

CORRECTING THE ERRORS: VOLATILITY FORECAST EVALUATION USING HIGH-FREQUENCY DATA AND REALIZED VOLATILITIES

BY TORBEN G. ANDERSEN, TIM BOLLERSLEV, AND NOUR MEDDAHI¹

We develop general model-free adjustment procedures for the calculation of unbiased volatility loss functions based on practically feasible realized volatility benchmarks. The procedures, which exploit recent nonparametric asymptotic distributional results, are both easy-to-implement and highly accurate in empirically realistic situations. We also illustrate that properly accounting for the measurement errors in the volatility forecast evaluations reported in the existing literature can result in markedly higher estimates for the true degree of return volatility predictability.

KEYWORDS: Continuous-time models, integrated volatility, realized volatility, high-frequency data, time series forecasting, Mincer–Zarnowitz regressions.

1. INTRODUCTION

THE BURGEONING LITERATURE on time-varying financial market volatility is abound with empirical studies in which competing models are evaluated and compared on the basis of their forecast performance. Contrary to the typical setting for economic forecast evaluation, the variable of interest in that context—the volatility—is not directly observable, but rather inherently latent. Consequently, any *ex post* assessment of forecast precision must contend with a fundamental errors-in-variable problem associated with the measurement of the realization of the forecasted variable. Growing recognition of the importance of this issue has led a number of recent studies to advocate the use of so-called realized volatilities, constructed from the summation of finely sampled squared high-frequency returns, as a practical method for improving the *ex post* volatility measures.

The use of realized volatility as the practical benchmark may be justified by standard continuous-time arguments. Assuming that the sampling frequency of the squared returns utilized in the realized volatility computations approaches zero, the realized volatility then consistently estimates the true (latent) integrated volatility under quite general conditions, where importantly, the latter concept corresponds to the realization of the (cumulative) instantaneous

¹This work was supported by a grant from the National Science Foundation to the NBER (Andersen and Bollerslev), and from FQRSC, IFM2, MITACS, NSERC, SSHRC, and Jean-Marie Dufour's Econometrics Chair of Canada (Meddahi). Some of this material was circulated earlier as part of the paper "Analytical Evaluation of Volatility Forecasts." Detailed comments by Neil Shephard, Tom McCurdy, two anonymous referees, and the editor have importantly improved the paper. We would also like to thank Bryan Campbell and Francis X. Diebold for many discussions on closely related ideas, as well as seminar participants at the 2002 NBER/NSF Time Series Conference and the 2003 CIRANO-CIREQ Financial Econometrics Conference. The third author thanks the Bendheim Center for Finance at Princeton University for its hospitality during his visit where part of the research was done.

variance process over the relevant horizon (see, e.g., Andersen and Bollerslev (1998), Andersen, Bollerslev, Diebold, and Labys (2001), Barndorff-Nielsen and Shephard (2001, 2002a, 2002b, 2004a), Comte and Renault (1998), along with the recent survey by Andersen, Bollerslev, and Diebold (2003)). Unfortunately, market microstructure frictions distort the measurement of returns at the highest frequencies so that, e.g., tick-by-tick return processes blatantly violate the theoretical semi-martingale restrictions implied by the no-arbitrage assumptions in continuous-time asset pricing models. These same features also bias empirical realized volatility measures constructed directly from the ultra high-frequency returns, so in practice the measures are instead typically constructed from intraday returns sampled at an intermediate frequency.² As such, the integrated volatility is invariably measured with error (see, e.g., the numerical calculations in Andreou and Ghysels (2002) and Bai, Russell, and Tiao (2000)). The exact form of the measurement error will, of course, depend on the assumed model structure (see, e.g., Meddahi (2002) and Barndorff-Nielsen and Shephard (2002a)), but it will generally result in a downward bias in the estimated degree of predictability obtained through any forecast evaluation criterion that simply uses the realized volatility in place of the true (latent) integrated volatility. Although this bias may be large (Andersen and Bollerslev (1998)), it is almost always ignored in empirical applications.

We address that issue by developing general model-free adjustment procedures that allow for the calculation of simple unbiased loss functions in realistic forecast situations. Moreover, the adjustments are simple to implement in practice. The derivation exploits the asymptotic (for increasing sampling frequency) distributional results in Barndorff-Nielsen and Shephard (2002a). While these results explicitly rule out so-called leverage effects, building on the recent insights in Barndorff-Nielsen and Shephard (2002b), we show that the same approximate adjustment procedures apply in the context of the general eigenfunction stochastic volatility class of models pioneered by Meddahi (2001), explicitly allowing for nonzero contemporaneous correlations between the separate shocks in the return and volatility processes. Following Andersen and Bollerslev (1998) and ABDL (2003), we focus our forecast comparisons on the value of the coefficient of multiple correlation, or R^2 , in the Mincer-Zarnowitz style regressions of the ex post realized volatility on the corresponding model forecasts,³ but our procedures are general and could be applied in the adjustment of other loss functions used in the evaluation of any ar-

²For instance, the daily realized volatilities in Andersen, Bollerslev, Diebold, and Labys (2003) (henceforth ABDL (2003)) discussed further below are based on the summation of squared half-hourly foreign exchange rate returns, but either 5-minute or 15-minute returns are other common choices in the literature.

³This particular loss function is directly inspired by the work of Mincer and Zarnowitz (1969), and we will refer to the corresponding regressions as such; see also the discussion in Chong and Hendry (1986).

bitrary set of volatility forecasts. On applying the procedures in the context of ABDL (2003), we obtain markedly higher estimates for the true degree of return-volatility predictability, with the adjusted R^2 's exceeding their unadjusted counterparts by up to forty percent.

We proceed as follows. Section 2.1 introduces the notions of integrated and realized volatility, along with the (feasible) asymptotic distribution theory due to Barndorff-Nielsen and Shephard (2002a). The development of the practical and easy-to-implement adjustment procedures based on this theory is then presented in Section 2.2, followed by our theoretical extensions explicitly allowing for leverage effects. Utilizing these results, the last section provides a reassessment of the empirical evidence in ABDL (2003) related to the fit of the Mincer-Zarnowitz style volatility regressions. The accuracy of the asymptotic approximations—which form the basis for our approach—is confirmed through Monte Carlo simulations for models calibrated to reflect empirically relevant and challenging specifications. The details of the simulations and some of the technical proofs are deferred to two appendices.

2. THEORY

We focus on a single asset traded in a liquid financial market. Assuming that the sample path of the logarithmic price process, $\{\log(S_t), 0 \leq t\}$, is continuous, the class of continuous-time stochastic volatility models employed in the finance literature is then conveniently expressed in terms of the following generic stochastic differential equation (sde),

$$(1) \quad d \log(S_t) = \mu_t dt + \sigma_t dW_t,$$

where W_t denotes a standard Brownian motion given an increasing filtration $\{\mathcal{F}_t, t \geq 0\}$, the drift term μ_t is (locally) predictable and of finite variation, and σ_t is a càdlàg process such that $\int_0^t \sigma_u^2 du < \infty$ a.s. for any $t > 0$. We also assume that the process (μ_t, σ_t, W_t) is adapted to the filtration $\{\mathcal{F}_t\}$. Consequently, $\int_0^t \sigma_u dW_u$ is a local martingale and $\log(S_t)$ a semi-martingale (see, for instance, Protter (2004)).

2.1. Integrated and Realized Volatility

Although the sde in equation (1) is very convenient from a theoretical arbitrage-free pricing perspective, practical return calculations and volatility measurements are invariably restricted to discrete time intervals. In particular,

focusing on the unit time interval, the one-period continuously compounded return for the price process in equation (1) is formally given by⁴

$$(2) \quad r_t \equiv \log(S_t) - \log(S_{t-1}) = \int_{t-1}^t \mu_u du + \int_{t-1}^t \sigma_u dW_u,$$

with the corresponding *integrated volatility*,

$$(3) \quad IV_t \equiv \int_{t-1}^t \sigma_u^2 du,$$

affording a natural measure of the inherent, or notional, return variability (see, e.g., Andersen, Bollerslev, and Diebold (2003), for further discussion of the integrated and notional volatility concepts).⁵

Of course, the integrated volatility is not directly observable. However, by the theory of quadratic variation (see, e.g., Protter (2004) for a general discussion), the corresponding *realized volatility* defined by the summation of the $1/h$ intra-period squared returns, $r_t^{(h)} \equiv \log(S_t) - \log(S_{t-h})$,

$$(4) \quad RV_t(h) \equiv \sum_{i=1}^{1/h} r_{t-1+ih}^{(h)2},$$

where $1/h$ is assumed to be an integer, converges uniformly in probability to IV_t as $h \rightarrow 0$.

The consistency of the realized volatility relies on the (conceptual) idea of an ever increasing number of finer sampled high-frequency returns, or $h \rightarrow 0$. However, as previously noted, the requisite semi-martingale property of returns invariably breaks down at ultra-high frequencies, so that in actual applications market microstructure frictions in effect put a limit on the number of return observations per unit time interval that may be used productively in the computation of the realized volatility measures. As such, the realized volatility will necessarily be subject to a finite-sample (nonzero h) measurement error vis-a-vis the true (latent) integrated volatility, say

$$(5) \quad U_t(h) \equiv RV_t(h) - IV_t.$$

This observation was the original motivation for the development of the Barndorff-Nielsen and Shephard (2002a) asymptotic theory, which gives us a tool with which to study the errors for finite h .

⁴For notational simplicity, we focus our discussion on one-period return and volatility measures, but the general results and associated measurement error adjustment extend in a straightforward manner to the multiperiod case.

⁵The integrated volatility also plays a crucial role in the pricing of options; see, e.g., Garcia, Ghysels, and Renault (2003).

Specifically, assuming that the mean, $\{\mu_u, u \geq 0\}$, and volatility, $\{\sigma_u, u \geq 0\}$, processes are jointly independent of $\{W_u, u \geq 0\}$, it follows from Barndorff-Nielsen and Shephard (2002a, 2004a) that,⁶

$$(6) \quad z_t \equiv \sqrt{h^{-1}} \frac{U_t(h)}{\sqrt{2IQ_t}} \xrightarrow{\text{law}} \mathcal{N}(0, 1),$$

where the *integrated quarticity*, IQ_t , is defined by

$$(7) \quad IQ_t \equiv \int_{t-1}^t \sigma_u^4 du.$$

Moreover, under these same assumptions the integrated quarticity may be consistently estimated by the (standardized) *realized quarticity*

$$(8) \quad RQ_t(h) \equiv \frac{1}{h} \frac{1}{3} \sum_{i=1}^{1/h} r_{t-1+ih}^{(h)4}.$$

These asymptotic results allow for general model-free approximations to the distribution of the realized volatility error. In particular no parametric or finite moment assumptions are made on the $\{\mu_u, u \geq 0\}$ and $\{\sigma_u, u \geq 0\}$ processes to derive this result.

2.2. Practical Measurement Error Adjustments

The results discussed in the previous section imply that the time $t + 1$ realized volatility error is (approximately) serially uncorrelated and orthogonal to any variables (volatility forecasts) in the time t information set.⁷ This justifies the common use of realized volatility as a convenient simple and unbiased, albeit potentially noisy, benchmark in ex post volatility forecast evaluations and model comparisons.

Specifically, consider the Mincer–Zarnowitz style regressions of the realized volatility on a set of predetermined regressors (volatility forecasts) employed

⁶Importantly, the same asymptotic distributional results have recently been shown by Barndorff-Nielsen and Shephard (2004b) to hold in the case of leverage, or dependence between the $\{\sigma_u, u \geq 0\}$ and $\{W_u, u \geq 0\}$ processes, under the assumption that $\int_0^t \mu_u^2 du < \infty$. However, the measurement error adjustment procedures developed here involve an additional covariance term in the leverage case that is not covered by this theory. Hence, we continue under the maintained assumption of no leverage effects, returning to the general leverage case in Section 2.3.

⁷In the case of zero drift, equations (9) and (11) in Meddahi (2002) imply that $U_{t+1}(h) = 2 \sum_{i=1}^{1/h} \int_{t+(i-1)h}^{t+ih} (\int_{t+(i-1)h}^u \sigma_s dW_s) \sigma_u dW_u$, and as a result $E[U_{t+1}(h)|\mathcal{F}_t] = 0$ for all $h > 0$. With a nonzero drift, this same orthogonality condition holds approximately for small values of h .

in ABDL (2003) among others.⁸ Assuming that the underlying continuous time process satisfies a weak uniform integrability condition so that the consistency of $RQ_i(h)$ for IQ_i also guarantees convergence in mean (see, e.g., Billingsley (1995, p. 338) and Hoffmann-Jørgensen (1994, Section 5.13, p. 376)), it follows directly from the definition in equation (5) that,

$$\text{var}[RV_i(h)] = \text{var}[IV_i] + \text{var}[U_i(h)] + 2\text{cov}(U_i(h), IV_i).$$

Meanwhile, equations (6)–(8) readily imply that $\text{var}[U_i(h)] = 2hE[RQ_i(h)] + o(h)$. Moreover, it follows from Meddahi (2002) that if $\{\mu_u, u \geq 0\}$, $\{\sigma_u, u \geq 0\}$, and $\{W_u, u \geq 0\}$ are jointly independent, the covariance term vanishes. Obviously, the condition that the mean and volatility processes are independent is unduly strict in general. However, the covariance between the conditional mean and variance will invariably be small (negligible) over short horizons (daily or weekly) for small values of h . Formally, the resulting covariance between the realized volatility error and the integrated volatility will be of order less than h , as argued at length in Section 4.2 of Andersen, Bollerslev, and Diebold (2003). Combining these results, we therefore have

$$(9) \quad \text{var}[IV_i] = \text{var}[RV_i(h)] - 2hE[RQ_i(h)] + o(h).$$

Hence, any MSE type forecast evaluation criteria based on a comparison of the volatility forecasts with the ex post $RV_i(h)$ in place of IV_i will on average overstate the true variability of the forecast errors by $2hE[RQ_i(h)]$. In particular, ignoring the $o(h)$ term, it follows that the (feasible) R^2 from the commonly employed Mincer–Zarnowitz regression will underestimate the true predictability as measured by the (infeasible) R^2 from the regression of the future (latent) integrated volatility on the same set of predetermined regressors (volatility forecasts) by the multiplicative factor: $\text{var}[RV_i(h)]/\{\text{var}[RV_i(h)] - 2hE[RQ_i(h)]\}$.⁹

Meanwhile, the predictive regressions and related loss functions reported in the extant volatility literature are often formulated in terms of the realized standard deviation, $RV_i(h)^{1/2}$, or the logarithmic standard deviation, $\log RV_i(h)^{1/2}$. To properly gauge the true predictability in those situations, the sample variances of the transformed realized volatilities may be similarly replaced by (feasible) expressions for the true (latent) variances, $\text{var}[IV_i^{1/2}]$ and

⁸Although this is not required for the Barndorff-Nielsen and Shephard (2002a, 2004a, 2004b) asymptotic theory discussed in the previous section, the Mincer–Zarnowitz regression implicitly assumes that the variable of interest, i.e., the integrated and realized volatility processes, have finite second-order moments. This in turn requires that the fourth moment of σ_t is finite. This holds for any affine and log-normal diffusion, and is also satisfied for the GARCH diffusions considered in the Monte Carlo experiment discussed in the Appendix.

⁹As previously noticed by Meddahi (2002), the approximation in (9) also allows for the construction of more efficient (in the sense of MSE) model-free integrated volatility estimates, by downweighting the realized volatility by the multiplicative factor $\{\text{var}[RV_i(h)] - 2hE[RQ_i(h)]\}/\text{var}[RV_i(h)]$ and adding the constant $\{E[RV_i(h)]2hE[RQ_i(h)]\}/\text{var}[RV_i(h)]$.

$\text{var}[\log IV_t^{1/2}]$, respectively.¹⁰ To this end, it follows from equation (6) and a second-order Taylor series expansion of the square-root function of $RV_t(h)$ around IV_t , that conditional on the sample-path realization of the (latent) point-in-time volatility process (see the Appendix),

$$(10) \quad RV_t(h)^{1/2} \approx IV_t^{1/2} + 2^{-1/2} h^{1/2} IV_t^{-1/2} IQ_t^{1/2} z_t - \frac{h}{4} IV_t^{-3/2} IQ_t z_t^2.$$

Thus, subject to the necessary uniform integrability conditions on the underlying continuous-time process ensuring convergence in mean of the relevant quantities (see also Barndorff-Nielsen and Shephard (2004a)),

$$E[IV_t^{1/2}] = E[RV_t(h)^{1/2}] + \frac{h}{4} E[RV_t(h)^{-3/2} RQ_t(h)] + o(h),$$

so that, in particular,

$$(11) \quad \text{var}[IV_t^{1/2}] \\ = \text{var}[RV_t(h)^{1/2}] - \frac{h}{2} E[RV_t(h)^{1/2}] E[RV_t(h)^{-3/2} RQ_t(h)] + o(h).$$

The variance of the square-root of the realized volatility, as used in a number of previous empirical studies, obviously exceeds the expression in (11) by the absence of the second (positive) term on the right-hand side of the equation. This in turn will result in a downward bias in the R^2 's from the (feasible) Mincer–Zarnowitz predictive regressions formulated in terms of $RV_t(h)^{1/2}$ in place of $IV_t^{1/2}$.

By similar arguments (see the Appendix),¹¹

$$(12) \quad \log RV_t(h) \approx \log IV_t + 2^{1/2} h^{1/2} IV_t^{-1} IQ_t^{1/2} z_t - h IV_t^{-2} IQ_t z_t^2$$

and

$$(13) \quad [\log RV_t(h)]^2 \approx [\log IV_t]^2 + 2^{3/2} h^{1/2} IV_t^{-1} [\log IV_t] IQ_t^{1/2} z_t \\ - 2h IV_t^{-2} (1 - \log IV_t) IQ_t z_t^2,$$

¹⁰Any transformed unbiased forecast for IV_{t+1} will generally not be unbiased for $IV_{t+1}^{1/2}$ or $\log IV_{t+1}^{1/2}$. However, allowing for a nonzero intercept and a slope coefficient different from unity in the Mincer–Zarnowitz regression of the future transformed realized volatilities on the transformed forecasts explicitly corrects this (unconditional) bias in the forecasts; see also the discussion in Andersen, Bollerslev, and Meddahi (2004).

¹¹Interestingly, the Monte Carlo evidence in Barndorff-Nielsen and Shephard (2003) also suggests that the asymptotic approximation obtained by equating z_t^2 to one in (12), i.e., $[\log RV_t(h) - \log IV_t + h IV_t^{-2} IQ_t] / [2^{1/2} h^{1/2} IV_t^{-1} IQ_t^{1/2}]$, is closer to a standard normal distribution than the approximation obtained by applying the delta rule directly to (6), i.e., $[\log RV_t(h) - \log IV_t] / [2^{1/2} h^{1/2} IV_t^{-1} IQ_t^{1/2}]$.

so that again, subject to the necessary integrability conditions,

$$\begin{aligned}
 (14) \quad \text{var}[\log IV_t] &= \text{var}[\log RV_t(h)] \\
 &\quad - 2hE[RV_t(h)^{-2}(1 - \log RV_t(h))RQ_t(h)] \\
 &\quad - 2hE[\log RV_t(h)E[RV_t(h)^{-2}RQ_t(h)]] + o(h).
 \end{aligned}$$

The accuracy of the distributional assumption and second-order Taylor series expansions underlying the (feasible) expressions for the latent variances in equations (9), (11), and (14) are underscored by the simulation results for the baseline models reported in Table II of the Appendix.¹² It is evident that the simulated medians and ninety-percent confidence intervals for the asymptotic approximations to $\text{var}[IV_t]$, $\text{var}[IV_t^{1/2}]$, and $\text{var}[\log(IV_t^{1/2})]$ are extremely close to the simulated sampling distributions for the true variances (labelled $h = 1/\infty$) as long as the frequency of the returns used in the calculation of the realized volatility and quarticity measures, $RV_t(h)$ and $RQ_t(h)$, respectively, exceeds half-an-hour, or $h \leq 1/48$.¹³

Similar arguments could, of course, be applied with any other twice continuously differentiable function of integrated volatility in order to obtain an approximate value for $\text{var}[f(IV_t)]$. This in turn would allow for simple model-free approximations to the true (infeasible) R^2 's that would obtain in the hypothetical regressions of $f(IV_t)$ on any forecasts by scaling the (feasible) R^2 's from the corresponding regressions based on $f(RV_t(h))$ by the multiplicative adjustment factor, $\text{var}[f(RV_t(h))]/\{\text{var}[f(IV_t(h))]\}$.

2.3. Leverage Effects

The assumptions underlying the adjustment procedures discussed in the previous section formally rule out leverage effects. This is especially problematic for equity index returns, which are often found to be negatively correlated with future volatility (e.g., Black (1976) and Nelson (1991)). Meanwhile, as noted above, Barndorff-Nielsen and Shephard (2004b) have recently shown that the same approximate Central Limit Theorem in equation (6) remains valid in this situation, so that in particular $\text{var}[U_t(h)] = 2hE[RQ_t(h)] + o(h)$. Of course, the calculation of $\text{var}[IV_t]$ from equation (5) still requires the

¹²The accuracy of (6) and the corresponding CLT for $RV_t(h)^{1/2}$ and $\log(RV_t(h))$ based on the Δ method has also previously been investigated by Barndorff-Nielsen and Shephard (2003).

¹³To highlight the practical importance of the asymptotic adjustments developed above, we also calculated the simple naive estimators for $\text{var}[IV_t]$, $\text{var}[IV_t^{1/2}]$, and $\text{var}[\log(IV_t^{1/2})]$ ignoring the h terms in equations (9), (11), and (14), given by $\text{var}[RV_t(h)]$, $\text{var}[RV_t(h)^{1/2}]$, and $\text{var}[\log(RV_t(h)^{1/2})]$, respectively. For the two-factor affine diffusion model in the second panel of Table II and $h = 1/48$ these variances overstate the true variability of the integrated volatility measures by 44.9, 40.9, and 41.3 percent, respectively; see also the related analytical calculations in Andersen, Bollerslev, and Meddahi (2004).

$\text{cov}(U_i(h), IV_i)$ term, which is complicated by the possible dependence between the $\{\sigma_u, u \geq 0\}$ and $\{W_u, u \geq 0\}$ processes. However, building on the powerful eigenfunction stochastic volatility (ESV) class of models introduced by Meddahi (2001),¹⁴ it is possible to show that the covariance term generally takes the form $2hE[r_t] \text{cov}[r_t, RV_i(h)] + o(h)$.¹⁵ Consequently,

$$(15) \quad \begin{aligned} \text{var}[IV_i] &= \text{var}[RV_i(h)] - \text{var}[U_i(h)] - 2 \text{cov}[U_i(h), IV_i] \\ &= \text{var}[RV_i(h)] - 2hE[RQ_i(h)] \\ &\quad - 4hE[r_t] \text{cov}[r_t, RV_i(h)] + o(h). \end{aligned}$$

Hence, relative to equation (9) in the previous section, the leverage effect introduces the additional $4hE[r_t] \text{cov}[r_t, RV_i(h)]$ term in the (feasible) asymptotic approximation to $\text{var}[IV_i]$.

In actual empirical applications $2E[r_t]$ and $\text{cov}[r_t, RV_i(h)]$ will both generally be orders of magnitude smaller than $E[RQ_i(h)]$ so that, invariably, the magnitude of the new adjustment term will be negligible relative to the $2hE[RQ_i(h)]$ term. To illustrate, consider the five-minute high-frequency S&P500 and U.S. T-Bond futures returns spanning the period from January, 1990 through December, 2002.¹⁶ The relative importance of the leverage adjustment term, as measured by the daily $2E[r_t] \text{cov}[r_t, RV_i(h)]/E[RQ_i(h)]$ ratios, equals -7.85×10^{-5} and -6.94×10^{-4} for each of the two markets. Also, for the DM/\$, Yen/\$, and Yen/DM half-hour returns underlying the empirical results in ABDL (2003) discussed below, these same daily ratios for the full December, 1986 through June, 1999 sample period equal 1.01×10^{-4} , -7.87×10^{-5} , and 3.54×10^{-4} , respectively, clearly an inconsequential addition to the approximation for $\text{var}[IV_i]$ in (9).

These empirical observations are further corroborated by the Monte Carlo simulation results for the leverage models with constant as well as time-varying drifts reported in Tables III and IV. The medians in the asymptotic approximations to $\text{var}[IV_i]$, $\text{var}[IV_i^{1/2}]$, and $\text{var}[\log(IV_i^{1/2})]$ in equations (9), (11), and (14), respectively, derived under the assumption of no leverage, are all right on the true medians (labelled $h = 1/\infty$). Moreover, as long as the frequency of the returns used in the calculation of the realized volatility and quarticity measures

¹⁴The use of eigenfunctions in modelling Markovian time series was pioneered by Chen, Hansen, and Scheinkman (2000). The ESV class of models is very general, encompassing all of the continuous-time volatility models most commonly employed in the literature, including the GARCH diffusion model of Nelson (1990), the log-normal diffusion model popularized by Hull and White (1987) and Wiggins (1987), and the square-root diffusion model of Heston (1993), along with multifactor extensions of all these models.

¹⁵For a formal proof, see the earlier Andersen, Bollerslev, and Meddahi (2003) working paper version of this note.

¹⁶These data have previously been analyzed in Andersen, Bollerslev, Diebold, and Vega (2003) from a very different perspective. We refer the reader to that study for a more detailed description of the data source and return construction.

exceeds half an hour, or $h \leq 1/48$, the simulated distributions for the leverage models are indistinguishable from the corresponding distributions for the same models without leverage reported in Table II.

In short, the realized volatility measurement error adjustment procedures developed in the preceding section remain highly accurate in empirically realistic situations, allowing for both leverage and time-varying drift. We next turn to a reinterpretation of the empirical evidence related to the Mincer–Zarnowitz volatility regressions reported in ABDL (2003) based on an application of these procedures.

3. ABDL (2003) REVISITED

The forecast comparisons in ABDL (2003) are based on daily realized volatilities constructed from high-frequency half-hourly, or $h = 1/48$, spot exchange rates for the U.S. dollar, the Deutschemark, and the Japanese yen, spanning twelve-and-a-half years.¹⁷ Separate forecast evaluation regressions are reported for the “in-sample” period comprised of the 2,449 “regular” trading days from December 1, 1986 through December 1, 1996, and the shorter “out-of-sample” forecast period consisting of the 596 days from December 2, 1996 through June 30, 1999. Separate results are also reported for one-day-ahead and ten-days-ahead forecasts. However, for all series and both sample periods and forecast horizons, a simple AR(5) model estimated directly from the realized volatilities generally performs as well or better than any of the many alternative models considered, including several GARCH type models estimated directly from the high-frequency data (both with and without corrections for the pronounced intradaily seasonal pattern in volatility). The representative R^2 's for the DM/\$, Yen/\$, and Yen/DM forecast regressions for $RV_{t+1}(1/48)$, $RV_{t+1}(1/48)^{1/2}$, $\log RV_{t+1}(1/48)^{1/2}$, $RV_{t+10,10}(1/48)$, $RV_{t+10,10}(1/48)^{1/2}$, and $\log RV_{t+10,10}(1/48)^{1/2}$, where $RV_{t+10,10}(1/48) \equiv RV_{t+1}(1/48) + RV_{t+2}(1/48) + \dots + RV_{t+10}(1/48)$, as reported in ABDL (2003) and the accompanying Appendix, are given in square brackets in Table I.¹⁸

By failing to account for the measurement errors in the future realized volatilities, these R^2 's understate the true degree of predictability in the (latent) integrated volatilities. This problem is rectified by the main entries in Table I, which report the adjusted R^2 's obtained by applying the (feasible) asymptotic approximations in equations (9), (11), and (14), along with the rel-

¹⁷The high-frequency data were generously provided by Olsen & Associates in Zürich, Switzerland; see Dacorogna, Gencay, Müller, Olsen, and Pictet (2001) for further discussion of the data capture, filtering, and return construction.

¹⁸The out-of-sample period contains a “once-in-a-generation” move in the Japanese yen on October 8, 1998. Somewhat higher R^2 's, but qualitatively similar results, were obtained by excluding this and the neighboring two days; see ABDL (2003) and the accompanying Appendix for further discussion and sensitivity analysis along these lines.

evant multiplicative adjustment factors.¹⁹ The results are quite striking. For some of the forecast horizons and rates, the “true” R^2 's exceed the standard predictive R^2 's, as reported in ABDL (2003), by up to *forty percent*. For instance, the in-sample, one-day-ahead R^2 for the DM/\$ series given in the very first entry in the table equals .219, whereas the true (albeit estimated) R^2 is substantially higher at .314. As such, the results clearly highlight the importance of appropriately adjusting for measurement error when assessing the quality of volatility forecasts in practical empirical applications.

Interestingly, the numerical values for the adjusted R^2 's for the DM-dollar series in Table I are quite close to the exact theoretical R^2 's implied by the specific two-factor affine diffusion discussed in Andersen, Bollerslev, and Meddahi (2004). This is especially noteworthy insofar as the parameter values for this model are based on the identical DM-dollar sample underlying the results reported in Table I. This suggests that the simple AR(5) models for the realized volatilities estimated in ABDL (2003)—when adjusted for the measurement error problem—capture a degree of predictability that is consis-

TABLE I
ABDL (2003) ADJUSTED PREDICTIVE R^2 'S

	IV	$IV^{1/2}$	$\log IV^{1/2}$
In-sample, one-day-ahead			
DM/\$.314 [.219]	.399 [.351]	.482 [.431]
Yen/\$.315 [.229]	.412 [.374]	.476 [.433]
Yen/DM	.450 [.361]	.559 [.499]	.630 [.567]
Out-of-sample, one-day-ahead			
DM/\$.200 [.158]	.296 [.246]	.350 [.285]
Yen/\$.230 [.197]	.366 [.338]	.419 [.373]
Yen/DM	.215 [.189]	.378 [.344]	.483 [.424]
In-sample, ten-days-ahead			
DM/\$.411 [.374]	.463 [.436]	.499 [.473]
Yen/\$.386 [.355]	.414 [.396]	.424 [.407]
Yen/DM	.536 [.513]	.606 [.589]	.653 [.637]
Out-of-sample, ten-days-ahead			
DM/\$.182 [.168]	.209 [.195]	.228 [.213]
Yen/\$.197 [.187]	.287 [.279]	.347 [.336]
Yen/DM	.186 [.178]	.301 [.293]	.401 [.390]

Note: The table reports the adjusted predictive R^2 's from the Mincer-Zarnowitz regressions of the realized volatilities on the AR(5) volatility forecasts in ABDL (2003), along with the corresponding unadjusted R^2 's (in square brackets). The realized volatility measures are constructed from high-frequency half-hour returns. The “in-sample” period covers December 1, 1986 through December 1, 1996, while the “out-sample” period spans December 2, 1996 through June 30, 1999.

¹⁹The adjustments are constructed separately for each series and for the in-sample and out-of-sample periods using the corresponding realized volatility and quarticity series.

tent with that implied by a conventional two-factor affine model. This type of benchmarking of the true predictive power of such reduced-form forecast procedures relative to that of a specific continuous-time volatility model would, of course, be impossible without the type of measurement error correction developed here.

4. CONCLUDING REMARKS

Building on the recent theoretical results of Barndorff-Nielsen and Shephard (2002a, 2004a, 2004b), this note develops a set of simple and practically feasible expressions for calculating true measures of return volatility predictability relative to that of the corresponding underlying (latent) integrated volatility. The procedures are general and could be applied in the evaluation of any volatility forecasts. The analytical results for the eigenfunction stochastic volatility class of models and accompanying simulation based evidence confirm that the procedures work equally well in situations with pronounced leverage effects. On specifically applying the procedures to the ex post forecast evaluation regressions reported in ABDL (2003), we document sizeable downward biases in terms of the previously reported predictive powers. More generally, the practical techniques developed here hold the promise for further development of new and improved easy-to-implement volatility forecasting procedures guided by proper benchmark comparisons. The techniques should also prove useful in more effectively calibrating the type of continuous-time models routinely employed in modern asset pricing theories.

Dept. of Finance, Kellogg School of Management, Northwestern University, Evanston, IL 60208, and NBER, U.S.A.; t-andersen@kellogg.northwestern.edu,
Dept. of Economics, Duke University, Durham, NC 27708, and NBER, U.S.A.;
boller@econ.duke.edu,

and

Département de Sciences Économiques, CIRANO, CIREQ, Université de Montréal, C.P. 6128, Succursale Centre-ville, Montréal (Québec), H3C 3J7, Canada;
nour.meddahi@umontreal.ca; <http://www.mapageweb.umontreal.ca/meddahin/>.

Manuscript received December, 2002; final revision received June, 2004.

APPENDIX A: MONTE CARLO SIMULATIONS

In order to assess the accuracy of the distributional assumptions and second-order Taylor series expansions underlying the asymptotic approximations in (9), (11), and (14) in empirically relevant specifications and sample sizes compatible with those of ABDL (2003), Tables II–IV report the simulated medians and ninety-percent confidence intervals (in square brackets) across 1,000 replications, each consisting of 2,500 “days.” We report the results for a total of nine different continuous-time models along with $1/h = 288, 96, 48,$ and $1,$ corresponding to the use of “5-minute,” “15-minute,” “half-hourly,” and “daily” returns.

The first three models reported in Table II fix the mean returns at zero, and assume that the volatility and the Brownian motion driving the price process are independent, i.e., no leverage

TABLE II
ASYMPTOTIC VARIANCE APPROXIMATIONS
BASELINE VOLATILITY MODELS

h	$\text{var}[IV_t]$	$\text{var}[IV_t^{1/2}]$	$\text{var}[\log(IV_t^{1/2})]$
GARCH(1, 1) diffusion			
$1/\infty$.170 [.117, .265]	.0647 [.0518, .0853]	.138 [.112, .168]
1/288	.170 [.116, .266]	.0647 [.0517, .0854]	.138 [.112, .168]
1/96	.171 [.116, .266]	.0648 [.0520, .0859]	.138 [.112, .168]
1/48	.170 [.115, .268]	.0650 [.0520, .0861]	.139 [.112, .169]
1	.167 [.0923, .313]	.208 [.175, .248]	1.19 [1.08, 1.30]
Two-factor affine			
$1/\infty$.0259 [.0222, .0316]	.0126 [.0111, .0145]	.0261 [.0235, .0290]
1/288	.0260 [.0222, .0316]	.0126 [.0111, .0145]	.0261 [.0234, .0291]
1/96	.0260 [.0221, .0315]	.0126 [.0111, .0146]	.0263 [.0235, .0294]
1/48	.0259 [.0219, .0315]	.0127 [.0112, .0148]	.0267 [.0238, .0302]
1	.0245 [.00617, .0462]	.136 [.125, .149]	1.07 [.973, 1.16]
Log-normal diffusion			
$1/\infty$.145 [.0640, .333]	.0544 [.0328, .0946]	.109 [.0764, .163]
1/288	.144 [.0643, .338]	.0543 [.0330, .0943]	.109 [.0762, .163]
1/96	.145 [.0642, .337]	.0546 [.0330, .0952]	.109 [.0766, .164]
1/48	.144 [.0635, .341]	.0547 [.0331, .0953]	.109 [.0769, .165]
1	.145 [.0529, .390]	.177 [.127, .252]	1.15 [1.05, 1.27]

Note: The table reports the simulated medians and ninety-percent confidence intervals (in square brackets) for the asymptotic approximations in equations (9), (11), and (14) across 1,000 replications, each consisting of 2,500 “days.”

effects. In the general equation

$$\begin{aligned} d \log(S_t) &= \mu dt + \sigma_t dW_t \\ &= \mu dt + \sigma_t [\rho_1 dW_{1,t} + \rho_2 dW_{2,t} + \sqrt{1 - \rho_1^2 - \rho_2^2} dW_{3,t}], \end{aligned}$$

one has $\mu = \rho_1 = \rho_2 = 0$, which leads to

$$d \log(S_t) = \sigma_t dW_{3,t}.$$

TABLE III
ASYMPTOTIC VARIANCE APPROXIMATIONS
VOLATILITY MODELS WITH LEVERAGE AND CONSTANT DRIFT

h	$\text{var}[IV_t]$	$\text{var}[IV_t^{1/2}]$	$\text{var}[\log(IV_t^{1/2})]$
GARCH(1, 1) diffusion			
$1/\infty$.170 [.117, .265]	.0647 [.0518, .0853]	.138 [.112, .168]
$1/288$.170 [.116, .262]	.0647 [.0518, .0849]	.138 [.112, .168]
$1/96$.170 [.116, .265]	.0647 [.0520, .0852]	.138 [.112, .168]
$1/48$.170 [.115, .268]	.0650 [.0520, .0853]	.138 [.113, .168]
1	.165 [.0964, .303]	.205 [.173, .247]	1.16 [1.07, 1.27]
Two-factor affine			
$1/\infty$.0259 [.0222, .0316]	.0126 [.0111, .0145]	.0261 [.0235, .0290]
$1/288$.0260 [.0221, .0317]	.0126 [.0111, .0145]	.0261 [.0234, .0292]
$1/96$.0261 [.0222, .0321]	.0127 [.0112, .0145]	.0263 [.0236, .0294]
$1/48$.0262 [.0221, .0323]	.0129 [.0112, .0150]	.0267 [.0239, .0301]
1	.0370 [.0155, .0654]	.139 [.128, .154]	1.07 [.973, 1.16]
Log-normal diffusion			
$1/\infty$.145 [.0640, .333]	.0544 [.0328, .0946]	.109 [.0764, .163]
$1/288$.144 [.0640, .336]	.0545 [.0329, .0941]	.109 [.0763, .162]
$1/96$.145 [.0637, .337]	.0545 [.0331, .0952]	.109 [.0763, .163]
$1/48$.146 [.0635, .340]	.0547 [.0335, .0943]	.110 [.0766, .162]
1	.145 [.0515, .375]	.177 [.127, .251]	1.15 [1.04, 1.27]

Note: See Table II.

The numbers in the first panel refer to the GARCH(1, 1) diffusion analyzed in Andersen and Bollerslev (1998),

$$d\sigma_t^2 = .035(.636 - \sigma_t^2) dt + .144\sigma_t^2 dW_{1,t}.$$

The second panel gives the results for the two-factor affine diffusion estimated by Bollerslev and Zhou (2002), $\sigma_t^2 = \sigma_{1,t}^2 + \sigma_{2,t}^2$, where

$$d\sigma_{1,t}^2 = .5708(.3257 - \sigma_{1,t}^2) dt + .2286\sigma_{1,t} dW_{1,t},$$

$$d\sigma_{2,t}^2 = .0757(.1786 - \sigma_{2,t}^2) dt + .1096\sigma_{2,t} dW_{2,t}.$$

TABLE IV
ASYMPTOTIC VARIANCE APPROXIMATIONS
VOLATILITY MODELS WITH LEVERAGE AND TIME-VARYING DRIFT

h	$\text{var}[IV_t]$	$\text{var}[IV_t^{1/2}]$	$\text{var}[\log(IV_t^{1/2})]$
GARCH(1, 1) diffusion			
$1/\infty$.170 [.117, .265]	.0647 [.0518, .0853]	.138 [.112, .168]
$1/288$.170 [.116, .262]	.0648 [.0518, .0850]	.138 [.112, .168]
$1/96$.171 [.116, .266]	.0647 [.0521, .0853]	.138 [.112, .168]
$1/48$.171 [.116, .270]	.0652 [.0521, .0855]	.138 [.113, .168]
1	.196 [.116, .398]	.225 [.189, .272]	1.17 [1.05, 1.27]
Two-factor affine			
$1/\infty$.0259 [.0222, .0316]	.0126 [.0111, .0145]	.0261 [.0235, .0290]
$1/288$.0261 [.0222, .0318]	.0126 [.0111, .0146]	.0262 [.0235, .0292]
$1/96$.0264 [.0225, .0324]	.0128 [.0113, .0147]	.0265 [.0238, .0296]
$1/48$.0268 [.0226, .0329]	.0131 [.0115, .0152]	.0272 [.0243, .0306]
1	.0661 [.0362, .106]	.163 [.150, .180]	1.09 [.998, 1.18]
Log-normal diffusion			
$1/\infty$.145 [.0640, .333]	.0544 [.0328, .0946]	.109 [.0764, .163]
$1/288$.145 [.0641, .336]	.0545 [.0329, .0942]	.109 [.0764, .163]
$1/96$.145 [.0639, .337]	.0546 [.0332, .0954]	.109 [.0764, .163]
$1/48$.146 [.0636, .342]	.0548 [.0336, .0948]	.110 [.0768, .163]
1	.164 [.0608, .543]	.192 [.136, .281]	1.15 [1.04, 1.27]

Note: See Table II.

These parameter values were obtained from estimation based on the identical DM-dollar sample used in ABDL (2003). The third set of numbers refer to the log-normal diffusion reported in Andersen, Benzoni, and Lund (2002) with volatility dynamics governed by

$$d \log(\sigma_t^2) = -.0136[.8382 + \log(\sigma_t^2)] dt + .1148 dW_{1,t}.$$

All of the models in Table II satisfy the Barndorff-Nielsen and Shephard (2002a, 2004a) regularity conditions discussed in Sections 2.1 and 2.2.

The results reported in Table III are based on the same three volatility specifications, but incorporate a positive drift and strong leverage effects. For the one-factor GARCH and log-normal diffusions,

$$d \log(S_t) = .0314 dt + \sigma_t[-.576 dW_{1,t} + \sqrt{1 - .576^2} dW_{3,t}],$$

where the values for the drift and leverage parameters are taken from Andersen, Benzoni, and Lund (2002). For the two-factor affine model the instantaneous return dynamic is governed by

$$d \log(S_t) = .0314 dt + \sigma_t[.9 dW_{1,t} - .4 dW_{2,t} + \sqrt{1 - .9^2 - .4^2} dW_{3,t}],$$

with the two leverage parameters adapted from the estimates reported in Chernov, Gallant, Ghysels, and Tauchen (2003).

In addition to the contemporaneous correlation between the return and volatility for the leverage models in Table III, the last set of models in Table IV also include a volatility feedback, or ARCH-in-mean, effect in the drift component. Specifically, for the two one-factor models,

$$d \log(S_t) = (.0314 + .3\sigma_t^2) dt + \sigma_t[-.576 dW_{1,t} + \sqrt{1 - .576^2} dW_{3,t}],$$

while for the two-factor model,

$$d \log(S_t) = (.0314 + .3\sigma_t^2) dt + \sigma_t[.9 dW_{1,t} - .4 dW_{2,t} + \sqrt{1 - .9^2 - .4^2} dW_{3,t}].$$

The value of the slope coefficient in the drift is taken from Chernov (2003).

APPENDIX B: TECHNICAL PROOFS

PROOF OF EQUATIONS (10), (11), (12), (13), AND (14): Let $f(\cdot)$ be a twice-differentiable function. By (6) and a second-order Taylor approximation of $f(\cdot)$ at the point $RV_t(h)$ around IV_t , it follows that

$$(A.1) \quad f(RV_t(h)) \approx f(IV_t) + f'(IV_t)\sqrt{2hIQ_t}z_t + \frac{1}{2}f''(IV_t)2hIQ_tz_t^2.$$

Consequently,

$$\begin{aligned} E[f(RV_t(h))] &= E[f(IV_t)] + E[f'(IV_t)\sqrt{2hIQ_t}z_t] + \frac{1}{2}E[f''(IV_t)2hIQ_tz_t^2] + o(h) \\ &= E[f(IV_t)] + E[f'(IV_t)\sqrt{2hIQ_t}E[z_t|\sigma_u, t-1 \leq u \leq t]] \\ &\quad + \frac{1}{2}E[f''(IV_t)2hIQ_tE[z_t^2|\sigma_u, t-1 \leq u \leq t]] + o(h) \\ &= E[f(IV_t)] + \frac{1}{2}E[f''(IV_t)2hIQ_t] + o(h), \end{aligned}$$

so that,

$$(A.2) \quad E[f(RV_t(h))] = E[f(IV_t)] + \frac{1}{2}E[f''(RV_t(h))2hRQ_t(h)] + o(h)$$

provided $E[f''(RV_t(h))RQ_t(h)] - E[f''(IV_t)IQ_t] = o(1)$. Equations (10), (12), and (13) follow by applying (A.1) to the functions $f_1(x) = x^{1/2}$, $f_2(x) = \log(x)$, and $f_3(x) = \log(x)^2$, where $f_1'(x) = 2^{-1}x^{-1/2}$, $f_1''(x) = -2^{-2}x^{-3/2}$, $f_2'(x) = x^{-1}$, $f_2''(x) = -x^{-2}$, $f_3'(x) = 2x^{-1}\log(x)$, and $f_3''(x) = 2x^{-2}(1 - \log(x))$. Applying (A.2) to the function $f_1(\cdot)$ results in (11). Similarly, applying (A.2) to the functions $f_2(\cdot)$ and $f_3(\cdot)$ yields (14). Q.E.D.

REFERENCES

- ANDERSEN, T. G., L. BENZONI, AND J. LUND (2002): "An Empirical Investigation of Continuous-Time Equity Return Models," *Journal of Finance*, 57, 1239–1284.
- ANDERSEN, T. G., AND T. BOLLERSLEV (1998): "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts," *International Economic Review*, 39, 885–905.
- ANDERSEN, T. G., T. BOLLERSLEV, AND F. X. DIEBOLD (2003): "Parametric and Nonparametric Volatility Measurement," in *Handbook of Financial Econometrics*, ed. by Y. Aït-Sahalia and L. P. Hansen. Amsterdam: North-Holland Press, forthcoming.
- ANDERSEN, T. G., T. BOLLERSLEV, F. X. DIEBOLD, AND P. LABYS (2001): "The Distribution of Exchange Rate Volatility," *Journal of the American Statistical Association*, 96, 42–55.
- (2003): "Modeling and Forecasting Realized Volatility," *Econometrica*, 71, 579–625.
- ANDERSEN, T. G., T. BOLLERSLEV, F. X. DIEBOLD, AND C. VEGA (2003): "The Evolving Effects of Macroeconomic News on Global Stock, Bond and Foreign Exchange Markets," Unpublished Manuscript, University of Pennsylvania.
- ANDERSEN, T. G., T. BOLLERSLEV, AND N. MEDDAHI (2003): "Correcting the Errors: Volatility Forecast Evaluation Using High-Frequency Data and Realized Volatilities," Unpublished Manuscript, Northwestern, Duke, and Montreal Universities.
- (2004): "Analytic Evaluation of Volatility Forecasts," *International Economic Review*, forthcoming.
- ANDREOU, E., AND E. GHYSELS (2002): "Rolling-Sampling Volatility Estimators: Some New Theoretical, Simulation and Empirical Results," *Journal of Business and Economic Statistics*, 20, 363–376.
- BAI, X., J. R. RUSSELL, AND G. C. TIAO (2000): "Beyond Merton's Utopia: Effects of Non-Normality and Dependence on the Precision of Variance Estimates Using High-Frequency Financial Data," Unpublished Manuscript, University of Chicago.
- BARNDORFF-NIELSEN, O. E., AND N. SHEPHARD (2001): "Non-Gaussian Ornstein–Uhlenbeck-Based Models and Some of Their Uses in Financial Economics," *Journal of the Royal Statistical Society, Ser. B*, 63, 167–207.
- (2002a): "Econometric Analysis of Realised Volatility and Its Use in Estimating Stochastic Volatility Models," *Journal of the Royal Statistical Society, Ser. B*, 64, 253–280.
- (2002b): "Estimating Quadratic Variation Using Realised Variance," *Journal of Applied Econometrics*, 17, 457–477.
- (2003): "How Accurate is the Asymptotic Approximation to the Distribution of Realised Volatility?" in *Identification and Inference for Econometric Models: A Festschrift for Thomas J. Rothenberg*, ed. by D. Andrews, J. Powell, P. Ruud, and J. Stock. Cambridge, U.K.: Cambridge University Press, forthcoming.
- (2004a): "Econometric Analysis of Realized Covariation: High Frequency Based Covariance, Regression, and Correlation in Financial Economics," *Econometrica*, 72, 885–925.
- (2004b): "A Feasible Central Limit Theory for Realised Volatility under Leverage," Unpublished Manuscript, Nuffield College, Oxford University.
- BILLINGSLEY, P. (1995): *Probability and Measure* (Third Ed.). New York: John Wiley & Sons.
- BLACK, F. (1976): "Studies in Stock Price Volatility Changes," in *Proceedings of the 1976 Business Meeting of the Business and Statistics Section*, American Statistical Association, 177–181.
- BOLLERSLEV, T., AND H. ZHOU (2002): "Estimating Stochastic Volatility Diffusion Using Conditional Moments of Integrated Volatility," *Journal of Econometrics*, 109, 33–65.
- CHEN, X., L. P. HANSEN, AND J. SCHEINKMAN (2000): "Principal Components and the Long Run," Unpublished Manuscript, University of Chicago.
- CHERNOV, M. (2003): "Empirical Reverse Engineering of the Pricing Kernel," *Journal of Econometrics*, 116, 329–364.
- CHERNOV, M., R. GALLANT, E. GHYSELS, AND G. TAUCHEN (2003): "Alternative Models for Stock Price Dynamics," *Journal of Econometrics*, 116, 225–257.

- CHONG, Y. Y., AND D. HENDRY (1986): "Econometric Evaluation of Linear Macro-Economic Models," *Review of Economic Studies*, 53, 671–690.
- COMTE, F., AND E. RENAULT (1998): "Long Memory in Continuous Time Stochastic Volatility Models," *Mathematical Finance*, 8, 291–323.
- DACOROGNA, M. M., R. GENÇAY, U. A. MÜLLER, R. B. OLSEN, AND O. V. PICTET (2001): *An Introduction to High-Frequency Finance*. San Diego: Academic Press.
- GARCIA, R., E. GHYSELS, AND E. RENAULT (2003): "Option Pricing Models," in *Handbook of Financial Econometrics*, ed. by Y. Aït-Sahalia and L. P. Hansen. Amsterdam: North-Holland Press, forthcoming.
- HESTON, S. L. (1993): "A Closed Form Solution for Options with Stochastic Volatility with Applications to Bond and Currency Options," *Review of Financial Studies*, 6, 327–344.
- HOFFMAN-JØRGENSEN, J. (1994): *Probability with a View Toward Statistics*, Vol. 1, Chapman and Hall Probability Series. New York: Chapman & Hall.
- HULL, J., AND A. WHITE (1987): "The Pricing of Options on Assets with Stochastic Volatilities," *The Journal of Finance*, 42, 281–300.
- MEDDAHI, N. (2001): "An Eigenfunction Approach for Volatility Modeling," CIRANO Working Paper, 2001s–70.
- (2002): "A Theoretical Comparison Between Integrated and Realized Volatility," *Journal of Applied Econometrics*, 17, 479–508.
- MINCER, J., AND V. ZARNOWITZ (1969): "The Evaluation of Economic Forecasts," *Economic Forecasts and Expectation*, ed. by J. Mincer. New York: Columbia University Press, 3–46.
- NELSON, D. B. (1990): "ARCH Models as Diffusion Approximations," *Journal of Econometrics*, 45, 7–39.
- (1991): "Conditional Heteroskedasticity in Asset Returns: A New Approach," *Econometrica*, 59, 347–370.
- PROTTER, P. (2004): *Stochastic Integration and Differential Equations* (Second Ed.). New York: Springer-Verlag.
- WIGGINS, J. B. (1987): "Options Values under Stochastic Volatility: Theory and Empirical Estimates," *Journal of Financial Economics*, 19, 351–372.