UC Santa Barbara
Departmental Working Papers

Title
Noise Reduced Realized Volatility: A Kalman Filter Approach

Permalink
https://escholarship.org/uc/item/4n80536m

Authors
Owens, John
Steigerwald, Douglas G

Publication Date
2009-06-01

Noise Reduced Realized Volatility: A Kalman Filter Approach

John P. Owens [*] and Douglas G. Steigerwald[†]

Abstract

Microstructure noise contaminates high-frequency estimates of asset price volatility. Recent work has determined a preferred sampling frequency under the assumption that the properties of noise are constant. Given the sampling frequency, the high-frequency observations are given equal weight. While convenient, constant weights are not necessarily efficient. We use the Kalman filter to derive more efficient weights, for any given sampling frequency. We demonstrate the efficacy of the procedure through an extensive simulation exercise, showing that our filter compares favorably to more traditional methods.

Keywords: Realized Volatility, Microstructure Noise, Kalman Filter

JEL Classification: C32, C52, G00

---

[*]School of Economics and Finance, Faculty of Commerce and Administration, Victoria University of Wellington, PO Box 600, Wellington, New Zealand, phone: +64-4-463 5451, fax: +64-4-463 5014, e-mail: John.Owens@vuw.ac.nz.

[†]Department of Economics, University of California, Santa Barbara

1 Introduction

Long-standing interest in asset price volatility, combined with recent developments in its estimation with high-frequency data, has provoked research on the correct use of such data. In this paper we offer a framework for high-frequency measurement of asset returns that provides a means of clarifying the impact of microstructure noise. Additionally, we provide Kalman filter based techniques for the efficient removal of such noise.

In a series of widely cited articles, Andersen, Bollerslev, Diebold, and Labys (2001a,b) and Barndorff-Nielsen and Shephard (2002a,b) lay out a theory of volatility estimation from high-frequency sample variances. According to the theory, realized volatility estimators can recover the volatility defined by the quadratic variation of the semimartingale for prices. Realized volatility estimators are constructed as the sums of squared returns, where each return is measured over a short interval of time.[1]

Realized volatility differs markedly from model-based estimation of volatility. The widely used class of volatility models derived from the ARCH specification of Engle (1982) place constraints on the parameters that correspond to the interval over which returns are measured. Empirical analyses of these models rarely support the constraints. In contrast, realized volatility estimators do not require a specified volatility model.

The asymptotic theory underpinning realized volatility estimators suggests that the estimators should be constructed from the highest frequency data available. One would then sum the squares of these high-frequency returns, giving each squared return equal weight.

---

[1]Andersen (2002) provides a survey of both theory and empirics for realized volatility.

In practice, however, very high frequency data is contaminated by noise arising from the microstructure of asset markets.

By now, it is widely accepted that market microstructure contamination obscures high-frequency returns through several channels. For example, transaction returns exhibit negative serial correlation due to what Roll (1984) terms the bid-ask bounce. When prices are observed at only regular intervals, or are treated as if this were the case, measured returns exhibit nonsynchronous trading biases as described in Cohen, Maier, Schwartz, and Whitcomb (1978, 1979), and Atchison, Butler, and Simonds (1987), and Lo and MacKinlay (1988, 1990). Because transaction prices are discrete and tend to cluster at certain fractional values, prices exhibit rounding distortions as described in Gottlieb and Kalay (1985), Ball (1988), and Cho and Frees (1988). Noise cannot be removed simply by working with the middle of specialist quotes; while mid-quotes are less impacted by asynchronous trade and the bid-ask bounce, mid-quotes are distorted by the inventory needs of specialists and by the regulatory requirements that they face.[2]

Simulations by Andersen and Bollerslev (1998) and Andreou and Ghysels (2002), among many others, illustrate the effects of finite sampling and microstructure noise on volatility estimates under a variety of specifications. Differences in model formulation and assumed frictions make drawing robust conclusions about the effects of specific microstructure features difficult. Nevertheless, from the cited work, it is clear that microstructure frictions, as a group, cannot be safely ignored.

---

[2]Surveys by O'Hara (1995), Hasbrouck (1996), Campbell, Lo, and MacKinlay (1997), and Madhavan (2000) document these and other microstructure frictions.

3

Essentially, three strands of research exist that treat the problem of microstructure noise in realized volatility estimation. The first attempts to remove the noise with a simple moving-average filter as in Zhou (1996). Andersen, Bollerslev, Diebold and Ebens (2001) and Corsi, Zumbach, Muller and Dacorogna (2001) who select a sample frequency of five minutes based on a volatility signature plots and then apply a moving-average filter. In contrast, Russell and Bandi (2004) work with an explicit model of microstructure noise. Rather than filtering the data to reduce the noise, they determine an optimal sampling frequency in the presence of noise. To do so, they construct a mean-squared error criterion that trades off the increase in precision against the corresponding increase in noise that arises as the sampling frequency increases. Although squared returns are given equal weight for a given asset, the optimal sampling interval that arises can vary across assets. Ait-Sahalia, Mykland, and Zhang (2003) and Oomen (2004) offer similar treatments. Finally, Hansen and Lunde (2004) derive a Newey and West (1987) type correction to account for spurious correlations in observed returns.

Theory suggests that noise volatility remains relatively constant. However, it is known that return volatility varies markedly. Thus, the relative contributions of noise and actual returns towards observed returns vary. During periods of high return volatility, return innovations tend to dominate the noise in size. In consequence we propose a somewhat different estimator in which the weight given to each return varies. Observed returns during periods of high volatility are given larger weight.

Our argument has three parts. First, we frame a precise definition of noise in terms

4

of market microstructure theory. Second, we show how the Kalman filter can be used to remove the microstructure noise. We pay particular attention to how the variability of the optimal return weights depends on high-frequency volatility. Third, we demonstrate the efficacy of the filter in removing the noise.

2 Model

To formalize, consider a sequence of fixed intervals (5-minute periods, for example) indexed by t. The log of the observed price at t is $\tilde{p}_t = p_t + \eta_t$, where p_t denotes true price and η_t denotes microstructure noise. The observed return is

$$\tilde{r}_t = r_t + \varepsilon_t, \tag{1}$$

where $r_t = p_t - p_{t-1}$ is the latent (true) return and $\varepsilon_t = \eta_t - \eta_{t-1}$ is the return noise.

We employ assumptions typical of the realized volatility literature. Our first assumption concerns the latent (true) price process.

Assumption 1. (Latent price process.)

The log true price process is a continuous local martingale. Specifically,

$$p_\tau = \int_0^\tau v_s dw_s$$

where w_s is standard Brownian motion and the spot volatility process v_s is a

strictly positive cadlag process such that the quadratic variation (or integrated

volatility) process, V_τ, obeys

$$V_\tau = \int_0^\tau v_s ds < \infty$$

with probability one for all τ.

Assumptions about noise dynamics must be selected with care. Close study of microstructure noise reveals strong positive correlation at high frequency.[3] The correlation declines sharply with the sampling frequency, due to intervening transactions. To understand these effects, we discuss three prominent sources of noise.

The bid-ask bounce, discussed by Roll (1984), arises because transactions cluster at quotes rather than the true price. Hasbrouck and Ho (1987) show that this source of noise may be positively correlated as a result of clustered trade at one quote (due to the break up of large block trades). However, the positive noise correlation due to trade clustering nearly vanishes between trades more than a few transactions apart. In similar fashion, positive noise correlation arising from the common rounding of adjacent transactions, vanishes at lower sampling frequencies.

The nonsynchronous trading effect, discussed by Lo and MacKinlay (1990), arises when transactions are relatively infrequent. If transactions are infrequent relative to the measurement of prices at regular intervals, then multiple price measurements refer to the

---

[3]Noise outcomes of adjacent price measurements are almost perfectly correlated when no transaction intervenes (they are not perfectly correlated because, although measured price remains constant in the absence of new transactions, the latent true price changes through time).

same transaction, inducing positive noise correlation. Again, the positive noise correlation vanishes as the sampling frequency declines.

To determine the sampling frequency at which noise is uncorrelated, Hasbrouck and Ho (1987) study a large sample of NYSE stocks. They find no significant correlation for observations sampled more than 10 transactions apart. Hansen and Lunde (2004) find supporting evidence in their recent study of Dow Jones Industrial Average stocks. In consequence, we assume that the sampling interval contain at least 10 transactions, to justify treating microstructure noise as an i.i.d. sequence.

Assumption 2. (Microstructure noise.)

The microstructure noise forms an i.i.d. sequence of random variables each with mean zero and variance $\sigma_\eta^2 < \infty$ and independent of the latent return process.

We do not make any distributional assumptions about microstructure noise. However, as the noise is composed of a sum of several largely independent features, and because these features tend to be symmetric, the assumption of normally distributed noise may be a plausible approximation. Consequently, we consider normally distributed microstructure noise in Section 3.

Under Assumption 2, it is clear that return noise forms an MA(1) process with a unit root. To determine the covariance structure of observed returns, we assume that latent returns form a weakly stationary martingale.

7

Lemma 1 *If in addition to Assumption 1 and Assumption 2, r_t forms a weakly stationary process with unconditional mean zero and unconditional variance σ_r^2, then the autocovariance function of observed returns obeys*

$$Cov\left(\tilde{r}_t, \tilde{r}_{t-k}\right) = \begin{cases} \sigma_r^2 + 2\sigma_\eta^2 & if \quad k = 0 \\[2mm] -\sigma_\eta^2 & if \quad k = 1 \\[2mm] 0 & if \quad k > 1 \end{cases} \cdot$$

Moreover, the first-order autocorrelation is given by

$$\rho = -\frac{\sigma_\eta^2}{\sigma_r^2 + 2\sigma_\eta^2}$$

For each day, which contains n intervals, define the following three volatility measures.

1. The integrated volatility, $V = \sum_{t=1}^n \sigma_t^2$.

2. An infeasible estimator, constructed from latent returns, $\bar{V} = \sum_{t=1}^n r_t^2$.

3. A feasible estimator, $\hat{V} = \sum_{t=1}^n \hat{r}_t^2$, where $\hat{r}_t = \tilde{r}_t$ in the absence of noise.

To motivate the form of the feasible estimator, decompose the estimation error as

$$V - \hat{V} = V - \bar{V} + \bar{V} - \hat{V}. \tag{2}$$

The behavior of $V - \bar{V}$ as a function of step length and the underlying volatility process has been studied by Barndorff-Nielsen and Shephard (2002a). If the step length is cho-

sen (hence n is fixed), then this part of the error is beyond the control of the researcher. Therefore, we focus on minimizing the the mean squared error $\mathsf{E}\left(\bar{V} - \hat{V}\right)^2$, where $\tilde{r}' = (\tilde{r}_1, \cdots, \tilde{r}_T)$ and $T = n \cdot J$ (J is the number of days in the sample). It is well known that the mean squared error is minimized by choosing

$$\hat{V} = \mathsf{E}\left(\bar{V}|\tilde{r}\right) = \mathsf{E}\left(\left.\sum_{t=1}^{n} r_t^2 \right| \tilde{r}\right) = \sum_{t=1}^{n} \mathsf{E}\left(r_t^2|\tilde{r}\right).$$

Thus, in order to minimize the effects of microstructure noise, we must extract expected squared latent returns from observed returns. The effectiveness with which the extraction can be achieved depends on the correct treatment of the microstructure noise.

2.1   Kalman Filter and Smoother

The Kalman filter provides a technique to separate (observed) contaminated returns into two components: the first corresponds to (latent) true returns and the second to microstructure noise. To construct the filter, we follow the notation in Hamilton (1994) (Harvey, 1989 also provides textbook treatment). The state vector consists of latent variables, $\xi_t' = \left(r_t, \eta_t, \eta_{t-1}\right)$. The observation equation relates the state vector to observed returns

$$\tilde{r}_t = \mathbf{H}'\xi_t \tag{3}$$

where $H' = (1, 1, -1)$. The state equation describes the dynamic evolution of the state vector

$$\boldsymbol{\xi}_{t+1} = \mathbf{F}\boldsymbol{\xi}_t + \mathbf{R}\mathbf{v}_{t+1} \tag{4}$$

where $v_t' = (r_t, \eta_t)$ with covariance matrix $Q_t = \begin{pmatrix} \sigma_t^2 & 0 \\ 0 & \sigma_\eta^2 \end{pmatrix}$ and the coefficient matrices

are $\mathbf{F} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 1 & 0 \end{pmatrix}$ and $\mathbf{R} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \\ 0 & 0 \end{pmatrix}$.

The Kalman filter delivers the linear projection of the state vector $\boldsymbol{\xi}_t$ given the sequence of observations $\{\tilde{r}_1, \cdots, \tilde{r}_t\}$. The Kalman smoother delivers the corresponding linear projection onto the extended sequence $\tilde{r}$. While these linear projections make efficient use of information, they may have larger MSE's than nonlinear projections. As conditional expectations need not be linear projections, we distinguish between linear projections and conditional expectations. Let $\hat{\mathsf{E}}_t$ represent linear projection onto $\mathcal{F}_t = \{\tilde{r}_t, \tilde{r}_{t-1}, \ldots, \tilde{r}_1, 1\}$. Let $\hat{\boldsymbol{\xi}}_{\tau|t} = \hat{\mathsf{E}}_t(\boldsymbol{\xi}_\tau)$ and let

$$\mathbf{P}_{\tau|t} = \mathsf{E}\left[\left(\boldsymbol{\xi}_\tau - \hat{\boldsymbol{\xi}}_{\tau|t}\right)\left(\boldsymbol{\xi}_\tau - \hat{\boldsymbol{\xi}}_{\tau|t}\right)'\right] \tag{5}$$

represent the mean squared error matrices of these projections. For example, the one-step-ahead mean squared error matrix $\mathbf{P}_{t|t-1}$ is a diagonal matrix with the first diagonal element equal to σ_t^2 and the second equal to σ_η^2. The third diagonal element we define as $c_t = Var\left(\hat{\eta}_{t-1|t-1}\right)$. The c_t are determined though a recursion described below. Let u_t

10

denote the one-step-ahead prediction error for the observed returns. Then it follows that the variance of u_t, which we denote by M_t, is given by $M_t = H'P_{t|t-1}H = \sigma_t^2 + \sigma_\eta^2 + c_t$.

The projections from the Kalman filter are given by the recursion

$$\hat{r}_{t|t}\left(\sigma_t^2\right) = \frac{\sigma_t^2}{M_t}\left(\tilde{r}_t + \hat{\eta}_{t-1|t-1}\right) \tag{6}$$

$$\hat{\eta}_{t|t} = \frac{\sigma_\eta^2}{M_t}\left(\tilde{r}_t + \hat{\eta}_{t-1|t-1}\right) \tag{7}$$

$$c_{t+1} = \frac{\sigma_\eta^2\left(\sigma_t^2 + c_t\right)}{M_t}. \tag{8}$$

The recursion and the boundary conditions $\hat{\eta}_{t|t} = 0$ and $c_1 = \sigma_\eta^2$ determine the sequence of filtered returns and filtered noise.

The projections from the Kalman smoother, which are the elements of $\hat{\boldsymbol{\xi}}_{\tau|T} = \hat{\mathsf{E}}_T\left(\boldsymbol{\xi}_\tau\right)$, are

$$\hat{r}_{t|T}\left(\sigma_t^2\right) = \hat{r}_{t|t}\left(\sigma_t^2\right) - \frac{\sigma_t^2}{\sigma_t^2 + c_t}\left(\hat{\eta}_{t|T} - \hat{\eta}_{t|t}\right) \tag{9}$$

$$\hat{\eta}_{t-1|T} = \hat{\eta}_{t-1|t} + \frac{c_t}{\sigma_t^2 + c_t}\left(\hat{\eta}_{t|T} - \hat{\eta}_{t|t}\right). \tag{10}$$

Smoothed quantities exhibit smaller variances than their filtered counterparts. For example, if we let $d_{t+1} = Var\left(\hat{\eta}_{t|T}\right)$, then it can be shown that

$$d_t = c_t\left(\frac{\sigma_t^2 + \sigma_\eta^2}{\sigma_t^2 + \sigma_\eta^2 + c_t}\right) - \left(\frac{c_t}{\sigma_t^2 + c_t}\right)^2\left(c_{t+1} - d_{t+1}\right). \tag{11}$$

As is easily verified from their definitions, $d_{T+1} = c_{T+1}$. Consequently, $d_T < c_T$, and,

11

by an induction argument, it follows that $d_t < c_t$ for $t = 1, 2, \ldots, T$, establishing that $Var\left(\hat{\eta}_{t|T}\right) < Var\left(\hat{\eta}_{t|t}\right)$ for $t = 1, 2, \ldots, T-1$.

The smoother estimates the latent returns as weighted averages of contemporaneous, lagged, and future observed returns. If variance is constant, so that $\sigma_t^2 = \sigma_r^2$ for all t, then the weights are nearly the same for all smoothed returns. To see the point clearly, consider a numerical example. If $\sigma_r^2 = 10$, $\sigma_\eta^2 = 1$, and $T = 7$, then, ignoring weights less than 0.001, we have that

$$\hat{r}_{4|7} = 0.006\tilde{r}_2 + 0.0709\tilde{r}_3 + 0.8452\tilde{r}_4 + 0.0709\tilde{r}_5 + 0.006\tilde{r}_6$$

while

$$\hat{r}_{3|7} = 0.0059\tilde{r}_1 + 0.0709\tilde{r}_2 + 0.8452\tilde{r}_3 + 0.0709\tilde{r}_4 + 0.006\tilde{r}_5.$$

Thus, an almost identical weighting scheme determines the third and fourth optimally estimated latent returns. For large samples, the weights are even more consistent. Except for a few returns at the beginning and end of the sample, the assumption of constant volatility leads to estimates of latent returns that are essentially a weighted average of the observed returns where the weights, for all practical purposes, are constants.

If, as is almost certainly the case in practice, latent returns do not exhibit constant volatility, then the optimal weights for estimating latent returns in (6)-(8) and (9)-(10) are not constant. Instead, during periods of high volatility the optimal weights are larger for the currently observed return and lower for the other returns.

2.2 Bias

With the estimated latent returns $\hat{r}_{t|T} = \hat{\mathsf{E}}\left(r_t|\tilde{r}\right)$, it seems natural to estimate realized volatility by $\sum_{t=1}^{n} \hat{r}_{t|T}^2$ (note, $\hat{r}_{t|T}^2$ stands for $\left(\hat{r}_{t|T}\right)^2$). Yet because filtering is a linear transformation, while squaring is not, $\hat{r}_t^2$ is a downward biased estimator of $\mathsf{E}\left(r_t^2|\tilde{r}\right)$. Fortunately, the size and direction of the bias is determined in the normal course of constructing the Kalman smoother. For the bias $\mathsf{E}\left(r_t^2 - \hat{r}_{t|T}^2\right)$, the properties of projection mappings imply[4]

$$\mathsf{E}\left[\left(r_t - \hat{r}_{t|T}^2\right)^2\right] = \mathsf{E}\left[\hat{\mathsf{E}}_T\left(r_t^2 - 2\hat{r}_{t|T}r_t + \hat{r}_{t|T}^2\right)\right] = \mathsf{E}\left[r_t^2 - \hat{r}_{t|T}^2\right]. \tag{12}$$

Thus, the bias equals the $(1,1)$ element of the mean squared error matrix for $\hat{r}_{t|T}$.

We find that the bias is

$$b_t\left(\sigma_t^2\right) = b_t^f\left(\sigma_t^2\right) - \left(\frac{\sigma_t^2}{\sigma_t^2 + c_t}\right)^2 \left(c_{t+1} - d_{t+1}\right), \tag{13}$$

where $b_t^f\left(\sigma_t^2\right)$ (the bias of the filtered return $\hat{r}_{t|t}$) is

$$b_t^f\left(\sigma_t^2\right) = \sigma_t^2 \left(\frac{\sigma_\eta^2 + c_t}{\sigma_t^2 + \sigma_\eta^2 + c_{t-1}}\right). \tag{14}$$

Recall that c_t is the element of the variance for the filtered prediction of η_{t-1} and d_t is the corresponding variance element for the smoothed prediction. As shown in (11), for $t < T$ the smoothed estimator has lower variance than the filtered estimator.[5] As a result, the

---

[4]Brockwell and Davis (1987, Proposition 2.3.2).
[5]If $t = T$, then the smoothed estimator is identical to the filtered estimator.

bias of the smoothed estimator is smaller than the bias of the filtered estimator and so we concentrate on the smoothed estimator in what follows.

To determine the magnitude of the bias, consider the simple case in which $\sigma_t^2 = \sigma_r^2$ (constant volatility). The bias, b_t, is well approximated by

$$\sigma_r^2 \left(1 - \frac{1}{\sqrt{1 + 4\sigma_\eta^2/\sigma_r^2}} \right).$$

In accord with intuition, the bias is a decreasing function of the return variance and an increasing function of the noise variance. If $\frac{\sigma_\eta^2}{\sigma_r^2} = .1$, then the bias is 15 percent of the return variance and 50 percent larger than the expected squared noise term.[6] If the noise variance dominates, so that $\sigma_r^2/\sigma_\eta^2 \simeq 0$, then the bias is approximately σ_r^2. If the return variance dominates, then the bias is near zero.

3 Multivariate Normal Approach

To analyze the multivariate normal case, it is convenient to work in vector form. Let $\mathbf{r} = (r_1, r_2, \ldots, r_T)'$ and $\boldsymbol{\eta} = (\eta_0, \eta_1, \ldots, \eta_T)'$ so that

$$\tilde{\mathbf{r}} = \mathbf{r} + \mathbf{B}\boldsymbol{\eta}.$$

---

[6]The bias of the filtered estimator is approximately $-2\rho\sigma_r^2$ (recall $\rho < 0$).

Here B is a selection matrix with first row $[-1, -1, 0, \cdots, 0]$. The covariance matrix of r

is $\Lambda = diag\left(\sigma_t^2\right)$. The Kalman smoother equations (in vector form) are

$$\hat{\mathbf{r}} = \mathbf{\Lambda} \left(\mathbf{\Lambda} + \sigma_\eta^2 \mathbf{B}\mathbf{B}'\right)^{-1} \tilde{\mathbf{r}} \quad \text{and} \quad \mathbf{\Sigma} = \sigma_\eta^2 \mathbf{B} \left(\mathbf{I} + \sigma_\eta^2 \mathbf{B}'\mathbf{\Lambda}^{-1}\mathbf{B}\right)^{-1} \mathbf{B}'. \tag{15}$$

From Assumption 1 it follows that $r_t|\sigma_t^2 \sim N\left(0, \sigma_t^2\right)$. If we extend the assumption to

$$\begin{pmatrix} \mathbf{r} \\ \boldsymbol{\eta} \end{pmatrix} \sim N_T \left(\mathbf{0}, \begin{pmatrix} \mathbf{\Lambda} & \mathbf{0} \\ \mathbf{0} & \sigma_\eta^2 \mathbf{I} \end{pmatrix}\right),$$

then it is simple to derive the conditional distribution of $\tilde{r}|r$. Specifically,

$$\mathbf{r}|\tilde{\mathbf{r}} \sim N_T\left(\hat{\mathbf{r}}, \mathbf{\Sigma}\right), \tag{16}$$

where $\hat{\mathbf{r}}$ and $\mathbf{\Sigma}$ are identical to the quantites from the Kalman smoother (15).

Under the assumption of joint normality $\hat{r} = E\left(r|\tilde{r}\right)$, so the smoothed estimator is the conditional expectation rather than simply the optimal linear projection. Similarly, $\mathbf{\Sigma} = Var\left(r|\hat{r}\right)$, rather than simply the MSE matrix of the linear projection. This is especially useful for understanding the source of the bias that arises from squaring filtered returns. Here

$$Var\left(r_t|\tilde{\mathbf{r}}\right) = \mathsf{E}\left(r_t^2|\tilde{\mathbf{r}}\right) - \mathsf{E}^2\left(r_t|\tilde{\mathbf{r}}\right). \tag{17}$$

The optimal estimator $\mathsf{E}\left(r_t^2|\tilde{\mathbf{r}}\right)$ exceeds the square of the optimal estimator for latent re-

turns $\mathsf{E}^2\left(r_t|\tilde{\mathbf{r}}\right)$. The correction term $Var\left(r_t|\tilde{\mathbf{r}}\right)$ does more than simply correct for the bias. Because the correction term corresponds to the conditional covariance matrix of r given the observed returns, the correction delivers the conditional expectation of squared returns. In consequence, we are able to form an optimal nonlinear estimator from an optimal linear estimator as

$$\mathsf{E}\left(r_t^2|\tilde{\mathbf{r}}\right) = \mathsf{E}^2\left(r_t|\tilde{\mathbf{r}}\right) + Var\left(r_t|\tilde{\mathbf{r}}\right).$$

4 Implementation

The above analysis, in which it is assumed that $\{\sigma_t^2\}$ and σ_η^2 are known, suggests the use of the bias-corrected estimator

$$\hat{V}\left(\sigma_t^2\right) = \sum_{t=1}^{n}\left[\hat{r}_{t|T}^2\left(\sigma_t^2\right) + b_t\left(\sigma_t^2\right)\right].$$

To implement the method, we need estimators of $\{\sigma_t^2\}$ and σ_η^2.

If the latent return variance is assumed constant, then $\sigma_t^2 = \sigma_r^2$ and the bias corrected estimator $\hat{V}\left(\sigma_r^2\right)$ is a function only of σ_r^2 and σ_η^2. From Lemma 1, the first two autocovariances of the observed returns series are sufficient for determining the variance of the noise and the expected variance of the true returns. (If one wishes to make further distributional assumptions, then ML estimators may be used in place of the method of moments estimators.) Andersen, Bollerslev, Diebold and Embens (2001a) employ an MA(1) estimator that, while similar to $\hat{V}\left(\sigma_r^2\right)$, does not contain smoothed estimates and makes no bias

16

correction.

For the case in which the return variances are not constant, we begin with $\hat{r}^2_{t|T}\left(\hat{\sigma}^2_r\right)$ and $b_t\left(\hat{\sigma}^2_r\right)$. We then estimate the time-varying variance with a rolling window[7]

$$\hat{\sigma}^2_{t|T} = \frac{1}{25}\sum_{k=t-12}^{t+12}\left(\hat{r}^2_{k|T}\left(\hat{\sigma}^2_r\right) + b_k\left(\hat{\sigma}^2_r\right)\right).\tag{18}$$

The estimated time-varying variances from (18) together with $\hat{\sigma}^2_\eta$, yield $\hat{r}^2_{t|T}\left(\hat{\sigma}^2_t\right)$ (from (9)) and $b_t\left(\hat{\sigma}^2_t\right)$ (from (13)).

For the case of constant variance, laws of large numbers ensure the consistency of $\hat{\sigma}^2_r$ and $\hat{\sigma}^2_\eta$. Similar results are derived for ML estimators in Ait-Sahalia, Mykland and Zhang (2003). To establish consistency if the return variance is not constant, it seems natural to specify a dynamic structure for $\{\sigma^2_t\}$. Rather than focus on this problem, we seek to recover latent realized volatility with a general purpose filter that minimizes mean squared error.

5 Performance

To test the performance of the suggested filter against realistic scenarios, we use a model for the simulated latent returns that is consistent with the return behavior of the S&P 500

---

[7]The rolling window width of 24 corresponds to two hours, which balances bias and variance in the presence of diurnal features.

stock index. A popular special case of Assumption 1 is

$$dp_t = \sigma_t dw_t$$

$$d\sigma_t^2 = \theta \left(\omega - \sigma_t \right) dt + (2\lambda\theta)^{1/2} dw_{\sigma_t}$$

(19)

where w_t and w_{σ_t} are independent Brownian motion. Drost and Werker (1996, Corollary 3.2) provide the map between (19) and the (discrete) GARCH(1,1),

$$p_t - p_{t-1/m} = r_{(m),t} = \sigma_{(m),t} \ z_{(m),t}$$

$$\sigma_{(m),t}^2 = \phi_{(m)} + \alpha_{(m)} r_{(m),t}^2 + \beta_{(m)} \sigma_{(m),t-1/m}^2$$

(20)

where $z_{(m),t}$ is (for the purposes of simulation) i.i.d. $N(0,1)$. Andreou and Ghysels (2002) find that 5-minute returns from the S&P 500 index are well approximated by the values

$$\phi_{(m)} = 0.0004, \ \alpha_{(m)} = 0.0037, \ \beta_{(m)} = 0.9963.$$

(21)

These parameters imply an unconditional return variance of $\sigma_r^2 = 7.9$ basis points over the 5-minute interval. While this unconditional variance is high (daily estimates of return variance are roughly 8 basis points), an appropriate rescaling by multiplying by 1/78 results in such small parameter values that simulation is difficult. As the relative mean squared error measurements that we report are invariant to such scaling, we follow Andreou and Ghysels and use the values in (21).

From (20) and (21) we simulate latent returns, r_t. We construct observed returns as

$\tilde{r}_t = r_t + \eta_t - \eta_{t-1}$, where η_t is generated as an i.i.d. $N\left(0, \sigma_\eta^2\right)$ random variable. To determine the noise variance, we invert the formula for ρ in Lemma 1 to obtain

$$\sigma_\eta^2 = \frac{-\rho}{1+2\rho}\sigma_r^2.$$

Hasbrouck and Ho (1987) report estimates of ρ between -.4 and -.1, so we allow ρ to take the values $[-.4, -.3, -.2, -.1]$. As decreasing the value of ρ increases σ_η^2, the resultant values of noise variance vary from $\sigma_\eta^2 = 1$ (for $\rho = -.1$) to $\sigma_\eta^2 = 15.8$ (for $\rho = -.4$).

To mirror trading days on the NYSE, which are 6.5 hours long, each simulated day contains 78 five-minute returns. We generate 10,000 trading days, a span that roughly corresponds to 50 years. For each day, j, we construct the latent realized volatility

$$\bar{V}_j = \sum_{t=(j-1)78+1}^{78j} r_t^2,$$

the feasible bias-corrected realized volatility estimator

$$\hat{V}_j\left(\hat{\sigma}_{t|T}^2\right) = \sum_{t=(j-1)78+1}^{78j} \left[\hat{r}_{t|T}^2\left(\hat{\sigma}_{t|T}^2\right) + b_t\left(\hat{\sigma}_{t|T}^2\right)\right],$$

and the infeasible bias-corrected estimator $\hat{V}_j\left(\sigma_t^2\right)$.

To compare this filter to methods that assume constant return variance, such as the MA(1) filter mentioned above, we construct $\hat{V}_j\left(\hat{\sigma}_r^2\right)$. To determine the gains from smooth-

19

ing, we also construct the estimator based on filtered (rather than smoothed) quantites

$$\hat{V}_j^f \left(\hat{\sigma}_{t|t}^2 \right) = \sum_{t=(j-1)78+1}^{78j} \left[\hat{r}_{t|t}^2 \left(\hat{\sigma}_{t|t}^2 \right) + b_t^f \left(\hat{\sigma}_{t|t}^2 \right) \right],$$

where $\hat{\sigma}_{t|t}^2$ is obtained from (18) with $\hat{r}_{k|t}^2 \left(\hat{\sigma}_r^2 \right)$ and $b_k^f \left(\hat{\sigma}_r^2 \right)$ in place of $\hat{r}_{k|T}^2 \left(\hat{\sigma}_r^2 \right)$ and $b_k \left(\hat{\sigma}_r^2 \right)$, respectively. Finally, for completeness, we construct $\hat{V}_j^f \left(\hat{\sigma}_r^2 \right)$.

To judge the quality of the realized volatility estimators, we measure the mean squared error (MSE) of each estimator relative to the infeasible (optimal) estimator. For example, the relative MSE for $\hat{V} \left(\hat{\sigma}_{t|T}^2 \right)$ is

$$\frac{MSE \left[\hat{V} \left(\hat{\sigma}_{t|T}^2 \right) \right]}{MSE \left[\hat{V} \left(\sigma_t^2 \right) \right]} = \frac{\sum_{j=1}^{10000} \left(\bar{V}_j - \hat{V}_j \left(\hat{\sigma}_{t|T}^2 \right) \right)^2}{\sum_{j=1}^{10000} \left(\bar{V}_j - \hat{V}_j \left(\sigma_t^2 \right) \right)^2}.$$

In Table 1, we present the relative efficiency calculations. Regardless of the degree of noise variance, or indeed of the decision to smooth, the gain from estimating a time-varying return variance is substantial. For the case with the smallest noise variance, the relative MSE for the smoothed estimator is reduced by more than half (from 3.5 to 1.4). As one would expect, increasing the noise variance renders the estimation problem more difficult, yet even for the highest noise variance the relative MSE for the smoothed estimator is substantially reduced (from 6.7 to 4.7). Moreover, while smoothing always leads to an efficiency gain, the magnitude of the efficiency gain resulting from smoothing is dominated by efficiency gain from allowing for time-varying volatility.

To determine the impact of diurnal patterns, we generate time-varying volatility that

	Constant Variance		Time-Varying Variance	
σ_η^2	$\hat{V}^f\left(\hat{\sigma}_r^2\right)$	$\hat{V}\left(\hat{\sigma}_r^2\right)$	$\hat{V}^f\left(\hat{\sigma}_{t|t}^2\right)$	$\hat{V}\left(\hat{\sigma}_{t|T}^2\right)$
15.8	8.4	6.7	5.7	4.7
5.9	7.6	6.0	3.9	3.4
2.6	6.1	5.1	2.4	2.2
1.0	3.9	3.5	1.5	1.4

Table 1: Relative Efficiency

	Constant Variance		Time-Varying Variance	
σ_η^2	$\hat{V}^f\left(\hat{\sigma}_r^2\right)$	$\hat{V}\left(\hat{\sigma}_r^2\right)$	$\hat{V}^f\left(\hat{\sigma}_{t|t}^2\right)$	$\hat{V}\left(\hat{\sigma}_{t|T}^2\right)$
15.8	8.1	6.5	5.5	4.6
5.9	7.4	5.9	3.8	3.3
2.6	6.0	5.0	2.4	2.2
1.0	3.9	3.5	1.5	1.4

Table 2: Relative efficiency with a diurnal pattern.

mirrors the U-shape pattern often observed in empirical returns. To do so, we construct a new sequence of return variances $\left\{\sigma_{(m),t}^{2*}\right\}$:

$$\sigma_{(m),t}^{2*} = \sigma_{(m),t}^2\left(1 + 1/3\cos\left(\frac{2\pi}{78}t\right)\right),$$

where $\sigma_{(m),t}^2$ is obtained from (20). Note that with the cyclic component, the expected variance doubles between the diurnal peak and trough. This process mimics the U-shape pattern as the maximum of the cosine term to corresponds to the beginning and ending of each day.

In Table 2, we find that the relative MSE measurements are surprisingly robust to the presence of diurnal patterns. When noise variance is about an order of magnitude smaller than the expected innovation variance (when $\sigma_\eta^2 = 1.0$), the MSE of the realized volatility estimator is about 4 percent larger when based on filtered returns. When noise variance

is roughly twice as large as the expected innovation variance (when $\sigma_\eta^2 = 15.8$), the filter based mean squared error is about 10 percent larger. Larger gains are achieved by the estimator based on the rolling volatility proxy, especially when noise volatility is relatively small. The mean squared errors based on the naive estimators are between 40 and 160 percent larger than corresponding mean squared errors based on the volatility proxy. The improvements from smoothing, relative to filtering, are shown in the last column.

Although currently used filters vary widely, we are aware of none that exploit the gains available from either smoothing or from the used of a high-frequency volatility proxy. Most filtering methods in uses are similar to the filtered naive estimator. Notice that the mean squared errors of the filtered naive estimators are more than double those of the smoothed estimators based on our proposed smoothed estimator based on the volatility proxy.

6   Conclusions

This article applies market microstructure theory to the problem of removing noise from a popular volatility estimate. The theory suggests that a Kalman smoother can optimally extract the latent squared returns, which are required for determining realized volatility from their noisy observable counterparts. However, the correct specification of the filter requires knowledge of a latent stochastic volatility state variable, and is therefore infeasible. We show that a feasible Kalman smoothing algorithm based on a simple rolling regression proxy for high-frequency volatility can improve realized volatility estimates. In simulations, the algorithm substantially reduces the mean squared error of realized volatil-

22

ity estimators even in the presence of strong diurnal patterns. The broad conclusion is that realized volatility estimators can be improved in an obvious way, by smoothing instead of merely filtering the data, and in a less obvious way, by bias correcting and using a straight-forward proxy of latent high-frequency volatility.

References

Ait-Sahalia, Y., Mykland, P., Zhang, L., 2003. How Often to Sample a Continuous-time Process in the Presence of Market Microstructure Noise. Working Paper w9611, NBER.

Andersen, T.G., 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., Ebens, H., 2001a. The Distribution of Realized Stock Return Volatility. Journal of Financial Economics 61, 43–76.

Andersen, T.G., Bollerslev, T., Diebold, F.X., Labys, P., 2001b. The Distribution of Realized Exchange Rate Volatility. Journal of the American Statistical Association 96, 42–55.

Andreou, E., Ghysels, E., 2002. Rolling Sample Volatility Estimators: Some New Theoretical, Simulation and Empirical Results. Journal of Business and Economic Statistics 20, 363–376.

Atchison, M., Butler, K., Simonds, R., 1987. Nonsynchronous Security Trading and Market Index Autocorrelation. Journal of Finance 42, 111–118.

Ball, C., 1988., Estimation Bias Induced by Discrete Security Prices. Journal of Finance 43, 841–865.

Russell, J., Bandi, F., 2004. Microstructure Noise, Realized Volatility, and Optimal Sampling. Econometric Society 2004 Latin American Meetings 220, Econometric Society.

Barndorff-Nielsen, O.E., Shephard, N., 2002a. Econometric Analysis of Realised Volatility and its Use in Estimating Stochastic Volatility Models. Journal of the Royal Statistical Society, Series B 64, 253–280.

Barndorff-Nielsen, O.E., Shephard, N., 2002b. Estimating Quadratic Variation Using Realised Variance. Journal of Applied Econometrics 17, 457–477.

Brockwell, P., Davis, R., 1987. Time Series: Theory and Methods. Springer-Verlag, New York; Berlin.

Campbell, A., Lo, W., MacKinlay, A., 1997. The Econometrics of Financial Markets. Princeton University Press, Princeton.

Cho, D., Frees, E., 1988. Estimating the Volatility of Discrete Stock Prices. Journal of Finance 43, 451–466.

Cohen, K., Maier, S., Schwartz, R., Whitcomb, D., 1978. The Returns Generation Process, Returns Variance, and the Effect of Thinness in Securities Markets. Journal of Finance 33, 149–167.

Cohen, K., Maier, S., Schwartz, R., Whitcomb, D., 1979. On the Existence of Serial Correlation in an Efficient Securities Market. TIMS Studies in the Management Sciences 11, 151–168.

Corsi, F., Zumbach, G., Müller, U., Dacorogna, M.M., 2001. Consistent High-Precision Volatility from High-Frequency Data. Economic Notes 30, 183–204.

Drost, F. Werker, B., 1996. Closing the GARCH Gap: Continuous Time GARCH Modeling. Journal of Econometrics 74, 31–57.

Engle, R. 1982., Autoregressive Conditional Heteroscedasticity with Estimates of the Variance of United Kingdom Inflation. Econometrica 50, 987–1007.

Gottlieb, G., Kalay, A., 1985. Implications of the Discreteness of Observed Stock Prices. Journal of Finance 40, 135–154.

Harvey, A., 1989. Forecasting, Structural Time Series Models and the Kalman Filter. Cambridge University Press, Cambridge.

Hamilton, J., 1994. Time Series Analysis. Princeton University Press, Princeton.

Hansen, P., Lunde, A., 2004. An Unbiased Measure of Realized Variance. Manuscript, Department of Economics, Brown University.

Hasbrouck, J., 1996. Modeling Market Microstructure Time Series. in: Maddala, G.S., Rao, C.R. (Eds.), Handbook of Statistics, Vol. 14. North-Holland, Amsterdam, pp. 647–692.

Hasbrouck, J., Ho, T., 1987. Order Arrival, Quote Behaviour, and the Return-Generating Process. Journal of Finance 42, 1035–1048.

Lo, A., MacKinlay, A.C., 1988. Stock Market Prices Do Not Follow Random Walks: Evidence from a Simple Specification Test. Review of Financial Studies 1, 41–66.

Lo, A., MacKinlay, A.C., 1990. An Econometric Analysis of Nonsynchronous-Trading. Journal of Econometrics 45, 181–212.

Madhavan, A., 2000. Market Microstructure: A Survey. Journal of Financial Markets 3, 205–258.

Newey,W., West, K., 1987. A Simple Positive Semi-definite, Heteroskedasticity and Auto-correlation Consistent Covariance Matrix. Econometrica 55, 703–708.

O'Hara, M., 1995. Market Microstructure Theory. Blackwell Publishers, Cambridge.

Oomen, R., 2004. Properties of Bias Corrected Realized Variance in Calendar Time and Business Time. Manuscript Warwick Business School.

Roll, R., 1984. A Simple Measure of the Effective Bid-Ask Spread in an Efficient Market. Journal of Finance 39, 1127–1139.

Zhou, B., 1996. High-Frequency Data and Volatility in Foreign Exchange Rates. Journal of Business and Economic Statistics 14, 45–52.