Fixed-k Inference for Volatility[*]

Tim Bollerslev[†]   Jia Li[‡]   Zhipeng Liao[§]

April 8, 2021

Abstract

We present a new theory for the conduct of nonparametric inference about the latent spot volatility of a semimartingale asset price process. In contrast to existing theories based on the asymptotic notion of an increasing number of observations in local estimation blocks, our theory treats the estimation block size k as fixed. While the resulting spot volatility estimator is no longer consistent, the new theory permits the construction of asymptotically valid and easy-to-calculate pointwise confidence intervals for the volatility at any given point in time. Extending the theory to a high-dimensional inference setting with a growing number of estimation blocks further permits the construction of uniform confidence bands for the volatility path. An empirically realistically calibrated simulation study underscores the practical reliability of the new inference procedures. An empirical application based on intraday data for the S&P 500 equity index reveals highly significant abrupt changes, or jumps, in the market volatility at FOMC news announcement times, validating recent uses of various high-frequency-based identification schemes in asset pricing finance and monetary economics.

Keywords: spot volatility, high-frequency identification, semimartingale, uniform inference.
JEL Codes: C14, C22, C32.

[†]Department of Economics, Duke University, Durham, NC 27708; e-mail: boller@duke.edu.
[‡]Department of Economics, Duke University, Durham, NC 27708; e-mail: jl410@duke.edu.
[§]Department of Economics, UCLA, Log Angeles, CA 90095; e-mail: zhipeng.liao@econ.ucla.edu.

1   Introduction

We propose new inference procedures for the latent spot volatility of a semimartingale asset price process. Our approach is decidedly nonparametric in nature and allows for the construction of pointwise and uniform confidence sets that exhibit reliable coverage properties in empirically realistic settings. Applying the new procedures in the study of aggregate equity market volatility reveals highly significant abrupt changes, akin to "volatility jumps," at the time of FOMC announcements, adding important new insights to the burgeoning literature in macroeconomics and finance on identification through heteroskedasticity and news announcement effects.

Virtually all financial and macroeconomic time series exhibit time-varying volatility. Proper characterization of this variation plays a critical role not only in asset pricing finance and risk management (see, e.g., the survey by Andersen, Bollerslev, Christoffersen, and Diebold (2013)), but increasingly so also in macroeconomic models of aggregate economic fluctuations and policy analysis (see, e.g., the survey by Fernandez-Villaverde and Rubio-Ramirez (2013)). Meanwhile, the true volatility is inherently latent, and the development of econometric procedures for making reliable inference on volatility, and financial market volatility in particular, has been among the most active and vibrant areas of research in econometrics over the past thirty years (see, e.g., the introductory chapter in Andersen and Bollerslev (2018) summarizing some of the most important works).

Most of the early work in the area relied on specific discrete-time GARCH or continuous-time stochastic volatility models, in which the volatility is inferred from the particular parametric structure. As demonstrated by Nelson (1992) and Foster and Nelson (1996), even if the underlying parametric model is formally misspecified, the resulting filtered volatilities may still be interpreted as nonparametric consistent estimates. However, the assumption of ever finer sampled observations over diminishing, or "local," time windows underpinning this interpretation (and the consistency of spot volatility estimators more generally, see, e.g., Comte and Renault (1996), Kristensen (2010) and Alvarez, Panloup, Pontier, and Savy (2012)), is very difficult to mimic in practice (see, e.g., the early discussion in Andersen and Bollerslev (1997)).

Instead, building on the ideas in Andersen, Bollerslev, Diebold, and Labys (2001) and Barndorff-Nielsen and Shephard (2002), a more recent strand of literature has focused not on the spot volatility at any given point in time, but rather on the integrated volatility over a non-trivial time interval, like a day or a month. In this case, it is possible to consistently estimate integrated functionals of the true latent volatility process via a single level in-fill asymptotic scheme that is much more amenable to practical empirical implementations (see also Andersen and Bollerslev (1998),

Andersen, Bollerslev, Diebold, and Labys (2003), and Aït-Sahalia and Jacod (2007)). While these nonparametric realized volatility measures constructed from high-frequency intraday data are now widely used for modeling and forecasting time-varying volatilities over daily and longer horizons, they are not well suited for pinpointing abrupt changes in the volatilities, or studying issues regarding volatility jumps and the underlying economic causes and mechanisms at work.[1]

As a case in point, Andersen, Bollerslev, Diebold, and Vega (2003) document that many macroeconomic news announcements are accompanied by short, but sustained periods of elevated volatility typically lasting for one to two hours, during which the information content of the news release is processed more fully and a new equilibrium price is established. This high-frequency empirical feature has also received a lot of attention in the recent literature concerned with identification of macroeconomic shocks. Nakamura and Steinsson (2018), in particular, in their estimation of a simultaneous equation model for policy news shocks rely explicitly on identification through heteroskedasticity (following the approach of Rigobon (2003)) and asset price volatilities being higher during short "treatment" windows around FOMC news announcements than during "control" non-announcement periods.[2] Similarly, Bollerslev, Li, and Xue (2018) in their estimation of the extent of heterogeneous beliefs among market participants, as embodied in estimates of the intraday volume-volatility elasticity, rely crucially on a regression-discontinuity design and elevated volatilities immediately following macroeconomic news announcements.

Set against this background, we propose a new nonparametric framework for conducting reliable inference about spot volatilities over relatively short time windows, including tests for volatility jumps as a by-product. In contrast to the aforementioned existing theory pertaining to nonparametric spot volatility inference, which is built upon the asymptotic thought experiment of a growing number (i.e., $k \to \infty$) of observations over shrinking time windows for establishing consistency and asymptotic Gaussian-based inferential procedures, our "fixed-k" theory treats the local window size as a fixed constant, rather than a carefully chosen tuning sequence, thereby alleviating the practical empirical concerns about mimicking a "double asymptotic" theoretical framework in which the number of observations in each shrinking local estimation window is assumed to grow to infinity. However, as a result of this change in the asymptotic embedding, under the fixed-k

---

[1]Empirical evidence in support of financial market volatility jumps traces back to Bakshi, Cao, and Chen (1997), Bates (2000), and Duffie, Pan, and Singleton (2000), all of whom rely on fairly tightly parameterized stochastic volatility models and daily data in showing the importance of allowing for volatility jumps when valuing financial options; for more recent evidence based on high-frequency intraday data and a more flexible stochastic volatility model see Andersen, Fusari, and Todorov (2015).

[2]The importance of volatility, or "uncertainty," shocks for understanding aggregate investment decisions and recessions more generally has been forcefully emphasized by Bloom (2009). Fernandez-Villaverde, Guerrón-Quintana, Rubio-Ramirez, and Uribe (2011) also further highlight the real effects of volatility shocks.

theory, the nonparametric estimator based on a small number of observations does *not* consistently estimate the true latent spot volatility, simply because the law of large numbers has nothing to say in this "small-number" setting. Similarly, a central limit theorem is no longer applicable for justifying standard asymptotic Gaussian-based inference.[3]

Instead, we derive valid inference about the spot volatility process via a novel "approximate finite-sample" approach. To help fix ideas, suppose that the returns in a small k-sized window are generated as increments of a Brownian motion with constant volatility. In this situation, the nonparametric spot volatility estimator, defined as the local average of the normalized squared returns, has an *exact* scaled chi-squared distribution with k degrees of freedom. Our new fixed-k theory extends this simple baseline case to the case of a general Itô semimartingale model, allowing for stochastic drift, stochastic volatility, as well as price and volatility jumps, by exploiting the local Gaussianity of Itô semimartingales. As such, the asymptotic distribution of the estimator under the fixed-k theory is essentially inherited from the finite-sample distribution of the approximating Gaussian model, rather than from the aggregation of a large number of martingale differences as in conventional high-frequency theory on volatility inference (see, e.g., Chapter 13.3 of Jacod and Protter (2012)). This idea in turn allows for the construction of pointwise confidence intervals (CI) for the spot volatility at a given point in time based on scaled chi-squared limiting distributions. While the width of these new fixed-k CIs generally exceed those of the conventional Gaussian-based CIs, our Monte Carlo simulations confirm their far superior finite-sample coverage properties, especially for smaller block sizes k.

Going one step further, we extend this pointwise inference theory to a high-dimensional setting, allowing for uniform inference about the entire volatility path over non-trivial time intervals. We arrive at this extension by casting the joint inference problem in terms of "many" local estimation blocks within a given time span. The issue of high-dimensionality arises from the fact that with the size of each estimation block fixed, the number of blocks has to grow at the same rate as the sample size. This presents a unique challenge that has not hitherto been addressed in the literature. However, we show that the same approximation underlying the pointwise inference holds uniformly across all estimation blocks, thus allowing for the construction of uniform CIs for the volatility process.

An empirical application pertaining to the U.S. equity market volatility demonstrates the practical usefulness of the new theoretical results. In particular, adding to the burgeon empirical litera-

---

[3]This lack of consistency together with the decidedly nonparametric nature of our approach also sets it apart from other procedures based on specific parametric stochastic volatility models coupled with various in-fill asymptotic sampling schemes, as in, e.g., Bandi and Reno (2016).

ture in monetary economics that relies on the short-term impact of FOMC news announcements on asset prices for the identification of policy shocks (for early contributions, see, e.g., Cochrane and Piazzesi (2002) and Bernanke and Kuttner (2005)), we enlist the new fixed-k inference procedures to provide an in-depth look at asset market volatility around FOMC news announcements. Our analysis is based on high-frequency intraday data for the S&P 500 futures price around the release times of FOMC announcements during the 2003–2017 sample period. We find that the vast majority of FOMC announcements result in significant abrupt increases in the volatility at the time of the announcement, followed by a slow decline over the subsequent hour. We also uncover a slight "up-tick" in the volatility in advance of the official release time for certain announcements, suggestive of possible leakage of information. This richer, and statistically more rigorous, characterization of the short-term volatility movements, afforded by the new fixed-k inference procedures, formally corroborate the on-average heightened volatility following FOMC news announcements asserted in the aforementioned recent literature (e.g., Nakamura and Steinsson (2018) and Bollerslev, Li, and Xue (2018)) that rely on discontinuity-in-volatility-based identification strategies.

The remainder of the paper is organized as follows. Section 2 presents the new theory. Section 3 summarizes results from a Monte Carlo experiment corroborating the practical applicability of the new inference procedures in empirically realistic settings. Section 4 applies the proposed procedures in analyzing the behavior of aggregate equity market volatility around FOMC announcements. Section 5 concludes. The Appendix collects all proofs. The Online Supplemental Appendix contains additional simulation and empirical results.

2   Theory

Section 2.1 describes the setting. Section 2.2 presents the fixed-k pointwise inference theory for spot volatility, which is further extended to accommodate uniform inference in Section 2.3. Below, we use $\xrightarrow{\mathbb{P}}$ and $\xrightarrow{d}$ to denote convergence in probability and convergence in distribution, respectively. All limits are for $n \to \infty$.

2.1   The setting

Suppose that the (log) price process X is an Itô semimartingale defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$ written as

$$X_t = X_0 + \int_0^t b_s ds + \int_0^t \sigma_s dW_s + J_t, \tag{2.1}$$

where the drift process b and the stochastic volatility process σ are optional, and W is a standard Brownian motion. The J process captures the price jumps, defined as the sum of a purely discontinuous local martingale with jump sizes not bigger than 1, and a pure-jump process with jump sizes bigger than 1, both of which are driven by a homogeneous Poisson random measure on $\mathbb{R}_+ \times \mathbb{R}$. Our econometric interest centers on the spot variance process $c_t \equiv \sigma_t^2$.

We consider an in-fill asymptotic framework that is now standard in the high-frequency econometrics literature (see, e.g., Andersen, Bollerslev, Diebold, and Labys (2001) and Barndorff-Nielsen and Shephard (2002), and the book by Aït-Sahalia and Jacod (2014)). In particular, suppose that the price process X is observed at discrete times $i\Delta_n$ for $i = 0, 1, \ldots, n$ within the fixed time interval $[0, T]$ with $\Delta_n = T/n$. Let the ith return of X be denoted by

$$\Delta_i^n X \equiv X_{i\Delta_n} - X_{(i-1)\Delta_n}, \quad i \in \{1, \ldots, n\}.$$

We study a standard spot variance estimator constructed over asymptotically shrinking blocks; see, for example, Foster and Nelson (1996), Comte and Renault (1998), Lee and Mykland (2008), Kristensen (2010), and Chapter 13 in Jacod and Protter (2012). Specifically, we divide the sample into m_n nonoverlapping blocks, each of which contains k returns. Let $\mathcal{I}_{n,j} \equiv \{(j-1)k+1, \ldots, jk\}$ denote the collection of indices in the jth block, which spans the time interval $\mathcal{T}_{n,j} \equiv [(j-1)k\Delta_n jk\Delta_n).$[4] We stress that our theory treats the block size k as a *fixed* number. This contrasts with existing work on nonparametric spot volatility estimation, where the block size is invariably required to grow to infinity. Correspondingly, to distinguish our theory from the traditional "large-k" asymptotic theory, we refer to it as the "fixed-k" setting.

To obtain jump-robust volatility estimates, we adopt a truncation technique (as proposed by Mancini (2001)) based on a sequence of truncation threshold u_n satisfying $u_n \asymp \Delta_n^\varpi$ for some $\varpi \in (0, 1/2)$.[5] With this truncation in place the spot variance estimator for the jth block is simply defined as

$$\hat{c}_{n,j} \equiv \frac{1}{k\Delta_n} \sum_{i \in \mathcal{I}_{n,j}} (\Delta_i^n X)^2 \, 1_{\{|\Delta_i^n X| \leq u_n\}}. \tag{2.2}$$

This collection of blockwise estimators $(\hat{c}_{n,j})_{1 \leq j \leq m_n}$ also serve as the functional estimator for the entire process $(c_t)_{t \in [0,T]}$. To emphasize this point, we further identify $(\hat{c}_{n,j})_{1 \leq j \leq m_n}$ with a t-indexed

---

[4]For ease of exposition, we assume that the blocks are equally sized. Our theory can be easily extended to allow for different block sizes.

[5]In practice, one may take $\varpi = 0.49$ and $u_n = C\bar{\sigma}\Delta_n^\varpi$ for some constant $C > 0$ and a preliminary estimator $\bar{\sigma}$ of the daily average volatility. For example, when $C = 5$, u_n may be roughly interpreted as a 5-standard-deviation rule. A popular choice of $\bar{\sigma}$ is the bipower variation estimator of Barndorff-Nielsen and Shephard (2004). If the price process X is assumed to be continuous, the truncation is not needed (i.e., $u_n = \infty$).

functional estimator $(\hat{c}_{n,t})_{t\in[0,T]}$ by defining

$$\hat{c}_{n,t} \equiv \hat{c}_{n,j},$$

for $t \in \mathcal{T}_{n,j}$ and $j \in \{1, \ldots, m_n\}$. We will use this notation interchangeably in the sequel.

2.2 Fixed-k pointwise inference

In this subsection, we describe the fixed-k pointwise inference theory for the spot variance c_t. This part of the theory concerns a finite collection of estimation blocks, and it is developed under very mild regularity conditions collected in the following assumption.

Assumption 1. *Suppose that X has the form (2.1) and there exists a sequence $(T_m)_{m\geq 1}$ of stopping times increasing to infinity and a sequence $(K_m)_{m\geq 1}$ of constants such that the following conditions hold for each $m \geq 1$: (i) for some constant $r \in [0, 2)$, $|b_t|+|\sigma_t|+|\sigma_t|^{-1}+\int(|x|^r\wedge 1)F_t(dx) \leq K_m$ for all $t \in [0, T_m]$, where F_t denotes the spot Lévy measure of J; (ii) $\mathbb{E}[|\sigma_{t\wedge T_m} -\sigma_{s\wedge T_m}|^2] \leq K_m|t-s|$ for all $t, s \in [0, T]$.*

Assumption 1 entails some very mild regularity conditions, allowing for essentially unrestricted price and volatility jumps, leverage effect, and intraday periodicity. Condition (i), in particular, imposes local boundedness on various processes, while condition (ii) states that the volatility process is locally $(1/2)$-Hölder continuous under the L_2 norm. The latter condition can be readily verified if σ is an Itô semimartingale, or a long-memory process driven by a fractional Brownian motion (see, e.g., Comte and Renault (1998)).

Theorem 1, below, describes the finite-dimensional asymptotic property of the spot variance estimator $\hat{c}_{n,t}$ under the fixed-k asymptotic setting. To simplify notation, we use $\bar{\chi}_k^2$ to denote the distribution of Z_k/k, where Z_k is a generic chi-squared random variable with k degrees of freedom (i.e., $Z_k \sim \chi_k^2$), and use $\bar{\chi}_k^{-2}$ to denote the distribution of k/Z_k. We will refer to $\bar{\chi}_k^2$ and $\bar{\chi}_k^{-2}$ as the scaled chi-squared and scaled inverse-chi-squared distributions, respectively.

Theorem 1. *Suppose that Assumption 1 holds. Then, for any finite subset $\mathcal{M} \subseteq \{1, \ldots, m_n\}$, there exists a collection of independent $\bar{\chi}_k^2$-distributed random variables $(\bar{S}_j)_{j\in\mathcal{M}}$ such that for any $j \in \mathcal{M}$ and $t \in \mathcal{T}_{n,j}$,*

$$\frac{\hat{c}_{n,t}}{c_t} - \bar{S}_j = O_p\left(\Delta_n^{(2-r)\varpi\wedge(1/2)}\right) = o_p(1). \tag{2.3}$$

COMMENT. Theorem 1 shows that $\hat{c}_{n,t}/c_t$ can be approximated by a $\bar{\chi}_k^2$-distributed random variable, given by $(k\Delta_n)^{-1}\sum_{i\in\mathcal{I}_{n,j}}(W_{i\Delta_n} - W_{(i-1)\Delta_n})^2$. As a consequence,

$$\frac{\hat{c}_{n,t}}{c_t} \xrightarrow{d} \bar{\chi}_k^2. \tag{2.4}$$

7

With k fixed, the estimation error in $\hat{c}_{n,t}$ does not vanish when $\Delta_n \to 0$, and hence, $\hat{c}_{n,t}$ is not a consistent estimator for c_t. This is very different from the standard "large-k" theory, under which $\hat{c}_{n,t} \xrightarrow{\mathbb{P}} c_t$ (see, e.g., Theorem 9.3.2 in Jacod and Protter (2012)). $\qquad\square$

Although the spot variance estimator $\hat{c}_{n,t}$ is no longer consistent under the fixed-k asymptotics, we can nevertheless use Theorem 1 to guide the conduct of valid feasible inference. In particular, for any $\alpha \in (0,1)$, pick an interval $[\bar{L}_\alpha, \bar{U}_\alpha]$ such that $\bar{\chi}_k^2([\bar{L}_\alpha, \bar{U}_\alpha]) = 1 - \alpha$. It follows then readily by (2.4) that $\mathbb{P}(\bar{L}_\alpha \leq \hat{c}_{n,t}/c_t \leq \bar{U}_\alpha) \to 1 - \alpha$, or equivalently, $\mathbb{P}(\bar{U}_\alpha^{-1}\hat{c}_{n,t} \leq c_t \leq \bar{L}_\alpha^{-1}\hat{c}_{n,t}) \to 1 - \alpha$. In other words, $[\bar{U}_\alpha^{-1}\hat{c}_{n,t}, \bar{L}_\alpha^{-1}\hat{c}_{n,t}]$ provides a confidence interval (CI) for c_t with asymptotic level $1 - \alpha$. Note that since the scaled chi-squared distribution $\bar{\chi}_k^2$ is not symmetric, a conventional symmetric CI centered at $\hat{c}_{n,t}$ will not have the smallest possible length (i.e., minimal $\bar{L}_\alpha^{-1} - \bar{U}_\alpha^{-1}$). To deduce the shortest CI, it is convenient to reparameterize $\tilde{L}_\alpha = \bar{U}_\alpha^{-1}$ and $\tilde{U}_\alpha = \bar{L}_\alpha^{-1}$, so that the CI can be rewritten as $[\tilde{L}_\alpha \hat{c}_{n,t}, \tilde{U}_\alpha \hat{c}_{n,t}]$. Then, since the $\bar{\chi}_k^2([\bar{L}_\alpha, \bar{U}_\alpha]) = 1 - \alpha$ coverage constraint is equivalent to $\bar{\chi}_k^{-2}([\tilde{L}_\alpha, \tilde{U}_\alpha]) = 1 - \alpha$, it is easy to see that the length of $[\tilde{L}_\alpha \hat{c}_{n,t}, \tilde{U}_\alpha \hat{c}_{n,t}]$ is minimized by taking $[\tilde{L}_\alpha, \tilde{U}_\alpha]$ as the highest density (HD) interval of the scaled inverse-chi-squared distribution $\bar{\chi}_k^{-2}$. We will denote this as $[\tilde{L}_\alpha^{HD}, \tilde{U}_\alpha^{HD}]$ below. Corollary 1 summarizes these results.

Corollary 1. *Let $[\tilde{L}_\alpha, \tilde{U}_\alpha]$ be an interval that assigns probability $1 - \alpha$ under the $\bar{\chi}_k^{-2}$ distribution for some $\alpha \in (0,1)$. Under the conditions of Theorem 1, we have*

$$\mathbb{P}(\tilde{L}_\alpha \hat{c}_{n,t} \leq c_t \leq \tilde{U}_\alpha \hat{c}_{n,t}) \to 1 - \alpha.$$

Moreover, the length of $[\tilde{L}_\alpha, \tilde{U}_\alpha]$ is minimized by $[\tilde{L}_\alpha^{HD}, \tilde{U}_\alpha^{HD}]$.

COMMENT. An interesting special case is obtained by setting $\alpha = 0.5$, $\tilde{L}_\alpha = -\infty$, and $\tilde{U}_\alpha$ to be the median of the $\bar{\chi}_k^{-2}$ distribution. In this case, we have $\mathbb{P}(c_t \leq \tilde{U}_\alpha \hat{c}_{n,t}) \to 0.5$, suggesting that $\tilde{U}_\alpha \hat{c}_{n,t}$ is an asymptotically median unbiased estimator of c_t.[6] $\qquad\square$

By comparison, a conventional Gaussian-based CI, as commonly employed in the literature, relies on the following central limit theorem

$$\frac{\sqrt{k}(\hat{c}_{n,t} - c_t)}{\sqrt{2}\hat{c}_{n,t}} \xrightarrow{d} \mathcal{N}(0,1), \tag{2.5}$$

where it is assumed that $k \to \infty$ and $k\Delta_n^{1/2} \to 0$ (see Theorem 2 in Foster and Nelson (1996) for an early contribution and Theorem 13.3.3 in Jacod and Protter (2012) for a more general result). The

---

[6]The median of $\bar{\chi}_k^{-2}$ is approximately $(1 - \frac{2}{9k})^{-3}$, with the relative approximation error bounded by 0.25% over $k \in \{5, \ldots, 200\}$. We are grateful to an anonymous referee for suggesting this asymptotically median unbiased estimator.

resulting CI takes the form $[L^*_\alpha \hat{c}_{n,t}, U^*_\alpha \hat{c}_{n,t}]$ with $L^*_\alpha \equiv 1 - \sqrt{2/k}z_{\alpha/2}$, $U^*_\alpha \equiv 1 + \sqrt{2/k}z_{\alpha/2}$, and $z_{\alpha/2}$ denoting the $1 - \alpha/2$ quantile of the standard normal distribution. A commonly used approach to improve the conventional asymptotic Gaussian approximation is to consider the log-transformed variance. By applying the delta method to (2.5), we obtain

$$\sqrt{k/2}(\log(\hat{c}_{n,t}) - \log(c_t)) \xrightarrow{d} \mathcal{N}(0,1). \tag{2.6}$$

Correspondingly, under the "large-k" asymptotic setting, the $1 - \alpha$ level two-sided symmetric CI for $\log(c_t)$ is given by

$$[\log(\hat{c}_{n,t}) - \sqrt{2/k}z_{\alpha/2}, \log(\hat{c}_{n,t}) + \sqrt{2/k}z_{\alpha/2}]. \tag{2.7}$$

The results in Theorem 1 combined with the continuous mapping theorem also speak directly to the inference about the log spot variance.[7] In particular, it follows readily that for each $j \in \mathcal{M}$ and $t_j \in \mathcal{T}_{n,j}$,

$$\left(\log(\hat{c}_{n,j}) - \log(c_{t_j})\right)_{j \in \mathcal{M}} \xrightarrow{d} \left(\log(\bar{S}_j)\right)_{j \in \mathcal{M}},$$

which further implies

$$\max_{j \in \mathcal{M}} \left|\log(\hat{c}_{n,j}) - \log(c_{t_j})\right| \xrightarrow{d} \max_{j \in \mathcal{M}} \left|\log(\bar{S}_j)\right|. \tag{2.8}$$

For $\alpha \in (0,1)$, let $\bar{z}_\alpha$ denote the $1 - \alpha$ quantile of $\max_{j \in \mathcal{M}} \left|\log(\bar{S}_j)\right|$ so that

$$\mathbb{P}\left(\max_{j \in \mathcal{M}} \left|\log(\hat{c}_{n,j}) - \log(c_{t_j})\right| \leq \bar{z}_\alpha\right) \longrightarrow 1 - \alpha.$$

A uniform $1 - \alpha$ level confidence band for $\log(c_{t_j})$ across $j \in \mathcal{M}$ may therefore be constructed as

$$[\log(\hat{c}_{n,j}) - \bar{z}_\alpha, \log(\hat{c}_{n,j}) + \bar{z}_\alpha], \quad j \in \mathcal{M}.$$

We summarize the property of this confidence band in the following corollary.

Corollary 2. *Under the same setting as Theorem 1, we have*

$$\mathbb{P}(\log(\hat{c}_{n,j}) - \bar{z}_\alpha \leq \log(c_{t_j}) \leq \log(\hat{c}_{n,j}) + \bar{z}_\alpha \text{ for all } j \in \mathcal{M}) \rightarrow 1 - \alpha,$$

for $t_j \in \mathcal{T}_{n,j}$ and each $j \in \mathcal{M}$.

---

[7] Since the spot variance estimator is not consistent under the fixed-k setting, the delta method is not applicable because it relies on a linear approximation in a shrinking neighborhood at the true value. Nevertheless, we can still use the continuous mapping theorem to obtain the convergence result for certain nonlinear transformations (such as the logarithm).

COMMENTS. (i) The symmetry of the $[\log{(\hat{c}_{n,j})} - \bar{z}_\alpha, \log{(\hat{c}_{n,j})} + \bar{z}_\alpha]_{j \in \mathcal{M}}$ confidence band, combined with the constant width for all time points, makes this band particularly appealing for empirical work. These features stem from the fact that volatility is a scaling factor and the log transformation effectively "descales" the process.

(ii) Note that when $\mathcal{M}$ is a singleton, the $\bar{z}_\alpha$ quantile can be solved numerically from

$$F_k\left(ke^{\bar{z}_\alpha}\right) - F_k\left(ke^{-\bar{z}_\alpha}\right) = \mathbb{P}(|\log{(\bar{S}_j)}| \le \bar{z}_\alpha) = 1 - \alpha, \qquad (2.9)$$

where $F_k(\cdot)$ denotes the χ_k^2 cumulative distribution function. More generally, if $\mathcal{M}$ contains m elements, it follows by the independence of the $\bar{S}_j$ random variables that $\mathbb{P}(\max_{j \in \mathcal{M}} |\log(\bar{S}_j)| \le \bar{z}_\alpha) = \mathbb{P}(F_k\left(ke^{\bar{z}_\alpha}\right) - F_k\left(ke^{-\bar{z}_\alpha}\right))^m$, so that $\bar{z}_\alpha$ can be solved from $F_k\left(ke^{\bar{z}_\alpha}\right) - F_k\left(ke^{-\bar{z}_\alpha}\right) = (1 - \alpha)^{1/m}$. The joint $1 - \alpha$ level critical value is thus the same as the marginal critical value at confidence level $(1 - \alpha)^{1/m}$. $\qquad \square$

To more clearly illustrate the differences between the fixed-k CIs described in Corollary 1 and the conventional Gaussian-based CIs, the left panel of Figure 1 depicts $[\tilde{L}_\alpha^{HD}, \tilde{U}_\alpha^{HD}]$ and $[L_\alpha^*, U_\alpha^*]$ for a range of block sizes $k \in \{5, \ldots, 30\}$ at confidence level 90%.[8] As the figure shows, not only are the fixed-k CIs asymmetric, they are also generally wider than their conventional Gaussian benchmarks, with the differences especially pronounced for small block sizes. As demonstrated in our Monte Carlo simulation experiments discussed in Section 3 below, these differences manifest in markedly different finite sample coverage probabilities. Similarly, the right panel of Figure 1 compares the critical values for the log variance: $[-\bar{z}_\alpha, \bar{z}_\alpha]$ for the fixed-k setting (recall (2.9)) and $[-\sqrt{2/k}z_{\alpha/2}, \sqrt{2/k}z_{\alpha/2}]$ for the conventional Gaussian setting (recall (2.7)). These critical values for the log variance are numerically closer than those for the variance itself. Needless to say, the apparent closeness in the former case does not automatically imply that the corresponding CIs exhibit better finite-sample coverage probabilities, in that the finite-sample distribution of $\log{(\hat{c}_{n,t})}$ and $\hat{c}_{n,t}$ are also very different. That said, in our Monte Carlo simulation experiment, we find that the log-transformation does indeed help mitigate the size distortions of the Gaussian-based CIs, but that the distortions remain larger than for the fixed-k based CIs.

A limitation of Corollary 2, which it inherits from Theorem 1, is that it concerns a *fixed* number of blocks, resulting in a subtlety regarding the scope of its applicability. In particular, directly applying Corollary 2 to the full sample period $[0, T]$ would require treating the total number of blocks m_n as a fixed number, but this would lead to a logical inconsistency as $m_n = n/k \to \infty$

---

[8]Here, and throughout the remainder of the paper, we focus on 90% CIs. The results for other confidence levels are qualitatively very similar, and hence, are omitted for brevity.
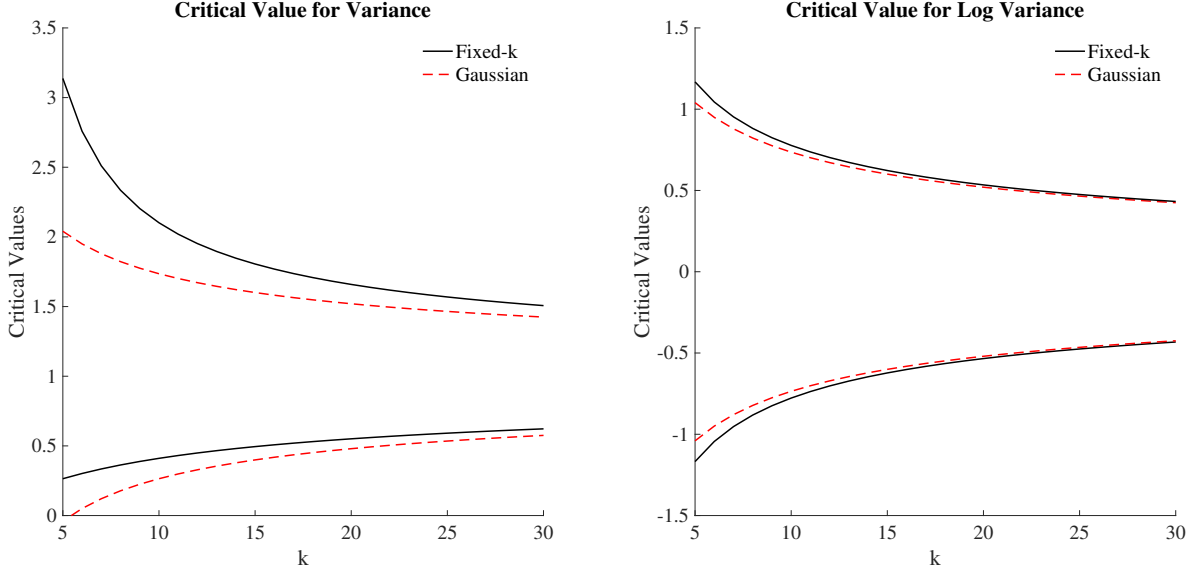
Figure 1: The figure plots critical values of fixed-k and Gaussian-based confidence intervals for the spot variance c_t (left) and its log transformation (right). The confidence level is fixed at $1 - \alpha = 90\%$. The block size k ranges from 5 to 30. On the left panel, the solid lines correspond to the fixed-k critical values $(\tilde{L}_\alpha^{HD}, \tilde{U}_\alpha^{HD})$ defined by the $1 - \alpha$ level highest density interval of the scaled inverse-chi-squared distribution $\bar{\chi}_k^{-2}$, and the dashed lines correspond to the benchmark Gaussian-based critical values $1 \pm \sqrt{2/k} z_{\alpha/2}$, where $z_{\alpha/2}$ denotes the $1 - \alpha/2$ quantile of the standard normal distribution. On the right panel, the solid lines correspond to the fixed-k critical values $\pm \bar{z}_\alpha$, where $\bar{z}_\alpha$ denotes the $1 - \alpha$ quantile of the $|\log \bar{\chi}_k^2|$ distribution, and the dashed lines correspond to the Gaussian-based critical values $\pm\sqrt{2/k} z_{\alpha/2}$ shown in equation (2.7).

under the maintained asymptotic framework.[9] The joint inference for a growing number of blocks is a much more challenging problem, to which we now turn.

2.3   Fixed-k inference in the high-dimensional case

This subsection is devoted to the uniform inference for the entire spot variance process $(c_t)_{t \in [0,T]}$ under the fixed-k asymptotic setting. Given the appealing "descaling" feature of the log transformation highlighted by Corollary 2, we focus exclusively on the log variance process $\log(c_t)$. The key technical challenge stems from the growing number of blocks in the sample span $[0, T]$. In particular, with k fixed, the number of blocks $m_n = n/k$ grows at the same rate as the sample size n, rendering the joint inference a high-dimensional problem. To address this challenge, we need to

---

[9]To put this into perspective, consider a representative empirical scenario with $\Delta_n = 1$-minute and a sample span $T = 1$ corresponding to a 6.5-hour trading day. For a fixed block size of say $k = 10$, there would be "only" $m_n = 39$ blocks.

strengthen Assumption 1 as follows.

Assumption 2. *Suppose that Assumption 1 with $r = 0$ holds true. Moreover, for any $p \geq 2$, assume that*

$$\max_{1 \leq j \leq m_n} \mathbb{E}\left[\sup_{t,s \in \mathcal{T}_{n,j}} |\sigma_{t \wedge T_m} - \sigma_{s \wedge T_m}|^p\right] \leq K_{m,p} |t - s|^{p/2}, \tag{2.10}$$

for all $m \geq 1$ and some finite constant $K_{m,p}$.

Assumption 2 strengthens Assumption 1 in two important ways. First, the restriction that $r = 0$ implies that the jumps in J have finite activity, whereas the activity of price jumps is essentially unrestricted in Assumption 1. This finite-activity condition is needed to ensure that the truncation is uniformly valid across the large number of $m_n = O(n)$ blocks.[10] Second, the condition in (2.10) further implies that the paths of the volatility process are Hölder continuous (with an index strictly smaller than $1/2$) on each subinterval $\mathcal{T}_{n,j}$. This condition holds if the volatility process behaves as a continuous Itô semimartingale, or a long-memory process within each subinterval. Importantly, it does not rule out jumps on the boundary time points between the $\mathcal{T}_{n,j}$ subintervals.[11] While this piecewise-continuity assumption restricts the form of volatility jumps, it readily accommodates volatility jumps induced by regularly scheduled macroeconomic and other precisely timed news announcements, for which it is natural to use the known announcement times to divide nearby estimation blocks.

Theorem 2, below, describes the uniform confidence band for the log variance process and its asymptotic validity.

Theorem 2. *Suppose that Assumption 2 holds and $k \geq 7$. Let $(\bar{S}_j)_{1 \leq j \leq k_n}$ be i.i.d. $\bar{\chi}_k^2$-distributed random variables, and $\bar{z}_{n,\alpha}$ be the $1 - \alpha$ quantile of $\sup_{1 \leq j \leq m_n} |\log(\bar{S}_j)|$. Then,*

$$\mathbb{P}\left(\sup_{t \in [0,T]} |\log(\hat{c}_{n,t}) - \log(c_t)| \leq \bar{z}_{n,\alpha}\right) \to 1 - \alpha.$$

Consequently, $[\log(\hat{c}_{n,t}) - \bar{z}_{n,\alpha}, \log(\hat{c}_{n,t}) + \bar{z}_{n,\alpha}]$ forms a uniform confidence band for the $\log(c_t)$ process on $[0,T]$ with asymptotic level $1 - \alpha$.

COMMENTS. (i) The confidence band described in Theorem 2 is constructed in exactly the same manner as the band in Corollary 2. Thus, the proposed confidence band is valid regardless of

---

[10]It might be possible to allow for infinitely active jumps if one sufficiently slows down the growth rate of the number of blocks in the joint inference problem. This extension, however, would be secondary to our main contribution, and hence, is not developed here.

[11]A similar piecewise continuity condition is also needed in the conventional "large-k" setting for the uniform estimation of the volatility process; see Remark 1 of Li, Todorov, and Tauchen (2017a).

whether the number of blocks is fixed or divergent, although in the latter case we need somewhat stronger assumptions to establish the validity.

(ii) Theorem 2 requires the block size $k \geq 7$, which is not needed for the pointwise inference discussed in Section 2.2. It is instructive to illustrate why this (seemingly peculiar) condition is needed. It follows from Theorem 1 that $\hat{c}_{n,t}/c_t \approx \bar{S}_j$ for each $t \in \mathcal{T}_{n,j}$. In order to get a similar approximation for $\log(\hat{c}_{n,t}/c_t)$, we can apply a Taylor expansion to show that

$$\log\left(\frac{\hat{c}_{n,t}}{c_t}\right) - \log\left(\bar{S}_j\right) \approx \frac{1}{\bar{S}_j}\left(\frac{\hat{c}_{n,t}}{c_t} - \bar{S}_j\right).$$

Although each $1/\bar{S}_j$ variable is $O_p(1)$, their maximum $\max_{1 \leq j \leq m_n}\left|1/\bar{S}_j\right|$ is divergent, and we need to "tame" the growth rate. Recall that the density of the scaled chi-squared distribution $\bar{\chi}_k^2$ behaves like an $O(x^{k/2-1})$ function when $x \approx 0$. Hence, a larger k results in a smaller probability that $\bar{S}_j$ is near zero. It turns out that when $k \geq 7$ this probability is sufficiently small to ensure that $\max_{1 \leq j \leq m_n}\left|1/\bar{S}_j\right|$ diverges slower than the uniform rate at which the $\hat{c}_{n,t}/c_t \approx \bar{S}_j$ approximation occurs, which in turn ensures that $\log\left(\hat{c}_{n,t}/c_t\right) \approx \log\left(\bar{S}_j\right)$ uniformly well. $\qquad\square$

To underscore the practical applicability of the new fixed-k theory and related inference procedures, we turn next to the results from an empirically realistically calibrated Monte Carlo simulation experiment.

3   Monte Carlo simulations

Following Bollerslev and Todorov (2011), we simulate the (log) price process from a two-factor stochastic volatility model. Specifically, with the unit time interval normalized to "one day," we generate the process X according to

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

where W, B_1, and B_2 denote independent standard Brownian motions.[12] The $\rho = -0.7$ parameter captures the well-documented negative correlation between price and volatility shocks (i.e., the "leverage" effect). The V_1 volatility factor is highly persistent with a half-life of 2.5 months,

---

[12]In Supplemental Appendix SA we consider an additional data generating process, in which the Brownian motions B_1 and B_2 are replaced by Lévy processes with infinitely active jumps. The simulation results are very similar to the findings discussed below.
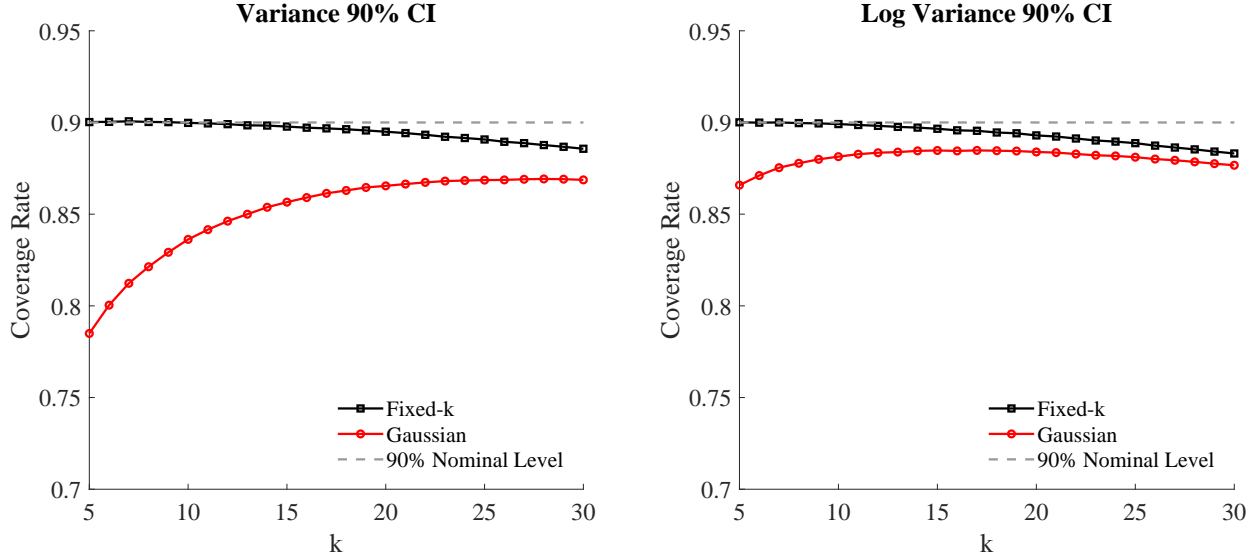
Figure 2: The figure plots the simulation coverage probabilities for the optimal fixed-k 90% CIs for the spot variance described in Corollary 1, the symmetric fixed-k 90% CIs for the log spot variance described in Corollary 2, together with the conventional large-k 90% Gaussian-based CIs. The block size k ranges from 5 to 30.

while the V_2 volatility factor is quickly mean-reverting with a half-life of only one day. We fix $V_{1,0} = V_{2,0} = 0.5$, so that $c_0 = 1$. We simulate the "continuous-time processes" using an Euler scheme on a 1-second mesh, with the "observed returns" actually used in the calculations sampled at $\Delta_n = 1$ minute intervals. All numerical results reported below are based on 100,000 Monte Carlo replications.

We begin by examining the performance of the fixed-k CIs for the spot variance c_t and its log transformation $\log(c_t)$. We consider the optimal CI for the spot variance described in Corollary 1, and the symmetric CI for the log spot variance described in Corollary 2. For comparison, we also include results for the conventional "large-k" CIs formed using the asymptotic Gaussian approximations in (2.5) and (2.6). We assess the coverage for the true latent volatility values at time $t = k\Delta_n$, corresponding to the endpoint of the first estimation block for block size k. We compute the CIs for a range of block sizes $k \in \{5, \ldots, 30\}$.

Figure 2 plots the resulting simulated coverage rates at nominal level $1 - \alpha = 90\%$. As the figure clearly shows, the coverage rates of the fixed-k CIs for both c_t and $\log(c_t)$ are notably closer to the nominal level than the Gaussian-based CIs. Quite remarkably, the fixed-k CIs have almost exactly 90% coverage when k is small (e.g., $k \leq 10$). In sharp contrast, the Gaussian-based CIs all suffer from nontrivial size distortions for these smaller block sizes, indicating that the asymptotic

Table 1: Monte Carlo Coverage Rates of Fixed-k Uniform Confidence Band

	Number of Blocks (m_n)					
	1	2	3	6	39	195
Corollary 2	0.896	0.897	0.898	0.897	0.898	0.851
Theorem 2	0.869	0.872	0.873	0.875	0.876	0.821

Note: The table reports the finite-sample coverage rates of the 90% uniform confidence band for the log spot variance $\log(c_t)$. The block size is fixed at $k = 10$. The two separate rows labeled Corollary 2 and Theorem 2 pertain to the uniform coverage for $(\log(c_{t(n,j)}))_{1 \leq j \leq m_n}$ and $(\log(c_t))_{t \in [0,T]}$, respectively.

Gaussian approximation that relies on $k \to \infty$ has not yet "kicked in." The log transformation alleviates some of the size distortion, although the coverage of the Gaussian-based CIs are still systematically below their nominal levels.[13]

We turn next to an analysis of the uniform confidence band for the log spot variance process. In view of the pointwise results discussed above, we fix the block size at $k = 10$, and focus our analysis on the effect of including additional blocks. Specifically, we consider the joint inference for $m_n \in \{1, 2, 3, 6, 39, 195\}$ blocks, corresponding to 10-minute, 20-minute, 30-minute, 1-hour, 1-day, and 1-week horizons, respectively. The results in Corollary 2 and Theorem 2 formally prescribe the same numerical confidence band,

$$\left[\log\left(\hat{c}_{n,t}\right) - \bar{z}_{n,\alpha}, \log\left(\hat{c}_{n,t}\right) + \bar{z}_{n,\alpha}\right], \quad \bar{z}_{n,\alpha} = 1 - \alpha \text{ quantile of } \max_{1 \leq j \leq m_n} \left|\log(\bar{S}_j)\right|.$$

However, the corollary makes the weaker assertion that the band covers the m_n-dimensional vector $\left(\log(c_{t_j})\right)_{1 \leq j \leq m_n}$ for $t_j \in \mathcal{T}_{n,j}$, whereas the theorem asserts uniform coverage for the entire process $(\log(c_t))_{t \in [0,T]}$. Accordingly, we report uniform coverage rates for two separate cases. The first case pertains to the setting of Corollary 2, for which we set t_j as the starting time point of the jth estimation block. The second case pertains to the setting of Theorem 2, for which we approximate the realized continuous-time sample path $(\log(c_t))_{t \in [0,T]}$ using the entire 1-second mesh from the Euler simulation scheme. As above, we focus on the simulated coverage rates at nominal level

[13]Gonçalves and Meddahi (2009) propose bootstrap inference methods for the integrated variance. Although their theory is not directly applicable in the present context concerning spot variances, motivated by their ideas we implement an alternative bootstrap-based version of the CIs. The results from these additional simulations, reported in Supplemental Appendix SA, show that the fixed-k CIs also control size better than the bootstrap CIs.

$1 - \alpha = 90\%$.

The results in the top row of Table 1 for the m_n-dimensional vector $\left(\log(c_{t_j})\right)_{1 \leq j \leq m_n}$ show coverage rates almost identical to the nominal level over daily and shorter time horizons (i.e., $m_n \leq 39$). However, at the longer 1-week horizon (i.e., $m_n = 195$) there is evidence of some under-coverage, reflecting the additional complications associated with this higher-dimensional scenario. For the more challenging case involving the entire process $(\log(c_t))_{t \in [0,T]}$, the simulated coverage rates reported in the second row are obviously lower, but still fairly close to the 90% nominal level.

Buttressed by these simulation results, we turn next to an empirical application of the new procedures.

4   FOMC policy announcements and asset market volatility

The ability of the new fixed-k asymptotic theory to reliably assess the spot volatility over relatively short estimation blocks renders the approach particularly appealing for situations in which the volatility process varies dramatically over short time-windows, possibly as a result of specific economic events. These types of empirical scenarios have also received considerable attention in the recent literature on high-frequency identification of macroeconomic shocks. Nakamura and Steinsson (2018), for example, rely on high-frequency bond futures data to identify monetary shocks and test for monetary non-neutrality via a simultaneous equation model identified by the assumption that asset price volatilities are higher during short "treatment" windows around FOMC announcements than during "control" windows void of any policy announcements. Bollerslev, Li, and Xue (2018) similarly rely on a regression-discontinuity, or jump-regression, strategy for estimating summary measures of investors' difference-of-opinion based on the premise that the volatilities of stock returns immediately following FOMC, and other macroeconomic, news announcements are "abnormally high" relative to their pre-announcement levels. Motivated by this recent literature and the identification schemes employed therein, we enlist the new inference procedures to more formally assess the short-term dynamics of aggregate stock market volatility around FOMC announcements.

There is already a large literature in macroeconomics studying the effect of monetary policy shocks, as identified through the impact of FOMC announcements on the level of asset prices, including early notable contributions by Cochrane and Piazzesi (2002), Rigobon and Sack (2004), and Bernanke and Kuttner (2005), and more recent work by Wright (2012) and Johnson and Paye (2019) specifically related to the Fed funds rate being at or near the zero lower bound. A closely related literature in asset pricing finance similarly seeks to associate large intraday asset price movements, in the form of statistically significant jumps, with specific economic news, like

FOMC announcements, as exemplified by Andersen, Bollerslev, and Diebold (2007), Lee and Mykland (2008), and Lee (2012). Meanwhile, other recent studies have documented that most of the equity risk premium appears to be earned in specific phases of the FOMC news release cycle; see, for example, Savor and Wilson (2014), Lucca and Moench (2015), and Cieslak, Morse, and Vissing-Jorgensen (2019). Related to this, Ai and Bansal (2018) exploit the unique environment created by FOMC announcements to help discriminating non-expected utility theory from standard Von Neumann-Morgenstern theory, and the role of uncertainty in determining asset prices in particular. It is not our intent to add to this burgeoning literature on the determinant of the equity risk premium, and the efficacy of monetary policy more generally. Instead, we simply recognize the unique position of FOMC announcements as the most important type of regularly scheduled macroeconomic news with an immediate impact on asset prices and asset market volatilities.

Our data consists of intraday transaction prices for the E-mini futures contract on the S&P 500 index, spanning July 1, 2003 to March 2, 2017. The data was obtained from Tick Data. To help mitigate the effect of market microstructure "noise," we follow standard practice in the literature to sparsely sample the data at a 1-minute sampling frequency (see, e.g., the discussion in Zhang, Mykland, and Aït-Sahalia (2005)). We rely on Bloomberg's Economic Calendar to pinpoint the exact announcement times for each of the 109 FOMC announcements that occurred during regular trading hours over our sample period.

Before considering the aggregate results for all of the FOMC announcements, it is instructive to detail the results for a few specific days. The most important episode in our sample vis-a-vis the role of the Fed is arguably the 2007–2008 financial crisis. We thus consider three FOMC announcements that occurred during the crisis, each of which is naturally associated with a clearcut narrative. The top-left panel in Figure 3, in particular, plots the time series of the S&P 500 E-mini futures prices over a 2-hour window on September 18, 2007. The window is centered at the exact time-of-day when the FOMC announced its decision to lower the target Fed funds rate from 5.25% to 4.75%, so as to forestall the adverse impact of the mortgage crisis on the U.S. economy. As the figure shows, this announcement resulted in a sharp price jump of more than one percent in the level of the S&P 500 index. Comparing the price paths in the one hour before and after the announcement, the FOMC announcement evidently also resulted in an increase in the intraday volatility. The top-middle panel, pertaining to the announcement on March 18, 2008, which occurred soon after the Fed guaranteed Bear Stearns' bad loans to facilitate its acquisition by J.P. Morgan Chase, tells a similar story in regards to the volatility. While the market index fluctuated within a fairly narrow price band in the one hour leading up to the announcement, the one hour following the announcement obviously witnessed much larger fluctuations. Moreover, although
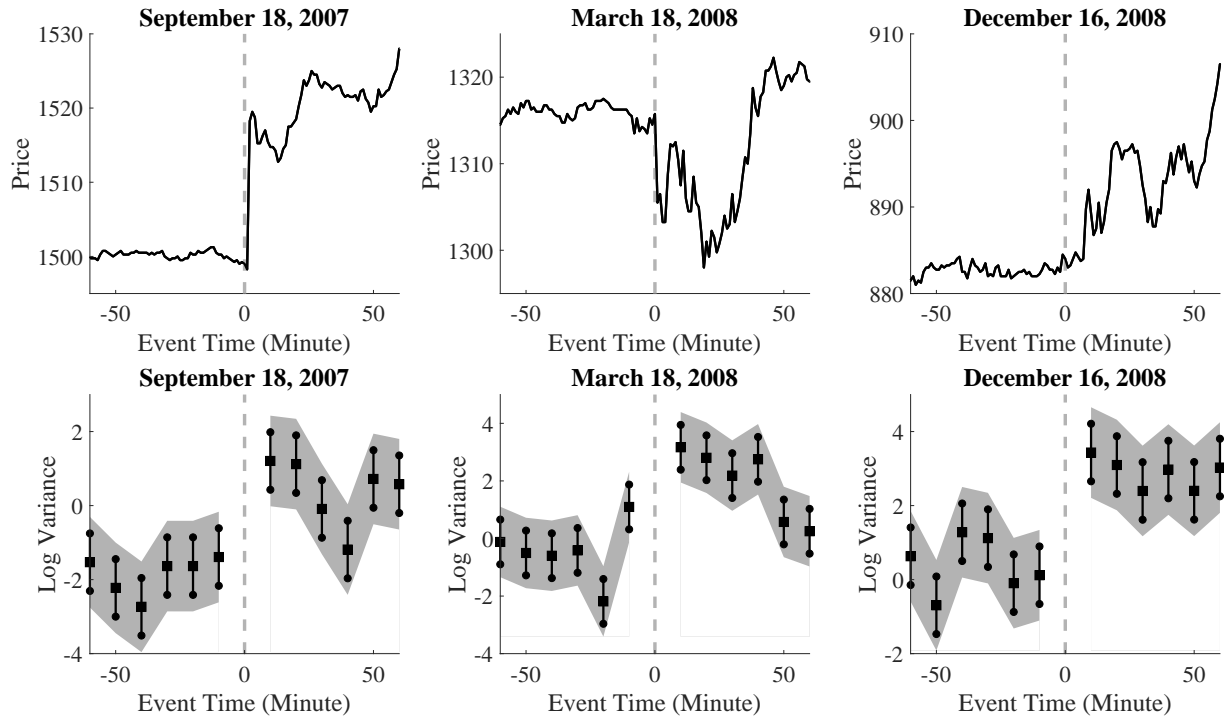
Figure 3: The top panel plots the S&P 500 E-mini futures prices one hour before and after selected FOMC announcements, with the announcement time normalized to zero. The bottom panel plots the corresponding log spot variance estimates formed on 10-minute estimation blocks (squares), together with 90% pointwise CIs (bars) and 90% uniform confidence bands (shaded areas).

this announcement triggered a negative price jump at the time of its release, this immediate price impact was reversed within the hour. Meanwhile, not all FOMC announcements necessarily result in immediate price jumps. As a case in point, the top-right panel shows the S&P 500 futures prices around the "zero-lower-bound" announcement on December 16, 2008. Nonetheless, even though the price itself did not jump at the time of that announcement, the volatility evidently increased quite sharply.

In an effort to more rigorously underpin these visual impressions, the bottom panels in Figure 3 display the log spot variance estimates computed over 10-minute blocks (i.e., $k = 10$) one hour before and after each of the three announcements. To prevent price jumps at announcement times from "contaminating" the spot volatility estimation, we remove the two (resp. three) 1-minute returns before (resp. after) the announcement time, resulting in a 5-minute gap between pre-announcement and post-announcement estimation blocks. In addition to the point estimates (squares), we also plot the corresponding 90% fixed-k pointwise CIs (bars), together with the 90% uniform confidence bands for the one hour before and one hour after the announcement times

18

(shaded areas).

Three interesting features emerge. Firstly, the lack of overlap in the CIs immediately before and after the announcements corroborate the visual impression gleaned from the price plots in the top panels, of significantly higher volatilities in the immediate aftermath of the announcements. Secondly, there is a tendency for the volatility to trend down over the post-announcement hour. However, it typically remains at elevated levels even after one hour of trading. Thirdly, for some days there is also evidence of a slight buildup in the volatility in advance of the announcement.[14] This pre-announcement "up-tick" may possibly be explained by informed trading steered by leaked information during the official press lock-up period.[15] Consistent with this, Bernile, Hu, and Tang (2016) and Kurov, Sancetta, Strasser, and Wolfe (2019) both report that some directional price-adjustments seemingly occurred in anticipation of the official FOMC news release.

We intentionally choose the three announcement days in Figure 3 because they all represent specific milestones in financial history. Nonetheless, the short-term volatility patterns that these specific events reveal are to a large extent representative of the full set of FOMC announcements. In particular, the finding that the announcements trigger positive volatility jumps, or $c_\tau > c_{\tau-}$, holds true almost uniformly across all FOMC announcements.

To formally test this hypothesis, we estimate spot volatilities immediately before and after all of the 109 announcements in our sample, say $\hat{c}_{1-}$ and $\hat{c}_{1+}$, respectively, based on 10-minute estimation blocks.[16] In parallel to the calculations above, we again exclude the two (resp. three) 1-minute returns immediately before (resp. after) the announcement as a conservative way to truncate announcement-induced price jumps. Referring to the distributional results in Theorem 1, it follows that under the null hypothesis that the spot volatilities are the same, the ratio $\hat{c}_{1+}/\hat{c}_{1-}$ converges in distribution to the ratio between two independent $\bar{\chi}_{10}^2$-distributed variables, that is,

---

[14]This is most notable for the announcement on March 18, 2008, and may also be gleaned from a closer look at the underlying price movements in Figure 3.

[15] The lock-up period refers to a short, often less than 30-minutes, period before the official announcement, when accredited journalists receive the information, thereby facilitating their preparation of media reports. However, under embargo agreements, the media cannot disclose the information until the scheduled release time. The FOMC explicitly expressed concerns about violations of this embargo in the transcript of its November 3, 2010 meeting. A report by the Office of Inspector General is available at https://oig.federalreserve.gov/reports/board-controls-sensitive-economic-information-apr2016.pdf.

[16]The choice of k reflects the usual "bias-variance" trade-off. Small values of k result in small biases and wide CIs, and vice versa. Our choice of $k = 10$ is mainly motivated by our simulation results summarized in Figure 2, which indicate that $k = 10$ provides excellent size control, while larger values of k may lead to size distortions. As an additional robustness check, we also repeated the empirical analysis with an even more conservative choice of $k = 5$. The results, detailed in Supplement Appendix SB.2, are qualitatively very similar to the ones based on $k = 10$ discussed here.

Table 2: 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	104	0	21	17	4	44

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) the announcements.

an $F_{(10,10)}$-distribution. Correspondingly, the $1 - \alpha$ quantile of this F-distribution may be used as the critical value to determine whether the $c_\tau = c_{\tau-}$ null hypothesis is rejected against the positive-jump alternative, or $c_\tau > c_{\tau-}$, at significance level α. Conversely, a test against the negative-jump alternative, or $c_\tau < c_{\tau-}$, may be carried out using the test statistic $\hat{c}_{1-}/\hat{c}_{1+}$ and the identical critical value.

The first two columns of Table 2 report the number of FOMC announcements, out of 109 in total, for which we reject the "no-jump" hypothesis against each of the two one-sided alternatives at the 10% significance level. For 104 of the days is the post-announcement spot variance estimate $\hat{c}_{1+}$ significantly higher than the pre-announcement estimate $\hat{c}_{1-}$, thus confirming that FOMC announcements generally induce positive volatility jumps.[17] Meanwhile, the tests never reject the null hypothesis against the negative-jump alternative, ruling out the possibility of downward volatility jumps, and a resolution of uncertainty, immediately following FOMC announcements. This perhaps is not surprising, as the new "lumpy" information conveyed by the FOMC statements may require additional interpretation by investors, even if the policy announcements may eventually help resolve uncertainty in the long-run.[18]

Further investigating the specificities of the patterns observed in Figure 3, we consider the volatilities calculated over the 10-minute estimation blocks immediately preceding (resp. following) the 10-minute blocks used in the estimation of $\hat{c}_{1-}$ (resp. $\hat{c}_{1+}$). Denoting these estimates by $\hat{c}_{2-}$ and

---

[17]The use of more stringent 5% and 1% significance levels slightly reduce the number of rejections to 102 and 95, respectively, but does not alter our main empirical message.

[18]Although the evidence is not as overwhelming, in Supplemental Appendix SB.1 we also document a tendency for upward volatility movements during post-announcement press conferences.
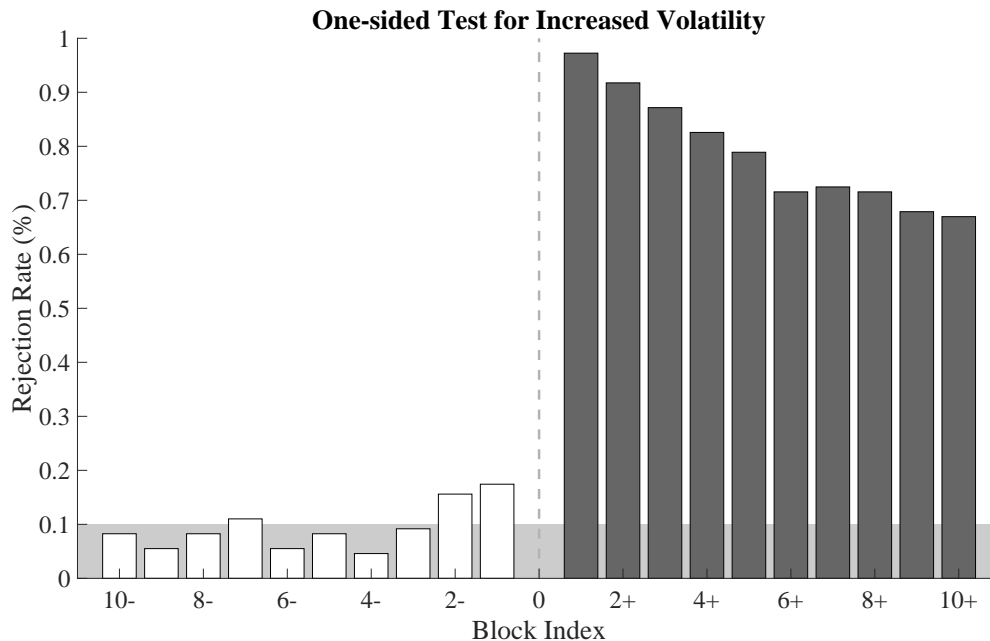
Figure 4: The figure reports the proportions of FOMC announcements for which the spot variance in each of the ten 10-minute estimation blocks before and after the announcement is significantly higher, at the 10% significance level, than the spot variance in the 11th 10-minute benchmark block before the announcement. The light-shaded area indicates 10%.

$\hat{c}_{2+}$, respectively, we test for a general buildup in the volatility in advance of FOMC announcements by comparing $\hat{c}_{2-}$ and $\hat{c}_{1-}$, and a decline in the volatility following the announcements by comparing $\hat{c}_{1+}$ and $\hat{c}_{2+}$. Column 3 (resp. 4) in Table 2 shows the number of instances in which $\hat{c}_{1-}$ is significantly higher (resp. lower) than $\hat{c}_{2-}$ at the 10% significance level. No obvious picture emerges: there is a significant pre-announcement increase in the volatility in 21 instances, while the volatility exhibits a significant decline before the announcement in 17 instances. Thus, the previously discussed pre-announcement "up-tick" on March 18, 2008 is by no means unique, but it is also not a widely shared phenomenon. Looking at the last two columns reveals a more systematic post-announcement pattern, with $\hat{c}_{2+}$ being significantly lower than $\hat{c}_{1+}$ for close to half of the days, suggesting that the volatility fairly quickly mean-reverts after the initial burst triggered by the announcement.

To gain additional insight on the strength of this mean-reversion, we consider a longer horizon comprised of the ten 10-minute estimation blocks before and after the announcement time. We refer to these estimates as $\hat{c}_{j-}$ and $\hat{c}_{j+}$, respectively, with $j \in \{1, \ldots, 10\}$ signifying the distance in 10-minute multiples from the announcement time. Since FOMC announcements typically take place at 2:15 PM (EST), the ten post-announcement blocks essentially span the remaining part

21

of the regular trading day. We use the 11th estimation block before the announcement, and the corresponding spot variance estimate $\hat{c}_{11-}$, as our benchmark, and perform one-sided tests at the 10% significance level to examine whether $\hat{c}_{j\pm} > \hat{c}_{11-}$ significantly for each $j \in \{1, \ldots, 10\}$.

Figure 4 summarizes the results by reporting the rejection rates for each of the twenty estimation blocks averaged across all of the 109 announcements in the sample. Looking at the pre-announcement window, the rejection rates for the first eight blocks are either below or very close to the 10% nominal level, indicating the lack of any systematic increase in the volatility up until twenty minutes before the announcement. However, there is a notable increase in the rejections for the two estimation blocks in the 20-minute window before the announcement time, echoing the previous findings in Table 2 for $\hat{c}_{1-} > \hat{c}_{2-}$, indicative of occasional pre-announcement "up-ticks" in the volatility. Meanwhile, the rejection rate for the first post-announcement block is close to 100%, providing unambiguous support for the notion that FOMC-induced monetary policy shocks result in heightened asset price volatility. The subsequent estimation blocks evidence a gradual decline in the average rejection rates the further apart in time the blocks are from the announcement time. In spite of this decline, the volatility usually remains at an elevated level in the trading hours after the announcement compared to the benchmark before the announcement.

In sum, not only do FOMC monetary policy announcements typically cause asset prices to jump, they almost always trigger highly statistically significant positive jumps in asset market volatility. This increase in the volatility is gradually reversed over the remainder of the trading day. In addition, there is also a significant buildup in the volatility shortly in advance of certain announcements. These "fine structures" are only formally revealed through the use of very short estimation blocks and accompanying reliable inference procedures, as afforded by our new fixed-k asymptotic theory. By contrast, conventional asymptotic Gaussian-based inference for spot volatility on this time scale is subject to nontrivial size distortions, and hence, would be difficult to interpret.

5 Conclusion

Conventional theory for nonparametric inference about spot volatility invariably relies on a "large-k" asymptotic scheme, under which local estimation windows are assumed to contain an increasing number of observations, while at the same time the length of the windows shrink to zero asymptotically. By "properly" choosing the block size of the windows to be "large enough but not too large," the asymptotic theory in turn delivers nonparametric consistent volatility estimates together with asymptotic Gaussian-based inference. However, the practical implementation of this conventional

double asymptotic framework is marred with difficulties.

Instead, we propose a novel fixed-k asymptotic theory and corresponding single asymptotic framework that is much more amenable to practical implementation. Our key insight is that even though the estimates themselves are no longer consistent, it is still possible to conduct valid nonparametric inference about the true latent spot volatilities. In order to do so, we derive easy-to-implement pointwise CIs determined by scaled chi-squared distributions. In contrast to the standard Gaussian-based CIs, which suffer from severe size distortions for small estimation block sizes, the new fixed-k CIs attain almost exact coverage. To allow for uniform inference about the entire volatility process, we further extend the fixed-k theory to a high-dimensional setting, in which the number of estimation blocks grows at the same rate as the sample size. This high-dimensional theory is also new to the literature on volatility estimation and inference.

We demonstrate the practical usefulness of the proposed new methods by studying the volatility of the S&P 500 aggregate equity index around FOMC news announcements. Our empirical results provide unambiguous support for positive volatility jumps at announcement times, thus formally corroborating the high-frequency identification-by-discontinuity strategies used in a number of recent studies. The new procedures also reveal other more subtle pre- and post-announcement drifts in the volatility, which may help shed new light on the way in which investors position themselves in advance of the announcements and process the information revealed by the announcements.

<center>APPENDIX: PROOFS</center>

Throughout the proofs, we use K to denote a positive constant that may change from line to line, and write K_p to emphasize its dependence on some parameter p. For $p \geq 1$, we use $\| \cdot \|_p$ to denote the L_p-norm of a random variable. For each j, we use $t(n, j)$ to denote the starting point of the $\mathcal{T}_{n,j}$ time interval, that is, $t(n, j) \equiv (j - 1)k\Delta_n$. In addition, by a standard localization procedure, we can strengthen Assumptions 1 and 2 by assuming that they hold with $T_1 = \infty$ without loss of generality; see Section 4.4.1 in Jacod and Protter (2012) for details on the localization procedure.

PROOF OF THEOREM 1. Denote $X' \equiv X - J$, that is, the continuous part of X. For each $j \in \{1, \ldots, m_n\}$, define $\hat{c}'_{n,j} \equiv (k\Delta_n)^{-1} \sum_{i \in \mathcal{I}_{n,j}} (\Delta_i^n X')^2$ and $S_{n,j} \equiv \Delta_n^{-1} \sum_{i \in \mathcal{I}_{n,j}} (\Delta_i^n W)^2$. We then decompose for each $t \in \mathcal{T}_{n,j}$,

$$k\hat{c}_{n,j} - c_t S_{n,j} = \left(c_{t(n,j)} - c_t\right) S_{n,j} + R_{n,j} + R'_{n,j}, \tag{A.1}$$

<center>23</center>

where $R_{n,j} \equiv k(\hat{c}_{n,j} - \hat{c}'_{n,j})$ and $R'_{n,j} \equiv k\hat{c}'_{n,j} - c_{t(n,j)}S_{n,j}$.

We rewrite $R_{n,j} = \sum_{i \in \mathcal{I}_{n,j}} \zeta_{n,i}$ where $\zeta_{n,i} \equiv \Delta_n^{-1}((\Delta_i^n X)^2 1_{\{|\Delta_i^n X| \leq u_n\}} - (\Delta_i^n X')^2)$. By some known estimates (see p. 1476 in Jacod and Rosenbaum (2013)),

$$\mathbb{E}\left[\left|(\Delta_i^n X)^2 1_{\{|\Delta_i^n X| \leq u_n\}} - (\Delta_i^n X')^2\right|\right] \leq K\Delta_n^{(2-r)\varpi+1},$$

which further implies $\mathbb{E}[|\zeta_{n,i}|] \leq K\Delta_n^{(2-r)\varpi}$. Hence,

$$R_{n,j} = O_p\left(\Delta_n^{(2-r)\varpi}\right). \tag{A.2}$$

For each $i \in \mathcal{I}_{n,j}$, let

$$\xi_{n,i} \equiv \Delta_i^n X' - \sigma_{t(n,j)}\Delta_i^n W = \int_{(i-1)\Delta_n}^{i\Delta_n} b_s ds + \int_{(i-1)\Delta_n}^{i\Delta_n} \left(\sigma_s - \sigma_{t(n,j)}\right) dW_s. \tag{A.3}$$

Note that by Itô's isometry and Assumption 1,

$$\mathbb{E}\left[\left|\int_{(i-1)\Delta_n}^{i\Delta_n} \left(\sigma_s - \sigma_{t(n,j)}\right) dW_s\right|^2\right] = \mathbb{E}\left[\int_{(i-1)\Delta_n}^{i\Delta_n} \left|\sigma_s - \sigma_{t(n,j)}\right|^2 ds\right] \leq K\Delta_n^2.$$

It is then easy to see that $\|\xi_{n,i}\|_2 \leq K\Delta_n$. Since $(\Delta_i^n X')^2 - (\sigma_{t(n,j)}\Delta_i^n W)^2 = \xi_{n,i}^2 + 2\xi_{n,i}\sigma_{t(n,j)}\Delta_i^n W$, we further deduce

$$\mathbb{E}\left[\left|(\Delta_i^n X')^2 - (\sigma_{t(n,j)}\Delta_i^n W)^2\right|\right] \leq \mathbb{E}\left[\xi_{n,i}^2\right] + 2\mathbb{E}\left[|\xi_{n,i}\sigma_{t(n,j)}\Delta_i^n W|\right] \leq K\Delta_n^{3/2}. \tag{A.4}$$

By the triangle inequality and (A.4),

$$\mathbb{E}\left[|R'_{n,j}|\right] \leq \Delta_n^{-1}\mathbb{E}\left[\sum_{i \in \mathcal{I}_{n,j}} \left|(\Delta_i^n X')^2 - (\sigma_{t(n,j)}\Delta_i^n W)^2\right|\right] \leq K\Delta_n^{1/2},$$

which implies that

$$R'_{n,j} = O_p\left(\Delta_n^{1/2}\right). \tag{A.5}$$

Finally, we note that $c_{t(n,j)} - c_t = O_p(\Delta_n^{1/2})$ and $S_{n,j} = O_p(1)$, and hence, $\left(c_{t(n,j)} - c_t\right)S_{n,j} = O_p(\Delta_n^{1/2})$. Combining this estimate with (A.1), (A.2), and (A.5), we deduce $k\hat{c}_{n,j} - c_t S_{n,j} = O_p(\Delta_n^{(2-r)\varpi \wedge (1/2)})$. Since $1/c_t = O_p(1)$, this further implies

$$\frac{\hat{c}_{n,j}}{c_t} - k^{-1}S_{n,j} = O_p(\Delta_n^{(2-r)\varpi \wedge (1/2)}).$$

Note that $k^{-1}S_{n,j}$ is $\bar{\chi}_k^2$-distributed. Moreover, the $S_{n,j}$ variables are clearly mutually independent. This finishes the proof of Theorem 1. $\qquad\qquad\square$

To prove Theorem 2, we need two technical lemmas. Recall that $X' \equiv X - J$ denote the continuous part of X. We consider another collection of index sets $(\widetilde{\mathcal{I}}_{n,j})_{1 \leq j \leq m_n}$ such that, for each j, $\widetilde{\mathcal{I}}_{n,j} \subseteq \mathcal{I}_{n,j}$ and $\widetilde{\mathcal{I}}_{n,j}$ contains k_j elements. We then set

$$\tilde{c}'_{n,j} \equiv (k\Delta_n)^{-1} \sum_{i \in \widetilde{\mathcal{I}}_{n,j}} \left(\Delta^n_i X'\right)^2.$$

Lemma A1 establishes a uniform approximation result for $(\tilde{c}'_{n,j})_{1 \leq j \leq m_n}$.

Lemma A1. *Suppose that Assumption 2 holds. Then, with $\widetilde{S}_{n,j} \equiv \Delta_n^{-1} \sum_{i \in \widetilde{\mathcal{I}}_{n,j}} (\Delta^n_i W)^2$, we have for any $\epsilon \in (0, 1/2)$,*

$$\max_{1 \leq j \leq m_n} \sup_{t \in \mathcal{T}_{n,j}} \left| \frac{\tilde{c}'_{n,j}}{c_t} - k^{-1}\widetilde{S}_{n,j} \right| = o_p(\Delta_n^\epsilon).$$

PROOF. For each $j \in \{1, \ldots, m_n\}$ and $t \in \mathcal{T}_{n,j}$, we decompose

$$k\tilde{c}'_{n,j} - c_t\widetilde{S}_{n,j} = \left(c_{t(n,j)} - c_t\right)\widetilde{S}_{n,j} + \Delta_n^{-1} \sum_{i \in \widetilde{\mathcal{I}}_{n,j}} \left((\Delta^n_i X')^2 - \left(\sigma_{t(n,j)}\Delta^n_i W\right)^2\right). \qquad (A.6)$$

By Assumption 2, we have for any $p \geq 1$, $1 \leq j \leq m_n$, and $t \in \mathcal{T}_{n,j}$,

$$\left\| \sup_{t \in \mathcal{T}_{n,j}} \left| c_{t(n,j)} - c_t \right| \right\|_p \leq K_p\Delta_n^{1/2}.$$

Hence, the L_p-norm of $\sup_{t \in \mathcal{T}_{n,j}} |(c_{t(n,j)} - c_t)\widetilde{S}_{n,j}|$ is also uniformly bounded by $K_p\Delta_n^{1/2}$. By a maximal inequality, we further deduce that

$$\max_{1 \leq j \leq m_n} \sup_{t \in \mathcal{T}_{n,j}} \left| c_{t(n,j)} - c_t \right| = O_p\left(\Delta_n^{1/2-1/p}\right). \qquad (A.7)$$

Define $\xi_{n,i}$ as in (A.3). Note that by the Burkholder–Davis–Gundy inequality, Hölder's inequality, and Assumption 2, we have for any $p \geq 2$,

$$\mathbb{E}\left[\left|\int_{(i-1)\Delta_n}^{i\Delta_n} \left(\sigma_s - \sigma_{t(n,j)}\right) dW_s\right|^p\right] \leq K_p\mathbb{E}\left[\left(\int_{(i-1)\Delta_n}^{i\Delta_n} \left|\sigma_s - \sigma_{t(n,j)}\right|^2 ds\right)^{p/2}\right]$$

$$\leq K_p\Delta_n^{p/2-1}\mathbb{E}\left[\int_{(i-1)\Delta_n}^{i\Delta_n} \left|\sigma_s - \sigma_{t(n,j)}\right|^p ds\right]$$

$$\leq K_p\Delta_n^p.$$

It is then easy to see that $\|\xi_{n,i}\|_p \leq K_p\Delta_n$. Consequently, for any $p \geq 1$,

$$\left\|(\Delta^n_i X')^2 - \left(\sigma_{t(n,j)}\Delta^n_i W\right)^2\right\|_p \leq \left\|\xi_{n,i}^2\right\|_p + 2\left\|\xi_{n,i}\sigma_{t(n,j)}\Delta^n_i W\right\|_p \leq K_p\Delta_n^{3/2}.$$

25

By the triangle inequality,

$$\left\| \Delta_n^{-1} \sum_{i \in \widetilde{\mathcal{I}}_{n,j}} \left((\Delta_i^n X')^2 - \left(\sigma_{t(n,j)} \Delta_i^n W \right)^2 \right) \right\|_p$$

$$\leq \Delta_n^{-1} \sum_{i \in \mathcal{I}_{n,j}} \left\| (\Delta_i^n X')^2 - \left(\sigma_{t(n,j)} \Delta_i^n W \right)^2 \right\|_p \leq K_p \Delta_n^{1/2}.$$

By a maximal inequality, we further deduce

$$\max_{1 \leq j \leq m_n} \left| \Delta_n^{-1} \sum_{i \in \widetilde{\mathcal{I}}_{n,j}} \left((\Delta_i^n X)^2 - \left(\sigma_{t(n,j)} \Delta_i^n W \right)^2 \right) \right| = O_p \left(\Delta_n^{1/2 - 1/p} \right). \tag{A.8}$$

In view of (A.6), (A.7), and (A.8), we can take p sufficiently large and deduce

$$\max_{1 \leq i \leq m_n} \sup_{t \in \mathcal{T}_{n,j}} \left| k \tilde{c}'_{n,j} - c_t \widetilde{S}_{n,j} \right| = o_p \left(\Delta_n^\epsilon \right).$$

From here, the assertion of the lemma readily follows. $\qquad \square$

Lemma A2, below, establishes a uniform bound for the density of the $\max_{1 \leq j \leq m_n} \left| \log(\bar{S}_j) \right|$ variable.

Lemma A2. *Let $k \geq 2$ and $p_n^*(\cdot)$ be the probability density function of $\max_{1 \leq j \leq m_n} \left| \log(\bar{S}_j) \right|$, where $(\bar{S}_j)_{1 \leq j \leq m_n}$ are i.i.d. $\bar{\chi}_k^2$-distributed random variables. Then, $p_n^*(x) \leq k$ for any $x \geq 0$.*

PROOF. Denote the probability density function and the cumulative distribution function of $|\log(\bar{S}_j)|$ by $g(\cdot)$ and $G(\cdot)$, respectively. Since the variables $(\bar{S}_j)_{1 \leq j \leq m_n}$ are i.i.d.,

$$\mathbb{P} \left(\max_{1 \leq j \leq m_n} \left| \log(\bar{S}_j) \right| \leq x \right) = \left(\mathbb{P} \left(\left| \log(\bar{S}_j) \right| \leq x \right) \right)^{m_n} = G(x)^{m_n},$$

which implies that

$$p_n^*(x) = m_n G(x)^{m_n - 1} g(x). \tag{A.9}$$

Since $G(\cdot)$ is an increasing function,

$$m_n G(x)^{m_n - 1} \int_x^\infty g(u) du \leq \int_x^\infty p_n^*(u) \, du = \mathbb{P} \left(\max_{1 \leq j \leq m_n} \left| \log(\bar{S}_j) \right| \geq x \right) \leq 1.$$

Therefore,

$$m_n G(x)^{m_n - 1} \leq \frac{1}{\int_x^\infty g(u) du}. \tag{A.10}$$

By (A.9) and (A.10),

$$p_n^*(x) \leq \frac{g(x)}{\int_x^\infty g(u) du}. \tag{A.11}$$

26

It remains to show that the right-hand side of (A.11) is bounded by k for all $x \geq 0$.

Let $p(\cdot)$ and $P(\cdot)$ be the density and distribution functions of $\log(\bar{S}_j)$, respectively. Note that $G(x) = P(x) - P(-x)$ and $g(x) = p(x) + p(-x)$. Since $k\bar{S}_j$ is χ_k^2-distributed,

$$P(x) = \mathbb{P}\left(k\bar{S}_j \leq ke^x\right) = \frac{1}{2^{k/2}\Gamma(k/2)} \int_0^{ke^x} u^{k/2-1} e^{-u/2} du,$$

where $\Gamma(\cdot)$ is the Gamma function, and hence,

$$p(x) = c_k \exp\left(\frac{k(x - e^x)}{2}\right), \quad \text{where} \quad c_k \equiv \frac{(k/2)^{k/2}}{\Gamma(k/2)}.$$

We further define $h(x) = x^{k/2-1} \exp(-kx/2)$ and rewrite $p(x) = c_k e^x h(e^x)$. By changing variables,

$$\int_x^\infty p(u)\, du = c_k \int_{e^x}^\infty h(u)\, du, \quad \int_x^\infty p(-u) du = c_k \int_0^{e^{-x}} h(u)\, du.$$

We can then rewrite

$$\frac{g(x)}{\int_x^\infty g(u)du} = \frac{p(x) + p(-x)}{\int_x^\infty p(u)\, du + \int_x^\infty p(-u)\, du} = \frac{e^x h(e^x) + e^{-x} h(e^{-x})}{\int_{e^x}^\infty h(u)\, du + \int_0^{e^{-x}} h(u)\, du}. \tag{A.12}$$

By (A.12),

$$\frac{g(x)}{\int_x^\infty g(u)du} \leq \frac{e^x h(e^x) + e^{-x} h(e^{-x})}{\int_0^{e^{-x}} h(u)\, du} = \frac{e^{-x} h(e^{-x})}{\int_0^{e^{-x}} h(u)\, du}\left(1 + e^{2x} \frac{h(e^x)}{h(e^{-x})}\right). \tag{A.13}$$

For ease of notation, we define

$$U(y) \equiv \frac{yh(y)}{\int_0^y h(u)\, du}, \quad V(x) \equiv \log\left(e^{2x} \frac{h(e^x)}{h(e^{-x})}\right). \tag{A.14}$$

The inequality (A.13) can then be rewritten as

$$\frac{g(x)}{\int_x^\infty g(u)du} \leq U\left(e^{-x}\right) \cdot \left(1 + e^{V(x)}\right). \tag{A.15}$$

Since $h(x) = x^{k/2-1} \exp(-kx/2)$, we can rewrite $V(x)$ more explicitly as

$$V(x) = kx - \frac{ke^x}{2} + \frac{ke^{-x}}{2}.$$

Note that $V(0) = 0$ and, for any x, $V'(x) = k(1 - e^x/2 - e^{-x}/2) \leq 0$. Hence,

$$\sup_{x \geq 0} V(x) \leq 0. \tag{A.16}$$

Turning to the $U(\cdot)$ function, we first compute its derivative

$$U'(y) = \frac{[h(y) + yh'(y)] \int_0^y h(u)\, du - yh(y)^2}{\left(\int_0^y h(u)\, du\right)^2}.$$

27

By direct calculation, we have

$$h(y) + yh'(y) = \frac{k}{2}(1-y)h(y). \qquad (A.17)$$

We can then rewrite

$$U'(y) = \frac{\widetilde{U}(y)h(y)}{\left(\int_0^y h(u)\,du\right)^2}, \quad \text{where} \quad \widetilde{U}(y) \equiv \frac{k}{2}(1-y)\int_0^y h(u)\,du - yh(y). \qquad (A.18)$$

Since $k \geq 2$, $\widetilde{U}(0) = 0$. Moreover, the derivative of $\widetilde{U}(y)$ satisfies, for $y > 0$,

$$\begin{aligned}
\widetilde{U}'(y) &= -\frac{k}{2}\int_0^y h(u)\,du + \frac{k}{2}(1-y)h(y) - h(y) - yh'(y) \\
&= -\frac{k}{2}\int_0^y h(u)\,du \leq 0,
\end{aligned}$$

where the first equality is by direct calculation, and the second equality is by (A.17). Therefore, $\widetilde{U}(y) \leq 0$ for all $y \geq 0$. By (A.18), we further see that $U(\cdot)$ is non-increasing on $[0,\infty)$. Recalling the definition of $U(\cdot)$ and applying L'Hôpital's rule, we have $U(0) = k/2$. Therefore, $U(y) \leq k/2$ for all $y \geq 0$, yielding

$$\sup_{x \geq 0} U\left(e^{-x}\right) \leq \frac{k}{2}. \qquad (A.19)$$

Combining (A.15), (A.16), and (A.19), we deduce $g(x)/\int_x^\infty g(u)du \leq k$. This inequality, together with (A.11), implies the assertion of the lemma. $\qquad\square$

We are now ready to prove Theorem 2.

PROOF OF THEOREM 2. Step 1. In this step, we derive a uniform approximation result for the truncated spot variance estimator. Let μ denote the standard Poisson random measure that drives the jumps in J, and $(\tau_q)_{q \geq 1}$ be the consecutive jumps of the Poisson process $t \mapsto \mu([0,t])$. Since the jumps of J has finite activity, the set $\{q \geq 1 : \tau_q \in [0,T]\}$ is almost surely finite. Denote $\mathcal{M}_n = \{1, \ldots, m_n\}$ and $\mathcal{M}'_n = \{j \in \mathcal{M}_n : \mathcal{T}_{n,j} \text{ does not contain any jump time } \tau_q\}$. We also define

$$\mathcal{J}_n = \{i : \Delta J_s \neq 0 \text{ for some } s \in ((i-1)\Delta_n, i\Delta_n]\},$$

and then set $\widetilde{\mathcal{I}}_{n,j} = \mathcal{I}_{n,j} \setminus \mathcal{J}_n$. Let k_j denote the number of elements in $\widetilde{\mathcal{I}}_{n,j}$. Note that for n sufficiently large, each time interval $\mathcal{T}_{n,j}$ can contain at most one jump time τ_q under the finite activity assumption. Therefore, $k_j = k - 1$ when $j \in \mathcal{M}_n \setminus \mathcal{M}'_n$ and $k_j = k$ when $j \in \mathcal{M}'_n$.

Recall that $X' = X - J$ is the continuous part of X. By Proposition 1 of Li, Todorov, and Tauchen (2017b), we have $\mathcal{J}_n = \{i : |\Delta_i^n X| > u_n\}$ with probability approaching 1, and hence,

$$\Delta_i^n X \mathbf{1}_{\{|\Delta_i^n X| \leq u_n\}} = \begin{cases} \Delta_i^n X' & \text{if } i \notin \mathcal{J}_n, \\ 0 & \text{if } i \in \mathcal{J}_n. \end{cases}$$

28

Consequently, with probability approaching 1,

$$\hat{c}_{n,j} = \tilde{c}_{n,j} \equiv (k\Delta_n)^{-1} \sum_{i \in \widetilde{\mathcal{I}}_{n,j}} \left(\Delta_i^n X'\right)^2, \quad \text{for all } j \in \mathcal{M}_n.$$

By Lemma A1, for any $\epsilon \in (0, 1/2)$,

$$\max_{1 \le j \le m_n} \sup_{t \in \mathcal{T}_{n,j}} \left| \frac{\hat{c}_{n,j}}{c_t} - k^{-1}\widetilde{S}_{n,j} \right| = o_p(\Delta_n^\epsilon), \tag{A.20}$$

where $\widetilde{S}_{n,j} \equiv \Delta_n^{-1} \sum_{i \in \widetilde{\mathcal{I}}_{n,j}} (\Delta_i^n W)^2$.

Step 2. In this step, we show that

$$\max_{1 \le j \le m_n} \sup_{t \in \mathcal{T}_{n,j}} \left| \log\left(\hat{c}_{n,t}\right) - \log\left(c_t\right) - \log\left(k^{-1}\widetilde{S}_{n,j}\right) \right| = o_p(1). \tag{A.21}$$

Denote $Y_{n,j} \equiv \hat{c}_{n,j}/c_{t(n,j)}$ and $D_{n,j} \equiv Y_{n,j} - k^{-1}\widetilde{S}_{n,j}$. Note that

$$\max_{1 \le j \le m_n} \sup_{t \in \mathcal{T}_{n,j}} \left| \log\left(\hat{c}_{n,t}\right) - \log\left(c_t\right) - \log\left(k^{-1}\widetilde{S}_{n,j}\right) \right|$$

$$\le \max_{1 \le j \le m_n} \left| \log\left(Y_{n,j}\right) - \log\left(k^{-1}\widetilde{S}_{n,j}\right) \right| + \max_{1 \le j \le m_n} \sup_{t \in \mathcal{T}_{n,j}} \left| \log\left(c_t\right) - \log\left(c_{t(n,j)}\right) \right|. \tag{A.22}$$

By localization, we can assume that $c_t \ge \eta$ for some constant $\eta > 0$ without loss of generality. Then, by the mean value theorem, $|\log\left(c_t\right) - \log(c_{t(n,j)})| \le K|c_t - c_{t(n,j)}|$. From (A.7), we further see that for any $\iota \in (0, 1/2)$,

$$\max_{1 \le j \le m_n} \sup_{t \in \mathcal{T}_{n,j}} \left| \log\left(c_t\right) - \log\left(c_{t(n,j)}\right) \right| = O_p\left(\Delta_n^{1/2 - \iota}\right). \tag{A.23}$$

By (A.20), there exists a positive real sequence $\delta_n = o(\Delta_n^{0.49})$ such that the sequence of events $\Omega_n \equiv \{|D_{n,j}| \le \delta_n, \text{ for all } 1 \le j \le m_n\}$ satisfies $\mathbb{P}(\Omega_n) \to 1$. By the mean value theorem, for each j, there exists some $\lambda_{n,j} \in [0, 1]$ such that

$$\left| \log\left(Y_{n,j}\right) - \log\left(k^{-1}\widetilde{S}_{n,j}\right) \right| = \frac{|D_{n,j}|}{k^{-1}\widetilde{S}_{n,j} + \lambda_{n,j}D_{n,j}}. \tag{A.24}$$

Let $S_{n,j} \equiv \Delta_n^{-1} \sum_{i \in \mathcal{I}_{n,j}} (\Delta_i^n W)^2$. For any constant $C_1 > 0$, we observe that

$$\mathbb{P}\left(\max_{j \in \mathcal{M}_n'} \frac{1}{\left|k^{-1}\widetilde{S}_{n,j} + \lambda_{n,j}D_{n,j}\right|} \ge C_1 n^{2/(k-2)}\right)$$

$$\le \mathbb{P}\left(\left\{\max_{1 \le j \le m_n} \frac{1}{|k^{-1}S_{n,j} + \lambda_{n,j}D_{n,j}|} \ge C_1 n^{2/(k-2)}\right\} \cap \Omega_n\right) + o(1)$$

$$\le \sum_{j=1}^{m_n} \mathbb{P}\left(\left\{|k^{-1}S_{n,j} + \lambda_{n,j}D_{n,j}| \le \frac{1}{C_1} n^{-2/(k-2)}\right\} \cap \Omega_n\right) + o(1)$$

$$\le \sum_{j=1}^{m_n} \mathbb{P}\left(S_{n,j} \le \frac{k}{C_1} n^{-2/(k-2)} + k\delta_n\right) + o(1), \tag{A.25}$$

29

where the first inequality follows from the fact that $\widetilde{S}_{n,j} = S_{n,j}$ when $j \in \mathcal{M}'_n$ and $\mathbb{P}(\Omega_n) \to 1$, the second inequality is obvious, and the last line holds because $|D_{n,j}| \leq \delta_n$ for all j in restriction to Ω_n.

Note that $S_{n,j} \sim \chi^2_k$, so its probability density is bounded by $Kx^{k/2-1}$ for x near zero. Therefore,

$$\sum_{j=1}^{m_n} \mathbb{P}\left(S_{n,j} \leq \frac{k}{C_1}n^{-2/(k-2)} + k\delta_n\right) \leq Km_n\left(\frac{1}{C_1}n^{-2/(k-2)} + \delta_n\right)^{k/2-1}$$

$$\leq KC_1^{1-(k/2)} + Km_n\delta_n^{k/2-1}. \tag{A.26}$$

Since $\delta_n = o(\Delta_n^{0.49})$ and $k \geq 7$, $m_n\delta_n^{k/2-1} = o(1)$. Then, by (A.25) and (A.26),

$$\limsup_{n\to\infty} \mathbb{P}\left(\max_{j\in\mathcal{M}'_n} \frac{1}{\left|k^{-1}\widetilde{S}_{n,j} + \lambda_{n,j}D_{n,j}\right|} \geq C_1n^{2/(k-2)}\right) \leq KC_1^{1-(k/2)}.$$

Therefore,

$$\max_{j\in\mathcal{M}'_n} \frac{1}{\left|k^{-1}\widetilde{S}_{n,j} + \lambda_{n,j}D_{n,j}\right|} = O_p\left(n^{2/(k-2)}\right). \tag{A.27}$$

Note that for $j \in \mathcal{M}_n \setminus \mathcal{M}'_n$, $\widetilde{S}_{n,j} \sim \chi^2_{k-1}$, and it is easy to see that $1/|k^{-1}\widetilde{S}_{n,j} + \lambda_{n,j}D_{n,j}| = O_p(1)$. Since the set $\mathcal{M}_n \setminus \mathcal{M}'_n$ is finite,

$$\max_{j\in\mathcal{M}_n\setminus\mathcal{M}'_n} \frac{1}{\left|k^{-1}\widetilde{S}_{n,j} + \lambda_{n,j}D_{n,j}\right|} = O_p\left(1\right),$$

which together with (A.27) implies

$$\max_{1\leq j\leq m_n} \frac{1}{\left|k^{-1}\widetilde{S}_{n,j} + \lambda_{n,j}D_{n,j}\right|} = O_p\left(n^{2/(k-2)}\right). \tag{A.28}$$

Since $k \geq 7$, $2/(k-2) \leq 2/5 < 1/2$. By (A.20), (A.24), and (A.28), we deduce that (by taking ϵ sufficiently close to $1/2$)

$$\max_{1\leq j\leq m_n} \left|\log(Y_{n,j}) - \log\left(k^{-1}\widetilde{S}_{n,j}\right)\right| = o_p(\Delta_n^{\epsilon-2/(k-2)}) = o_p(1). \tag{A.29}$$

Combining (A.22), (A.23), and (A.29), we deduce (A.21) as claimed.

Step 3. We prove the assertion of the theorem in this step. By (A.21),

$$\sup_{t\in[0,T]} |\log(\hat{c}_{n,t}) - \log(c_t)| - \max_{1\leq j\leq m_n} \left|\log\left(k^{-1}\widetilde{S}_{n,j}\right)\right| = o_p(1). \tag{A.30}$$

30

We note that the variables $(\widetilde{S}_{n,j})_{1\le j\le m_n}$ are independent. Moreover, $\widetilde{S}_{n,j} = S_{n,j} \sim \chi_k^2$ for all but finitely many j's (which correspond to the $\mathcal{T}_{n,j}$ intervals containing jump times), and $\widetilde{S}_{n,j} \sim \chi_{k-1}^2$ in the exceptional cases. Since $m_n \to \infty$, it is easy to see that

$$\mathbb{P}\left(\max_{1\le j\le m_n}\left|\log\left(k^{-1}\widetilde{S}_{n,j}\right)\right| \ne \max_{1\le j\le m_n}\left|\log\left(k^{-1}S_{n,j}\right)\right|\right) = o(1). \tag{A.31}$$

By (A.30) and (A.31),

$$\sup_{t\in[0,T]}\left|\log\left(\hat{c}_{n,t}\right) - \log\left(c_t\right)\right| - \max_{1\le j\le m_n}\left|\log\left(k^{-1}S_{n,j}\right)\right| = o_p(1).$$

Hence, there exists a positive real sequence $\delta_n = o(1)$ such that with probability approaching 1,

$$\left|\sup_{t\in[0,T]}\left|\log\left(\hat{c}_{n,t}\right) - \log\left(c_t\right)\right| - \max_{1\le j\le m_n}\left|\log\left(k^{-1}S_{n,j}\right)\right|\right| < \delta_n. \tag{A.32}$$

Consequently,

$$\mathbb{P}\left(\sup_{t\in[0,T]}\left|\log\left(\hat{c}_{n,t}\right) - \log\left(c_t\right)\right| \le \bar{z}_{n,\alpha}\right)$$
$$\le \mathbb{P}\left(\max_{1\le j\le m_n}\left|\log\left(k^{-1}S_{n,j}\right)\right| \le \bar{z}_{n,\alpha} + \delta_n\right) + o(1)$$
$$= \mathbb{P}\left(\max_{1\le j\le m_n}\left|\log\left(k^{-1}S_{n,j}\right)\right| \le \bar{z}_{n,\alpha}\right)$$
$$+ \mathbb{P}\left(\bar{z}_{n,\alpha} < \max_{1\le j\le m_n}\left|\log\left(k^{-1}S_{n,j}\right)\right| \le \bar{z}_{n,\alpha} + \delta_n\right) + o(1)$$
$$= \mathbb{P}\left(\max_{1\le j\le m_n}\left|\log\left(k^{-1}S_{n,j}\right)\right| \le \bar{z}_{n,\alpha}\right) + o(1)$$
$$= 1 - \alpha + o(1), \tag{A.33}$$

where the first inequality is by (A.32), the first equality holds obviously, the second equality holds because the density of $\max_{1\le j\le m_n}\left|\log\left(k^{-1}S_{n,j}\right)\right|$ is bounded (see Lemma A2) and $\delta_n = o(1)$, and the last line follows from the definition of $\bar{z}_{n,\alpha}$. By a similar argument, we can also show that

$$\mathbb{P}\left(\sup_{t\in[0,T]}\left|\log\left(\hat{c}_{n,t}\right) - \log\left(c_t\right)\right| \le \bar{z}_{n,\alpha}\right) \ge 1 - \alpha - o(1). \tag{A.34}$$

The assertion of the theorem then readily follows from (A.33) and (A.34). $\qquad\square$

References

AI, H., AND R. BANSAL (2018): "Risk Preferences and the Macroeconomic Announcement Premium," *Econometrica*, 86(4), 1383–1430.

AÏT-SAHALIA, Y., AND J. JACOD (2007): "Volatility Etimators for Discretely Sampled Levy Processes," *Annals of Statistics*, 37(1), 355–392.

——— (2014): *High-Frequency Financial Econometrics*. Princeton University Press.

ALVAREZ, A., F. PANLOUP, M. PONTIER, AND N. SAVY (2012): "Estimation of the Instantaneous Volatility," *Statistical Inference and Stochastic Processes*, 15, 27–59.

ANDERSEN, T., AND T. BOLLERSLEV (1997): "Intraday Periodicity and Volatility Persistence in Financial Markets," *Journal of Empirical Finance*, 4(2-3), 115–158.

——— (1998): "Answering the Skeptics: Yes, Standard Volatility Models do Provide Accurate Forecasts," *International Economic Review*, 39(4), 885–905.

ANDERSEN, T., T. BOLLERSLEV, F. DIEBOLD, AND P. LABYS (2001): "The Distribution of Realized Exchange Rate Volatility," *Journal of the American Statistical Association*, 96, 42–55.

ANDERSEN, T. G., AND T. BOLLERSLEV (2018): *Volatility*. Edward Elgar Publishing.

ANDERSEN, T. G., T. BOLLERSLEV, P. F. CHRISTOFFERSEN, AND F. X. DIEBOLD (2013): "Financial Risk Measurement for Financial Risk Management," in *Handbook of Economics of Finance, Volume 2B*, ed. by G. Constantinides, M. Harris, and R. Stulz, pp. 1127–1220. Elsevier, Amsterdam.

ANDERSEN, T. G., T. BOLLERSLEV, AND F. X. DIEBOLD (2007): "Roughing it Up: Disentangling Continuous and Jump Components in Measuring, Modeling and Forecasting Asset Return Volatility," *Review of Economics and Statistics*, 89(4), 701–720.

ANDERSEN, T. G., T. BOLLERSLEV, F. X. DIEBOLD, AND P. LABYS (2003): "Modeling and Forecasting Realized Volatility," *Econometrica*, 71(2), pp. 579–625.

ANDERSEN, T. G., T. BOLLERSLEV, F. X. DIEBOLD, AND C. VEGA (2003): "Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange," *American Economic Review*, 93(1), 251 – 277.

ANDERSEN, T. G., N. FUSARI, AND V. TODOROV (2015): "Parametric Inference and Dynamic State Recovery From Option Panels," *Econometrica*, 83(3), 1081–1145.

BAKSHI, G., C. CAO, AND Z. CHEN (1997): "Empirical Performance of Alternative Option Pricing Models," *Journal of Finance*, 52(5), 2003–2049.

BANDI, F., AND R. RENO (2016): "Price and Volatility Co-Jumps," *Journal of Financial Economics*, 119(1), 107–146.

BARNDORFF-NIELSEN, O., AND N. SHEPHARD (2002): "Econometric Analysis of Realized Volatility and its Use in Estimating Stochastic Volatility Models," *Journal of the Royal Statistical Society*, Series B, 64, 253–280.

BARNDORFF-NIELSEN, O. E., AND N. SHEPHARD (2004): "Power and bipower variation with stochastic volatility and jumps (with discussion)," *Journal of Financial Econometrics*, 2, 1–48.

BATES, D. S. (2000): "Post-'87 Crash Fears in the S&P 500 Futures Option Market," *Journal of Econometrics*, 94(1-2), 181–238.

BERNANKE, B. S., AND K. N. KUTTNER (2005): "What Explains the Stock Market's Reaction to Federal Reserve Policy?," *The Journal of Finance*, 60(3), 1221–1257.

BERNILE, G., J. HU, AND Y. TANG (2016): "Can Information Be Locked Up? Informed Trading Ahead of Macro-News Announcements," *Journal of Financial Economics*, 121(3), 496 – 520.

BLOOM, N. (2009): "The Impact of Uncertainty Shocks," *Economerrica*, 77(3), 623–685.

BOLLERSLEV, T., J. LI, AND Y. XUE (2018): "Volume, Volatility, and Public News Announcements," *Review of Economic Studies*, 85(4), 2005–2041.

BOLLERSLEV, T., AND V. TODOROV (2011): "Estimation of Jump Tails," *Econometrica*, 79(6), 1727–1783.

CIESLAK, A., A. MORSE, AND A. VISSING-JORGENSEN (2019): "Stock Returns over the FOMC Cycle," *Journal of Finance*, 74(5), 2201–2248.

COCHRANE, J. H., AND M. PIAZZESI (2002): "The Fed and Interest Rates - A High-Frequency Identification," *American Economic Review*, 92(2), 90–95.

COMTE, F., AND E. RENAULT (1996): "Long Memory Continuous Time Models," *Journal of Econometrics*, 73(1), 101–149.

——— (1998): "Long Memory in Continuous Time Stochastic Volatility Models," *Mathematical Finance*, 8(4), 291–323.

DUFFIE, D., J. PAN, AND K. SINGLETON (2000): "Transform Analysis and Asset Pricing for Affine Jump-Diffusions," *Econometrica*, 68(6), 1343–1376.

FERNANDEZ-VILLAVERDE, J., P. GUERRÓN-QUINTANA, J. F. RUBIO-RAMIREZ, AND M. URIBE (2011): "Risk Matters: The Real Effects of Volatility Shocks," *American Economic Review*, 101(6), 2530–2561.

FERNANDEZ-VILLAVERDE, J., AND J. F. RUBIO-RAMIREZ (2013): "Macroeconomics and Volatility: Data, Models, and Estimation," in *Advances in Economics and Econometrics, Tenth World Congress*, ed. by D. Acemoglu, M. Arellano, and E. Dekel, pp. 137–183. Cambridge University Press.

FOSTER, D. P., AND D. B. NELSON (1996): "Continuous Record Asymptotics for Rolling Sample Variance Estimators," *Econometrica*, 64(1), 139–174.

GONÇALVES, S., AND N. MEDDAHI (2009): "Bootstrapping realized volatility," *Econometrica*, 77, 283–306.

JACOD, J., AND P. PROTTER (2012): *Discretization of Processes*. Springer.

Jacod, J., and M. Rosenbaum (2013): "Quarticity and Other Functionals of Volatility: Efficient Estimation," *The Annals of Statistics*, 41(3), 1462–1484.

Johnson, J. A., and B. Paye (2019): "Asset Price Reactions to Unconventional Monetary Policy Surprises," *Working paper*.

Kristensen, D. (2010): "Nonparametric Filtering of the Realized Spot Volatility: A Kernel-Based Approach," *Econometric Theory*, 26(1), 60–93.

Kurov, A., A. Sancetta, G. Strasser, and M. H. Wolfe (2019): "Price Drift Before U.S. Macroeconomic News: Private Information about Public Announcements?," *Journal of Financial and Qunatitative Analysis*, 54(1), 449–479.

Lee, S. (2012): "Jumps and Information Flow in Financial Markets," *Review of Financial Studies*, 25(2), 439–479.

Lee, S., and P. Mykland (2008): "Jumps in Financial Markets: A New Nonparametric Test and Jump Dynamics," *Review of Financial Studies*, 21(6), 2535–2563.

Li, J., V. Todorov, and G. Tauchen (2017a): "Adaptive Estimation of Continuous-Time Regression Models using High-Frequency Data," *Journal of Econometrics*, 200(1), 36–47.

——— (2017b): "Jump Regressions," *Econometrica*, 85(1), 173–195.

Lucca, D. O., and E. Moench (2015): "The Pre-FOMC Announcement Drift," *Journal of Finance*, 70(1), 329–371.

Mancini, C. (2001): "Disentangling the Jumps of the Diffusion in a Geometric Jumping Brownian Motion," *Giornale dell'Istituto Italiano degli Attuari*, LXIV, 19–47.

Nakamura, E., and J. Steinsson (2018): "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect," *Quarterly Journal of Economics*, 133(3), 1283–1330.

Nelson, D. B. (1992): "Filtering and Forecasting with Misspecified ARCH Models I: Getting the Right Variance with the Wrong Model," *Journal of Econometrics*, 52(1-2), 61–90.

Rigobon, R. (2003): "Identification Through Heteroskedasticity," *Review of Economics and Statistics*, 85(4), 777–792.

Rigobon, R., and B. Sack (2004): "The Impact of Monetary Policy on Asset Prices," *Journal of Monetary Economics*, 51(8), 1553 – 1575.

Savor, P., and M. Wilson (2014): "Asset Pricing: A Tale of Two Days," *Journal of Financial Economics*, 113, 171–201.

Wright, J. H. (2012): "What Does Monetary Policy do at the Zero Lower Bound?," *Economic Journal*, 122(564), 447–466.

ZHANG, L., P. A. MYKLAND, AND Y. AÏT-SAHALIA (2005): "A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data," *Journal of the American Statistical Association*, 100(472), 1394–1411.