

Range-Based Estimation of Stochastic Volatility Models

or

Exchange Rate Dynamics are More Interesting Than You Think

Sassan Alizadeh

Bear, Stearns & Co. Inc.

Michael W. Brandt

Wharton School
University of Pennsylvania
and NBER

Francis X. Diebold

Stern School, NYU
University of Pennsylvania
and NBER

This draft / print: December 20, 1999

Abstract: We propose using the price range, a recently-neglected volatility proxy with a long history in finance, in the estimation of stochastic volatility models. We show both theoretically and empirically that the log range is approximately Gaussian, in sharp contrast to popular volatility proxies, such as log absolute or squared returns. Hence Gaussian quasi-maximum likelihood estimation based on the range is not only simple, but also highly efficient. We illustrate and enrich our theoretical results with a Monte Carlo study and a substantive empirical application to daily exchange rate volatility. Our empirical work produces sharp conclusions. In particular, the evidence points strongly to the inadequacy of one-factor volatility models, favoring instead two-factor models with one highly persistent factor and one quickly mean reverting factor.

Acknowledgments: This work was supported by the National Science Foundation. Siem Koopman graciously shared both his wisdom and his Ox routines. We thank Steve Brown, Freddy Delbaen, Rob Engle, Joel Hasbrouck, Chris Rogers, Steve Satchell, George Tauchen, and the participants at the 6th Annual CAP Workshop on Mathematical Finance at Columbia University and the Wharton Micro Finance Brown Bag Luncheon for their comments and discussions. Canlin Li provided excellent research assistance. All remaining errors are ours.

Range-Based Estimation of Stochastic Volatility Models

or

Exchange Rate Dynamics are More Interesting Than You Think

Abstract: We propose using the price range, a recently-neglected volatility proxy with a long history in finance, in the estimation of stochastic volatility models. We show both theoretically and empirically that the log range is approximately Gaussian, in sharp contrast to popular volatility proxies, such as log absolute or squared returns. Hence Gaussian quasi-maximum likelihood estimation based on the range is not only simple, but also highly efficient. We illustrate and enrich our theoretical results with a Monte Carlo study and a substantive empirical application to daily exchange rate volatility. Our empirical work produces sharp conclusions. In particular, the evidence points strongly to the inadequacy of one-factor volatility models, favoring instead two-factor models with one highly persistent factor and one quickly mean reverting factor.

1. Introduction

Volatility is a central concept in finance, whether in asset pricing, portfolio choice, or risk management. Not long ago, theoretical models routinely assumed constant volatility (e.g., Merton, 1969; Black and Scholes, 1973). Today, however, we widely acknowledge that volatility is both time-varying and predictable (e.g., Andersen and Bollerslev, 1997). *Stochastic volatility* models are now central; they have emerged as the paradigm of choice for modeling time-varying and predictable volatility. Discrete- and continuous-time stochastic volatility models are extensively used in theoretical finance, empirical finance, and financial econometrics, both in academe and industry (e.g., Hull and White, 1987; Heston, 1993; Bates, 1996; Ghysels, Harvey, and Renault, 1996; Jarrow, 1998).

Unfortunately, the estimation of stochastic volatility models has proven quite difficult. The Gaussian quasi-maximum likelihood estimation (QMLE) approach of Ruiz (1994) and Harvey, Ruiz, and Shephard (1994), which initially seemed appealing because of its simplicity, fell by the wayside as it became apparent that stochastic volatility models are highly non-Gaussian. The problem is that standard volatility proxies, such as log absolute or squared returns, are contaminated with highly non-Gaussian measurement error (e.g., Andersen and Sorensen, 1996). Unfortunately, highly non-Gaussian volatility proxies produce highly inefficient Gaussian quasi-maximum likelihood estimators.

The literature therefore turned toward alternative estimators. In particular, attention turned to variants of the generalized method of moments (GMM) that use moment conditions obtained by integrating out volatility, either through simulations (e.g., Duffie and Singleton, 1993) or analytically (e.g., Singleton, 1997). These estimators, however, can also be highly inefficient,

depending on the choice of moment conditions and weighting matrix. Although recent work has tried to maximize the efficiency of these GMM estimators through the optimal choice of moment conditions, the empirical implementation of this approach remains challenging (Gallant, Hsieh, and Tauchen, 1997; Gallant, Hsu, and Tauchen, 1999; Chernoff and Ghysels, 1999).

Another literature focuses on likelihood-based estimation using importance sampling or Markov chain Monte Carlo methods, whether in a Bayesian setting (e.g., Jacquier, Polson, and Rossi, 1994) or in a classical setting (e.g., Danielsson 1994; Kim, Shephard, and Chib, 1998; Sandmann and Koopman, 1998). Such simulation methods can in principle deliver highly accurate approximations to the exact maximum likelihood estimator, but they are not widely adopted due to practical considerations. In particular, all of these methods are computationally intense and rely on assumptions that are hard to check in practice, such as the accuracy of mixture approximations to non-Gaussian distributions and the convergence of simulated Markov chains to their steady state.

Motivated both by the popularity of stochastic volatility models and by the difficulties associated with estimating them, we propose a simple yet highly efficient estimation method based on the range, defined here as the difference between the highest and lowest log asset price during a discrete sampling interval. The range is a volatility proxy with a long and colorful history in finance (e.g., Garman and Klass, 1980; Parkinson, 1980; Beckers, 1983; Ball and Torous 1984; Rogers and Satchell, 1991; Anderson and Bollerslev, 1998). Curiously, however, it has been neglected in the recent stochastic volatility literature.¹

¹ Gallant, Hsu, and Tauchen (1999) also make use of the log range, albeit with a very different estimator. Although they are aware of the efficiency of the log range, they are unaware of and do not exploit its normality.

We set the stage for the paper in Section 2, in which we describe a general class of continuous-time stochastic volatility models and the particular discretization that we exploit. In Section 3 we use both analytical and numerical methods to motivate and establish the remarkable normality of the log range. In addition, we note that the log range is a highly efficient volatility measure, a fact known at least since Parkinson (1980) and recently formalized by Andersen and Bollerslev (1998). The approximate normality and high efficiency of the log range suggest its use in Gaussian quasi-maximum likelihood estimation. We pursue this idea first in the Monte Carlo study of Section 4, which reveals huge efficiency gains from our approach relative to traditional methods, and then in a substantive empirical analysis of exchange rate volatility in Section 5, which delivers sharp new insights. In Section 6 we summarize, conclude, and sketch directions for future research.

2. Stochastic Volatility

A Continuous-Time Stochastic Volatility Model

In a generic continuous-time stochastic volatility model, the price S of a security evolves as a diffusion with instantaneous drift μ and volatility σ . Both the drift and volatility depend on a latent state variable v , which itself evolves as a diffusion. Formally, we write:

$$\begin{aligned} dS_t &= \mu(S_t, v_t)dt + \sigma(S_t, v_t)dW_{S_t} \\ dv_t &= \alpha(S_t, v_t)dt + \beta(S_t, v_t)dW_{v_t}, \end{aligned} \tag{1}$$

where W_{S_t} and W_{v_t} are two Wiener processes with correlation $E_t[dW_{S_t} dW_{v_t}] = \theta(S_t, v_t)$. The functions α and β govern the drift and volatility of the state variable process.

The stochastic volatility literature contains numerous variations on the generic model (1). In this paper we work with a first-order parameterization, which is rich enough to be interesting,

yet simple enough to permit a streamlined exposition:

$$\frac{dS_t}{S_t} = \mu dt + \sigma_t dW_{St} \quad (2)$$

$$d \ln \sigma_t = \alpha (\ln \bar{\sigma} - \ln \sigma_t) dt + \beta dW_{vt}.$$

The simple stochastic volatility model (2) emerges from the general model (1) when

$\sigma(S_t, v_t) = \sigma_t S_t$, $\sigma_t = \exp(v_t)$, $\alpha(S_t, v_t) = \alpha (\ln \bar{\sigma} - v_t)$, $\beta(S_t, v_t) = \beta$, and $\theta(S_t, v_t) = 0$. In this

parameterization, the log volatility $\ln \sigma$ of returns dS/S is the latent state variable. It evolves as a mean-reverting Ornstein-Uhlenbeck process, with mean $\ln \bar{\sigma}$ and mean reversion parameter $\alpha > 0$.

The instantaneous drift of returns and the instantaneous drift and standard deviation of log volatility are assumed constant. Furthermore, the return innovations are independent of the log volatility innovations.²

Discretization of the Continuous-Time Model

In practice, we have to rely on N discrete-time price observations to draw inference about the continuous-time model. Thus, we divide the sample period $[0, T]$ into N intervals, each of length $H = T/N$, corresponding to the discrete-time data.³ We then replace the continuous volatility dynamics with a piecewise constant process, where within each interval i , that is between times iH and $(i+1)H$, for $i = 1, 2, \dots, N$, volatility is assumed constant at $\sigma_t = \sigma_{iH}$, but from one interval to the next, volatility is allowed to be stochastic.

²Harvey and Shephard (1996) show how to relax this zero correlation assumption in a setting that is very similar to ours, in order to capture the leverage effect in equity volatilities (e.g., Schwert, 1989; Nelson, 1991; Engle and Ng, 1993). Since our empirical work focuses on exchange rates, not equities, we maintain the zero correlation assumption.

³The assumption of equally spaced observations is made for notational convenience and can be relaxed.

This piecewise constant approximation implies that within each interval i the security price evolves as a geometric Brownian motion:

$$\frac{dS_t}{S_t} = \mu dt + \sigma_{iH} dW_{S_t}, \quad \text{for } iH < t \leq (i+1)H, \quad (3)$$

and, by Ito's lemma, that the log security price $s_t = \ln S_t$ evolves as a Brownian motion:

$$ds_t = \left(\mu - \frac{1}{2}\sigma_{iH}^2\right)dt + \sigma_{iH}dW_{S_t}, \quad \text{for } iH < t \leq (i+1)H. \quad (4)$$

Log volatility varies from one interval to the next according to its Ornstein-Uhlenbeck dynamics. For small interval lengths H , the conditional distribution of log volatility is approximately:⁴

$$\ln \sigma_{(i+1)H} | \ln \sigma_{iH} \sim \mathbf{N}\left[\ln \bar{\sigma} + \rho_H(\ln \sigma_{iH} - \ln \bar{\sigma}), \beta^2 H\right]. \quad (5)$$

In words, the discretized log volatility follows a Gaussian first-order autoregressive process with mean $\ln \bar{\sigma}$, autoregressive parameter $\rho_H = 1 - \alpha H$, and variance $\beta^2 H$.

3. Econometric Approach

Measuring Volatility

Even the discretized stochastic volatility model is difficult to estimate because the sample path of the asset price within each interval is not fully observed. If it were observed, we could infer the diffusion coefficients σ_{iH} with arbitrary precision.⁵ In practice, we are forced to use discretely observed statistics of the sample paths, such as the absolute or squared returns over each interval,

⁴ This conditional distribution is only an approximation for small H . The exact conditional distribution of $\ln \sigma_{(i+1)H}$ is normal with mean $\ln \bar{\sigma} + \exp(-\alpha H)(\ln \sigma_{iH} - \ln \bar{\sigma})$ and variance $\beta^2 [1 - \exp(-2\alpha H)] / (2\alpha)$. The approximation follows from Taylor series expansions of $\exp(-\alpha H)$ and $\exp(-2\alpha H)$ around $H=0$.

⁵ See, for example, Merton (1980).

to draw inferences about the discretized log volatilities and their dynamics.

To formalize this idea, consider a statistic $f(s_{iH,(i+1)H})$ of the continuous sample path $s_{iH,(i+1)H}$ of the log asset price between times iH and $(i+1)H$. If the statistic is homogeneous in some power γ of volatility, then we can write it as:

$$f(s_{iH,(i+1)H}) = \sigma_{iH}^\gamma f(s_{iH,(i+1)H}^*), \quad (6)$$

which implies that:

$$\ln |f(s_{iH,(i+1)H})| = \gamma \ln \sigma_{iH} + \ln |f(s_{iH,(i+1)H}^*)|, \quad (7)$$

where $s_{iH,(i+1)H}^*$ denotes the continuous sample path of a standardized diffusion generated by the same innovations as $s_{iH,(i+1)H}$, but with volatility $\sigma_{iH}^*=1$.

Equation (7) makes clear that the statistic $f(\cdot)$ is a noisy volatility proxy: the first term is proportional to log volatility, but the second term is a measurement error. Other things the same, the variability of the measurement error reduces the informational content of the volatility proxy. The more variable the measurement error, the less precise is our inference about discretized log volatility and its dynamics.

Linear State Space Representation

Following Ruiz (1994) and Harvey, Ruiz, and Shephard (1994), we recognize that equations (5) and (7) form a linear state space system:

$$\ln \sigma_{(i+1)H} = \ln \bar{\sigma} + \rho_H (\ln \sigma_{iH} - \ln \bar{\sigma}) + \beta \sqrt{H} v_{(i+1)H} \quad (8a)$$

$$\ln |f(s_{iH,(i+1)H})| = \gamma \ln \sigma_{iH} + \mathbb{E} \left[|f(s_{iH,(i+1)H}^*)| \right] + \epsilon_{(i+1)H}. \quad (8b)$$

The transition equation (8a) follows from the conditional distribution of log volatility. It describes

the dynamics of the unobserved log volatility. The transition errors v are i.i.d. $N[0,1]$, which follows from the fact that $\ln \sigma_t$ is a diffusion. The measurement equation (8b) makes precise the way in which the log volatility proxy $\ln |f(\cdot)|$ is related to the true log volatility $\ln \sigma_{iH}$; it follows from equation (7) with the projection $\ln |f(\cdot)| \equiv E[\ln |f(\cdot)|] + \epsilon$. The expectation of $|f(s_{iH,(i+1)H}^*)|$ depends on s_{iH}^* , the functional form of $f(\cdot)$, and interval length H , but it is by construction independent of the log volatility $\ln \sigma_{iH}$. The projection errors ϵ have a zero mean, but are not necessarily Gaussian.

Quasi-Maximum Likelihood Estimation

If the projection errors in the measurement equation are Gaussian, exact maximum likelihood estimation of the stochastic volatility model is straightforward. Consistent and asymptotically efficient estimates of the model parameters θ are obtained by maximizing the Gaussian log likelihood function:

$$\ln \mathcal{L}(\ln |f(s_{0,H})|, \ln |f(s_{H,2H})|, \dots, \ln |f(s_{(N-1)H,NH})|; \theta) = c - \frac{1}{2} \sum_{i=1}^N \ln \eta_i - \frac{1}{2} \sum_{i=1}^N \frac{e_i^2}{\eta_i}, \quad (9)$$

where we can use the Kalman filter to evaluate the one-step ahead forecast errors and their conditional variances:⁶

$$e_i = \ln |f(s_{(i-1)H,iH})| - E_{i-1}[\ln |f(s_{(i-1)H,iH})|] \quad (10a)$$

$$\eta_i = \text{Var}_{i-1}[e_i]. \quad (10b)$$

When the projection errors in the measurement equation are not Gaussian, maximum likelihood

⁶ For a recent overview of the Kalman filter, see Hamilton (1994).

estimation is more involved. In that case, a tidy closed-form expression for the likelihood such as equation (9) does not exist in general, which makes the evaluation and maximization of the likelihood extremely challenging. Related, in the non-Gaussian case the prediction errors e_i produced by the Kalman filter are merely linear projection errors, not a conditional expectation errors, since the linear projections produced by the Kalman filter do not in general coincide with conditional expectations in non-Gaussian settings.

Nevertheless, maximizing the Gaussian likelihood function (9) can yield consistent parameter estimates even when the projection errors are not Gaussian. This approach is called Gaussian quasi-maximum likelihood estimation (QMLE).⁷ The benefits of Gaussian quasi-maximum likelihood estimation are its simplicity and consistency. Its drawbacks are that the estimates are inefficient, even asymptotically, and that its small-sample properties are suspect. Intuitively, the further the distribution of the projection errors is from a Gaussian density, the more severe are the problems with Gaussian quasi-maximum likelihood estimation.

Properties of Alternative Volatility Proxies

a. Log Absolute or Squared Returns

The stochastic volatility literature routinely uses absolute or squared returns as volatility proxies.⁸

The continuously compounded return over the i th interval is just the difference between the log asset prices at times $(i+1)H$ and iH . Thus, the traditional log volatility proxy is:

$$\ln |f(s_{iH,(i+1)H})| = \gamma \ln |s_{(i+1)H} - s_{iH}| = \gamma \ln \sigma_{iH} + \gamma \ln |s_{(i+1)H}^* - s_{iH}^*|, \quad (11)$$

⁷ Watson (1989) provides sufficient conditions for consistency of the Gaussian quasi-maximum likelihood estimator in linear state space models, and Ruiz (1993) verifies these conditions for stochastic volatility models.

⁸ See Ghysels, Harvey, and Renault (1996) for a survey.

where $\gamma=1$ or $\gamma=2$, depending on whether we consider absolute or squared returns. Because γ only scales the volatility proxy, and hence does not affect the distribution of the projection errors in the measurement equation, we focus exclusively, but without loss of generality, on absolute returns. That is, throughout the remainder of the paper, we set $\gamma=1$.

The second equality of equation (11) formally requires that the log price is a martingale, so that it is homogeneous in volatility. This assumption is not too troubling because over sufficiently small sampling intervals H , such as a day or even a week, the price drift of most securities is negligible in practice. In fact, from a statistical perspective, the assumption is likely to be helpful. By using a drift estimator that always takes the value zero we inject only a small bias, to the extent that the true drift differs slightly from zero, but we greatly reduce the variance relative to other estimators. This results in a very small mean squared error forecast of the drift.

It is by now well known that the conditional distribution of log absolute or squared returns is far from Gaussian. Jacquier, Polson, and Rossi (1994), Andersen and Sorensen (1996), and Kim, Shephard, and Chib (1998) argue that, as a result, quasi-maximum likelihood estimation with these traditional volatility proxies is highly inefficient and often severely biased in finite samples. Indeed, the relevant subset of our own Monte Carlo results, which we present in the next section, confirm their conclusions.

To deepen our theoretical understanding of why the conditional normality assumption for log absolute or squared returns fails, we examine the distribution of the log absolute value of a driftless Brownian motion x , with origin $x_0=0$ and constant diffusion coefficient σ , over an interval of finite length τ .⁹ Karatzas and Shreve (1991) characterize the distribution of the

⁹ The assumption $x_0=0$ allows us to interpret x_t directly as a continuously compounded return.

absolute value of a Brownian motion. A simple transformation of their result reveals that the distribution of the log absolute value is:

$$\text{Prob}[\ln |x_\tau| \in dy] = \frac{2e^y}{\sigma\sqrt{\tau}} \phi\left(\frac{e^y}{\sigma\sqrt{\tau}}\right) dy, \quad (12)$$

where ϕ denotes a standard normal density.

From this distribution, we can compute the mean, standard deviation, skewness, and kurtosis of $\ln|x_\tau|$, which we present in the first row of Table 1. Notice that different values of σ and τ affect only the mean, not the variance, skewness, or kurtosis of log absolute returns. In other words, those parameters determine the location, but not the shape, of the distribution. Without loss of generality then, we graph in Figure 1a the distribution of $\ln|x_\tau|$ with both σ and τ set to one. For comparison, we also show a Gaussian density with matching mean and variance.

Table 1 and Figure 1a clearly demonstrate that the distribution of log absolute returns is far from Gaussian. The skewness and kurtosis of $\ln|x_\tau|$ are -1.5 and 6.9, instead of zero and three for a normal random variable. The intuition of this result is that both positive and negative returns close to zero, observations that are “inliers” of the return distribution, become large negative outliers of the distribution of log absolute returns.¹⁰

¹⁰ The use of log absolute returns is even more problematic in empirical work, where some returns can be exactly zero because of the discreteness in prices. In that case, the logarithm of absolute returns is undefined and the quasi-maximum likelihood approach fails. Various ad hoc procedures have been devised to skirt this problem. For example, Kim, Shephard, and Chib (1998) suggest adding a small constant to absolute returns.

b. The Log Range

Now consider using the range as volatility proxy, where the range over the i th interval is defined as the difference between the security's highest and lowest log price between times iH and $(i+1)H$. Formally, consider use of the log volatility proxy:

$$\begin{aligned} \ln |f(s_{iH,(i+1)H})| &= \ln \left(\sup_{iH < t \leq (i+1)H} s_t - \inf_{iH < t \leq (i+1)H} s_t \right) \\ &= \ln \sigma_{iH} + \ln \left(\sup_{iH < t \leq (i+1)H} s_t^* - \inf_{iH < t \leq (i+1)H} s_t^* \right). \end{aligned} \tag{13}$$

For the second equality we require again that the log price is homogeneous in volatility (i.e., that it is a martingale).¹¹ We drop the absolute value signs because the range cannot be negative.

The log range is a superior volatility proxy, relative to log absolute or squared returns, for two reasons. First, it is more efficient, in the sense that the variance of the measurement errors associated with the log range is far less than the variance of the measurement errors associated with log absolute or squared returns. Second – and this is a central insight exploited throughout this paper – the log range is Gaussian, to a very good approximation. On both counts, the log range is an attractive volatility proxy for Gaussian quasi-maximum likelihood estimation of stochastic volatility models.

Let us first discuss in more detail the superior efficiency of the log range. The intuition for

¹¹ Instead of assuming a zero drift, we can perform a change of variable from the Brownian motion to a Brownian bridge (e.g., Doob, 1949; Feller, 1951). The distribution of the log range of the Brownian bridge is nearly identical to that of the log range of the corresponding Brownian motion (i.e., it is virtually Gaussian). However, the Brownian bridge is by construction independent of the drift. See Alizadeh (1998) for details.

this result is simple: on days with substantial price reversals, return-based measures underestimate the daily volatility because the closing price is not very different from the opening price, despite the large intraday price fluctuations. The range, in contrast, reflects the intraday price fluctuations.

The mathematics underlying the superior efficiency of the log range is less simple, but nevertheless standard. Specifically, consider the log range of a driftless Brownian motion x , with origin $x_0=0$ and constant diffusion coefficient σ , over an interval of finite length τ . Feller (1951) derives the distribution of the range of a Brownian motion. A simple transformation of his result reveals that the distribution of the log range is:

$$\text{Prob} \left[\ln \left(\sup_{0 < t \leq \tau} x_t - \inf_{0 < t \leq \tau} x_t \right) \in dy \right] = 8 \sum_{k=1}^{\infty} (-1)^{k-1} \frac{k^2 e^y}{\sigma \sqrt{\tau}} \phi \left(\frac{ke^y}{\sigma \sqrt{\tau}} \right) dy. \quad (14)$$

Although this distribution is expressed as an infinite series, it is straightforward to compute its moments after suitably truncating the sum. In the second row of Table 1 we report the mean and standard deviation. The superior efficiency of the log range, relative to the log absolute return, emerges clearly. Both proxies move one-for-one with log volatility on average, but the standard deviation of the log range is a quarter of the standard deviation of the log absolute return!

The efficiency of the range as a volatility measure has been appreciated implicitly for decades in the business press, which routinely reports high and low prices and sometimes displays high-low-close, or so-called “candlestick,” plots. Range-based volatility estimation has featured in the academic literature at least since Parkinson (1980), who proposes and rigorously analyzes the use of the range for estimating volatility in an i.i.d. setting. Since then, Parkinson’s estimator has been improved by combining the range with opening and closing prices (e.g., Garman and

Klass, 1980; Beckers, 1983; Ball and Torous, 1984) and correcting the downward bias in the range induced by discrete sampling (Rogers and Satchell, 1991).¹² More recently, Andersen and Bollerslev (1998) formalize the efficiency of the range in the context of estimating the integrated volatility of a diffusion.

Let us now discuss in more detail the approximate normality of the log range, or equivalently, the approximate log-normality of the range. This aspect of the range is not particularly intuitive, and it is certainly not widely appreciated. Nevertheless, it is a fact. The second row of Table 1 shows that the skewness and kurtosis of the log range are 0.17 and 2.80, respectively. These values are very close to the corresponding values of zero and three for a normal random variable, and they represent a sharp contrast to the earlier-presented skewness and kurtosis of the log absolute return. Figure 1b plots the density (14), with σ and τ set to one, together with a Gaussian density with matching mean and variance. Again, the distribution of the log range is remarkably Gaussian, except that it is slightly skewed to the right and has somewhat thinner tails.

4. Monte Carlo Analysis

The diffusion theory sketched above reveals that the log range is a much less noisy volatility proxy than log absolute or squared returns and that the distribution of the log range is approximately Gaussian, in contrast to the skewed and leptokurtic distribution of log absolute or squared

¹² Although including the opening and closing prices can improve the estimation of volatility in principle, the gains are not necessarily realized in practice. In particular, Brown (1990) argues against the inclusion of the opening and closing prices on the grounds that they are highly influenced by microstructure effects, such as the lack of trading at the close or “market on the close” orders that have a disproportionate effect on the closing price. Furthermore, experimentation by Alizadeh (1998) reveals little theoretical efficiency gain from combining the range with the opening and closing prices in the estimation of stochastic volatility modes. Thus, we shall not pursue the idea in this paper.

returns. Both of these findings suggests that Gaussian quasi-maximum likelihood estimation with the log range as volatility proxy is highly efficient, not only relative to Gaussian quasi-maximum likelihood estimation with the traditional volatility proxies, but also relative to the corresponding exact maximum likelihood estimation.

Using a Monte Carlo study, we now compare the small-sample properties of Gaussian quasi-maximum likelihood estimation with the log range as volatility proxy to those of Gaussian quasi-maximum likelihood and exact maximum likelihood estimation with the log absolute return as volatility proxy. For each of two data-generating processes, which we describe below, we generate 5000 return samples of length $T = 500, 1000,$ and 1500 “days,” where each daily return is generated by 1000 intraday price moves. For every Monte Carlo sample, we perform Gaussian quasi-maximum likelihood estimation of the diffusion parameters with either the log range or the log absolute return as volatility proxy. For comparison, we also perform exact maximum likelihood estimation using the simulated likelihood method of Sandmann and Koopman (1998).

Constant Intraday Volatility

We first use a data-generating process for which volatility is in fact constant throughout the day. In particular, we simulate the data from the following Euler discretization of the stochastic volatility model (4)-(5):

$$s_t = s_{t-\Delta t} + \sigma_{iH} \epsilon_{st} \sqrt{\Delta t} \quad (15)$$

$$\ln \sigma_{(i+1)H} = \ln \bar{\sigma} + \rho_H (\ln \sigma_{iH} - \ln \bar{\sigma}) + \beta \epsilon_{vi} \sqrt{H}, \quad (16)$$

for $iH < t \leq (i+1)H$, where ϵ_{st} and ϵ_{vi} are independent $N[0,1]$ innovations. The discrete time increment Δt , a small fraction of the discrete sampling interval H , approximates the continuous

time dt . We set $H=1/257$ and $\Delta t=H/1000$, which corresponds to using daily data generated by 1000 trades per day. We set $\alpha=3.855$, $\ln\bar{\sigma}=-2.5$ and $\beta=0.75$, which implies a volatility process with a daily autocorrelation of $\rho_H=0.985$ and an annualized average volatility of 8.51 percent. These volatility dynamics are broadly consistent with our empirical results for five major currencies, presented in the next section, as well as with the literature.

Table 2 summarizes the sampling distributions of the three estimators of the diffusion parameters ρ_H , β , and $\ln\bar{\sigma}$ for the benchmark case of $T=1000$ observations. Using the absolute return as the volatility proxy, the average Gaussian quasi-maximum likelihood estimate of ρ_H is 0.95, compared to an average estimate of 0.98 using the range and the true value of 0.985. The standard deviations of the estimates are 0.13 and 0.01, respectively. Thus, using the range instead of the absolute return as volatility proxy produces quasi-maximum likelihood estimates that are less biased and dramatically less variable.

The performance difference between the two quasi-maximum likelihood estimators is even more striking for the volatility β of innovations to log volatility. The average estimate using the log absolute return is 1.07, with a standard error of 1.10. In contrast, the average estimate using the log range is 0.8, close to the true value of 0.75, with a standard error of only 0.11!

Interestingly, the performances of the quasi-maximum likelihood estimators of the mean volatility $\ln\bar{\sigma}$ are basically identical. Both estimators appear unbiased and equally efficient.

In Figure 2, we illustrate graphically the very different sampling properties of the two Gaussian quasi-maximum likelihood estimators. The first three plots of the first two rows show the empirical distributions of the parameter estimates using the log absolute return and the log range as volatility proxy, respectively. The drastic efficiency gains from using the range are

immediately apparent.¹³ Notice also that the sampling distribution of the estimates of ρ and β for the log absolute return are severely skewed, which implies that the usual Gaussian inferences based on asymptotic standard errors are not trustworthy. In contrast, the distribution of the corresponding estimates using the log range are very close to Gaussian.

The results thus far indicate that Gaussian quasi-maximum likelihood estimation with the log range as volatility proxy is highly efficient relative to quasi-maximum likelihood estimation with the log absolute return as volatility proxy. This efficiency gain stems from the fact that the log range is approximately Gaussian, as well as the fact that the range is a much less noisy volatility measure than absolute or squared returns. To separate these two effects, we now compare the range-based quasi-maximum likelihood estimator to the exact maximum likelihood estimator for absolute returns, which we compute using the simulation method of Sandmann and Koopman (1998). If the only benefit from using the range is its approximate normality, the sampling properties of the range-based quasi-maximum likelihood estimator should be very similar to the sampling distribution of the exact maximum likelihood estimator for absolute returns. If, however, the information about intraday volatility that is revealed by the range but not by absolute or squared returns is useful in the estimation of the model, the sampling properties of the range-based quasi-maximum likelihood estimator could well dominate the sampling properties of the exact maximum likelihood estimator for absolute returns.¹⁴

Comparing the third row of each panel in Table 2, which summarizes the sampling

¹³ Notice different scales of the second plots in each row. The horizontal axis of the second plots of the second and third row corresponds to the region between the two vertical lines in the second plot of the first row.

¹⁴ Alternatively, we could compare the sampling properties of the range-based quasi-maximum likelihood estimator to the sampling properties of the exact maximum likelihood estimator for the range. However, given the near-normality of the log range, the difference in performance between these two estimators would be minimal.

properties of the exact maximum likelihood estimator for absolute returns, to the first two rows of each panel, reveals that much but not all of the efficiency gain from using the log range as volatility proxy is attributed to its approximate normality (for a graphical representation of this result, see also the third row of Figure 2). In terms of bias, the range-based quasi-maximum likelihood estimator and the exact maximum likelihood estimator for absolute returns perform equally well. However, the standard deviations of the range-based estimates of ρ and β are less than half of the standard deviations of the corresponding exact maximum likelihood estimates. This demonstrates that the information about intraday volatility contained in the range is an important aspect of the success of our range-based estimation of the stochastic volatility model.

Once the model has been estimated, the Kalman filter can be used to extract the latent stochastic volatility series. When using a Gaussian volatility proxy, the extraction of the latent volatilities is best (i.e., minimum variance) unbiased, whereas the extraction is merely best *linear* unbiased when using a non-Gaussian volatility proxy. Hence there are two reasons to expect volatility extraction with the log range as volatility proxy to dominate the extraction with the log absolute return as volatility proxy. First, the range-based parameter estimates are more accurate. Second, even for the same parameter values, the approximate normality of the log range yields a more efficient volatility extraction.

With this in mind, we summarize in the last two panels of Table 2 (and in the last column of Figure 2) the sampling distributions of the average extraction error $\frac{1}{T} \sum_{t=1}^T (\ln \hat{\sigma}_t - \ln \sigma_t)$, which is an estimator of the expected extraction error $E[\ln \hat{\sigma}_t - \ln \sigma_t]$, and the average squared extraction error $\frac{1}{T} \sum_{t=1}^T (\ln \hat{\sigma}_t - \ln \sigma_t)^2$, which is an estimator of the expected squared extraction error

$E[\ln \hat{\sigma}_t - \ln \sigma_t]^2$.¹⁵ Not surprisingly, among the quasi-maximum likelihood estimators the range-based estimator is superior: both extractions of $\ln \sigma_t$ appear unbiased, but the range-based estimator is much more efficient. Using the log absolute return as volatility proxy, the average squared extraction error is 0.05, with a 90 percent confidence interval of [0.03, 0.08], whereas using the log range as volatility proxy, the average squared extraction error is only 0.02, with a confidence interval of [0.01, 0.05]. Furthermore, comparing the range-based quasi-maximum likelihood extractions to the exact maximum likelihood extractions for absolute returns, we again notice that the information about intraday volatility contained in the range is important.

It is of interest to see how the Monte Carlo results vary with the sample size T . We have already discussed in detail the results for $T = 1000$ observations, and in particular the clear superiority of our range-based estimator. Now we discuss the results for a smaller sample size of $T = 500$ observations and a larger sample size of $T = 5000$ observations.

We show the results for $T = 500$ in Table 3; they are qualitatively identical to those in Table 2. Quantitatively, however, the relative performance of the quasi-maximum likelihood estimator with the log absolute return as volatility proxy, which was already poor with $T = 1000$ observations, is much worse with $T = 500$ observations. This is largely due to dramatic increases in the biases of the estimators of ρ and β , which then translate into larger extraction errors.

We present the results for $T = 5000$ in Table 4. Qualitatively, they are again identical to the results in Table 2; quantitatively, the comparative performance of the quasi-maximum likelihood estimator with the log absolute return as volatility proxy is improved in some respects, but it remains poor in others. On the up-side, the estimators' biases largely vanish, and their

¹⁵ When we write $\ln \hat{\sigma}$, we are referring to an estimator of $\ln \sigma$, not the log of an estimator of σ .

standard deviations are close to those of the exact maximum likelihood estimator, except for β . On the down-side, however, the mean squared volatility extraction error remains poor.

In summary, the clear superiority of range-based estimation evident with $T = 1000$ observations is amplified with $T = 500$ observation and remains significant even with $T = 5000$ observations. In particular, it generates substantially smaller mean squared volatility extraction error regardless of the sample size, both because of the strong optimality of the Kalman filter in the Gaussian case, in contrast to the weak optimality in the non-Gaussian case, and because of the information about intraday volatility contained in the range.

Stochastic Intraday Volatility

The assumption of constant intraday volatility is perhaps controversial. Therefore, we repeat our Monte Carlo study assuming this time that log volatility evolves continuously, just as it does in the continuous time stochastic volatility model. We simulate data from the following Euler discretization of the bivariate diffusion (2):

$$s_t = s_{t-\Delta t} + \sigma_{t-\Delta t} \epsilon_{st} \sqrt{\Delta t} \quad (17)$$

$$\ln \sigma_t = \ln \bar{\sigma} + \rho_{\Delta t} (\ln \sigma_{t-\Delta t} - \ln \bar{\sigma}) + \beta \epsilon_{vt} \sqrt{\Delta t}, \quad (18)$$

where $\rho_{\Delta t} = (1 - \alpha \Delta t)$ and ϵ_{vt} is a $N[0,1]$ innovation independent of ϵ_{st} . All other details are the same as in the case of constant intraday volatility.¹⁶

Tables 5-7 summarize the second set of Monte Carlo results. Most of the statistics are virtually identical to those in Tables 2-4, which suggests that the discretization of the volatility

¹⁶ Except that the exact maximum likelihood estimator is now for the data-generating process with time-varying intraday volatility (17)-(18), in contrast to the exact maximum likelihood estimator examined earlier, which was for the data-generating process with constant intraday volatility (15)-(16).

process has minimal effects on the small sample properties of the estimators.¹⁷

5. Empirical Analysis of Exchange Rate Volatility Dynamics

We estimate stochastic volatility models for the U.S. dollar price of five actively-traded currencies: the British pound, Canadian dollar, Deutsche Mark, Japanese yen, and Swiss franc. The volatility proxies are constructed from daily high, low, and closing futures prices. We first tabulate some statistics describing salient aspects of the volatility proxies, and then we estimate and interpret both one- and two-factor stochastic volatility models.

Data

We use daily high, low, and closing (3pm EST) prices of currency futures contracts traded on the International Monetary Market, a subsidiary of the Chicago Mercantile Exchange, from January 1978 through December 1998.¹⁸ A futures contract represents delivery of the currency on the second Wednesday of the following March, June, September, or December. Each day there are at least three futures contracts with different quarterly delivery dates traded on each currency. We use futures prices from the front-month contract, which is the one closest to delivery with at least ten days to delivery. This front-month contract is typically the most actively traded contract.

There are several advantages to using futures, as opposed to spot, exchange rate data. First, all futures prices (including the daily high and low) result from open outcry, so that all transactions are open to the market and orders are filled at the best price. Currency spot market trading, in contrast, is bilateral between banks, and any one particular executed price is not

¹⁷ Our results, as with all Monte Carlo results, are of course specific to the particular data-generating process studied, and perhaps in particular to the high degree of persistence in the log volatility process.

¹⁸ The data source is FAME Information Services.

necessarily representative of overall market conditions. Second, the closing, or “settlement,” futures price is based on the best sentiment of the market at the time of close (3pm EST, after which spot market trading declines) and is widely scrutinized, since it is used to mark account balances to market. Therefore, the futures closing price is likely to be a very accurate measure of the “true” market price at that time. Finally, futures returns are the actual returns from investing in a foreign currency, whereas spot “returns” are less meaningful, unless one accounts for the interest rate differential between the two countries.

Empirical Description of the Volatility Proxies

In Table 8 we present statistics summarizing the distributions of log absolute returns and the log range for each of the five currencies. The superior efficiency of the log range as a volatility proxy emerges not only in terms of its smaller standard deviation stressed thus far, but also in terms of its time-series dynamics. In particular, the large and slowly-decaying autocorrelations of the log range clearly reveal strong volatility persistence for each exchange rate, in sharp contrast to the spuriously small autocorrelations of log absolute returns. Clearly, the measurement errors associated with the log absolute returns masks the persistence in volatility.

Estimation of the One-Factor Stochastic Volatility Model With the Log Range as Volatility Proxy

We report estimates of the one-factor stochastic volatility model described by equations (4)-(5) for the five currencies in the left panel of Table 9. For now, we focus on the estimates obtained using our preferred volatility proxy, the log range. The most striking result concerns the volatility persistence parameter ρ , which tends to be small compared to typical values reported in the literature (ranging typically from 0.80 to 0.99). The estimates range from 0.62 to 0.85, with four

of the five estimates less than 0.75.

However, the residual diagnostics presented in Table 10 indicate serious problems with the one-factor model. While the residuals are clearly less persistent than the log range itself, substantial residual serial correlation remains. Effectively, the one-factor stochastic volatility model adequately accounts for the volatility correlation at lag one, but not at longer lags, which results in a humped-shaped residual autocorrelation function.

The misspecification of the one-factor model can be seen in another way. To obtain the estimates in Table 9, we set the standard deviation of the measurement equation disturbances to 0.29, per the results in Table 1. Alternatively, however, we can estimate the standard deviation of the measurement errors along with the other parameters, and when we do so, we typically obtain a much larger estimate of ρ . Consider, for example, the British pound. When we set the standard deviation of the measurement errors to 0.29, we obtain $\hat{\rho}=0.66$, as recorded in Table 9, but when we estimate the standard deviation of the measurement errors along with the other parameters, we obtain $\hat{\rho}=0.97$ and an estimate of the standard deviation of 0.42. The difference in maximized log likelihoods, moreover, is greater than two hundred. Hence, the measurement errors of the one-factor model are much more variable than expected if the one-factor model were correct – a standard deviation of 0.42 vs. 0.29 – which again suggests that the model is *not* correct.¹⁹

It should be noted that except for the remaining residual autocorrelation, the residual diagnostics in Table 10 are encouraging. The measurement equation residuals display little

¹⁹ In fact, the sum of the unconditional variance of the measurement errors and the unconditional variance of the latent log volatility process exceeds the unconditional variance of the log range (from Table 8), which suggests a negative correlation between log volatility and the measurement errors. In theory, of course, the measurement errors are uncorrelated with log volatility.

skewness or excess kurtosis, and more generally, the histograms and quantile-quantile (QQ) plots in Figure 3 illustrate that the residuals are virtually indistinguishable from Gaussian.²⁰

Estimation of a Two-Factor Stochastic Volatility Model With the Log Range as Volatility Proxy

In light of the severe deficiencies of the one-factor stochastic volatility model revealed by our range-based estimation and analysis, we move to a two-factor model, with transition equation:

$$\ln \sigma_{(i+1)H} = \ln \sigma_{1,(i+1)H} + \ln \sigma_{2,(i+1)H}, \quad (19)$$

where

$$\begin{aligned} \ln \sigma_{1,(i+1)H} &= \ln \bar{\sigma}_1 + \rho_{1,H} (\ln \sigma_{1,iH} - \ln \bar{\sigma}_1) + \beta_1 \sqrt{H} v_{1,(i+1)H} \\ \ln \sigma_{2,(i+1)H} &= \ln \bar{\sigma}_2 + \rho_{2,H} (\ln \sigma_{2,iH} - \ln \bar{\sigma}_2) + \beta_2 \sqrt{H} v_{2,(i+1)H}. \end{aligned} \quad (20)$$

The model is similar in spirit to the two-factor stochastic volatility models of Chacko and Viceira (1999), Gallant, Hsu, and Tauchen (1999), and Chernov, Gallant, Ghysels, and Tauchen (1999). It also resembles the component GARCH approach of Engle and Lee (1999). The difference is that ours is a truly two-factor stochastic volatility model, whereas theirs is a one-factor GARCH(2,2) model decomposed tautologically into a component structure.

When we estimate the above two-factor stochastic volatility model, the results of which we report in the right panel of Table 9, we obtain a persistent and a transient factor. Each factor is responsible for about half the long-run (unconditional) variance of log volatility, but the transient factor responsible for much more of the short-run variance. This result is intuitively appealing and in line with properties of volatilities estimated using very different procedures, such as the realized volatilities of Andersen, Bollerslev, Diebold, and Labys (1999), which display slow

²⁰ A Gaussian QQ plot is simply a graph of the quantiles of a standardized distribution against the corresponding quantiles of a $N[0,1]$ distribution. Hence if a volatility proxy is normally distributed, its Gaussian QQ plot is a straight line with a unit slope, which enables simple visual assessment of closeness to normality.

persistent movement of log volatility with high-frequency noise superimposed.

The residual diagnostics in Table 11 and Figure 4 indicate that the two-factor models are adequate. The measurement equation residuals are serially uncorrelated and virtually Gaussian.

An interesting feature of our results is that the estimated one-factor volatility persistence parameter is an average of the estimated persistence parameters from the two-factor model. To understand this finding, assume that, in the spirit of equation (19):

$$\ln\sigma_t = c + \ln\sigma_{1,t} + \ln\sigma_{2,t}, \quad (21)$$

where

$$\begin{aligned} \ln\sigma_{1,t} &= \rho_1 \ln\sigma_{1,t-1} + v_{1,t} \\ \ln\sigma_{2,t} &= \rho_2 \ln\sigma_{2,t-1} + v_{2,t} \end{aligned} \quad (22)$$

and the volatility innovations $v_{1,t}$ and $v_{2,t}$ are independent. Suppose, however, that although the two-factor model is true, we fit a one-factor model, which captures only the sum of the components $\ln\sigma_t$, instead of the components $\ln\sigma_{1,t}$ and $\ln\sigma_{2,t}$ themselves. Then, the first autocovariance of $\ln\sigma_t$ is:

$$\begin{aligned} \text{Cov}[\ln\sigma_t, \ln\sigma_{t-1}] &= \text{Cov}[\ln\sigma_{1,t} + \ln\sigma_{2,t}, \ln\sigma_{1,t-1} + \ln\sigma_{2,t-1}] \\ &= \text{Cov}[\ln\sigma_{1,t}, \ln\sigma_{1,t-1}] + \text{Cov}[\ln\sigma_{2,t}, \ln\sigma_{2,t-1}] \\ &= \rho_1 \text{Var}[\ln\sigma_{1,t}] + \rho_2 \text{Var}[\ln\sigma_{2,t}], \end{aligned} \quad (23)$$

where of course the variances are unconditional. Hence the first autocorrelation of $\ln\sigma_t$ in the one-factor model is simply:

$$\rho = \frac{\text{Cov}[\ln\sigma_t, \ln\sigma_{t-1}]}{\text{Var}[\ln\sigma_t]} = \frac{\rho_1 \text{Var}[\ln\sigma_{1,t}] + \rho_2 \text{Var}[\ln\sigma_{2,t}]}{\text{Var}[\ln\sigma_{1,t}] + \text{Var}[\ln\sigma_{2,t}]}. \quad (24)$$

In words, the first autocorrelation in the one-factor model is a relative variance weighted average of the first autocorrelations of the two factors. This is approximately true in the estimates. The

fact that we can successfully predict the outcome of estimating a one-factor model on the basis of our estimates of the two-factor model is further evidence in favor of the two-factor model.

The range-based estimates of the one- and two-factor models are also consistent in terms of implied unconditional variances of log volatility. Obviously, given the independence of $\ln\sigma_1$ and $\ln\sigma_2$, the unconditional variance of $\ln\sigma_t$ in the one-factor model should be equal to the sum of the variances of the two factors in the two-factor model. This is approximately true.

Estimation With the Log Absolute Return as Volatility Proxy: Comparison and Reconciliation

At first sight, the estimates of the one-factor model based on the log absolute return, in the left panel of Table 9, and the corresponding residual diagnostics, in the left panel of Table 10, show no evidence of model misspecification.²¹ In particular, the single factor that emerges resembles closely the persistent factor of our two-factor model, instead of the variance-weighted average of the two factors. Apparently, the choice of volatility proxy drastically effects our inferences.

The explanation of this result is that the additional noise in the log absolute return, compared to the log range, masks the presence of the second, less persistent factor. More mechanically, both estimators choose the parameters $\ln\bar{\sigma}$, ρ_H , and β to match two features of the data: the autocorrelation of the volatility proxy and the difference between the unconditional variance of the volatility proxy and the corresponding unconditional variance of the measurement errors from Table 1. The relative importance of those features, however, differs across the volatility proxies. Specifically, the latent volatility dynamics explain less than 10 percent of the unconditional variance of log absolute returns, but more than 70 percent of the variance of the log

²¹This explains the current state of the literature, which centers primarily on one-factor models.

range, which is just another manifestation of the informational efficiency of the log range. The quasi-maximum likelihood estimator for log absolute returns therefore chooses parameters that explain entirely the autocorrelation of the volatility proxy, but leave half of the variance of log absolute returns that is attributed to the volatility dynamics (which is very little relative to the total variance of log absolute returns) unexplained. In contrast, the estimator for the log range chooses parameters that explain all of the variance of the volatility proxy, but leave a significant amount of autocorrelation (about half) unexplained.

In summary, upon closer inspection we notice that both sets of one-factor models are equally misspecified, but that the misspecification is revealed along different dimensions. The misspecification of the range-based models is immediately apparent from the autocorrelations of the residuals. The misspecification of the absolute return-based models, in contrast, is more subtle. It appears as a violation of the adding up constraint $\text{Var}[\ln|f_t|] = \text{Var}[\ln\sigma_t] + \text{Var}[\epsilon_t]$.

6. Summary, Conclusions and Directions for Future Research

The range has a long history in finance, from the stock charts in business newspapers to highbrow academics. We have used the log range to develop a simple, yet highly efficient method for estimating stochastic volatility models.

We argued both theoretically and empirically that the log range is nearly Gaussian, and moreover, that it is a much less noisy measure of volatility than popular volatility proxies, such as log absolute or squared returns. This translates into a simple range-based Gaussian quasi-maximum likelihood estimator that is highly efficient, both in small and large samples, which we established through a Monte Carlo study.

We conclude that range-based Gaussian quasi-maximum likelihood estimation of

stochastic volatility models provides the best of both worlds: simplicity and efficiency. Daily high and low prices are available over long historical periods for a variety of assets, in contrast to high-frequency intra-day data, which can be used under certain conditions to produce volatility proxies even more efficient than the range, but which have become available only recently and only for selected highly-developed markets.²²

We also performed an empirical analysis of volatility dynamics for a set of five major dollar exchange rates, which not only illustrated the simplicity and flexibility of our method, but also produced sharp and substantive results. In particular, for each exchange rate the analysis points emphatically to a two-factor stochastic volatility model, with one highly persistent and one transient factor. We are not the first to suggest that a two-factor stochastic volatility model is desirable; see Chacko and Viceira (1999), Gallant, Hsu, and Tauchen (1999), Chernov, Gallant, Ghysels, and Tauchen (1999), and Engle and Lee (1999). Our results do, however, contribute importantly to what may be an emerging consensus.

One appealing direction for future research is an extension of our methods to the multivariate case, through the construction and study of measures of covariation based on daily ranges. One might consider, for example, the daily maximum divergence $\text{Range}(s_i - s_j)$, effectively a “cross-range.” A key question is whether this cross range, for which we typically do not have data, is related to $\text{Range}(s_i)$ and $\text{Range}(s_j)$, for which we do have data. It is obvious that $\text{Range}(s_i - s_j) \neq \text{Range}(s_i) - \text{Range}(s_j)$, but a more subtle relationship may nevertheless be operative.

A second direction for future research concerns the robustness of the range to market microstructure effects. If high-frequency intra-day data are available, then under certain

²² See Andersen and Bollerslev (1998) and Andersen, Bollerslev, Diebold, and Labys (1999).

conditions daily realized volatilities constructed from this data are even more efficient than the range as volatility proxy (e.g., Andersen and Bollerslev, 1998; Andersen, Bollerslev, Diebold, and Labys, 1999). We conjecture, however, that market microstructure effects may plague daily realized volatilities constructed from high-frequency intra-day data more than they plague the daily range, depending on the specifics of the market.

References

- Alizadeh, S. (1998), "Essays in Financial Econometrics," Ph.D. Dissertation, Department of Economics, University of Pennsylvania.
- Andersen, T.G., and Bollerslev, T. (1997), "Heterogeneous Information Arrivals and Return Volatility Dynamics: Uncovering the Long-Run in High Frequency Returns," *Journal of Finance*, 52, 975-1005.
- Andersen, T.G., and Bollerslev, T. (1998), "Answering the Skeptics: Yes, Standard Volatility Models Do Provide Accurate Forecasts," *International Economic Review*, 39, 885-905.
- Andersen, T.G., Bollerslev, T., Diebold, F.X., and Labys, P. (1999), "The Distribution of Exchange Rate Volatility," Wharton Financial Institutions Center Working Paper 99-08 and NBER Working Paper 6961.
- Andersen, T.G., and Sorensen, B.E. (1997), "GMM and QML Asymptotic Standard Deviations in Stochastic Volatility Models," *Journal of Econometrics*, 76, 397-403.
- Ball, C.A., and Torous, W. (1984), "The Maximum Likelihood Estimation of Security Price Volatility: Theory, Evidence, and Application to Option Pricing," *Journal of Business*, 57, 97-112.
- Bates, D.S. (1996), "Jumps and Stochastic Volatility: Exchange Rate Processes Implicit in Deutschemark Options," *Review of Financial Studies*, 9, 69-107.
- Beckers, S. (1983), "Variance of Security Price Returns Based on High, Low, and Closing Prices," *Journal of Business*, 56, 97-112.
- Black, F., and Scholes, M. (1973), "The Pricing of Options and Corporate Liabilities," *Journal of Political Economy*, 81, 637-654.
- Bollerslev, T., Chou, R.Y., and Kroner, K.F. (1992), "ARCH Modeling in Finance: A Selective Review of the Theory and Empirical Evidence," *Journal of Econometrics*, 52, 5-59.
- Brown, S. (1990), "Estimating Volatility," in S. Figlewski, W. Silber, and M. Subrahmanyam (eds.), *Financial Options*. Homewood, Ill.: Business One Irwin.
- Chernoff, M., and Ghysels, E. (1999), "A Study Toward a Unified Approach to the Joint Estimation of Objective and Risk Neutral Measures for Options Valuation," *Journal of Financial Economics*, forthcoming.
- Chacko, G., and Viceira, L.M. (1999), "Dynamics Consumption and Portfolio Choice with Stochastic Volatility in Incomplete Markets," Manuscript, Graduate School of Business, Harvard University.

- Chernov, M., Ghysels, E., Gallant, A.R., and Tauchen, G.E. (1999), "A New Class of Stochastic Volatility Models with Jumps: Theory and Estimation," Manuscript, Departments of Economics, Penn State University, University of North Carolina, and Duke University.
- Danielsson, J. (1994), "Stochastic Volatility in Asset Prices: Estimation with Simulated Maximum Likelihood," *Journal of Econometrics* 64, 375-400.
- Doob, J. (1949), "Heuristic Approach to Kolmogorov-Smirnov Theorems," *Annals of Mathematical Statistics*, 20, 393-403.
- Duffie, D., and Singleton, K. (1993), "Simulated Moments Estimation of Markov Models of Asset Prices," *Econometrica*, 61, 929-952.
- Engle, R.F., and Lee, G.G.J. (1999), "A Permanent and Transitory Model of Stock Return Volatility," in R. Engle and H. White (eds.), *Cointegration, Causality, and Forecasting: A Festschrift in Honor of Clive W.J. Granger*, 475-497. Oxford: Oxford University Press.
- Engle, R.F., and Ng, V.K. (1993), "Measuring and Testing the Impact of News on Volatility," *Journal of Finance*, 48, 1749-1801.
- Feller, W. (1951), "The Asymptotic Distribution of the Range of Sums of Independent Random Variables," *Annals of Mathematical Statistics*, 22, 427-432.
- Galant, A.R., Hsieh, D., and Tauchen, G.E. (1997), "Estimation of Stochastic Volatility Models with Diagnostics," *Journal of Econometrics*, 81, 159-192.
- Gallant, A.R., Hsu, C.T., and Tauchen, G.E. (1999), "Calibrating Volatility Diffusions and Extracting Integrated Volatility," *Review of Economics and Statistics*, forthcoming.
- Garman, M.B., and Klass, M.J. (1980), "On the Estimation of Price Volatility from Historical Data," *Journal of Business*, 53, 67-78.
- Ghysels, E., Harvey, A., and Renault, E. (1996), "Stochastic Volatility," in G.S. Maddala and C.R. Rao (eds.), *Statistical Methods in Finance* (Handbook of Statistics, Volume 14). Amsterdam: North-Holland.
- Hamilton, J.D. (1994), *Time Series Analysis*. Princeton: Princeton University Press.
- Harvey, A., Ruiz, E., and Shephard, N. (1994), "Multivariate Stochastic Variance Models," *Review of Economic Studies*, 61, 247-264.
- Harvey, A., Ruiz, E., and Shephard, N. (1996), "Estimation of an Asymmetric Stochastic Volatility Model for Asset Returns," *Journal of Business and Economic Statistics*, 14, 429-434.

- Hull, J., and White, A. (1987), "The Pricing of Options on Assets with Stochastic Volatilities," *Journal of Finance*, 42, 381-400.
- Jacquier, E., Polson, N.G., and Rossi, P.E. (1994), "Bayesian Analysis of Stochastic Volatility Models," *Journal of Business and Economic Statistics*, 12, 413-417.
- Jarrow, R. (ed.) (1998), *Volatility*. London: Risk Publications.
- Kim, S., Shephard, N., and Chib, S. (1998), "Stochastic Volatility: Likelihood Inference and Comparison with ARCH Models," *Review of Economic Studies*, 65, 361-393.
- Merton, R.C. (1969), "Lifetime Portfolio Selection Under Uncertainty: The Continuous-Time Case," *Review of Economics and Statistics*, 51, 247-257.
- Merton, R.C. (1980), "On Estimating the Expected Return on the Market: An Exploratory Investigation," *Journal of Financial Economics*, 8, 323-361.
- Nelson, D.B. (1991), "Conditional Heteroskedasticity in Asset Pricing: A New Approach," *Econometrica*, 59, 347-370.
- Parkinson, M. (1980), "The Extreme Value Method for Estimating the Variance of the Rate of Return," *Journal of Business*, 53, 61-65.
- Rogers, L.C.G., and Satchell, S.E. (1991), "Estimating Variance from High, Low and Closing Prices," *Annals of Applied Probability*, 1, 504-512.
- Ruiz, E. (1994), "Quasi-Maximum Likelihood Estimation of Stochastic Variance Models," *Journal of Econometrics*, 63, 289-306.
- Sandmann, G., and Koopman, S. (1998), "Estimation of Stochastic Volatility Models via Monte Carlo Maximum Likelihood," *Journal of Econometrics*, 87, 271-301.
- Schwert, G.W. (1989), "Why Does Stock Market Volatility Change Over Time," *Journal of Finance*, 44, 1115-1153.
- Singleton, K. (1997), "Estimation of Affine Asset Pricing Models Using the Empirical Characteristic Function," Manuscript, Graduate School of Business, Stanford University.

Table 1
Moments of Alternative Volatility Proxies

We consider a driftless Brownian motion x , with origin $x_0=0$ and constant diffusion coefficient σ , over an interval of finite length τ . The table shows the first four moments of two volatility proxies: the log absolute return $\ln|x_\tau|$ and the log range $\ln|\sup x_t - \inf x_t|$.

Volatility Proxy	Mean	Standard Deviation	Skewness	Kurtosis
Log Absolute Return	$-0.64 + \frac{1}{2}\ln\tau + \ln\sigma$	1.11	-1.53	6.93
Log Range	$0.43 + \frac{1}{2}\ln\tau + \ln\sigma$	0.29	0.17	2.80

Table 2
Monte Carlo Analysis of Alternative Estimators
Constant Intraday Volatility, $T = 1000$

We report statistics summarizing the sampling distribution of several estimators of the parameters and the latent volatilities in the stochastic volatility model:

$$s_t = s_{t-\Delta t} + \sigma_{iH} \epsilon_{st} \sqrt{\Delta t}$$

$$\ln \sigma_{(i+1)H} = \ln \bar{\sigma} + \rho_H (\ln \sigma_{iH} - \ln \bar{\sigma}) + \beta \epsilon_{vi} \sqrt{H},$$

$iH < t \leq (i+1)H$, where ϵ_{st} and ϵ_{vi} are independent $N[0,1]$ random variables. The discrete time increment Δt , a fraction of the discrete sampling interval H , approximates the continuous time dt . We set $H=1/257$ and $\Delta t=H/1000$, which corresponds to using daily data generated by 1000 trades per day. We set $\alpha=3.855$, $\ln \bar{\sigma} = -2.5$ and $\beta=0.75$, which implies a volatility process with daily autocorrelation of $\rho_H=0.985$ and an annualized average volatility of 8.51 percent. “QML (Log Abs. Return)” denotes the Gaussian quasi-maximum likelihood estimator with the log absolute return as volatility proxy. “QML (Log Range)” denotes the Gaussian quasi-maximum likelihood estimator with the log range as volatility proxy. “Exact ML” denotes the estimator that maximizes the exact likelihood of log absolute returns, evaluated by simulation methods. All results are based on 5000 Monte Carlo replications.

Estimator	Mean	Std. Dev.	1%	5%	25%	50%	75%	95%	99%
$\rho=0.985$									
QML (Log Abs. Return)	0.95	0.13	0.14	0.84	0.97	0.98	0.99	0.99	1.00
QML (Log Range)	0.98	0.01	0.95	0.96	0.97	0.98	0.99	0.99	0.99
Exact ML	0.98	0.02	0.92	0.95	0.97	0.98	0.99	0.99	1.00
$\beta=0.750$									
QML (Log Abs. Return)	1.07	1.10	0.22	0.37	0.61	0.82	1.09	2.37	7.53
QML (Log Range)	0.80	0.11	0.55	0.63	0.72	0.79	0.87	0.98	1.07
Exact ML	0.80	0.22	0.34	0.49	0.64	0.77	0.91	1.16	1.44
$\ln \bar{\sigma} = -2.50$									
QML (Log Abs. Return)	-2.49	0.10	-2.74	-2.66	-2.56	-2.49	-2.42	-2.33	-2.27
QML (Log Range)	-2.53	0.09	-2.75	-2.68	-2.59	-2.53	-2.46	-2.38	-2.31
Exact ML	-2.51	0.10	-2.73	-2.66	-2.58	-2.50	-2.44	-2.35	-2.27
$E[\ln \hat{\sigma}_t - \ln \sigma_t]$									
QML (Log Abs. Return)	0.00	0.09	-0.22	-0.16	-0.06	0.00	0.07	0.15	0.22
QML (Log Range)	0.00	0.09	-0.22	-0.15	-0.06	0.00	0.07	0.15	0.22
Exact ML	0.00	0.09	-0.22	-0.15	-0.06	0.00	0.07	0.16	0.23
$E[\ln \hat{\sigma}_t - \ln \sigma_t]^2$									
QML (Log Abs. Return)	0.05	0.02	0.02	0.03	0.04	0.04	0.05	0.08	0.11
QML (Log Range)	0.02	0.01	0.01	0.01	0.01	0.02	0.02	0.05	0.07
Exact ML	0.03	0.01	0.01	0.02	0.02	0.02	0.03	0.06	0.08

Table 3
Monte Carlo Analysis of Alternative Estimators
Constant Intraday Volatility, $T = 500$

We report statistics summarizing the sampling distribution of several estimators of the parameters and the latent volatilities in the stochastic volatility model:

$$s_t = s_{t-\Delta t} + \sigma_{iH} \epsilon_{st} \sqrt{\Delta t}$$

$$\ln \sigma_{(i+1)H} = \ln \bar{\sigma} + \rho_H (\ln \sigma_{iH} - \ln \bar{\sigma}) + \beta \epsilon_{vi} \sqrt{H},$$

$iH < t \leq (i+1)H$, where ϵ_{st} and ϵ_{vi} are independent $N[0,1]$ random variables. The discrete time increment Δt , a fraction of the discrete sampling interval H , approximates the continuous time dt . We set $H=1/257$ and $\Delta t=H/1000$, which corresponds to using daily data generated by 1000 trades per day. We set $\alpha=3.855$, $\ln \bar{\sigma}=-2.5$ and $\beta=0.75$, which implies a volatility process with daily autocorrelation of $\rho_H=0.985$ and an annualized average volatility of 8.51 percent. “QML (Log Abs. Return)” denotes the Gaussian quasi-maximum likelihood estimator with the log absolute return as volatility proxy. “QML (Log Range)” denotes the Gaussian quasi-maximum likelihood estimator with the log range as volatility proxy. “Exact ML” denotes the estimator that maximizes the exact likelihood of log absolute returns, evaluated by simulation methods. All results are based on 5000 Monte Carlo replications.

Estimator	Mean	Std. Dev.	1%	5%	25%	50%	75%	95%	99%
$\rho=0.985$									
QML (Log Abs. Return)	0.86	0.27	0.00	0.02	0.91	0.97	0.98	0.99	1.00
QML (Log Range)	0.97	0.02	0.90	0.94	0.97	0.98	0.98	0.99	1.00
Exact ML	0.96	0.06	0.72	0.89	0.96	0.98	0.99	0.99	1.00
$\beta=0.750$									
QML (Log Abs. Return)	1.58	1.90	0.11	0.28	0.63	0.93	1.50	6.42	9.36
QML (Log Range)	0.82	0.16	0.47	0.56	0.71	0.81	0.92	1.09	1.26
Exact ML	0.88	0.43	0.19	0.38	0.62	0.81	1.03	1.56	2.49
$\ln \bar{\sigma} = -2.50$									
QML (Log Abs. Return)	-2.49	0.14	-2.83	-2.72	-2.59	-2.49	-2.40	-2.26	-2.16
QML (Log Range)	-2.52	0.13	-2.83	-2.74	-2.61	-2.53	-2.45	-2.32	-2.21
Exact ML	-2.51	0.13	-2.82	-2.72	-2.59	-2.50	-2.41	-2.28	-2.18
$E[\ln \hat{\sigma}_t - \ln \sigma_t]$									
QML (Log Abs. Return)	0.00	0.13	-0.30	-0.21	-0.08	0.00	0.09	0.21	0.32
QML (Log Range)	0.00	0.13	-0.29	-0.21	-0.08	0.00	0.09	0.21	0.31
Exact ML	0.00	0.13	-0.29	-0.20	-0.07	0.01	0.10	0.22	0.32
$E[\ln \hat{\sigma}_t - \ln \sigma_t]^2$									
QML (Log Abs. Return)	0.06	0.03	0.02	0.03	0.04	0.05	0.07	0.12	0.17
QML (Log Range)	0.03	0.02	0.01	0.01	0.01	0.02	0.03	0.08	0.13
Exact ML	0.04	0.03	0.01	0.01	0.02	0.03	0.04	0.08	0.14

Table 4
Monte Carlo Analysis of Alternative Estimators
Constant Intraday Volatility, $T = 5000$

We report statistics summarizing the sampling distribution of several estimators of the parameters and the latent volatilities in the stochastic volatility model:

$$s_t = s_{t-\Delta t} + \sigma_{iH} \epsilon_{st} \sqrt{\Delta t}$$

$$\ln \sigma_{(i+1)H} = \ln \bar{\sigma} + \rho_H (\ln \sigma_{iH} - \ln \bar{\sigma}) + \beta \epsilon_{vi} \sqrt{H},$$

$iH < t \leq (i+1)H$, where ϵ_{st} and ϵ_{vi} are independent $N[0,1]$ random variables. The discrete time increment Δt , a fraction of the discrete sampling interval H , approximates the continuous time dt . We set $H=1/257$ and $\Delta t=H/1000$, which corresponds to using daily data generated by 1000 trades per day. We set $\alpha=3.855$, $\ln \bar{\sigma}=-2.5$ and $\beta=0.75$, which implies a volatility process with daily autocorrelation of $\rho_H=0.985$ and an annualized average volatility of 8.51 percent. “QML (Log Abs. Return)” denotes the Gaussian quasi-maximum likelihood estimator with the log absolute return as volatility proxy. “QML (Log Range)” denotes the Gaussian quasi-maximum likelihood estimator with the log range as volatility proxy. “Exact ML” denotes the estimator that maximizes the exact likelihood of log absolute returns, evaluated by simulation methods. All results are based on 5000 Monte Carlo replications.

Estimator	Mean	Std. Dev.	1%	5%	25%	50%	75%	95%	99%
$\rho=0.985$									
QML (Log Abs. Return)	0.98	0.01	0.96	0.97	0.98	0.98	0.99	0.99	0.99
QML (Log Range)	0.98	0.00	0.97	0.98	0.98	0.98	0.99	0.99	0.99
Exact ML	0.98	0.00	0.97	0.98	0.98	0.98	0.99	0.99	0.99
$\beta=0.750$									
QML (Log Abs. Return)	0.77	0.15	0.48	0.55	0.67	0.76	0.86	1.02	1.19
QML (Log Range)	0.78	0.05	0.68	0.71	0.75	0.78	0.81	0.86	0.89
Exact ML	0.73	0.08	0.56	0.61	0.69	0.77	0.80	0.89	0.96
$\ln \bar{\sigma} = -2.50$									
QML (Log Abs. Return)	-2.49	0.05	-2.60	-2.57	-2.53	-2.49	-2.46	-2.42	-2.39
QML (Log Range)	-2.53	0.04	-2.63	-2.60	-2.56	-2.53	-2.50	-2.46	-2.43
Exact ML	-2.50	0.05	-2.60	-2.57	-2.52	-2.49	-2.47	-2.43	-2.39
$E[\ln \hat{\sigma}_t - \ln \sigma_t]$									
QML (Log Abs. Return)	0.00	0.04	-0.10	-0.07	-0.03	0.00	0.03	0.07	0.10
QML (Log Range)	0.00	0.04	-0.10	-0.07	-0.03	0.00	0.03	0.08	0.10
Exact ML	0.00	0.04	-0.10	-0.07	-0.02	0.00	0.04	0.08	0.11
$E[\ln \hat{\sigma}_t - \ln \sigma_t]^2$									
QML (Log Abs. Return)	0.04	0.00	0.03	0.03	0.04	0.04	0.04	0.05	0.05
QML (Log Range)	0.01	0.00	0.01	0.01	0.01	0.01	0.01	0.02	0.02
Exact ML	0.02	0.00	0.02	0.02	0.02	0.02	0.02	0.03	0.03

Table 5
Monte Carlo Analysis of Alternative Estimators
Stochastic Intraday Volatility, $T = 1000$

We report statistics summarizing the sampling distribution of several estimators of the parameters and the latent volatilities in the stochastic volatility model:

$$s_t = s_{t-\Delta t} + \sigma_{t-\Delta t} \epsilon_{st} \sqrt{\Delta t}$$

$$\ln \sigma_t = \ln \bar{\sigma} + \rho_{\Delta t} (\ln \sigma_{t-\Delta t} - \ln \bar{\sigma}) + \beta \epsilon_{vt} \sqrt{\Delta t},$$

where ϵ_{st} and ϵ_{vt} are independent $N[0,1]$ random variables. The discrete time increment Δt , a fraction of the discrete sampling interval H , approximates the continuous time dt . We set $H=1/257$ and $\Delta t=H/1000$, which corresponds to using daily data generated by 1000 trades per day. We set $\alpha=3.855$, $\ln \bar{\sigma}=-2.5$ and $\beta=0.75$, which implies a volatility process with daily autocorrelation of $\rho_H=0.985$ and an annualized average volatility of 8.51 percent. “QML (Log Abs. Return)” denotes the Gaussian quasi-maximum likelihood estimator with the log absolute return as volatility proxy. “QML (Log Range)” denotes the Gaussian quasi-maximum likelihood estimator with the log range as volatility proxy. “Exact ML” denotes the estimator that maximizes the exact likelihood of log absolute returns, evaluated by simulation methods. All results are based on 5000 Monte Carlo replications.

Estimator	Mean	Std. Dev.	1%	5%	25%	50%	75%	95%	99%
$\rho=0.985$									
QML (Log Abs. Return)	0.95	0.13	0.13	0.84	0.96	0.98	0.99	0.99	1.00
QML (Log Range)	0.98	0.10	0.95	0.96	0.97	0.98	0.99	0.99	0.99
Exact ML	0.98	0.11	0.92	0.95	0.97	0.98	0.99	0.99	1.00
$\beta=0.750$									
QML (Log Abs. Return)	1.07	1.08	0.24	0.39	0.61	0.82	1.11	2.46	7.10
QML (Log Range)	0.79	0.11	0.56	0.62	0.72	0.78	0.86	0.97	1.05
Exact ML	0.79	0.22	0.37	0.47	0.68	0.77	0.92	1.17	1.43
$\ln \bar{\sigma} = -2.50$									
QML (Log Abs. Return)	-2.50	0.10	-2.73	-2.66	-2.56	-2.50	-2.43	-2.34	-2.27
QML (Log Range)	-2.53	0.09	-2.75	-2.68	-2.59	-2.53	-2.47	-2.38	-2.32
Exact ML	-2.50	0.09	-2.73	-2.66	-2.56	-2.50	-2.44	-2.35	-2.28
$E[\ln \hat{\sigma}_t - \ln \sigma_t]$									
QML (Log Abs. Return)	0.00	0.09	-0.22	-0.16	-0.06	0.00	0.06	0.15	0.21
QML (Log Range)	0.00	0.09	-0.22	-0.15	-0.06	0.00	0.06	0.15	0.22
Exact ML	0.00	0.09	-0.22	-0.15	-0.06	0.00	0.07	0.15	0.21
$E[\ln \hat{\sigma}_t - \ln \sigma_t]^2$									
QML (Log Abs. Return)	0.05	0.02	0.02	0.03	0.04	0.04	0.05	0.07	0.10
QML (Log Range)	0.02	0.01	0.01	0.01	0.01	0.02	0.02	0.04	0.07
Exact ML	0.03	0.01	0.01	0.01	0.02	0.02	0.03	0.05	0.08

Table 6
Monte Carlo Analysis of Alternative Estimators
Stochastic Intraday Volatility, $T = 500$

We report statistics summarizing the sampling distribution of several estimators of the parameters and the latent volatilities in the stochastic volatility model:

$$s_t = s_{t-\Delta t} + \sigma_{t-\Delta t} \epsilon_{st} \sqrt{\Delta t}$$

$$\ln \sigma_t = \ln \bar{\sigma} + \rho_{\Delta t} (\ln \sigma_{t-\Delta t} - \ln \bar{\sigma}) + \beta \epsilon_{vt} \sqrt{\Delta t},$$

where ϵ_{st} and ϵ_{vt} are independent $N[0,1]$ random variables. The discrete time increment Δt , a fraction of the discrete sampling interval H , approximates the continuous time dt . We set $H=1/257$ and $\Delta t=H/1000$, which corresponds to using daily data generated by 1000 trades per day. We set $\alpha=3.855$, $\ln \bar{\sigma}=-2.5$ and $\beta=0.75$, which implies a volatility process with daily autocorrelation of $\rho_H=0.985$ and an annualized average volatility of 8.51 percent. “QML (Log Abs. Return)” denotes the Gaussian quasi-maximum likelihood estimator with the log absolute return as volatility proxy. “QML (Log Range)” denotes the Gaussian quasi-maximum likelihood estimator with the log range as volatility proxy. “Exact ML” denotes the estimator that maximizes the exact likelihood of log absolute returns, evaluated by simulation methods. All results are based on 5000 Monte Carlo replications.

Estimator	Mean	Std. Dev.	1%	5%	25%	50%	75%	95%	99%
$\rho=0.985$									
QML (Log Abs. Return)	0.85	0.27	0.00	0.00	0.91	0.97	0.98	0.99	1.00
QML (Log Range)	0.97	0.02	0.90	0.94	0.96	0.98	0.98	0.99	0.99
Exact ML	0.97	0.05	0.82	0.90	0.96	0.98	0.98	0.99	1.00
$\beta=0.750$									
QML (Log Abs. Return)	1.54	1.86	0.11	0.25	0.61	0.91	1.49	6.10	9.37
QML (Log Range)	0.81	0.16	0.46	0.57	0.70	0.80	0.91	1.08	1.23
Exact ML	0.85	0.42	0.23	0.37	0.59	0.78	1.01	1.53	2.19
$\ln \bar{\sigma} = -2.50$									
QML (Log Abs. Return)	-2.49	0.14	-2.81	-2.72	-2.59	-2.49	-2.40	-2.27	-2.18
QML (Log Range)	-2.53	0.13	-2.82	-2.74	-2.61	-2.53	-2.44	-2.32	-2.24
Exact ML	-2.51	0.13	-2.80	-2.71	-2.59	-2.50	-2.41	-2.29	-2.20
$E[\ln \hat{\sigma}_t - \ln \sigma_t]$									
QML (Log Abs. Return)	0.00	0.13	-0.30	-0.21	-0.08	0.00	0.09	0.21	0.29
QML (Log Range)	0.00	0.13	-0.29	-0.21	-0.08	0.00	0.09	0.21	0.28
Exact ML	0.00	0.13	-0.29	-0.20	-0.08	0.00	0.10	0.22	0.30
$E[\ln \hat{\sigma}_t - \ln \sigma_t]^2$									
QML (Log Abs. Return)	0.05	0.03	0.02	0.03	0.04	0.05	0.06	0.11	0.16
QML (Log Range)	0.03	0.02	0.01	0.01	0.01	0.02	0.03	0.07	0.11
Exact ML	0.04	0.03	0.01	0.01	0.02	0.03	0.04	0.08	0.13

Table 7
Monte Carlo Analysis of Alternative Estimators
Stochastic Intraday Volatility, $T = 5000$

We report statistics summarizing the sampling distribution of several estimators of the parameters and the latent volatilities in the stochastic volatility model:

$$s_t = s_{t-\Delta t} + \sigma_{t-\Delta t} \epsilon_{st} \sqrt{\Delta t}$$

$$\ln \sigma_t = \ln \bar{\sigma} + \rho_{\Delta t} (\ln \sigma_{t-\Delta t} - \ln \bar{\sigma}) + \beta \epsilon_{vt} \sqrt{\Delta t},$$

where ϵ_{st} and ϵ_{vt} are independent $N[0,1]$ random variables. The discrete time increment Δt , a fraction of the discrete sampling interval H , approximates the continuous time dt . We set $H=1/257$ and $\Delta t=H/1000$, which corresponds to using daily data generated by 1000 trades per day. We set $\alpha=3.855$, $\ln \bar{\sigma}=-2.5$ and $\beta=0.75$, which implies a volatility process with daily autocorrelation of $\rho_H=0.985$ and an annualized average volatility of 8.51 percent. “QML (Log Abs. Return)” denotes the Gaussian quasi-maximum likelihood estimator with the log absolute return as volatility proxy. “QML (Log Range)” denotes the Gaussian quasi-maximum likelihood estimator with the log range as volatility proxy. “Exact ML” denotes the estimator that maximizes the exact likelihood of log absolute returns, evaluated by simulation methods. All results are based on 5000 Monte Carlo replications.

Estimator	Mean	Std. Dev.	1%	5%	25%	50%	75%	95%	99%
$\rho=0.985$									
QML (Log Abs. Return)	0.98	0.01	0.96	0.97	0.98	0.98	0.99	0.99	0.99
QML (Log Range)	0.98	0.00	0.98	0.98	0.98	0.98	0.99	0.99	0.99
Exact ML	0.98	0.00	0.97	0.98	0.98	0.98	0.99	0.99	0.00
$\beta=0.750$									
QML (Log Abs. Return)	0.77	0.15	0.46	0.54	0.67	0.76	0.86	1.03	1.17
QML (Log Range)	0.77	0.05	0.67	0.70	0.74	0.77	0.80	0.85	0.88
Exact ML	0.74	0.09	0.56	0.61	0.68	0.74	0.80	0.89	0.94
$\ln \bar{\sigma} = -2.50$									
QML (Log Abs. Return)	-2.50	0.05	-2.60	-2.57	-2.53	-2.50	-2.46	-2.42	-2.39
QML (Log Range)	-2.53	0.04	-2.63	-2.61	-2.56	-2.53	-2.50	-2.46	-2.43
Exact ML	-2.50	0.04	-2.61	-2.57	-2.53	-2.50	-2.47	-2.43	-2.41
$E[\ln \hat{\sigma}_t - \ln \sigma_t]$									
QML (Log Abs. Return)	-0.00	0.04	-0.10	-0.08	-0.03	-0.00	0.03	0.07	0.10
QML (Log Range) avg	-0.00	0.04	-0.10	-0.08	-0.03	-0.00	0.03	0.07	0.10
Exact ML	-0.00	0.04	-0.10	-0.07	-0.02	-0.00	0.04	0.07	0.10
$E[\ln \hat{\sigma}_t - \ln \sigma_t]^2$									
QML (Log Abs. Return)	0.04	0.00	0.03	0.03	0.03	0.04	0.04	0.05	0.05
QML (Log Range)	0.01	0.00	0.01	0.01	0.01	0.01	0.01	0.02	0.02
Exact ML	0.02	0.00	0.02	0.02	0.02	0.02	0.02	0.03	0.03

Table 8
Distributions and Dynamics of Volatility Proxies
Five Dollar Exchange Rates

We report statistics summarizing both the unconditional moments and the autocorrelations of two volatility proxies for five dollar exchange rates, measured daily from 1 January 1978 through 31 December 1998. The underlying data used to compute the log absolute return and the log range are daily high, low, and settlement prices of front-month futures contracts traded on the International Monetary Market.

Volatility Proxy	Unconditional Moments				Autocorrelations				
	Mean	Std.	Skew.	Kurt.	1 st	2 nd	5 th	10 th	20 th
British Pound									
Log Absolute Return	-5.82	1.19	-0.86	3.64	0.09	0.06	0.10	0.07	0.05
Log Range	-4.87	0.53	0.09	3.10	0.39	0.33	0.30	0.27	0.22
Canadian Dollar									
Log Absolute Return	-6.67	1.10	-0.62	2.89	0.11	0.09	0.12	0.08	0.05
Log Range	-5.70	0.53	-0.02	3.39	0.49	0.46	0.41	0.35	0.31
Deutsche Mark									
Log Absolute Return	-5.77	1.17	-0.88	3.62	0.06	0.06	0.09	0.07	0.05
Log Range	-4.83	0.52	-0.07	3.12	0.40	0.37	0.35	0.30	0.23
Japanese Yen									
Log Absolute Return	-5.77	1.17	-0.88	3.62	0.10	0.05	0.08	0.07	0.07
Log Range	-4.88	0.58	0.03	3.19	0.41	0.34	0.32	0.26	0.20
Swiss Franc									
Log Absolute Return	-5.60	1.16	-0.95	3.77	0.05	0.02	0.06	0.05	0.04
Log Range	-4.67	0.48	0.04	3.13	0.32	0.29	0.30	0.25	0.19

Table 9
Quasi-Maximum Likelihood Estimates
One-Factor and Two-Factor Stochastic Volatility Models
Five Dollar Exchange Rates

We report estimates of one-factor and two-factor stochastic volatility models fit to five dollar exchange rates, using daily data from 1 January 1978 through 31 December 1998. Asymptotic standard errors are in parentheses. See text for model descriptions.

Volatility Proxy	One-Factor Model			Two-Factor Model				
	$\ln\bar{\sigma}$	ρ	β	$\ln\bar{\sigma}$	ρ_1	β_1	ρ_2	β_2
British Pound								
Log Absolute Return	-2.42 (0.06)	0.99 (0.01)	0.91 (0.18)	-2.42 (0.08)	0.99 (0.00)	0.60 (0.12)	0.06 (0.07)	7.44 (0.51)
Log Range	-2.51 (0.01)	0.65 (0.02)	5.33 (0.12)	-2.50 (0.04)	0.98 (0.00)	0.94 (0.09)	0.19 (0.03)	5.14 (0.10)
Canadian Dollar								
Log Absolute Return	-3.29 (0.06)	0.98 (0.01)	1.12 (0.16)	-3.29 (0.06)	0.98 (0.00)	1.03 (0.15)	0.24 (0.39)	3.26 (0.93)
Log Range	-3.34 (0.02)	0.85 (0.01)	3.69 (0.14)	-3.34 (0.05)	0.98 (0.00)	1.20 (0.10)	0.16 (0.04)	4.26 (0.11)
Deutsche Mark								
Log Absolute Return	-2.38 (0.04)	0.97 (0.01)	1.37 (0.25)	-2.38 (0.05)	0.98 (0.01)	1.07 (0.21)	-0.11 (0.16)	6.57 (0.61)
Log Range	-2.47 (0.02)	0.72 (0.02)	4.77 (0.14)	-2.47 (0.04)	0.97 (0.01)	1.23 (0.09)	0.05 (0.04)	4.64 (0.11)
Japanese Yen								
Log Absolute Return	-2.37 (0.04)	0.97 (0.01)	1.47 (0.28)	-2.38 (0.05)	0.98 (0.01)	0.94 (0.21)	0.17 (0.10)	7.31 (0.53)
Log Range	-2.53 (0.02)	0.62 (0.02)	6.20 (0.12)	-2.53 (0.04)	0.97 (0.01)	1.43 (0.13)	0.15 (0.03)	5.68 (0.12)
Swiss Franc								
Log Absolute Return	-2.22 (0.04)	0.98 (0.10)	0.74 (0.15)	-2.22 (0.05)	0.99 (0.00)	0.59 (0.13)	0.02 (0.02)	6.29 (0.58)
Log Range	-2.32 (0.01)	0.63 (0.02)	4.78 (0.13)	-2.32 (0.03)	0.97 (0.01)	1.05 (0.08)	0.03 (0.03)	4.50 (0.11)

Table 10
Residual Diagnostics
One-Factor Stochastic Volatility Models
Five Dollar Exchange Rates

We report statistics summarizing both the unconditional moments and the autocorrelations of measurement equation residuals from one-factor stochastic volatility models fit to five dollar exchange rates, using daily data from 1 January 1978 through 31 December 1998.

Volatility Proxy	Unconditional Moments			Autocorrelations				
	Std.	Skew.	Kurt.	1 st	2 nd	5 th	10 th	20 th
British Pound								
Log Absolute Return	1.17	-1.29	5.87	-0.00	-0.02	0.02	-0.00	-0.02
Log Range	0.30	0.15	3.06	0.18	0.20	0.22	0.20	0.16
Canadian Dollar								
Log Absolute Return	1.10	-1.19	5.46	0.10	-0.01	0.02	-0.02	-0.03
Log Range	0.26	0.17	3.25	-0.02	0.07	0.12	0.10	0.09
Deutsche Mark								
Log Absolute Return	1.16	-1.46	7.53	-0.03	-0.02	0.02	0.004	0.00
Log Range	0.28	0.06	3.07	0.10	0.16	0.21	0.18	0.15
Japanese Yen								
Log Absolute Return	1.14	-1.08	4.67	0.02	-0.03	0.01	0.12	0.03
Log Range	0.32	0.09	3.15	0.23	0.23	0.25	0.20	0.16
Swiss Franc								
Log Absolute Return	1.15	-1.25	5.59	-0.00	-0.04	0.01	0.01	-0.00
Log Range	0.29	0.12	3.08	0.12	0.16	0.21	0.18	0.15

Table 11
Residual Diagnostics
Two-Factor Stochastic Volatility Models
Five Dollar Exchange Rates

We report statistics summarizing both the unconditional moments and the autocorrelations of measurement equation residuals from two-factor stochastic volatility models fit to five dollar exchange rates, using daily data from 1 January 1978 through 31 December 1998.

Volatility Proxy	Unconditional Moments			Autocorrelations					
	Std.	Skew.	Kurt.	1 st	2 nd	5 th	10 th	20 th	
	British Pound								
Log Absolute Return	1.18	-1.26	5.72	0.01	-0.01	0.03	0.00	-0.01	
Log Range	0.37	0.24	3.17	0.11	0.03	0.02	0.02	0.01	
	Canadian Dollar								
Log Absolute Return	1.09	-1.17	5.39	0.01	-0.01	0.02	-0.01	-0.03	
Log Range	0.33	0.23	3.38	0.07	0.05	0.01	-0.01	0.01	
	Deutsche Mark								
Log Absolute Return	1.17	-1.28	5.88	-0.02	-0.01	0.03	0.01	0.00	
Log Range	0.37	0.19	3.09	0.04	0.00	0.02	0.02	0.02	
	Japanese Yen								
Log Absolute Return	1.16	-1.34	6.72	0.03	-0.01	0.02	0.02	0.02	
Log Range	0.39	0.26	3.27	0.09	0.01	0.03	0.01	0.02	
	Swiss Franc								
Log Absolute Return	1.16	-1.30	5.91	0.01	-0.03	0.02	0.01	0.00	
Log Range	0.37	0.27	3.17	0.02	-0.01	0.03	0.02	0.03	

Figure 1a
Distribution of Log Absolute Return

We consider a driftless Brownian motion, with zero origin and unit diffusion coefficient, over an interval of unit length. We plot the distribution of the log absolute return, with the best-fitting normal distribution superimposed for visual reference.




Figure 1b
Distribution of Log Range

We consider a driftless Brownian motion, with zero origin and unit diffusion coefficient, over an interval of unit length. We plot the distribution of the log range, with the best-fitting normal distribution superimposed for visual reference.

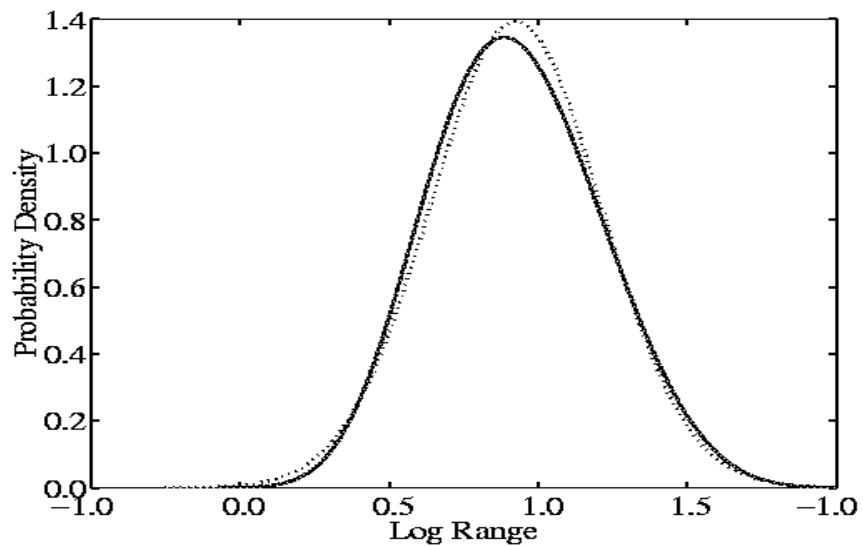


Figure 2
Monte Carlo Distributions of Parameter Estimates


We show estimates of the sampling distributions of several estimators of the parameters and the latent volatilities in the stochastic volatility model:

$$s_t = s_{t-\Delta t} + \sigma_{iH} \epsilon_{st} \sqrt{\Delta t}$$


$$\ln \sigma_{(i+1)H} = \ln \bar{\sigma} + \rho_H (\ln \sigma_{iH} - \ln \bar{\sigma}) + \beta \epsilon_{vi} \sqrt{H},$$

$iH < t \leq (i+1)H$, where ϵ_{st} and ϵ_{vi} are independent $N[0,1]$ random variables. The discrete time increment Δt , a fraction of the discrete sampling interval H , approximates the continuous time dt . We set $H=1/257$ and $\Delta t=H/1000$, which corresponds to using daily data generated by 1000 trades per day. We set $\alpha=3.855$, $\ln \bar{\sigma}=-2.5$ and $\sigma_v=0.75$, which implies a volatility process with daily autocorrelation of $\rho_H=0.985$ and an annualized average volatility of 8.51 percent. “QML (Log Abs. Return)” denotes the Gaussian quasi-maximum likelihood estimator with the log absolute return as volatility proxy. “QML (Log Range)” denotes the Gaussian quasi-maximum likelihood estimator with the log range as volatility proxy. “Exact ML” denotes the estimator that maximizes the exact likelihood of log absolute returns, evaluated by simulation methods. All results are based on 5000 Monte Carlo replications, sample size $T=1000$, and constant intraday volatility. Reading across the rows, we show the sampling distributions of the estimators of ρ , β , $\log \bar{\sigma}$, and $E[\ln \hat{\sigma}_t - \ln \sigma_t]^2$. The two vertical lines in the second plot of the first row mark the range of the same plots in the second and third row.

QML (Log Abs. Return)



QML (Log Range)



Exact ML




Figure 3
Measurement Equation Residual Distributions
One-Factor Stochastic Volatility Models, Five Dollar Exchange Rates

We show histograms of the measurement equation residuals for stochastic volatility models estimated using either log absolute returns or the log range as volatility proxy with the best-fitting normal imposed for visual reference, and the corresponding quantile-quantile (QQ) plot, which is a graph of the quantiles of the standardized residual distribution against the corresponding quantiles of a $N[0,1]$ distribution. If the residual is normally distributed, its Gaussian QQ plot is a straight line with a unit slope. The rows correspond to the five currencies examined: the British pound, Canadian dollar, Deutsche Mark, Japanese yen, and Swiss franc.

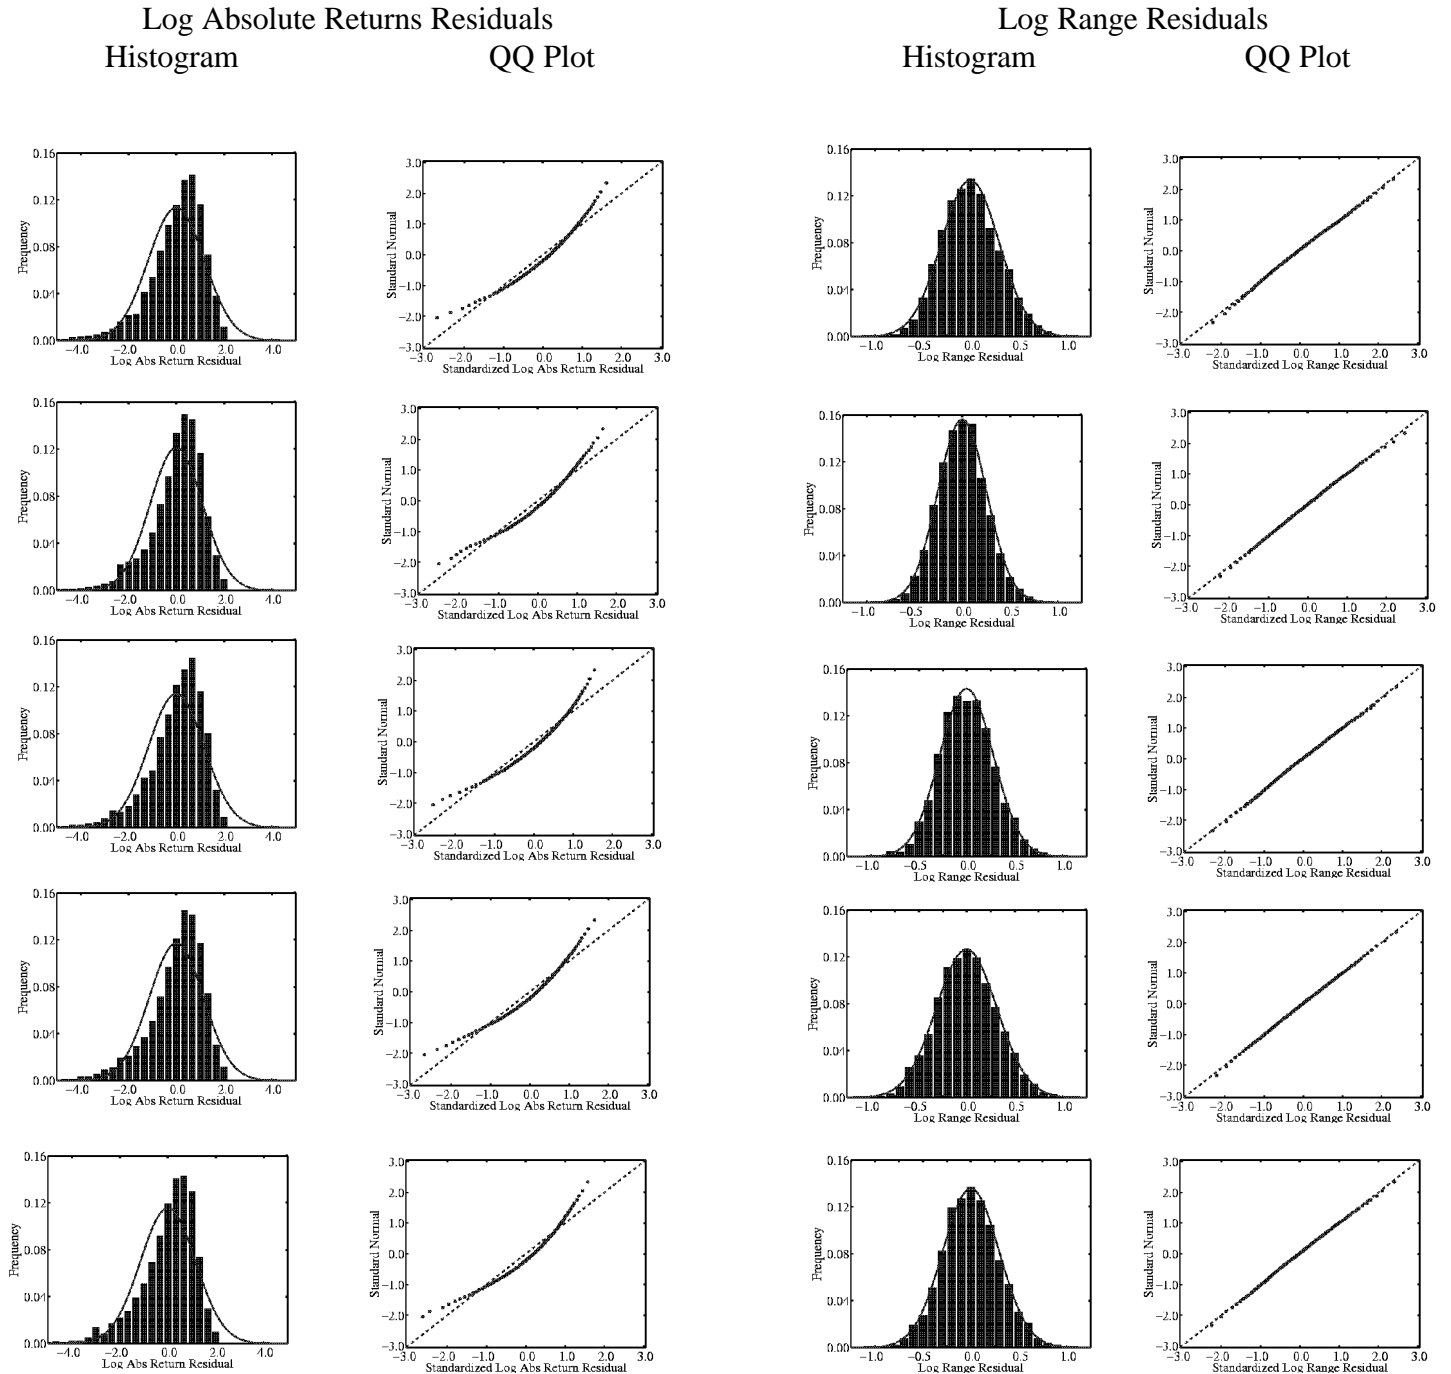


Figure 4 Measurement Equation Residual Distributions Two-Factor Stochastic Volatility Models, Five Dollar Exchange Rates

We show histograms of the measurement equation residuals for stochastic volatility models estimated using either log absolute returns or the log range as volatility proxy with the best-fitting normal imposed for visual reference, and the corresponding QQ plot, which is a graph of the quantiles of the standardized residual distribution against the corresponding quantiles of a $N[0,1]$ distribution. If the residual is normally distributed, its Gaussian QQ plot is a straight line with a unit slope. The rows correspond to the five currencies examined: the British pound, Canadian dollar, Deutsche Mark, Japanese yen, and Swiss franc.

