

Supplement to “Fixed- k inference for volatility”
(*Quantitative Economics*, Vol. 12, No. 4, November 2021, 1053–1084)

TIM BOLLERSLEV
Department of Economics, Duke University

JIA LI
Department of Economics, Duke University

ZHIPENG LIAO
Department of Economics, UCLA

SA. ADDITIONAL MONTE CARLO SIMULATIONS

In this section, we extend the simulation study in the main text in two ways. First, we consider another data generating process (DGP) in which the stochastic volatility is driven not only by the Brownian motion but also by Lévy processes with infinite-activity jumps. Second, we include confidence intervals (CI) constructed using a bootstrap method similar to [Gonçalves and Meddahi \(2009\)](#).

We start by introducing the DGP with Lévy shocks, under which the process X is generated according to

$$\begin{aligned}dX_t &= \sqrt{c_t} dW_t, & c_t &= V_{1,t} + V_{2,t}, & \rho &= -0.7, \\dV_{1,t} &= 0.0128(0.4068 - V_{1,t}) dt + 0.0954\sqrt{V_{1,t}}(\rho dW_t + \sqrt{1 - \rho^2} dL_{1,t}), \\dV_{2,t} &= 0.6930(0.4068 - V_{2,t}) dt + 0.7023\sqrt{V_{2,t}}(\rho dW_t + \sqrt{1 - \rho^2} dL_{2,t}),\end{aligned}$$

where W is a standard Brownian motion, and L_1, L_2 are independent Lévy processes. We simulate each Lévy process as a β -stable process with $\beta = 1.5$. To avoid unrealistic sample paths, we truncate the β -stable distribution on the $[-30, 30]$ interval. Note that the DGP used in Section 3 of the main text and the new DGP above share the same structure and parameter values, except that (B_1, B_2) are now replaced by (L_1, L_2) . For ease of discussion, we refer to them as the Brownian DGP and the Lévy DGP, respectively. We maintain the other Monte Carlo settings as in the main text.

Next, we introduce the bootstrap procedures for computing two-sided symmetric CIs of c_t and $\log(c_t)$, which are inspired by [Gonçalves and Meddahi \(2009\)](#). We clarify from the outset that [Gonçalves and Meddahi \(2009\)](#) study the bootstrap inference

Tim Bollerslev: boller@duke.edu
Jia Li: j1410@duke.edu
Zhipeng Liao: zhipeng.liao@econ.ucla.edu

for integrated variance instead of the spot variance, so their procedure is not directly applicable here. The bootstrap procedure considered below is constructed only in a similar spirit, and may be improved in future research on bootstrapping the spot variance.

Turning to the details, we recall that under the conventional “large- k ” asymptotics, the t-statistic associated with the spot variance estimator satisfies

$$\frac{\sqrt{k}(\hat{c}_{n,t} - c_t)}{\sqrt{2\hat{c}_{n,t}}} \xrightarrow{d} \mathcal{N}(0, 1). \quad (\text{SA.1})$$

The two-sided symmetric CI for c_t based on the conventional Gaussian approximation is given by

$$\left[\left(1 - \sqrt{\frac{2}{k}} z_{\alpha/2}\right) \hat{c}_{n,t}, \left(1 + \sqrt{\frac{2}{k}} z_{\alpha/2}\right) \hat{c}_{n,t} \right],$$

where the critical value $z_{\alpha/2}$ is the $1 - \alpha/2$ quantile of $\mathcal{N}(0, 1)$. The bootstrap method employs a different critical value. Specifically, we perform an i.i.d. resampling (with replacement) of the returns within the estimation block, and then use the resampled returns to compute the bootstrap spot variance estimator $\hat{c}_{n,t}^*$. We then set z_α^* as the $1 - \alpha$ data-conditional quantile of

$$\left| \frac{\sqrt{k}(\hat{c}_{n,t}^* - \hat{c}_{n,t})}{\sqrt{2\hat{c}_{n,t}^*}} \right|.$$

Note that z_α^* is the bootstrap analogue of $z_{\alpha/2}$ (which is the $1 - \alpha$ quantile of $|\mathcal{N}(0, 1)|$). The bootstrap CI of c_t is then given by

$$\left[\left(1 - \sqrt{\frac{2}{k}} z_\alpha^*\right) \hat{c}_{n,t}, \left(1 + \sqrt{\frac{2}{k}} z_\alpha^*\right) \hat{c}_{n,t} \right]$$

The bootstrap CI for the log spot variance can be constructed in a similar way. Applying the delta method on (SA.1) yields

$$\sqrt{\frac{k}{2}} (\log(\hat{c}_{n,t}) - \log(c_t)) \xrightarrow{d} \mathcal{N}(0, 1).$$

Let \tilde{z}_α^* be the $1 - \alpha$ data-conditional quantile of

$$\left| \frac{\sqrt{k}(\log(\hat{c}_{n,t}^*) - \log(\hat{c}_{n,t}))}{\sqrt{2}} \right|.$$

The bootstrap CI for $\log(c_t)$ is then given by

$$\left[\log(\hat{c}_{n,t}) - \sqrt{\frac{2}{k}} \tilde{z}_\alpha^*, \log(\hat{c}_{n,t}) + \sqrt{\frac{2}{k}} \tilde{z}_\alpha^* \right].$$

Figure S1 plots, the finite-sample coverage rates of the fixed- k , Gaussian-based, and bootstrap CIs under the Brownian DGP. This figure is constructed in the same way as

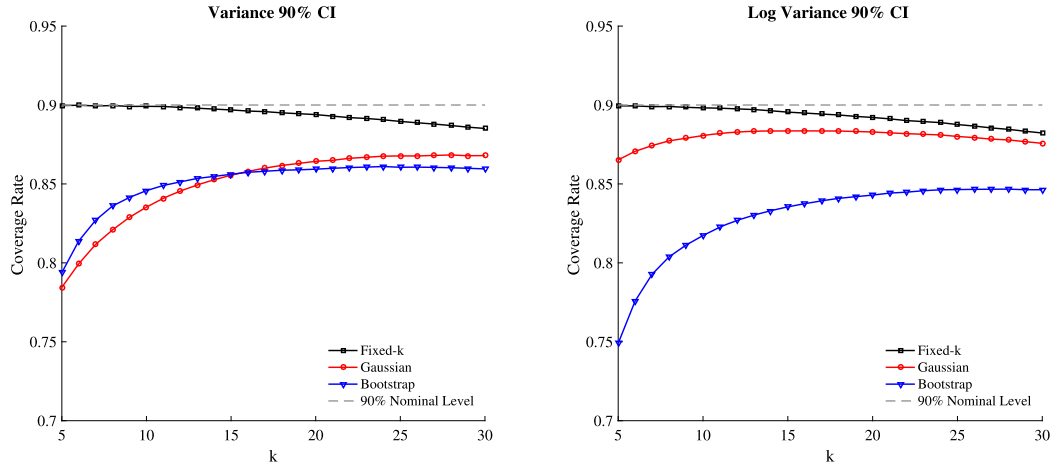


FIGURE S1. Coverage rates under the Brownian DGP. The figure plots the simulation coverage probabilities for the optimal fixed- k 90% CIs for the spot variance described in Corollary 1, the symmetric fixed- k 90% CIs for the log spot variance described in Corollary 2, together with the large- k 90% Gaussian-based and bootstrap CIs. The block size k ranges from 5 to 30.

Figure 2 in the main text, except that we now also include the coverage rates of the bootstrap CIs. From the left panel, we see that for the spot variance, the bootstrap CI has less size distortion than the Gaussian-based CI when k is small ($k < 15$). But the former underperforms the latter when k is large. However, the bootstrap CI appears to suffer from larger size distortion than the Gaussian-based CI for the log spot variance, as shown in the right panel. On the other hand, the fixed- k CIs always have better coverage than the other two alternatives. Figure S2 shows the results for the Lévy DGP, exhibiting essentially the same pattern as Figure S1.

Overall, these additional simulation results suggest that the proposed fixed- k CIs have better size control than the Gaussian-based and bootstrap-based alternatives. These findings are robust with respect to the inclusion of Lévy shocks in the volatility dynamics.

SB. ADDITIONAL EMPIRICAL RESULTS

SB.1 Volatility jumps near FOMC press conference

In this subsection, we study the behavior of spot volatility around the time of the press conference after each FOMC announcement. Specifically, we repeat the analysis underlying Table 2 in the main text, but setting the event time to be 30 minutes after the announcement, when the press conference typically starts. Table S1 reports the results. From the middle panel of the table, we see that \hat{c}_{1-} is statistically different from \hat{c}_{2-} 25 times in favor of the $\hat{c}_{1-} < \hat{c}_{2-}$ alternative. This provides further evidence for the decay of volatility after the initial announcement. In contrast, the same null hypothesis is rejected only 3 times in favor of $\hat{c}_{1-} > \hat{c}_{2-}$. Interestingly, as shown in the left panel of the

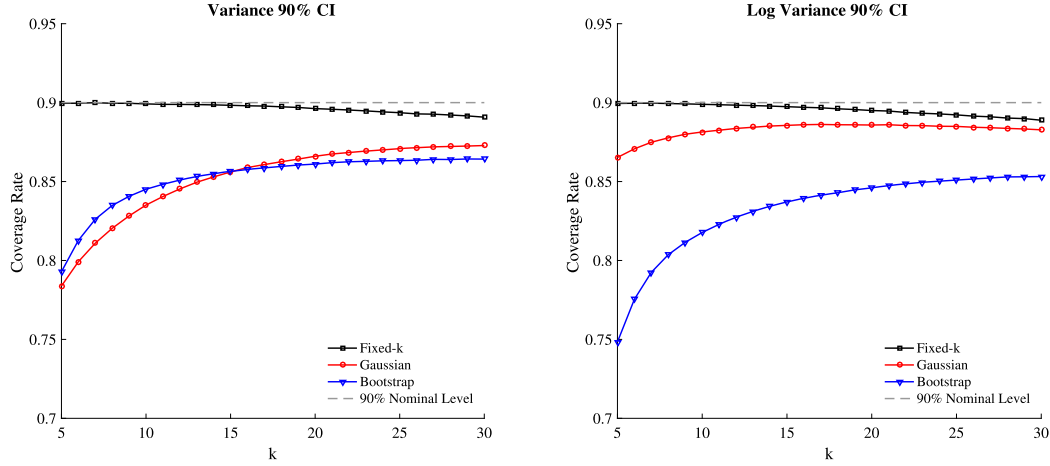


FIGURE S2. Coverage rates under the Lévy DGP. The figure plots the simulation coverage probabilities for the optimal fixed- k 90% CIs for the spot variance described in Corollary 1, the symmetric fixed- k 90% CIs for the log spot variance described in Corollary 2, together with the large- k 90% Gaussian-based and bootstrap CIs. The block size k ranges from 5 to 30.

TABLE S1. Significant volatility movements around FOMC press conference.

Null	$\hat{c}_{1+} - \hat{c}_{1-} = 0$		$\hat{c}_{1-} - \hat{c}_{2-} = 0$		$\hat{c}_{2+} - \hat{c}_{1+} = 0$	
Alternative	> 0	< 0	> 0	< 0	> 0	< 0
No. of rejections	11	16	3	25	13	22

Note: The table reports the number of FOMC announcements (out of a total of 109) for which the null hypothesis indicated in the top row is rejected by a one-sided test at the 10% significance level. \hat{c}_{2-} and \hat{c}_{1-} (resp., \hat{c}_{1+} and \hat{c}_{2+}) refer to the spot variance estimates in the two 10-minute estimation blocks before (resp., after) $\tau + 30$ minutes, where τ denotes the announcement time.

table, \hat{c}_{1-} is statistically different from \hat{c}_{1+} 11 times in favor of the alternative $\hat{c}_{1-} < \hat{c}_{1+}$, and \hat{c}_{1+} is statistically different from \hat{c}_{2+} 13 times in favor of $\hat{c}_{1+} < \hat{c}_{2+}$. These latter findings suggest that the volatility may increase during *some* press conferences, although the evidence is much weaker than that for the positive volatility jump at the announcement time.

SB.2 Robustness checks

Our empirical analysis in the main text is based on $k = 10$, corresponding to 10-minute windows for spot estimation. In this subsection, we repeat the analysis underlying Table 2 and Figure 4 in the main text with $k = 5$. Table S2 and Figure S3 present the results, which are qualitatively very similar to those in the main text.

TABLE S2. Significant volatility movements around FOMC announcements.

Null	$\hat{c}_{1+} - \hat{c}_{1-} = 0$		$\hat{c}_{1-} - \hat{c}_{2-} = 0$		$\hat{c}_{2+} - \hat{c}_{1+} = 0$	
Alternative	> 0	< 0	> 0	< 0	> 0	< 0
No. of rejections	98	1	19	11	6	33

Note: The table reports the number of FOMC announcements (out of a total of 109) for which the null hypothesis indicated in the top row is rejected by a one-sided test at the 10% significance level. \hat{c}_{2-} and \hat{c}_{1-} (resp., \hat{c}_{1+} and \hat{c}_{2+}) refer to the spot variance estimates in the two 5-minute estimation blocks before (resp., after) the announcements.

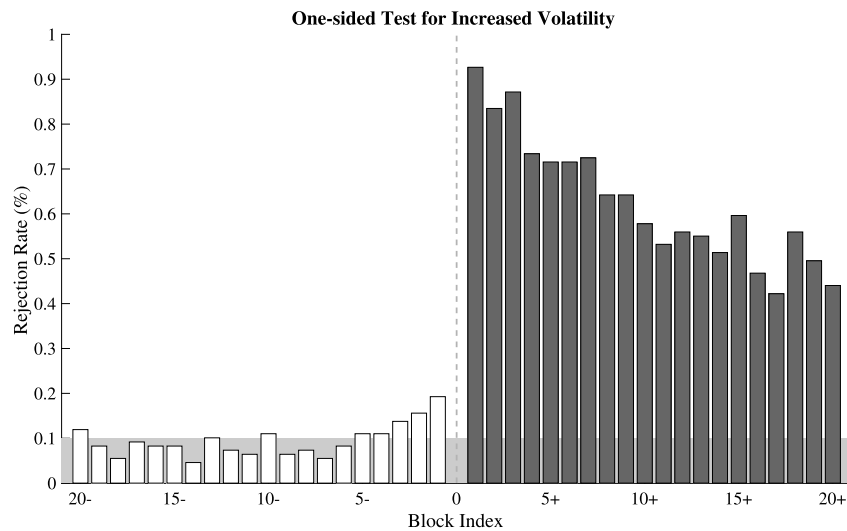


FIGURE S3. The figure reports the proportions of FOMC announcements for which the spot variance in each of the twenty 5-minute estimation blocks before and after the announcement is significantly higher, at the 10% significance level, than the spot variance in the 21th 5-minute benchmark block before the announcement. The light-shaded area indicates 10%.

REFERENCE

Gonçalves, S. and N. Meddahi (2009), “Bootstrapping realized volatility.” *Econometrica*, 77, 283–306. [1]

Co-editor Tao Zha handled this manuscript.

Manuscript received 16 October, 2020; final version accepted 16 March, 2021; available online 8 April, 2021.