

Supplemental Appendix to:
Realized Semicovariances

This version: 17 February, 2020

Tim Bollerslev, Jia Li, Andrew J. Patton and Rogier Quaadvlieg

S1. Additional theoretical results

S1.1. Joint convergence of realized semicovariance matrices

Theorem 2 in the main text is restricted to the joint stable convergence of \widehat{P}_{12} and \widehat{N}_{12} . In this appendix, we describe the joint convergence of all components of the realized semicovariance measures in the bivariate setting (i.e., $d = 2$). Since $\widehat{M}_{11} = \widehat{M}_{22} = 0$ by construction, it suffices to study the following 7-dimensional statistic:

$$\widehat{\mathbf{S}} \equiv \left(\widehat{P}_{11}, \widehat{P}_{12}, \widehat{P}_{22}, \widehat{N}_{11}, \widehat{N}_{12}, \widehat{N}_{22}, \widehat{M}_{12} \right).$$

From Theorem 1 of the main text, we recall that the probability limit of $\widehat{\mathbf{S}}$ is

$$\mathbf{S} \equiv (P_{11}, P_{12}, P_{22}, N_{11}, N_{12}, N_{22}, M_{12}).$$

The proof for this general setting follows essentially the same steps as that of Theorem 2, but it is far more tedious because of the increased number of asymptotic covariances between these statistics. For brevity, we only describe the theorem.

We define various components in the limit theorem following the same structure of Section 2.2. It is convenient to define a few functions. Recall from the main text that

$$\begin{aligned} \psi(\rho) &\equiv \frac{\rho \arccos(-\rho) + \sqrt{1 - \rho^2}}{2\pi}, \\ \Psi(\rho) &\equiv \frac{3\rho\sqrt{1 - \rho^2} + (1 + 2\rho^2) \arccos(-\rho)}{2\pi}. \end{aligned}$$

Here, we also define

$$\begin{aligned} \Xi(\rho) &\equiv \frac{(3 + 2\rho^2) \sqrt{1 - \rho^2} + 5\rho \arccos(-\rho)}{4\pi}, \\ \Phi(\rho) &\equiv \frac{2\rho \arccos(\rho) - (1 + \rho^2) \sqrt{1 - \rho^2}}{2\pi}. \end{aligned}$$

We now define the limiting components in turn.

Bias components due to price drift, \mathbf{B} . This bias term is given by

$$\mathbf{B} \equiv (B_{11}, B_{12}, B_{22}, -B_{11}, -B_{12}, -B_{22}, 0)^\top,$$

S.1

where

$$\begin{aligned} B_{11} &\equiv \sqrt{\frac{2}{\pi}} \int_0^T b_{1,s} v_{1,s} ds, & B_{22} &\equiv \sqrt{\frac{2}{\pi}} \int_0^T b_{2,s} v_{2,s} ds, \\ B_{12} &\equiv \frac{1}{2\sqrt{2\pi}} \int_0^T (b_{1,s} v_{2,s} + b_{2,s} v_{1,s}) (1 + \rho_s) ds. \end{aligned}$$

Bias components due to continuous price-volatility covariation, \mathbf{L} . We denote for $x \in \mathbb{R}^2$,

$$\begin{aligned} f_{11}(x) &\equiv 2 \max\{x_1, 0\}, & f_{22}(x) &\equiv 2 \max\{x_2, 0\}, \\ f_{121}(x) &\equiv 1_{\{x_1 \geq 0\}} \max\{x_2, 0\}, & f_{122}(x) &= \max\{x_1, 0\} 1_{\{x_2 \geq 0\}} \end{aligned}$$

and then denote, for any 2×2 matrix A ,

$$\begin{aligned} F_{11}(A) &\equiv \mathbb{E} \left[f_{11}(AW_1) \int_0^1 W_s dW_s^\top \right], \\ F_{22}(A) &= \mathbb{E} \left[f_{22}(AW_1) \int_0^1 W_s dW_s^\top \right], \\ F_{12j}(A) &= \mathbb{E} \left[f_{12j}(AW_1) \int_0^1 W_s dW_s^\top \right], \quad j = 1, 2. \end{aligned}$$

The bias term \mathbf{L} is defined as $(L_{11}, L_{12}, L_{22}, -L_{11}, -L_{12}, -L_{22}, 0)^\top$, where

$$\begin{aligned} L_{11} &\equiv \int_0^T \text{Trace} [\tilde{\sigma}_{1,s} F_{11}(\sigma_s)] ds, \\ L_{22} &\equiv \int_0^T \text{Trace} [\tilde{\sigma}_{2,s} F_{22}(\sigma_s)] ds, \\ L_{12} &\equiv \sum_{j=1}^2 \int_0^T \text{Trace} [\tilde{\sigma}_{j,s} F_{12j}(\sigma_s)] ds. \end{aligned}$$

Diffusive sampling error spanned by price risk, ζ . We set

$$\begin{aligned} \gamma_{11,t} &\equiv \sqrt{\frac{2}{\pi}} v_{1,t}^2 \begin{pmatrix} v_{1,t} \\ \rho_t v_{2,t} \end{pmatrix}, & \gamma_{22,t} &\equiv \sqrt{\frac{2}{\pi}} v_{2,t}^2 \begin{pmatrix} \rho_t v_{1,t} \\ v_{2,t} \end{pmatrix}, \\ \gamma_{12,t} &\equiv \frac{(1 + \rho_t)^2 v_{1,t} v_{2,t}}{2\sqrt{2\pi}} \begin{pmatrix} v_{1,t} \\ v_{2,t} \end{pmatrix}, \end{aligned}$$

We then set, for $(j, k) \in \{(1, 1), (1, 2), (2, 2)\}$,

$$\zeta_{jk} \equiv \int_0^T (c_s^{-1} \gamma_{jk,t})^\top (\sigma_s dW_s).$$

We then set $\zeta \equiv (\zeta_{11}, \zeta_{12}, \zeta_{22}, -\zeta_{11}, -\zeta_{12}, -\zeta_{22}, 0)^\top$. The quadratic covariation matrix of the local martingale ζ equals $\int_0^T \Gamma_s ds$, where

$$\Gamma_t = \begin{pmatrix} \Gamma_t^* & -\Gamma_t^* & 0 \\ -\Gamma_t^* & \Gamma_t^* & 0 \\ 0 & 0 & 0 \end{pmatrix},$$

and Γ_t^* is a symmetric matrix process given by

$$\Gamma_t^* \equiv \begin{pmatrix} \gamma_{11,t}^\top c_t^{-1} \gamma_{11,t} & \gamma_{11,t}^\top c_t^{-1} \gamma_{12,t} & \gamma_{11,t}^\top c_t^{-1} \gamma_{22,t} \\ \bullet & \gamma_{12,t}^\top c_t^{-1} \gamma_{12,t} & \gamma_{12,t}^\top c_t^{-1} \gamma_{22,t} \\ \bullet & \bullet & \gamma_{22,t}^\top c_t^{-1} \gamma_{22,t} \end{pmatrix}.$$

Diffusive sampling error orthogonal to price risk, $\tilde{\zeta}$. This limit variable may be represented by its \mathcal{F} -conditional distribution as

$$\tilde{\zeta} = \int_0^T \bar{\gamma}_s^{1/2} d\widetilde{W}_s,$$

where \widetilde{W} is a 7-dimensional standard Brownian motion that is independent of the σ -field \mathcal{F} , and the $\bar{\gamma}$ process is defined by

$$\bar{\gamma}_t \equiv \bar{\Gamma}_t - \Gamma_t,$$

where

$$\bar{\Gamma}_t \equiv \begin{pmatrix} A_{1,t} & A_{2,t} & A_{3,t} \\ A_{2,t} & A_{1,t} & A_{3,t} \\ A_{3,t}^\top & A_{3,t}^\top & A_{4,t} \end{pmatrix},$$

and

$$\begin{aligned} A_{1,t} &\equiv \begin{pmatrix} \frac{5}{4}v_{1,t}^4 & v_{1,t}^3 v_{2,t} \Xi(\rho_t) & v_{1,t}^2 v_{2,t}^2 (\Psi(\rho_t) - \frac{1}{4}) \\ \bullet & v_{1,t}^2 v_{2,t}^2 (\Psi(\rho_t) - \psi(\rho_t)^2) & v_{1,t} v_{2,t}^3 \Xi(\rho_t) \\ \bullet & \bullet & \frac{5}{4}v_{2,t}^4 \end{pmatrix}, \\ A_{2,t} &\equiv \begin{pmatrix} -\frac{1}{4}v_{1,t}^4 & -\frac{1}{2}v_{1,t}^3 v_{2,t} \psi(\rho_t) & v_{1,t}^2 v_{2,t}^2 (\Psi(-\rho_t) - \frac{1}{4}) \\ \bullet & -v_{1,t}^2 v_{2,t}^2 \psi(\rho_t)^2 & -\frac{1}{2}v_{1,t} v_{2,t}^3 \psi(\rho_t) \\ \bullet & \bullet & -\frac{1}{4}v_{2,t}^4 \end{pmatrix}, \\ A_{3,t} &\equiv \begin{pmatrix} \Phi(\rho_t) v_{1,t}^3 v_{2,t} \\ 2v_{1,t}^2 v_{2,t}^2 \psi(\rho_t) \psi(-\rho_t) \\ \Phi(\rho_t) v_{1,t} v_{2,t}^3 \end{pmatrix}, \quad A_{4,t} \equiv v_{1,t}^2 v_{2,t}^2 [2\Psi(-\rho_t) - 4\psi(-\rho_t)^2]. \end{aligned}$$

Jump-induced sampling error, ξ . For ease of discussion, we recall some notations from

the main text. Let \mathcal{T}_j for $j \in \{1, 2\}$ denote the collection of jump times of $(X_{j,t})_{t \in [0, T]}$, with the corresponding “signed” subsets denoted by

$$\mathcal{T}_{j+} \equiv \{\tau \in \mathcal{T}_j : \Delta X_{j,\tau} > 0\}, \quad \mathcal{T}_{j-} \equiv \{\tau \in \mathcal{T}_j : \Delta X_{j,\tau} < 0\}.$$

For each $\tau \in \mathcal{T}_1 \cup \mathcal{T}_2$ associate the variables $(\kappa_\tau, \tilde{\xi}_{\tau-}, \tilde{\xi}_{\tau+})$ that are, conditionally on \mathcal{F} , mutually independent with the following conditional distributions: $\kappa_\tau \sim \text{Uniform}[0, 1]$, $\tilde{\xi}_{\tau-} \sim \mathcal{MN}(0, c_{\tau-})$, and $\tilde{\xi}_{\tau+} \sim \mathcal{MN}(0, c_\tau)$. Further define $\tilde{\eta}_\tau = (\tilde{\eta}_{1,\tau}, \tilde{\eta}_{2,\tau})^\top \equiv \sqrt{\kappa_\tau} \tilde{\xi}_{\tau-} + \sqrt{1 - \kappa_\tau} \tilde{\xi}_{\tau+}$. The limiting variable $\boldsymbol{\xi}$ has the form

$$\boldsymbol{\xi} = (\xi_{P,11}, \xi_{P,12}, \xi_{P,22}, \xi_{N,11}, \xi_{N,12}, \xi_{N,22}, \xi_{M,12}),$$

where, for $j \in \{1, 2\}$,

$$\begin{aligned} \xi_{P,jj} &\equiv \sum_{\tau \in \mathcal{T}_{j+}} 2\Delta X_{j,\tau} \tilde{\eta}_{j,\tau}, & \xi_{N,jj} &\equiv \sum_{\tau \in \mathcal{T}_{j-}} 2\Delta X_{j,\tau} \tilde{\eta}_{j,\tau}, \\ \xi_{P,12} &\equiv \sum_{\tau \in \mathcal{T}_{1+} \cap \mathcal{T}_{2+}} (\Delta X_{1,\tau} \tilde{\eta}_{2,\tau} + \Delta X_{2,\tau} \tilde{\eta}_{1,\tau}) \\ &\quad + \sum_{\tau \in \mathcal{T}_{1+} \setminus \mathcal{T}_2} \Delta X_{1,\tau} p(\tilde{\eta}_{2,\tau}) + \sum_{\tau \in \mathcal{T}_{2+} \setminus \mathcal{T}_1} \Delta X_{2,\tau} p(\tilde{\eta}_{1,\tau}), \\ \xi_{N,12} &\equiv \sum_{\tau \in \mathcal{T}_{1-} \cap \mathcal{T}_{2-}} (\Delta X_{1,\tau} \tilde{\eta}_{2,\tau} + \Delta X_{2,\tau} \tilde{\eta}_{1,\tau}) \\ &\quad + \sum_{\tau \in \mathcal{T}_{1-} \setminus \mathcal{T}_2} \Delta X_{1,\tau} n(\tilde{\eta}_{2,\tau}) + \sum_{\tau \in \mathcal{T}_{2-} \setminus \mathcal{T}_1} \Delta X_{2,\tau} n(\tilde{\eta}_{1,\tau}), \\ \xi_{M,12} &\equiv \sum_{\tau \in (\mathcal{T}_{1+} \cap \mathcal{T}_{2-}) \cup (\mathcal{T}_{1-} \cap \mathcal{T}_{2+})} (\Delta X_{1,\tau} \tilde{\eta}_{2,\tau} + \Delta X_{2,\tau} \tilde{\eta}_{1,\tau}) \\ &\quad + \sum_{\tau \in \mathcal{T}_{1+} \setminus \mathcal{T}_2} \Delta X_{1,\tau} n(\tilde{\eta}_{2,\tau}) + \sum_{\tau \in \mathcal{T}_{2-} \setminus \mathcal{T}_1} \Delta X_{2,\tau} p(\tilde{\eta}_{1,\tau}) \\ &\quad + \sum_{\tau \in \mathcal{T}_{1-} \setminus \mathcal{T}_2} \Delta X_{1,\tau} p(\tilde{\eta}_{2,\tau}) + \sum_{\tau \in \mathcal{T}_{2+} \setminus \mathcal{T}_1} \Delta X_{2,\tau} n(\tilde{\eta}_{1,\tau}). \end{aligned}$$

We are now ready to state the extension of Theorem 2 for all realized semicovariance components.

Theorem 1 *Under the same conditions as Theorem 2,*

$$\Delta_n^{-1/2} \left(\widehat{\mathbf{S}} - \mathbf{S} \right) \xrightarrow{\mathcal{L}\text{-s}} \mathbf{B} + \mathbf{L} + \boldsymbol{\zeta} + \tilde{\boldsymbol{\zeta}} + \boldsymbol{\xi}.$$

S1.2. Comparison with prior results on semivariances

Theorem 1 extends the theory of Barndorff-Nielsen, Kinnebrock, and Shephard (2010) to a bivariate case while allowing for general price and volatility jumps. Below, we special-

ize Theorem 1 to the univariate case and describe the stable convergence of $\Delta_n^{-1/2}(\widehat{P}_{11} - P_{11}, \widehat{N}_{11} - N_{11})$.

In the univariate case, $v_{1,t} = \sigma_t$ and we can compute the function $F_{11}(\cdot)$ explicitly. By Itô's formula, $\int_0^1 W_s dW_s = (W_1^2 - 1)/2$. Then, by direct integration,

$$F_{11}(a) = \mathbb{E} [\max\{aW_1, 0\} (W_1^2 - 1)] = \frac{a}{\sqrt{2\pi}}.$$

Then it is easy to see that

$$\begin{aligned} B_{11} &= \sqrt{\frac{2}{\pi}} \int_0^T b_{1,s} \sigma_s ds, & L_{11} &= \frac{1}{\sqrt{2\pi}} \int_0^T \tilde{\sigma}_s \sigma_s ds, \\ \gamma_{11,t} &= \sqrt{\frac{2}{\pi}} \sigma_t^3, & \zeta_{11} &= \sqrt{\frac{2}{\pi}} \int_0^T \sigma_s^2 dW_s, \\ \Gamma_t^* &= \frac{2}{\pi} \sigma_t^4, & \mathbf{\Gamma}_t &= \frac{2}{\pi} \sigma_t^4 \begin{pmatrix} 1 & -1 \\ -1 & 1 \end{pmatrix}, \\ A_{1,t} &= \frac{5}{4} \sigma_t^4, & A_{2,t} &= -\frac{1}{4} \sigma_t^4, & \bar{\mathbf{\Gamma}}_t &= \frac{1}{4} \sigma_t^4 \begin{pmatrix} 5 & -1 \\ -1 & 5 \end{pmatrix}, \\ \xi_{P,11} &\equiv \sum_{\tau \in \mathcal{T}_{1+}} 2\Delta X_{1,\tau} \tilde{\eta}_{1,\tau}, & \xi_{N,11} &\equiv \sum_{\tau \in \mathcal{T}_{1-}} 2\Delta X_{1,\tau} \tilde{\eta}_{1,\tau}. \end{aligned}$$

Consequently, Theorem 1 implies that $\Delta_n^{-1/2}(\widehat{P}_{11} - P_{11}, \widehat{N}_{11} - N_{11})^\top$ converges stably in law to,

$$\begin{aligned} &\frac{1}{\sqrt{2\pi}} \int_0^T (2b_{1,s} \sigma_s + \tilde{\sigma}_s \sigma_s) ds \begin{pmatrix} 1 \\ -1 \end{pmatrix} + \sqrt{\frac{2}{\pi}} \int_0^T \sigma_s^2 dW_s \begin{pmatrix} 1 \\ -1 \end{pmatrix} \\ &+ \int_0^T \bar{\gamma}_s^{1/2} d\widetilde{W}_s + \begin{pmatrix} \sum_{\tau \in \mathcal{T}_{1+}} 2\Delta X_{1,\tau} \tilde{\eta}_{1,\tau} \\ \sum_{\tau \in \mathcal{T}_{1-}} 2\Delta X_{1,\tau} \tilde{\eta}_{1,\tau} \end{pmatrix} \end{aligned}$$

where

$$\bar{\gamma}_t = \bar{\mathbf{\Gamma}}_t - \mathbf{\Gamma}_t = \sigma_t^4 \begin{pmatrix} \frac{5}{4} - \frac{2}{\pi} & -\frac{1}{4} + \frac{2}{\pi} \\ -\frac{1}{4} + \frac{2}{\pi} & \frac{5}{4} - \frac{2}{\pi} \end{pmatrix}.$$

In the case without price jumps (i.e., $\Delta X_1 = 0$), the limiting variable coincides with that in Proposition 2 of Barndorff-Nielsen, Kinnebrock, and Shephard (2010).

S2. Monte Carlo simulations

In this appendix, we examine finite-sample properties of the JCSD and DCSD tests using simulated data. Our data generating process is based on the one-factor stochastic volatility model used in Barndorff-Nielsen, Hansen, Lunde, and Shephard (2011), aug-

mented with intraday volatility periodicity and jumps. We generate the price process $X_t = (X_{1,t}, X_{2,t})$ for the two assets according to the following jump-diffusion model:

$$\begin{aligned} dX_{j,s} &= b_j ds + dV_{j,s} + dF_{j,s} + dJ_{j,s}, \\ dV_{j,s} &= \sqrt{1-\rho} \sigma_{j,s} dB_{j,s}, \\ dF_{j,s} &= \sqrt{\rho} \sigma_{j,s} dW_s, \end{aligned} \tag{S2.1}$$

where B_1 , B_2 , and W are independent Brownian motions, F_j is the diffusive factor component of asset j , and J_j is a pure-jump process. Note that the spot correlation coefficient between the two assets is ρ . We consider $\rho \in \{0, 0.5, 0.7\}$ in our experiments below. The spot volatility process $\sigma_{j,t}$ is generated as $\sigma_{j,s} = \varsigma_s \exp(\beta_0 + \beta_1 \varrho_{j,s})$, with $d\varrho_{j,s} = \alpha \varrho_{j,s} ds + dB_{j,s}$, and $\varsigma_s = C + Ae^{-ps} + Be^{-q(1-s)}$. The latter ς process captures the diurnal volatility pattern of Andersen, Dobrev, and Schaumburg (2012). We calibrate the parameters $(\beta_0, \beta_1, \alpha)$ at $(-5/16, 1/8, -1/40)$ as in Barndorff-Nielsen, Hansen, Lunde, and Shephard (2011), and (A, B, C, p, q) at $(0.75, 0.25, 0.88929198, 10, 10)$ as in Andersen, Dobrev, and Schaumburg (2012).

The stationary distribution $\varrho(0) \sim N(0, (-2\alpha)^{-1})$ is used to start the process in each simulation. The parameter choice ensures that $\mathbb{E}[\int_0^1 \sigma_{j,s}^2 ds] = 1$. Returns are generated at the 1-second frequency using a standard Euler discretization scheme, and aggregated to the 1-minute frequency ($\Delta_n = 1/390$), to match the empirical setup of Section 3.

As in the main text, we follow Bollerslev and Todorov (2011a,b) to adjust for intraday periodicity in the jump detection. In order to estimate the time-of-day effect in our simulations, for each iteration we generate 1,000 days of data which are exclusively used to estimate the truncation threshold. In particular, from these days we nonparametrically estimate a time-of-day (*TOD*) factor as well as the average bi-power variation \overline{BPV} . Subsequently these estimates are used to determine the truncation threshold as $u_n = 3 \times \overline{BPV} \times \overline{TOD}_n \Delta_n^{0.49}$.¹

To examine the size and power properties of the JCSD test, we set $b_j = 0$ and simulate the jumps as follows. For the ease of interpretation, we deterministically add positive co-jumps of size $5\Delta_n^{1/2}$ to observation 130 of both the series, which roughly corresponds to 5 local standard deviations of the diffusive returns. On observation 260, we add negative jumps of size $5\vartheta\Delta_n^{1/2}$, with $\vartheta \in \{1, 1.5, 2\}$. Note that the $\vartheta = 1$ case corresponds to the JCSD null hypothesis $P^\dagger = N^\dagger$. When $\vartheta > 1$, we obtain the alternative $P^\dagger < N^\dagger$.

Table S.1 reports the rejection rates of the JCSD test based on 10,000 Monte Carlo replications. Results with and without adjusting for the intraday periodicity are reported on the left and right panels, respectively. The first row of the table reports the rejection rates under the null hypothesis. With the periodicity adjustment, the size of the test

¹In the empirical section we use the full history of a stock to estimate the *TOD*, and replace \overline{BPV} with BPV_{t-1} , as in Bollerslev and Todorov (2011a,b).

Table S.1: Rejection Frequencies of the JCSD Test

	Periodicity Adjusted			Periodicity Unadjusted		
	$\rho = 0$	$\rho = 0.5$	$\rho = 0.7$	$\rho = 0$	$\rho = 0.5$	$\rho = 0.7$
$\vartheta = 1.0$	0.053	0.060	0.063	0.067	0.085	0.102
$\vartheta = 1.5$	0.872	0.786	0.775	0.877	0.793	0.774
$\vartheta = 2.0$	0.986	0.977	0.969	0.988	0.972	0.958

Note: The table reports the rejection frequencies of the JCSD test at 5% nominal level, based on 10,000 Monte Carlo replications. The data are generated according to (S2.1), aggregated to one-minute frequency, for various levels of ρ . The negative co-jump is ϑ times the magnitude of the positive co-jump. The $\vartheta = 1$ case and the $\vartheta > 1$ cases correspond to the null and alternative hypotheses, respectively. The intraday volatility periodicity is adjusted using the method of Bollerslev and Todorov (2011a,b).

is quite close to the 5% nominal level. At the same time, we observe some non-trivial overrejection when the intraday periodicity is unadjusted, especially when the correlation between the two assets is high. This comparison highlights the importance of taking into account the intraday volatility pattern. Under the alternative (i.e., $\vartheta > 1$), the JCSD test rejects the null with high probabilities, which approach 1 as ϑ increases. The power of the test is higher when the two assets are less correlated.

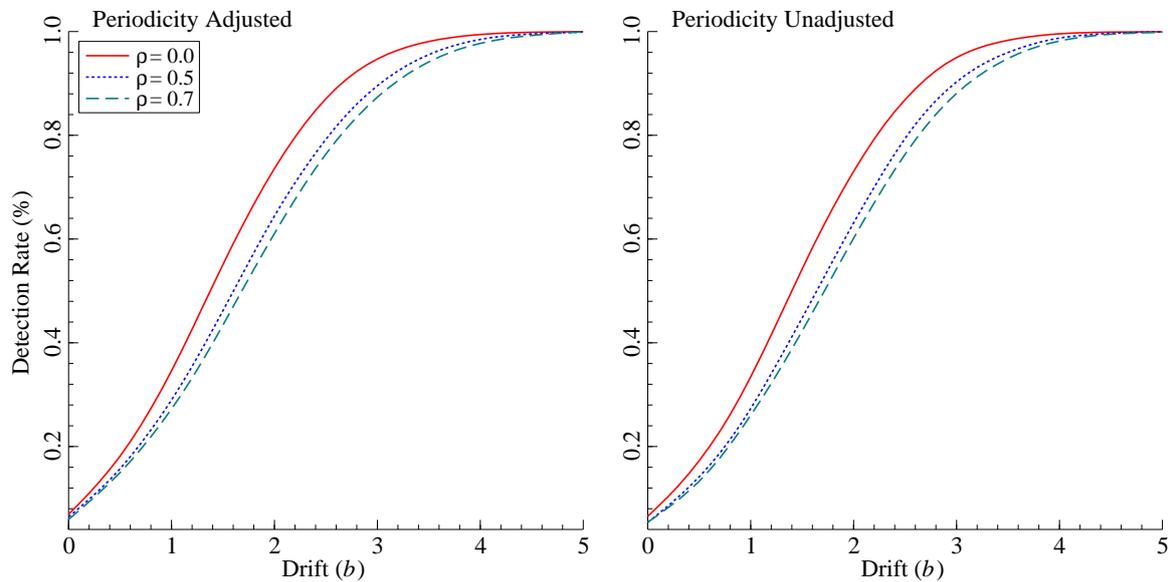
We next turn to the DCSD detection (recall Algorithm 2). As explained in main text, we interpret the DCSD t-statistic as a signal-to-noise ratio in the general case. To complement the empirical findings in Table 1 and Figure 3, here, we examine how the DCSD detection depends on assets' drift, which has a clear economic interpretation through the lens of Kyle's model. To this end, we set $\vartheta = 1$ and consider $b_j = b$ with $b \in [0, 5]$ and report the detection rate of large DCSD (i.e., $\hat{P}^* - \hat{N}^*$) in Figure S.1, with or without adjusting for intraday periodicity of volatility. From the figure, we see that as the co-drift effect (parameterized by b) becomes stronger, the detection rate increases from around 5% to close to 100%. Like in the JCSD test, the detection rates decreases when the correlation ρ is large.

Overall, the simulation results clearly corroborate the validity of the asymptotic results for making finite-sample inference, and the practical usefulness of the tests for detecting significant differences in the concordant semicovariance elements.

S3. JCSD and DCSD events

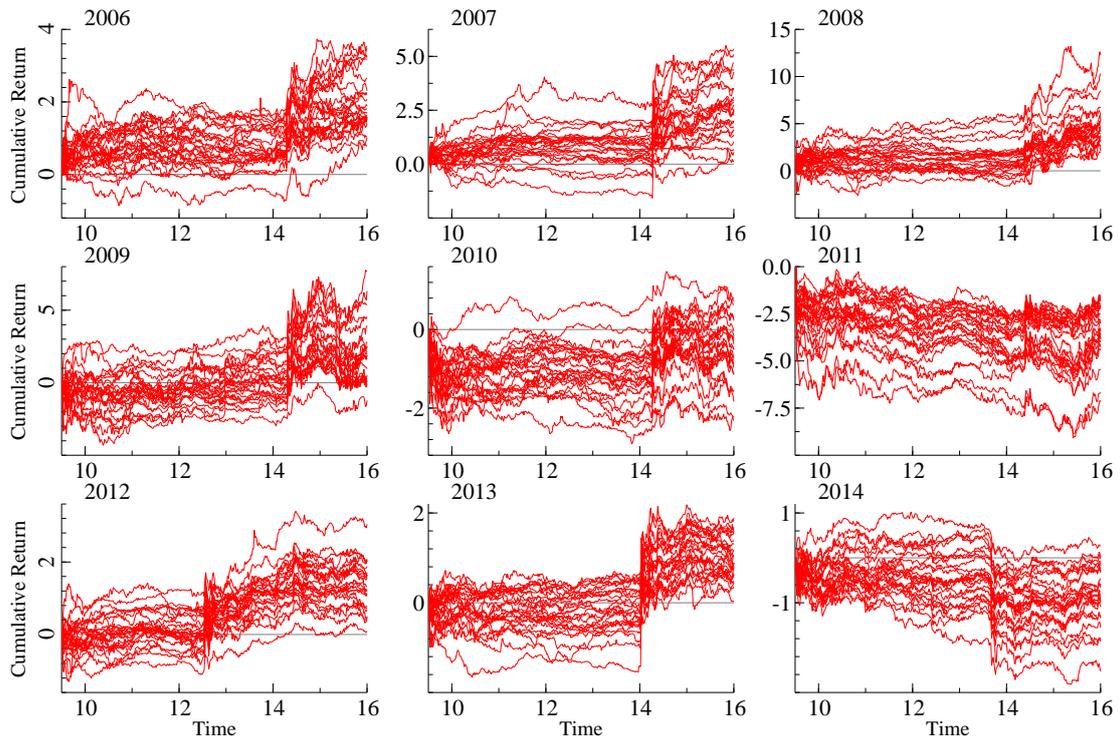
This section plots the cumulative intraday returns of the 30 DJIA stocks for all of the most significant event days listed in Table 1 of the main text. Figure S.2 and Figure S.3 show the results for JCSD and DCSD events, respectively.

Figure S.1: DCSD Detection Frequencies



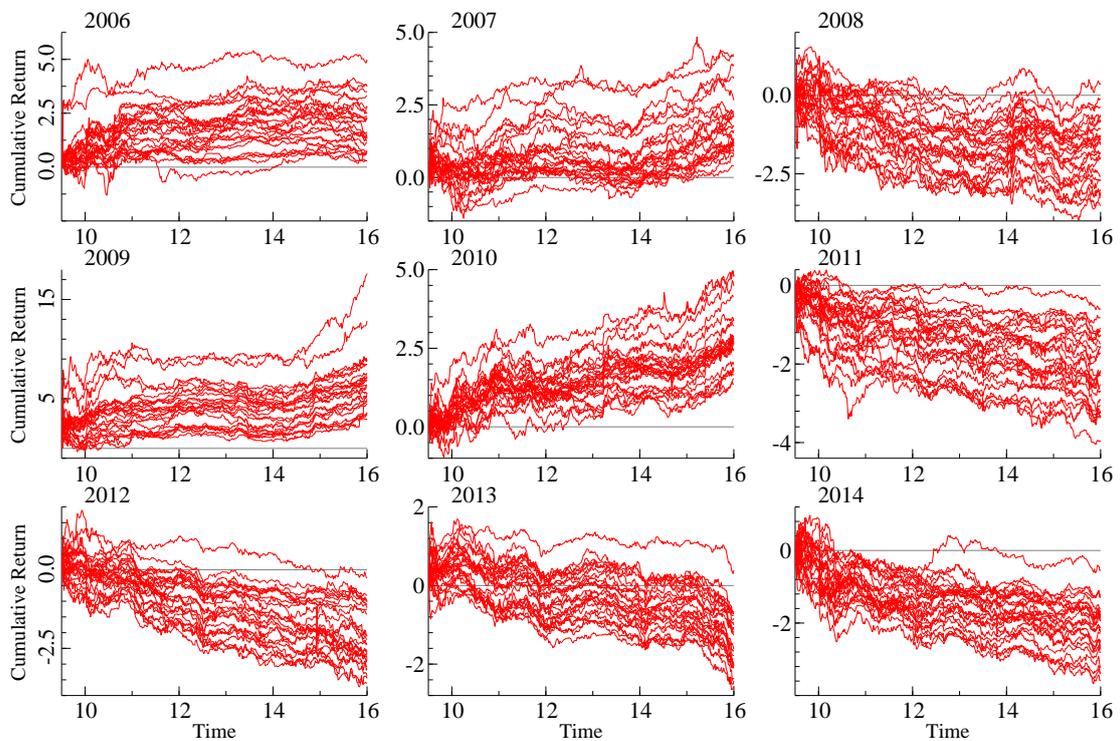
Note: The figure plots DCSD detection rates at the 5% level, based on 10,000 Monte Carlo Replications. The data are generated according to (S2.1), aggregated to one-minute frequency, for various levels of the correlation ρ , and drift $b_j = b$. The left-panel estimates the periodicity for jump-detection, while the right hand panel disregards the periodicity.

Figure S.2: DJIA Cumulative Returns on Jump CSD Event Days



Note: The figure plots the cumulative returns of the 30 DJIA stocks on the nine jump CSD event days listed in Panel A of Table 1.

Figure S.3: DJIA Cumulative Returns on Diffusive CSD Event Days



Note: The figure plots the cumulative returns of the 30 DJIA stocks on the nine diffusive CSD event days listed in Panel B of Table 1.

S4. Semicovariance dynamics

As noted in the main text, the within day differences in the price dynamics on DCSD and JCSD event days, also translate into different dynamic dependencies in the semicovariance components across days. To illustrate, Table S.2 reports for each pair of stocks, the average first-order autocorrelation between the realized semicovariance measures on day $t + 1$ and day t , conditional on the DCSD or JCSD tests being significant on day t .

Table S.2: Conditional First-Order Autocorrelations

	Conditioning Information				
	Unconditional	$\widehat{P}_{jk}^* \gg \widehat{N}_{jk}^*$	$\widehat{P}_{jk}^* \ll \widehat{N}_{jk}^*$	$\widehat{P}_{jk}^\dagger \gg \widehat{N}_{jk}^\dagger$	$\widehat{P}_{jk}^\dagger \ll \widehat{N}_{jk}^\dagger$
\widehat{C}_{jk}	0.641	0.732	0.757	0.642	0.687
\widehat{P}_{jk}	0.510	0.769	0.574	0.536	0.560
\widehat{N}_{jk}	0.658	0.608	0.773	0.659	0.704
\widehat{M}_{jk}	0.407	0.447	0.451	0.444	0.410

Note: The table reports the autocorrelations between the realized measures at time $t + 1$ and t , conditional on the rejection of CSD-based tests at time t . The first column presents the unconditional autocorrelations, while columns two and three (resp. four and five) report the autocorrelations conditional on the DCSD (resp. JCSD) test, where we use the \gg and \ll symbols to denote statistically significant inequalities.

Looking first at the unconditional autocorrelations reported in the first column, we see that \widehat{M} has the weakest unconditional autocorrelation (0.41), followed by positive semicovariance, \widehat{P} , (0.51), and negative semicovariance, \widehat{N} , has the strongest unconditional autocorrelation (0.66). This is a multivariate equivalent of the univariate result of Patton and Sheppard (2015) that negative realized semivariances tend to be more persistent than their positive counterparts. The mixed semicovariance term has no univariate analog, and so its relative lack of persistence cannot be foreshadowed in univariate work.

Meanwhile, conditioning on the outcome of the CSD-based tests can lead to notably different autocorrelations. In particular, \widehat{C} generally appears to be more persistent after a (positive or negative) diffusive event detected by the DCSD algorithm, with autocorrelation rising from 0.64 unconditionally to around 0.75 following a diffusive event. The effect of co-jumps detected using the JCSD test is smaller and depends on the direction of the rejection: persistence is slightly higher after negative co-jumps.

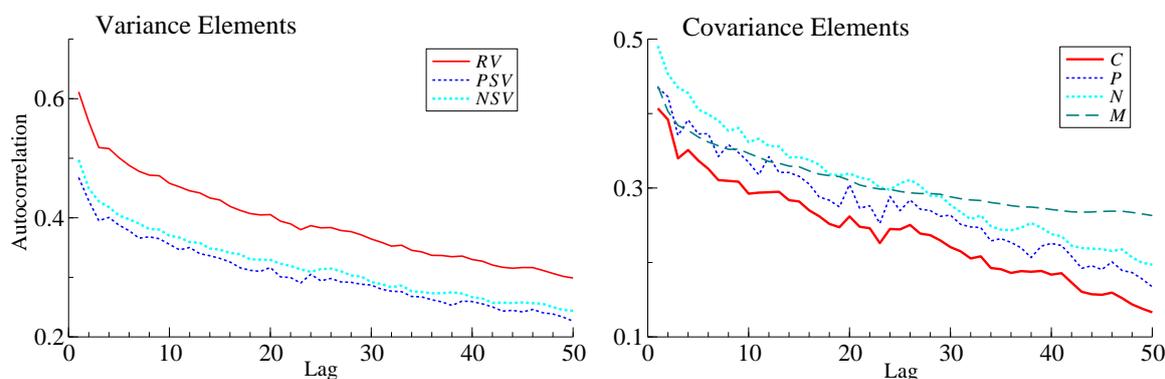
The impact of the CSD-based tests on the realized semicovariance dynamics is somewhat more complex. The persistence of \widehat{P} , in particular, is higher following a positive DCSD event on the previous day, with a much smaller increase after a negative DCSD event. Mirroring this, the persistence of \widehat{N} is higher following a *negative* DCSD event on the previous day, with a much smaller increase after a positive DCSD event. This is consistent with the idea that certain types of “soft” news is processed over multiple days. By comparison, the dynamic dependency after a JCSD jump event shows a different

pattern: the persistence of both the positive and the negative realized semicovariances slightly increases only after negative co-jumps; the effect of positive co-jumps is small.

S5. Unadjusted autocorrelograms

The magnitude and ordering of the first-order autocorrelations for the 1,000 randomly selected pairs of S&P 500 stocks shown in the main text differ somewhat from the results for the 30 DJIA stocks reported in the first column in the previous Table S.2. These differences are mostly driven by the instrumental variable approach of Hansen and Lunde (2014) that corrects for the inherent measurement error, and not by the expansion of the dataset from the 30 DJIA stocks to the broader set of S&P 500 constituent stocks. To underscore this, Figure S.4 shows the standard unadjusted empirical autocorrelation functions averaged across the S&P 500 stocks. The realized variance now appears more persistent than the semivariances. However, the calculation of the realized variances obviously also rely on twice as many observations as the semivariances, thus resulting in less downward-biased autocorrelation estimates. The ordering of the realized semicovariances is more or less in line with Figure 4 in the main text, albeit more erratic, but the magnitudes of the autocorrelations are obviously lower.

Figure S.4: Unadjusted Autocorrelations



Note: The graph plots the empirical autocorrelation functions for each of the individual elements of the realized semicovariance matrix. The autocorrelations are averaged over 1,000 random pairs of S&P 500 stocks.

S6. Semicovariances and volatility forecasting

S6.1. Model confidence sets

In order to assess whether the differences in forecast loss for the models discussed in Section 4.2 of the main text are statistically significant, Table S.3 below reports additional results from implementing the Model Confidence Set (MCS) of Hansen, Lunde, and Nason (2011), for each horizon considered in the main paper (i.e., 1, 5 and 22 days). The MCS procedure selects a set of models that contains the best model with a certain probability, here set to 80%. As the results in the tables below show, the SCHAR-r model is included in the MCS with a very high degree of confidence for both the daily, weekly and monthly forecasts.

The final column presents the results of the multi-horizon uMCS of Quaadvlieg (2020). This version of the MCS selects a set of models that have superior predictive ability across all horizons individually, rather than just on average. Correspondingly, the uMCS is applied to quadratic variation forecast paths, that is, RV_{t+h} , with $h = 1, \dots, 22$. A model that is removed from the set has significantly worse forecasting performance than one of the remaining models for at least one horizon. This more stringent criterion naturally decreases the size of the uMCS relative to each of the MCS. Most notably, the relative performance of the unrestricted SCHAR model deteriorates relative to the single-horizon MCS results. The restricted SCHAR-r is again included in the uMCS for the vast majority of randomly formed portfolios.

S6.2. Alternative forecasting models

This section complements the univariate portfolio forecasting results of Section 4.2 in the main text by including the results from a set of alternative specifications. We consider the following set of models for forecasting RV^w (i.e., the portfolio realized variance), succinctly expressed as functions of the specific realized measures used in the construction of the forecasts:

$$\begin{aligned}
 (1a) \quad RV_{t+1|t}^w &= f(\widehat{RV}_t^w, \widehat{RV}_{t-1:t-4}^w, \widehat{RV}_{t-5:t-21}^w) \\
 (1b) \quad RV_{t+1|t}^w &= f(\widehat{PSV}_t^w, \widehat{NSV}_t^w, \widehat{RV}_{t-1:t-4}^w, \widehat{RV}_{t-5:t-21}^w) \\
 (2) \quad RV_{t+1|t}^w &= f(\widehat{P}_t^w, \widehat{N}_t^w, \widehat{M}_t^w, \widehat{RV}_{t-1:t-4}^w, \widehat{RV}_{t-5:t-21}^w) \\
 (3) \quad RV_{t+1|t}^w &= f(\widehat{P}_t^w, \widehat{P}_{t-1:t-4}^w, \widehat{P}_{t-5:t-21}^w, \widehat{N}_t^w, \widehat{N}_{t-1:t-4}^w, \widehat{N}_{t-5:t-21}^w, \widehat{M}_t^w, \widehat{M}_{t-1:t-4}^w, \widehat{M}_{t-5:t-21}^w) \\
 (4) \quad RV_{t+1|t}^w &= f(\widehat{N}_t^w, \widehat{N}_{t-1:t-4}^w, \widehat{N}_{t-5:t-21}^w, \widehat{M}_t^w, \widehat{M}_{t-1:t-4}^w, \widehat{M}_{t-5:t-21}^w) \\
 (5) \quad RV_{t+1|t}^w &= f(\widehat{N}_t^w, \widehat{N}_{t-1:t-4}^w, \widehat{N}_{t-5:t-21}^w, \widehat{M}_{t-5:t-21}^w) \\
 (6) \quad RV_{t+1|t}^w &= f(\widehat{N}_t^w, \widehat{N}_{t-1:t-4}^w, \widehat{N}_{t-5:t-21}^w).
 \end{aligned}$$

Models 1a, 1b, 3, and 5 are referred to as HAR, SHAR, SCHAR and SCHAR-r, respectively, in the main paper. In addition to the models for directly forecasting RV^w , we

Table S.3: Formal Statistical Comparison of Forecasting Models.

	Daily Horizon			Weekly Horizon			Monthly Horizon			Path
	Avg	Ratio	MCS	Avg	Ratio	MCS	Avg	Ratio	MCS	uMCS
<i>Panel A. MSE Small Portfolio Case ($d = 10$)</i>										
HAR	1.849	1.000	0.296	1.314	1.000	0.268	1.407	1.000	0.192	0.246
SHAR	1.671	0.966	0.568	1.222	0.958	0.354	1.372	0.982	0.202	0.240
SCHAR	1.643	0.955	0.692	1.075	0.900	0.760	1.121	0.906	0.624	0.228
SCHAR-r	1.567	0.908	0.996	1.062	0.874	0.948	1.135	0.867	0.868	0.902
<i>Panel B. QLIKE Small Portfolio Case ($d = 10$)</i>										
HAR	0.141	1.000	0.420	0.086	1.000	0.238	0.108	1.000	0.166	0.322
SHAR	0.139	0.986	0.642	0.083	0.967	0.548	0.106	0.982	0.334	0.408
SCHAR	0.210	1.318	0.402	0.151	1.596	0.466	0.160	1.339	0.494	0.236
SCHAR-r	0.139	0.979	0.910	0.085	0.974	0.860	0.111	1.000	0.788	0.824
<i>Panel C. MSE Large Portfolio Case ($d = 100$)</i>										
HAR	0.048	1.000	0.188	0.028	1.000	0.422	0.031	1.000	0.246	0.408
SHAR	0.045	0.935	0.480	0.026	0.917	0.486	0.030	0.954	0.264	0.412
SCHAR	0.045	0.976	0.454	0.024	0.881	0.782	0.028	0.926	0.498	0.314
SCHAR-r	0.041	0.862	0.990	0.023	0.823	0.982	0.025	0.816	0.934	0.938
<i>Panel D. QLIKE Large Portfolio Case ($d = 100$)</i>										
HAR	0.119	1.000	0.390	0.070	1.000	0.536	0.071	1.000	0.240	0.430
SHAR	0.115	0.957	0.754	0.065	0.936	0.712	0.068	0.956	0.256	0.516
SCHAR	0.236	1.495	0.678	0.068	0.981	0.800	0.060	0.859	0.740	0.504
SCHAR-r	0.111	0.925	0.996	0.062	0.884	0.982	0.058	0.837	0.804	0.876

Note: The table reports the loss for forecasting the portfolio variance for portfolios of size $d = 10$ and $d = 100$, for the daily, weekly and monthly horizons, corresponding 1, 5 and 22 days, respectively. The reported numbers are based on 500 randomly selected portfolios. The Average column provides the average loss over time and all portfolios. The Ratio column gives the time-average ratio of losses across all sets of portfolios relative to the HAR model. The MCS column gives the fraction of pairs for which the model is in the 80% Model Confidence Set Hansen, Lunde, and Nason (2011). The final column presents the fraction of times the model is in the multi-horizon uniform Model Confidence Set of Quaedvlieg (2020), based on spot-volatility paths from $h = 1$ to 22 days ahead.

also consider a model in which we construct separate forecasts for \widehat{P}^w , \widehat{N}^w , and \widehat{M}^w , and then add up the forecasts to arrive at a forecast for RV^w , that is,

$$(7) \quad RV_{t+1|t}^w = f_P(\widehat{P}_t^w, \widehat{P}_{t-1:t-4}^w, \widehat{P}_{t-5:t-21}^w) \\ + f_N(\widehat{N}_t^w, \widehat{N}_{t-1:t-4}^w, \widehat{N}_{t-5:t-21}^w) + f_M(\widehat{M}_t^w, \widehat{M}_{t-1:t-4}^w, \widehat{M}_{t-5:t-21}^w)$$

where f_P , f_N , and f_M are models analogous to those in Models 1a to 6 above, but with the dependent variable set to \widehat{P}^w , \widehat{N}^w , and \widehat{M}^w , respectively.

We follow the forecasting setup discussed in the main text. To help assess statistical significance, we also include the results from implementing the 80% MCS of Hansen, Lunde, and Nason (2011). The detailed results for portfolios of size $d = 1, 2, 5, 10, 50, 100$ are reported in Table S.4. For each of the columns and each value of d we report the best model in boldface, corresponding to the smallest loss and/or the highest frequency of being in the MCS. For the univariate case, \widehat{M} is always zero and \widehat{P} and \widehat{N} are simply the realized semivariances, so for $d = 1$ we only report the loss based on Models 1a, 1b, 6 and 7.

The table clearly shows that for portfolios comprised of a large number of stocks, utilizing the information in the realized semicovariances can result in significantly better performance than the benchmark models that only use realized variances (or realized semivariances) estimated from the univariate portfolio returns. For $d \geq 50$, the best-performing model under both MSE and QLIKE loss is always Model 5 (the SCHAR-r model) which uses all three lags (daily, weekly and monthly) of \widehat{N}^w , but only the monthly lag of \widehat{M}^w , and excludes all lags of \widehat{P}_t^w altogether. For d between 2 and 10 the best-performing model under MSE loss is a more parsimonious version of that model, which includes only the three lags of \widehat{N}^w , while under QLIKE the preferred model is Model 2, which only decomposes the daily lag of \widehat{C}^w .

The MCS approach generally exhibits quite low power, failing to distinguish among the models in many of the cases. That being said, Model 5 (the SCHAR-r model) ends up in the MCS for almost all of the randomly selected collections of assets, suggesting that this model does indeed provide the best overall forecasting scheme. It is natural that the parsimonious specifications of Model 5 and 6 both do well, but the results indicate that the monthly lag of the mixed realized semicovariance \widehat{M}^w truly helps in predicting the future volatility, consistent with its strong persistence documented in Figure 4.

S6.3. Overnight returns

All of the results in the paper are based on open-to-close returns only. This section presents the results for predicting the portfolio return variance including the overnight variation. Although the actual values of the losses that include the overnight variation

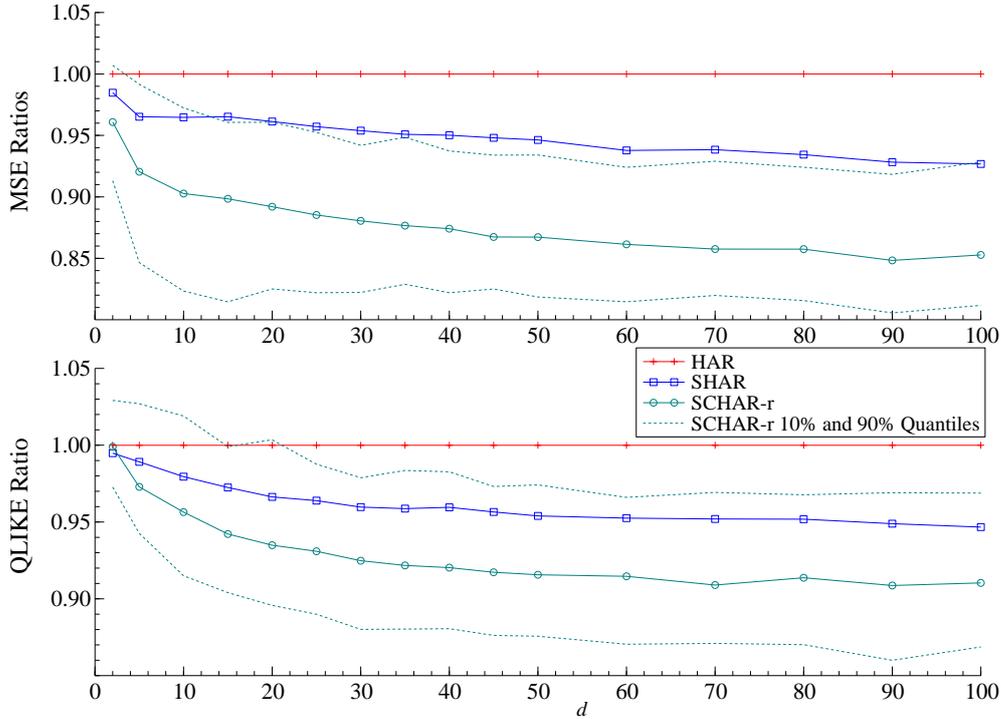
Table S.4: Additional Portfolio Variance Forecasting Models

Model	MSE			QLIKE		
	Average	Ratio	MCS	Average	Ratio	MCS
$d = 1$						
1a	35.112	1.000	0.886	0.239	1.000	0.728
1b	34.981	0.997	0.914	0.238	0.998	0.864
6	35.209	1.001	0.882	0.243	1.022	0.250
7	35.326	1.020	0.740	0.247	1.036	0.034
$d = 2$						
1a	8.094	1.000	0.760	0.184	1.000	0.652
1b	7.951	0.988	0.868	0.184	0.998	0.656
2	7.946	0.986	0.914	0.183	0.995	0.880
3	8.009	0.998	0.854	0.446	1.165	0.372
4	7.899	0.982	0.892	0.191	1.056	0.376
5	7.833	0.972	0.910	0.188	1.043	0.442
6	7.800	0.966	0.962	0.185	1.005	0.650
7	8.064	1.012	0.714	0.187	1.020	0.138
$d = 5$						
1a	3.014	1.000	0.466	0.151	1.000	0.544
1b	2.815	0.966	0.698	0.150	0.994	0.582
2	2.804	0.958	0.828	0.150	0.994	0.736
3	2.810	0.969	0.770	0.206	1.322	0.360
4	2.721	0.934	0.806	0.153	1.014	0.516
5	2.679	0.922	0.848	0.156	1.034	0.654
6	2.662	0.916	0.962	0.151	1.003	0.704
7	2.935	1.010	0.408	0.154	1.021	0.194
$d = 10$						
1a	1.849	1.000	0.306	0.141	1.000	0.442
1b	1.671	0.966	0.554	0.139	0.986	0.612
2	1.662	0.949	0.700	0.142	0.994	0.686
3	1.643	0.955	0.690	0.210	1.318	0.402
4	1.602	0.922	0.760	0.141	0.996	0.604
5	1.567	0.908	0.862	0.139	0.979	0.816
6	1.560	0.911	0.892	0.141	1.021	0.714
7	1.778	1.025	0.228	0.145	1.033	0.172
$d = 50$						
1a	0.335	1.000	0.254	0.130	1.000	0.462
1b	0.284	0.951	0.428	0.125	0.963	0.788
2	0.288	0.954	0.566	0.133	1.021	0.638
3	0.284	0.944	0.674	0.140	1.087	0.682
4	0.284	0.933	0.600	0.135	1.135	0.750
5	0.273	0.878	0.986	0.121	0.931	0.990
6	0.272	0.885	0.896	0.126	0.963	0.822
7	0.312	1.050	0.134	0.136	1.049	0.240
$d = 100$						
1a	0.048	1.000	0.282	0.119	1.000	0.492
1b	0.045	0.935	0.550	0.115	0.957	0.806
2	0.047	0.964	0.538	0.125	1.041	0.552
3	0.045	0.976	0.538	0.236	1.495	0.766
4	0.045	0.960	0.458	0.155	1.677	0.798
5	0.041	0.862	0.988	0.111	0.925	0.994
6	0.042	0.866	0.934	0.112	0.933	0.862
7	0.050	1.056	0.130	0.126	1.055	0.286

Note: This table provides aggregate forecasting results for the portfolio variance for random selected equally weighted portfolios of size $d = 1, \dots, 100$. Each number is based on 500 randomly selected sets of S&P500 stocks. The Average column provides the average loss over time and sets of series. The Ratio column provides the average over all sets of stocks of the ratio of time-averaged loss of model k relative to model 1. Finally, the MCS column provides the fraction of sets each respective model was part of the 80% Model Confidence Set. Boldface denotes the best model per dimension d and column.

are somewhat higher,² the relative losses depicted in Figure S.5 confirms the findings in Figure 5 in the main part of the paper: the SCHAR-r model systematically outperform the other models for all portfolio dimensions d .

Figure S.5: Median Loss Ratios, Overnight Returns Included



Note: The graph plots the median loss ratios as a function of the number of stocks in the portfolio, d . The ratio is calculated as the average loss of the models divided by the average loss of of the standard HAR, based on 500 random samples of d stocks.

S6.4. Multivariate Forecasting

In this section we forecast the full $(d \times d)$ realized covariance matrix, rather than the portfolio variance considered in the main paper. We consider a set of seven different models. All of the models rely on standard HAR-type specifications, along the lines of the vech-HAR model of Chiriac and Voev (2011), in which we estimate scalar autoregressive coefficients, but individual constant terms, based on standard OLS procedures. That is, each individual element has identical dynamics, but different unconditional means. The models mimic the extended set of portfolio variance model of Section S6.2.

²This is in line with earlier work that demonstrates that the overnight variation is difficult to predict; see, for example, Andersen, Bollerslev, and Huang (2011), who treat the overnight returns as jumps, requiring a separate specification.

- 1) $\widehat{C}_{t+1|t} = f(\widehat{C}_d, \widehat{C}_w, \widehat{C}_m)$
- 2) $\widehat{C}_{t+1|t} = f(\widehat{P}_d, \widehat{N}_d, \widehat{M}_d, \widehat{C}_w, \widehat{C}_m)$
- 3) $\widehat{C}_{t+1|t} = f(\widehat{P}_d, \widehat{P}_w, \widehat{P}_m, \widehat{N}_d, \widehat{N}_w, \widehat{N}_m, \widehat{M}_d, \widehat{M}_w, \widehat{M}_m)$
- 4) $\widehat{C}_{t+1|t} = f(\widehat{N}_d, \widehat{N}_w, \widehat{N}_m, \widehat{M}_d, \widehat{M}_w, \widehat{M}_m)$
- 5) $\widehat{C}_{t+1|t} = f(\widehat{N}_d, \widehat{N}_w, \widehat{N}_m, \widehat{M}_m)$
- 6) $\widehat{C}_{t+1|t} = f(\widehat{N}_d, \widehat{N}_w, \widehat{N}_m)$

As a final model we also consider the sum of the forecasts from three separate vech-HAR models for \widehat{P} , \widehat{N} , and \widehat{M} .

$$7) \widehat{C}_{t+1|t} = f_{\widehat{P}}(\widehat{P}_d, \widehat{P}_w, \widehat{P}_m) + f_{\widehat{N}}(\widehat{N}_d, \widehat{N}_w, \widehat{N}_m) + f_{\widehat{M}}(\widehat{M}_d, \widehat{M}_w, \widehat{M}_m)$$

Similar to the univariate case, for each d , ranging from 2 to 100, we take 500 random sets of stocks, and compute the semicovariances on their overlapping sample. We evaluate the performance of the forecasts using the multivariate analogs to MSE and QLIKE (see Laurent, Rombouts, and Violante, 2013):

$$Frobenius(\widehat{H}_t, \widehat{C}_t) = \text{Tr}((\widehat{C}_t - \widehat{H}_t)'(\widehat{C}_t - \widehat{H}_t))/d^2 \quad (\text{S6.2})$$

$$QLIKE(\widehat{H}_t, \widehat{C}_t, \widetilde{C}_t) = \text{Tr}(\widehat{H}_t^{-1}\widetilde{C}_t) - \log |\widehat{H}_t^{-1}\widetilde{C}_t| - d \quad (\text{S6.3})$$

where \widehat{H}_t refers to the forecast, and \widehat{C}_t denotes the ex-post Realized Covariance estimate.

The *QLIKE* loss function in equation (S6.3) is a modification of the original *QLIKE* loss function, also known as James and Stein (1961) loss, which is obtained if $\widetilde{C}_t = \widehat{C}_t$. The original loss function is not well defined if \widehat{C}_t is only positive semidefinite. We define \widetilde{C}_t as:

$$\widetilde{C}_t \equiv (1 - \lambda_t) C_t^* + \lambda_t \text{Diag} \{C_t^*\} \quad (\text{S6.4})$$

$$\text{where } \lambda_t = \arg \min(\lambda \in \{0, 0.1, \dots, 1\} : |\widetilde{C}_t| > 0) \quad (\text{S6.5})$$

$$\text{and } C_t^* \equiv I_d \odot \max \left\{ \widehat{C}_t, 0.001 \cdot J_d \right\} + (J_d - I_d) \odot \widehat{C}_t \quad (\text{S6.6})$$

In words, equations (S6.4)-(S6.5) shrink the off-diagonal elements of the matrix towards zero, and equation (S6.6) ensures that the diagonals are at least 0.001. This modification affects only the normalization of this loss function, not the ranking of competing forecasts.

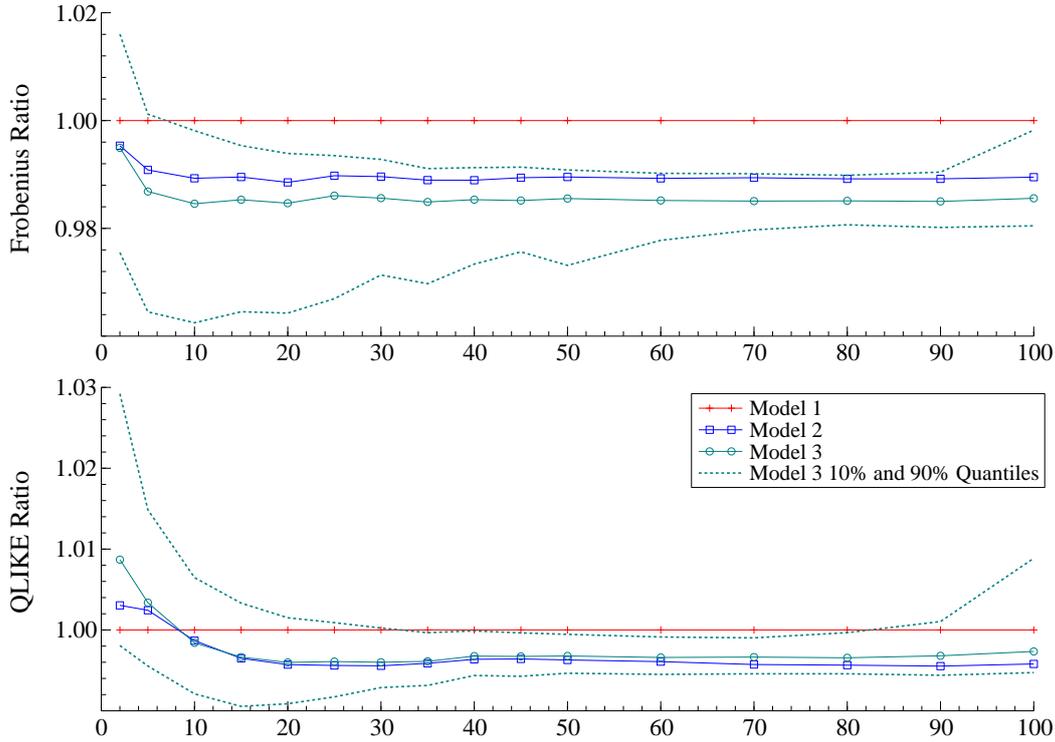
The results for matrices of size $d = 2, 5, 10, 50, 100$ are reported in Table S.5, in the same format as Table S.4. For all dimensions d , the standard HAR model is *never* preferred under Frobenius loss, while under *QLIKE* loss it is never preferred for $d \geq 10$. In vast dimensions ($d = 100$), Model 2 (where the daily lag of \widehat{C} is decomposed into semicovariances) and Model 3 (where all lags of \widehat{C} are decomposed into semicovariances) almost always outperform the other models. Whereas in the univariate setting parsimonious specifications appear to be preferred, in the multivariate setting the large cross-section

Table S.5: Performance Comparison for Covariance Forecasts

Model	Frobenius			QLIKE		
	Average	Ratio	MCS	Average	Ratio	MCS
$d = 2$						
1	18.087	1.000	0.882	0.581	1.000	0.928
2	17.971	0.995	0.924	0.584	1.005	0.576
3	17.983	0.996	0.918	0.590	1.017	0.416
4	18.111	1.001	0.794	0.599	1.032	0.032
5	18.104	1.001	0.800	0.599	1.031	0.038
6	18.310	1.013	0.166	0.627	1.080	0.006
7	18.206	1.017	0.696	0.595	1.025	0.038
$d = 5$						
1	9.546	1.000	0.586	2.154	1.000	0.738
2	9.397	0.988	0.852	2.161	1.004	0.420
3	9.367	0.985	0.948	2.166	1.005	0.482
4	9.417	0.992	0.656	2.202	1.023	0.022
5	9.433	0.993	0.620	2.210	1.026	0.020
6	9.818	1.047	0.062	2.777	1.283	0.014
7	9.563	1.010	0.572	2.169	1.008	0.314
$d = 10$						
1	5.943	1.000	0.280	6.333	1.000	0.386
2	5.815	0.986	0.678	6.329	1.000	0.384
3	5.786	0.982	0.952	6.327	0.999	0.572
4	5.806	0.992	0.454	6.424	1.015	0.006
5	5.815	0.994	0.406	6.443	1.018	0.002
6	6.222	1.151	0.020	8.382	1.508	0.006
7	5.926	1.005	0.380	6.328	0.999	0.456
$d = 50$						
1	1.372	1.000	0.026	155.347	1.000	0.096
2	1.347	0.988	0.424	154.762	0.996	0.820
3	1.351	0.984	0.986	154.842	0.997	0.460
4	1.337	0.994	0.138	155.535	1.002	0.002
5	1.341	0.996	0.124	155.831	1.002	0.002
6	1.575	1.307	0.006	169.305	1.106	0.014
7	1.363	1.000	0.038	155.506	1.001	0.010
$d = 100$						
1	0.677	1.000	0.116	431.189	1.000	0.260
2	0.674	0.994	0.398	430.952	0.999	0.804
3	0.671	0.990	0.928	433.053	1.002	0.224
4	0.704	1.041	0.086	444.891	1.082	0.008
5	0.744	1.099	0.044	449.751	1.055	0.016
6	0.948	1.419	0.004	440.220	1.023	0.014
7	0.679	1.002	0.066	433.445	1.006	0.000

Note: This table provides aggregate forecasting results for covariance matrices of size $d = 1, \dots, 100$. Each number is based on 500 randomly selected sets out of the 749 S&P500 stocks. We report the results for the models (1) to (7) defined above. The Average column provides the average loss over time and sets of series. The Ratio column provides the average over all sets of stocks of the ratio of time-averaged loss of model k relative to model 1. Finally, the MCS column provides the fraction of sets each respective model was part of the 80% Model Confidence Set. Boldface denotes the best model per dimension d and column.

Figure S.6: Median loss ratios: Covariance prediction



Note: The graph plots the median loss ratios as a function of the number of stocks in the portfolio, d . The ratio is calculated as the average loss of the models divided by the average loss of of the standard HAR, based on 500 random samples of d stocks.

allows us to pin down the parameters precisely, and the large models utilizing the full decomposition lead to the best forecasts.

While the reductions in loss reported in Table S.5 may appear numerically small, especially compared to those in the univariate case, they are statistically significant.³ In particular, for $d = 100$ the best models are in the MCS 93% and 85% of the samples for Frobenius and QLIKE loss, respectively. The next best model is part of the MCS for only 38% and 21% of the samples, respectively.

To help visualize the performance of the models across different matrix dimensions, Figure S.6 provides the multivariate analogue to Figure 5 in the main paper. It plots the median loss ratios across the 500 samples for d ranging from 2 to 100. To avoid cluttering the figure we only include the results for Models 1, 2 and 3. In addition, to illustrate the cross-sectional variation in out-performance, we also include the cross-sectional 10% and 90% quantiles of loss ratios for Model 3. The figure reveals that

³The estimation error in the realized covariances used in computing average losses means that the reductions in average loss are underestimated. Andersen, Bollerslev, and Meddahi (2005) show how to correct this underestimation in univariate applications. A similar correction could in principle be employed here, although it would not affect the relative ordering of the models.

the gains in forecast performance obtained from decomposing the realized covariance matrix into its semicovariance components appear even for relatively small values of d . Interestingly, the gains remain roughly constant for d ranging between 20 and 100. The 10% and 90% quantiles further show that these gains are present for almost all samples: under Frobenius loss, the 90% quantile is below one for all $d \geq 5$, while for QLIKE it is below one for d between 30 and 80.

S7. Multi-step forecast results

Table S.6 below complements the results in Table 3 and Figure 5 of the main paper by showing results for 5-day and 22-day ahead forecasts, corresponding to forecast horizons of one week and one month. For simplicity we rely on “direct” forecasts constructed by simply replacing the left-hand-side variable in the estimation of the different models with the realized portfolio variance over the relevant horizon.⁴ As the table shows, if anything the forecast improvements obtained by the SCHAR-r model vis-a-vis the standard HAR model appear even larger over longer forecast horizons; as a case in point the 18.4% relative improvement (measured by MSE) for the large portfolio at the monthly horizon is especially impressive.

S8. Realized semicovariances and market microstructure noise

We report signature plots of semicovariance estimates at different frequencies in Figures S.7. The top two plots show the variance and covariance estimates for the sample of 30 Dow Jones stocks over the period 2006-2014, which is the sample we use in Section 3 of the main paper. The bottom two panels show the plots for the 749 SP500 stocks over the period 1993-2014, which is the sample we use in Section 4 of the main paper. We start at the 1-minute frequency and increasingly aggregate returns up to half-hour returns, and report the estimates averaged over stocks and time. For the full sample we report the average over 500 random pairs, while for the DJIA sample we consider all pairs.

The signature plots were used by Hansen and Lunde (2006) to highlight the dangers of applying standard Realized Variance estimators to (too) high-frequency data. First consider the Dow Jones Sample for which we use the 1-minute frequency in our main paper. While the variance estimates based on 1-minute data are marginally higher than those based on 5- or 10-minute data, the increases are moderate relative to those found in Hansen and Lunde (2006), where the differences were a factor of 2 to 4 or more. Indeed,

⁴The construction of “iterated” forecasts for the SHAR and SCHAR models would necessitate additional modeling assumptions about the realized semi(co)variances; for further discussion of “direct” versus “iterated” volatility forecasts see, for example, Bollerslev, Hood, Huss, and Pedersen (2018) and the references therein.

Table S.6: Long-Horizon Results

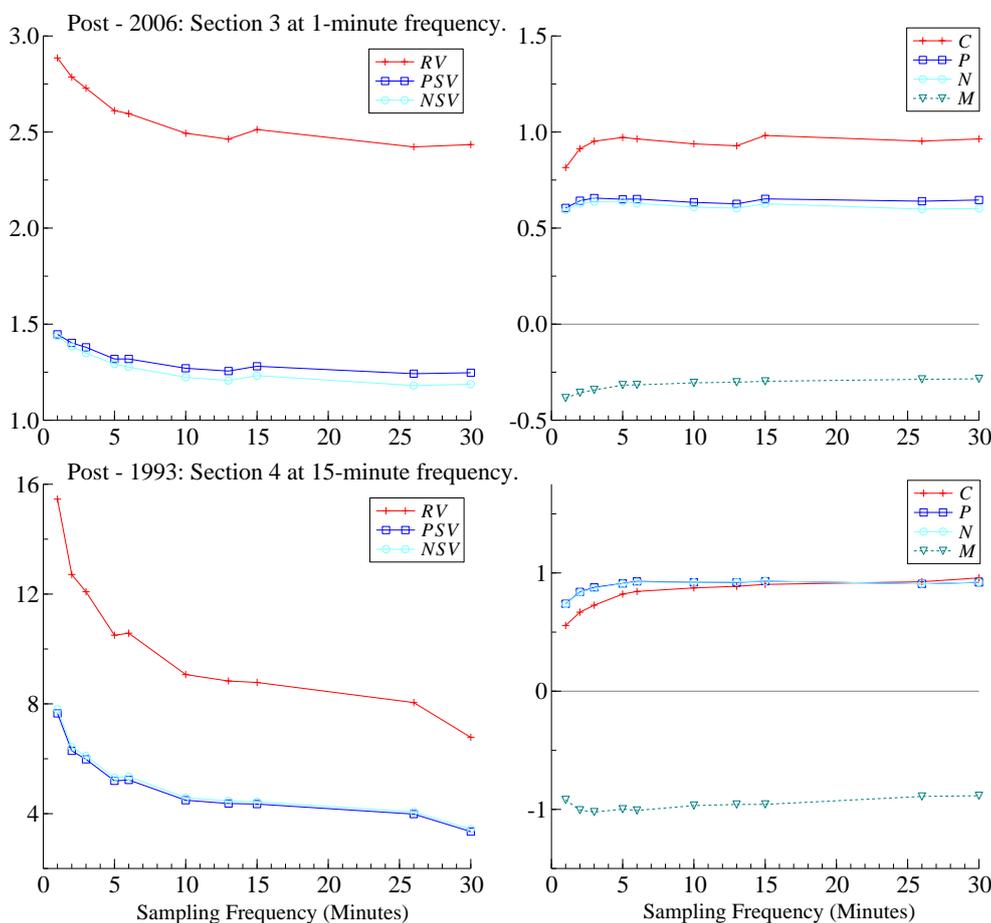
Model	MSE		QLIKE	
	Average	Ratio	Average	Ratio
<i>Panel A. Weekly Forecasts, Small Portfolio Case ($d = 10$)</i>				
HAR	1.314	1.000	0.086	1.000
SHAR	1.222	0.958	0.083	0.967
SCHAR	1.075	0.900	0.151	1.596
SCHAR-r	1.062	0.874	0.085	0.974
<i>Panel B. Weekly Forecasts, Large Portfolio Case ($d = 100$)</i>				
HAR	0.028	1.000	0.070	1.000
SHAR	0.026	0.917	0.065	0.936
SCHAR	0.024	0.881	0.068	0.981
SCHAR-r	0.023	0.823	0.062	0.884
<i>Panel C. Monthly Forecasts, Small Portfolio Case ($d = 10$)</i>				
HAR	1.407	1.000	0.108	1.000
SHAR	1.372	0.982	0.106	0.982
SCHAR	1.121	0.906	0.160	1.339
SCHAR-r	1.135	0.867	0.111	1.000
<i>Panel D. Monthly Forecasts, Large Portfolio Case ($d = 100$)</i>				
HAR	0.031	1.000	0.071	1.000
SHAR	0.030	0.954	0.068	0.956
SCHAR	0.028	0.926	0.060	0.859
SCHAR-r	0.025	0.816	0.058	0.837

Note: The table reports the loss for forecasting the portfolio variance for portfolios of size $d = 10$ and 100 for each of the different forecasting models, for weekly and monthly horizons. The reported numbers are based on 500 randomly selected portfolios. The Average column provides the average loss over time and all portfolios. The Ratio column gives the time-average ratio of losses across all sets of portfolios relative to the HAR model.

biases in 1-minute returns have substantially decreased since the writing of that paper, which is the foremost reason we start our sampling period in 2006. Next consider the full S&P500 sample, which we use in our forecasting section. Here, a long time-span was the primary consideration in our sampling frequency. As the bottom-left panel shows, 1-minute sampling here does lead to an almost two-fold increase in the realized variance compared to 15-minute sampling, which is why for this section and sample, we settled at 15-minute sampling, from which the signature plots remain mostly stable.

We test for the impact of microstructure noise on the 1-minute RV estimates in two ways. First, we compare our estimates with a realized kernel estimate (Barndorff-Nielsen, Hansen, Lunde, and Shephard (2008)). The average estimate is 2.68, compared to 2.89 for the Realized Variance, and the realized variance is within the 95% confidence interval of the realized kernel on 97.9% of the stock-days. Second, we test for the presence of microstructure noise using the test developed by Aït-Sahalia and Xiu (2019). We use the “H3n” statistic that is robust to stochastic volatility, non-Gaussian noise and jumps, and which came out as most favorable in their simulations. We reject the null of no noise in only 11.8% of the stock-days at the 5% level. Overall, we conclude that for this set of highly liquid stocks, on this relatively recent sample, measures based on a one-minute sampling frequency are not likely to be overly affected by microstructure noise.

Figure S.7: Signature Plots



Note: The figures shows signature plots for the datasets used in Sections 3 and Section 4 of the main paper. The top panel plots average realized (semi)variances across all DJIA stocks, and realized (semi)covariances across all possible pairs, at different sampling frequencies. The bottom panel shows average realized (semi)variances across all 749 S&P 500 stocks, and average realized (semi)covariances over 500 randomly selected pairs of the S&P 500 stocks, at different sampling frequencies.

References

- AÏT-SAHALIA, Y., AND D. XIU (2019): “A Hausman test for the presence of market microstructure noise in high frequency data,” *Journal of econometrics*, 211(1), 176–205.
- ANDERSEN, T. G., T. BOLLERSLEV, AND X. HUANG (2011): “A reduced form framework for modeling volatility of speculative prices based on realized variation measures,” *Journal of Econometrics*, 160(1), 176–189.
- ANDERSEN, T. G., T. BOLLERSLEV, AND N. MEDDAHI (2005): “Correcting the errors: Volatility forecast evaluation using high-frequency data and realized volatilities,” *Econometrica*, 73(1), 279–296.
- ANDERSEN, T. G., D. DOBREV, AND E. SCHAUMBURG (2012): “Jump-robust volatility estimation using nearest neighbor truncation,” *Journal of Econometrics*, 169(1), 75–93.
- BARNDORFF-NIELSEN, O. E., P. R. HANSEN, A. LUNDE, AND N. SHEPHARD (2008): “Designing Realized Kernels to Measure the ex post Variation of Equity Prices in the Presence of Noise,” *Econometrica*, 76, 1481–1536.
- BARNDORFF-NIELSEN, O. E., P. R. HANSEN, A. LUNDE, AND N. SHEPHARD (2011): “Multivariate Realised Kernels: Consistent Positive Semi-Definite Estimators of the Covariation of Equity Prices with Noise and Non-Synchronous Trading,” *Journal of Econometrics*, 162(2), 149–169.
- BARNDORFF-NIELSEN, O. E., S. KINNEBROCK, AND N. SHEPHARD (2010): “Measuring Downside Risk: Realised Semivariance,” in *Volatility and Time Series Econometrics: Essays in Honor of Robert F. Engle*, ed. by T. Bollerslev, J. R. Russell, and M. W. Watson, pp. 117–136. Oxford University Press, Oxford.
- BOLLERSLEV, T., B. HOOD, J. HUSS, AND L. H. PEDERSEN (2018): “Risk Everywhere: Modeling and Managing Volatility,” *Review of Financial Studies*, 31(7), 2729–2773.
- BOLLERSLEV, T., AND V. TODOROV (2011a): “Estimation of Jump Tails,” *Econometrica*, 79(6), 1727–1783.
- (2011b): “Tails, Fears, and Risk Premia,” *Journal of Finance*, 66(6), 2165–2211.
- CHIRIAC, R., AND V. VOEV (2011): “Modelling and forecasting multivariate realized volatility,” *Journal of Applied Econometrics*, 26(6), 922–947.
- HANSEN, P. R., AND A. LUNDE (2006): “Realized Variance and Market Microstructure Noise,” *Journal of Business and Economic Statistics*, 24(2), 127–161.
- (2014): “Estimating the Persistence and the Autocorrelation Function of a Time Series that is Measured with Error,” *Econometric Theory*, 30(1), 60–93.
- HANSEN, P. R., A. LUNDE, AND J. M. NASON (2011): “The Model Confidence Set,” *Econometrica*, 79(2), 453–497.
- JAMES, W., AND C. STEIN (1961): “Estimation with quadratic loss,” in *Proceedings of the fourth Berkeley symposium on mathematical statistics and probability*, vol. 1, pp. 361–379.

- LAURENT, S., J. V. K. ROMBOUTS, AND F. VIOLANTE (2013): “On loss functions and ranking forecasting performances of multivariate volatility models,” *Journal of Econometrics*, 173(1), 1–10.
- PATTON, A. J., AND K. SHEPPARD (2015): “Good Volatility, Bad Volatility: Signed Jumps and the Persistence of Volatility,” *Review of Economics and Statistics*, 97(3), 683–697.
- QUAEDVLIEG, R. (2020): “Multi-horizon forecast comparison,” *Journal of Business and Economic Statistics*, *Forthcoming*.