be inaccurate to say that the noise is small.

3.4 THE NOISE CHANGED OVER TIME

We fully agree. For instance, the size of the noise component in the observed mid-quotes and transaction prices, as summarized by the magnitude of the noise second moment, has been decreasing throughout the years (Bandi and Russell, 2005b; Oomen, 2005a).

Markets are becoming more liquid and more decentralized. More liquid markets will be characterized by a larger number of transactions. While the increased transaction rate will likely reduce the extent of the dependence $\mathbf{E}_M(r^e_{j\delta}\eta_{j\delta})$ between the efficient return process and microstructure noise over any time interval, the resulting effect on the bias component $M\mathbf{E}_M(r^e_{j\delta}\eta_{j\delta})$ in Eq. (9) is not obvious. The increased decentralization will likely induce less stickiness in the observed prices. This effect will probably contribute to render both the dependence between noise and efficient price and the noise persistence relatively less important.

4 FINAL REMARKS and SOME PERSPECTIVES

Remark 1. We argued that some of the dynamic features of the noise, namely the negative correlation between efficient returns and noise components and the persistence of the noise, may be explained by a *mechanical* relationship induced by a combination of sluggishness in the adjustments to the observed prices and an ever-evolving efficient price.

Remark 2. Are the inevitable negative correlation between noise and efficient return process and the noise persistence first-order effects in the data? It depends on the market, on the type of price measurement, and on the sampling scheme.

Consider the dependence between noise and efficient prices. The signature plots in Hansen and Lunde's paper nicely speak to the strong negative dependence between efficient price and noise component when the observed price ($\widetilde{p}$) is particularly sticky, i.e., in the case of calendar-time midquotes on the NYSE. However, different market structures, different sampling methods, and different price measurements can have drastically different noise features and a different nature of dependence between noise and efficient price. When the observed prices are relatively less sticky (i.e., in the case of foreign exchange markets, NASDAQ and combinations of markets, for example, but also when using transaction prices and/or sampling in event time rather than in calendar time), the evidence and logic in favor of a negative dependence between efficient price and noise component at high-frequencies, as well as the evidence and logic in favor of noise persistence, are weaker in general.

Remark 3 (The importance of finite sample performance.) Should the literature on integrated variance estimation be concerned with the inevitable negative correlation between noise and efficient return process and with noise persistence? Naturally, the answer depends on whether these features are first-order effects in the data.

Our own work has emphasized the need for differentiating between asymptotic and finite sample methods in the context of integrated variance estimation by virtue of noisy high-frequency data and the importance of finite sample criteria (Bandi and Russell, 2005a,d). Irrespective of the estimator used, the effective implementation of finite sample criteria requires the imposition of sensible restrictions on the properties of the noise. For example, one can optimize the finite sample MSE of a variety of integrated variance estimators in the presence of dependent noise (see Bandi and Russell, 2005a, for a study of the dependent noise case when dealing with realized variance). However, dependence between the noise and the efficient price is a more delicate issue. While, in general, one can express the MSE of integrated variance estimators in the presence of dependence between the noise and the efficient price, empirical estimation of the additional terms that would arise by virtue of this dependence is a very complicated matter.

Depending on the data (market structure and type of price measurement) and sampling scheme, empirically justifiable compromises need to be made for the finite sample properties of alternative integrated variance estimators to be optimized. Compromises are made in Hansen and Lunde's work (the variance of their first-order bias corrected estimator $\widehat{V}_1$ is optimized under an i.i.d. noise model with noise independent of the efficient price). Similarly, compromises are likely to be needed when optimizing the finite sample variance properties of the more general bias-corrected estimator advocated by Hansen and Lunde ($\widehat{V}_{q_M}$ in Eq. (3)).

Remark 4. What data do we suggest using?

(1) If interest is placed on the integrated variance of the underlying price process (as in Hansen and Lunde's work), then it is sensible to use high-frequency data that are less contaminated by residual market microstructure noise components, i.e., midquotes.

(2) As we showed, sampling in calendar time can generate severe distortions at high frequencies. The presence of stale quotes due to sampling between quote updates can bias the estimates downward substantially. Hence, sampling at high frequencies for the purpose of estimating midquote noise moments, as suggested in Bandi and Russell (2005a,b) and Zhang et al. (2005), should be conducted in event time.

(3) If the properties of the noise are not expected to change wildly across exchanges, information provided by multiple exchanges can be employed. Bandi and Russell (2005a,b) use event-time midquotes from the NYSE and the MIDWEST to estimate the noise moments. Fig. 11 presents event-time midquote volatility plots for IBM and GE using NYSE and the MIDWEST. At high frequencies the IBM variance doubles while the GE variance is five times as large as the corresponding value at low frequencies. These increases provide important information about the second moment of the noise.

(4) If interest is placed on the noise variance and η is interpreted as a transaction cost, then transaction prices sampled in event time should be used.

Remark 5. (The importance of economic metrics.) Numerous statistical refinements have been made for integrated variance estimation in the presence of market microstructure noise (Bandi and Russell, 2005c, and Barndorff-Nielsen and Shephard, 2005, summarize the current state of the literature). Hansen and Lunde's work is central to this exciting line of research. Even though we expect important statistical advances to continue to be made, we think it is now necessary to ask what are the economic implications of statistical refinements.

One way to assess these implications is to use the utility-based method advocated by Fleming et al. (2001) to evaluate alternative variance forecasts. This is the approach taken by Bandi and Russell (2005b). However, there are numerous alternative economic metrics that could (and should) be considered. For example, given an option pricing model linking option prices to integrated variance, the forecasting power of alternative integrated variance measures can be assessed. While proposing and implementing sensible economic metrics is a difficult task in general, this is, we believe, an important hurdle to overcome for the literature on integrated variance estimation.

Acknowledgments. We thank the Graduate School of Business of the University of Chicago for financial support. We are grateful to the conference participants at the 2005 Annual JBES Invited Lecture (American Statistical Association Meetings, Minneapolis, August 7-11) for discussions.

References

[1] Andersen, T.G., T. Bollerslev, F.X. Diebold and P. Labys (2000). Great realizations. *Risk,* (13)3, 105-108.

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

[3] Bandi, F.M., and J.R. Russell (2005a). Microstructure noise, realized variance, and optimal sampling. *Working paper.*

[4] Bandi, F.M., and J.R. Russell (2005b). Separating microstructure noise from volatility. *Journal of Financial Economics.* Forthcoming.

[5] Bandi, F.M., and J.R. Russell (2005c). Volatility. In J.R. Birge and V. Linetski (Eds.) *Handbook of Financial Engineering.* Elsevier North-Holland. Forthcoming.

[6] Bandi, F.M., and J.R. Russell (2005d). On the finite sample properties of kernel-based integrated variance estimators. *Working paper.*

[7] 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.

[8] Barndorff-Nielsen, O.E., and N. Shephard (2005). Variation, jumps, market frictions, and high-frequency data in financial econometrics. *Working paper.*

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

[10] Hansen, P.R. and A. Lunde (2006). Realized variance and market microstructure noise. *Journal of Business and Economic Statistics - this issue.*

[11] Hasbrouck, J. (2006). *Empirical Market Microstructure.* In press.

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

[13] Oomen, R. (2005a). Properties of realized variance for pure-jump processes: calendar time sampling versus business time sampling. *Working paper.*

[14] 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.* Forthcoming.

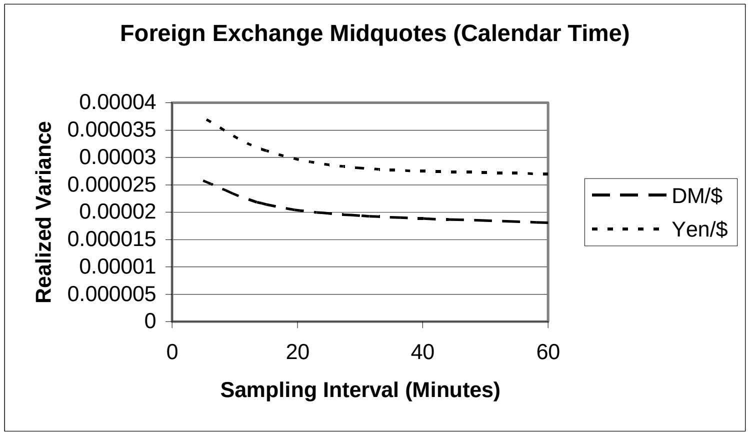

Figure 1.

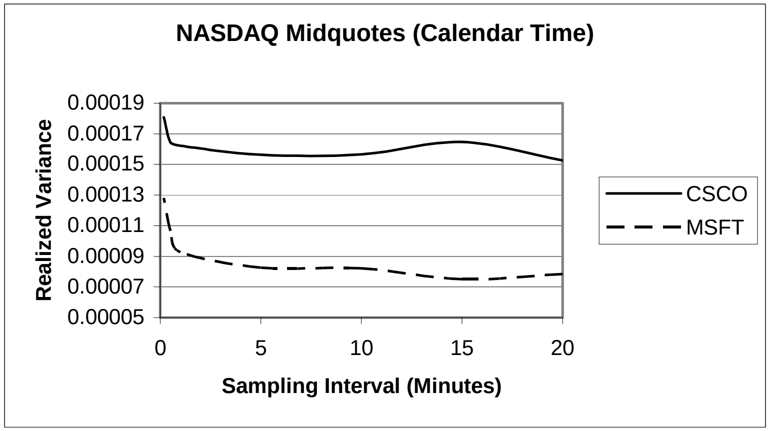

Figure 2.

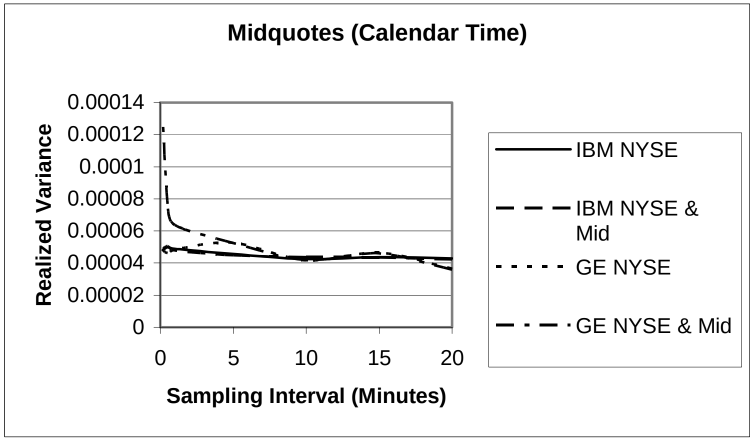

Figure 3.

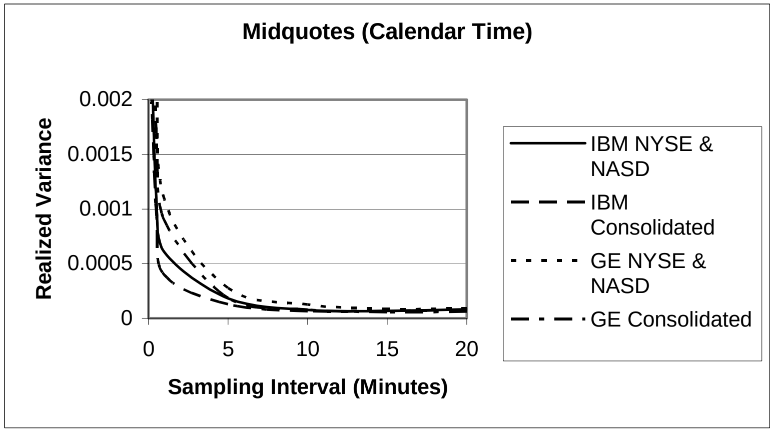

Figure 4.

Figure 1 presents midquote variance signature plots for Deutchmark/Dollar and Yen/Dollar Exchange rate data for all Tuesdays, Wednesdays, and Thursdays from 1996.
Figure 2 presents calendar-time midquote variance signature plots for the NASDAQ stocks Cisco Systems and Microsoft.
Figure 3 presents calendar-time midquote variance signature plots for IBM and GE using i) NYSE and ii) NYSE and Midwest.
Figure 4 presents calendar-time midquote variance signature plots for IBM and GE using i) NYSE and NASDAQ and ii) the consolidated market.

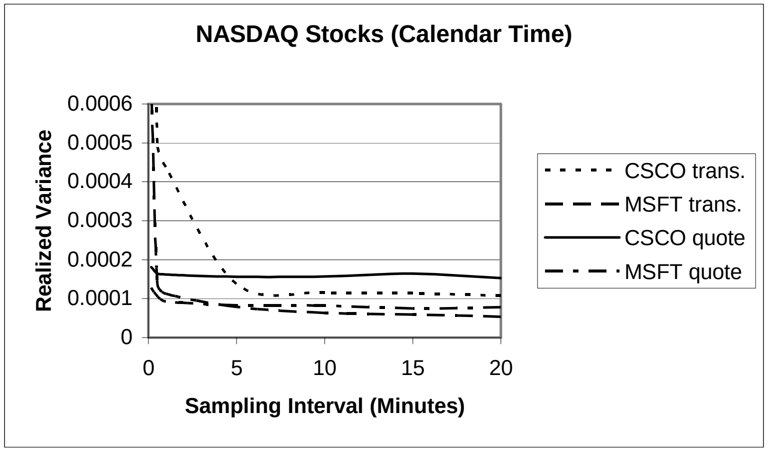

Figure 5.

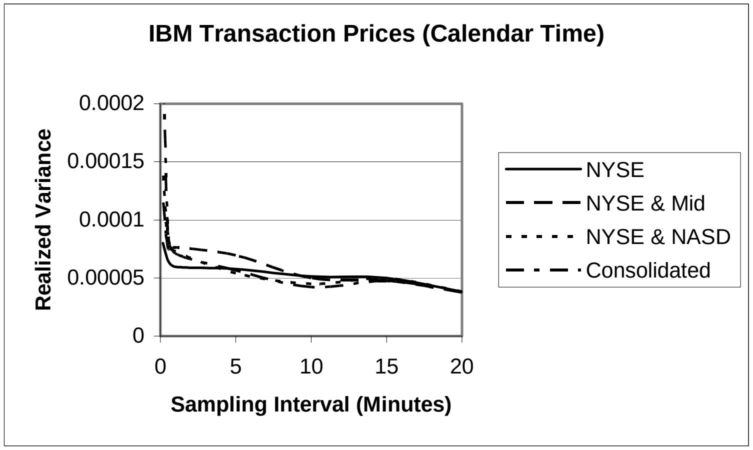

Figure 6.

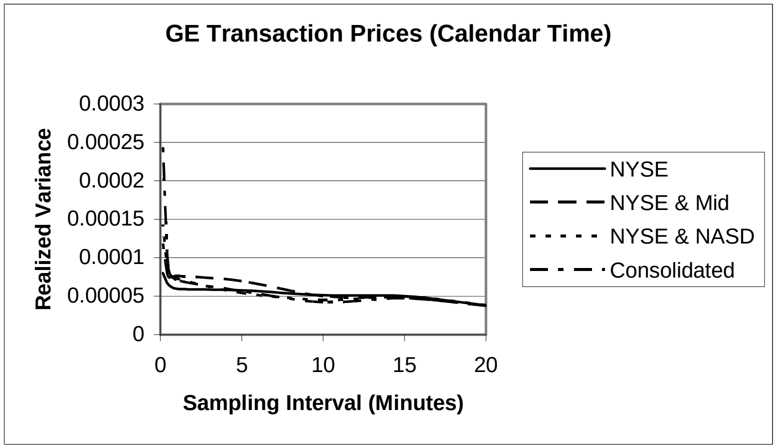

Figure 7.

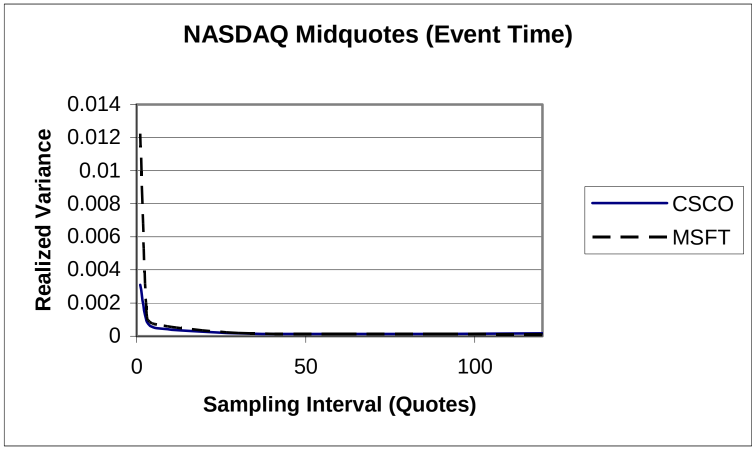

Figure 8.

Figure 5 presents calendar-time variance signature plots for Cisco Systems and Microsoft using i) transaction prices and ii) midquotes.
Figure 6 presents calendar-time transaction-price variance signature plots for IBM using i) NYSE, ii) NYSE and Midwest, iii) NYSE and NASDAQ, and iv) the consolidated market.
Figure 7 presents calendar-time transaction-price variance signature plots for GE using i) NYSE, ii) NYSE and Midwest, iii) NYSE and NASDAQ, and iv) the consolidated market.
Figure 8 presents event-time midquote variance signature plots for the NASDAQ stocks Cisco Systems and Microsoft.

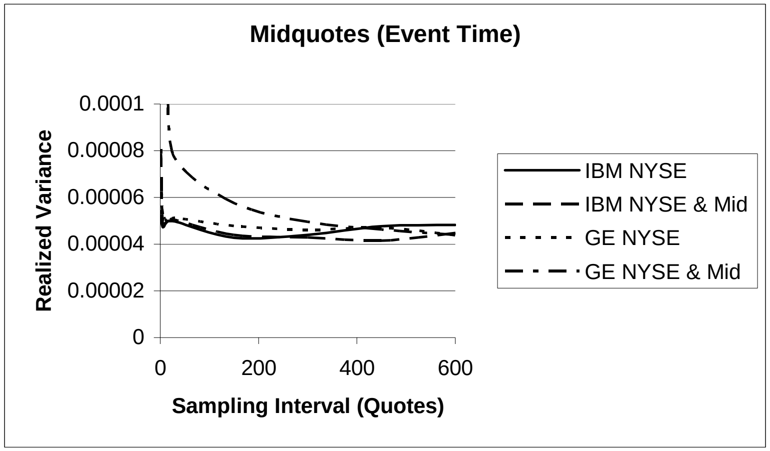

Figure 9.

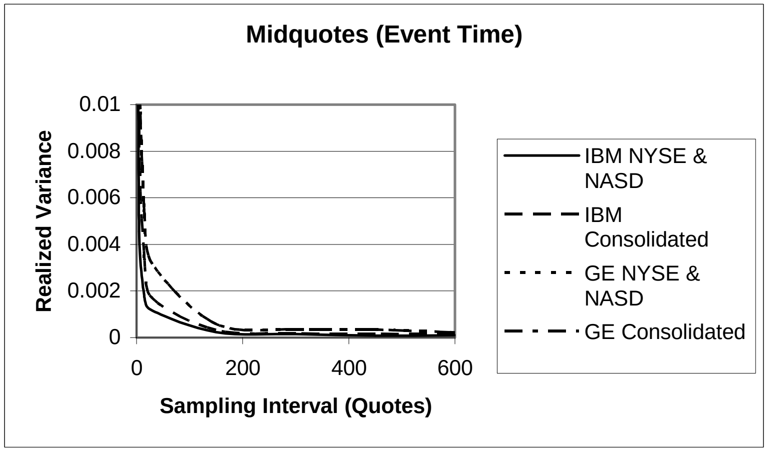

Figure 10.

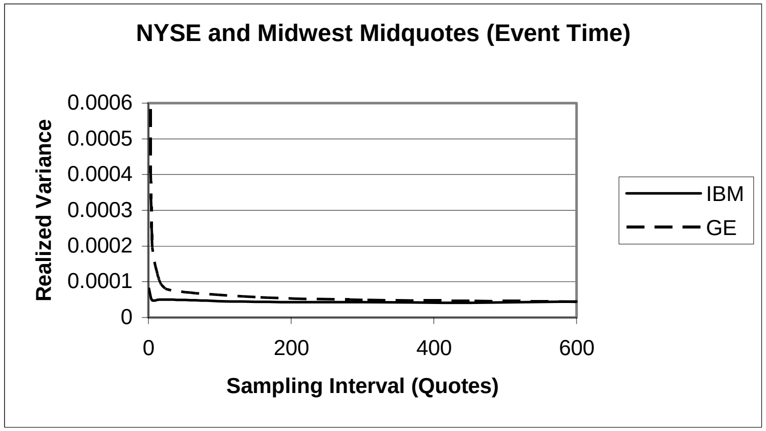

Figure 11.

Figure 9 presents event-time midquote variance signature plots for IBM and GE using i) NYSE and ii) NYSE and Midwest.
Figure 10 presents event-time midquote variance signature plots for IBM and GE using i) NYSE and NASDAQ and ii) the consolidated market.
Figure 11 presents NYSE and Midwest midquote variance signature plots for IBM and GE.