

TESTING THE DIMENSIONALITY OF POLICY SHOCKS

Jia Li, Viktor Todorov, and Qiushi Zhang*

Abstract—This paper provides a nonparametric test for deciding the dimensionality of a policy shock as manifest in the abnormal change in asset returns' stochastic covariance matrix, following the release of a macroeconomic announcement. We use high-frequency data in local windows before and after the event to estimate the covariance jump matrix and then test its rank. We find a one-factor structure in the covariance jump matrix of the yield curve resulting from the Federal Reserve's monetary policy shocks before the 2007–2009 financial crisis. The dimensionality of policy shocks increased afterwards because of the use of unconventional monetary policy tools.

I. Introduction

THE movements of financial asset prices are driven by underlying economic shocks. Large market-wide shocks are often caused by unanticipated components of economic policies, such as central banks' decisions on interest rates and governments' emergency aids during a crisis. Understanding the impact of policy shocks on financial markets and the macroeconomy is evidently of great interest for both academics and policymakers. Disentangling such shocks from the other confounding factors, however, presents a difficult empirical challenge. The conventional approach based on structural vector autoregressive (VAR) models (Christiano, Eichenbaum, and Evans, 1999) often relies on strong orthogonalization assumptions and may incorrectly interpret anticipated actions to be shocks (Rudebusch, 1998; Cochrane & Piazzesi, 2002).

An alternative approach to achieve identification, which has become increasingly popular in the recent literature, is to study the behavior of asset prices in short time windows around certain news announcements, so that high-frequency price movements can be plausibly attributed to the announced information. A prominent example in this regard concerns the reaction of asset prices to monetary policy shocks triggered by Federal Open Market Committee (FOMC) announcements. For instance, Kuttner (2001), Cochrane and Piazzesi (2002), and Nakamura and Steinsson (2018) use moves in bond futures prices around FOMC announcements to identify monetary shocks; Rigobon and Sack (2003, 2004) and Bernanke and Kuttner (2005) use information in asset price data around FOMC announcements to study the impact of monetary shocks on asset prices; and Beechey and Wright (2009) study the effect of macroeco-

nomic news on long-term yields and compare it with that implied by structural models.¹

A key implication of having one type of shock driving multiple asset prices around the news event (e.g., FOMC announcements) is that the resulting abnormal change, or “jump,” in assets' stochastic covariance matrix should have a one-factor structure. A case in point is the conventional monetary policy, which focuses solely on the short-term interest rate. But some policy shocks may be multidimensional. For example, during the Great Recession of 2007–2009 and its aftermath, the Federal Reserve and many other central banks were impelled to employ new monetary policy tools such as forward guidance and large-scale asset purchases; see Bernanke (2020) for a recent review. It is conceivable that the complex mix of policy tools may have resulted in multidimensional policy shocks in the era of unconventional monetary policy. Nonetheless, if the dimensionality of the policy shock is lower than the number of assets, which is likely the case, the announcement-induced jump in the covariance matrix may still be of reduced rank. Pinning down the dimensionality of policy shocks is a natural starting point for the further understanding of their impact on asset prices and other economic quantities. The dimensionality, however, cannot be directly read off from the actual announcement because the shock stems from the unanticipated policy component, rather than the policy itself.

Set against this background, our goal in this paper is to provide a formal test to rigorously uncover the dimensionality of policy shocks as manifest in asset prices. The method is decidedly nonparametric and is based on intraday high-frequency observations of assets' returns around news events such as prescheduled macroeconomic announcements. Using the high-frequency data, we form nonparametric estimates of the instantaneous (or “spot”) covariance matrices of asset returns immediately before and after the event. Their difference is our estimate of the covariance jump matrix. With the underlying confounding factors differenced out, the covariance jump matrix captures the effect of the policy shock on asset prices, and its rank is equivalent to the shock's dimensionality. Econometrically, we carry out the rank test by checking whether all but the r largest (in magnitude) estimated eigenvalues are statistically equal to zero, where r is the dimensionality being tested. The corresponding test statistic is defined as the sum of squared estimated eigenvalues associated with the zero eigenvalues under the null hypothesis.

Received for publication May 21, 2020. Revision accepted for publication November 9, 2021. Editor: Xiaoxia Shi.

*Li: Singapore Management University; Todorov: Northwestern University; Zhang: University of International Business and Economics.

Part of this research was conducted while Li was visiting Yale University and the Cowles Foundation. We would like to thank a coeditor and two anonymous referees for useful comments and suggestions.

A supplemental appendix is available online at https://doi.org/10.1162/rest_a_01139.

¹A large literature in finance also documents the equity risk premium earned during different phases of the FOMC news release cycle; see, for example, Savor and Wilson (2014), Lucca and Moench (2015), and Cieslak et al. (2019). Ai and Bansal (2018) provide a nonexpected utility theory that explains the risk premium related to macro news announcement.

We establish the asymptotic properties of the test in an in-fill asymptotic framework with the sampling interval of high-frequency observations going to zero asymptotically. The limiting distribution of the test statistic is nonstandard: it takes the form of a weighted mixture of chi-squared random variables, with the random weights depending on the stochastic volatilities before and after each event. Importantly, the event-specific heterogeneity is fully reflected in these random weights without being “averaged out” as in conventional event-study method based on classical time-series tools. This desirable feature is achievable here because we draw inference from the “large number” of intraday observations around individual events, rather than by pooling a large number of distinct events based on a (likely fragile) homogeneity presumption.² For the same reason, the proposed high-frequency econometric method can in fact be implemented on an event-by-event basis. Since the limiting distribution of the test statistic is nonstandard, we propose and theoretically justify an easy-to-implement bootstrap method for constructing critical values, which exhibits good finite-sample performance in empirically realistically calibrated Monte Carlo experiments. Our analysis of the asymptotic behavior of the test statistic and the associated bootstrap method is new to the literature, which is the main theoretical contribution of the present paper.

On the empirical side, we apply the proposed inference procedure to study the dimensionality of policy shocks generated by FOMC announcements through testing the rank of the resulting stochastic covariance jump matrix of the yield curve. Our sample spans the period from October 2001 till August 2018 and covers 135 scheduled FOMC announcements. We study the responses of Eurodollar and Treasury security futures with maturities ranging from one year to thirty years. The testing results show that, before the financial crisis of 2007–2009, the covariance jump matrix of bond futures around FOMC announcements had a one-factor structure, which suggests that monetary policy shocks were one-dimensional under the conventional monetary policy regime. However, from the onset of the use of unconventional monetary policies by the central bank, the dimensionality of the policy shock increased to three in the 2007–2009 crisis period and largely stayed at two from then on. These findings formally demonstrate the complex impact of monetary policy on asset prices during the crisis and its aftermath. In particular, they are consistent with the fact that news about forward guidance and large-scale asset purchases can generate shocks not only to investors’ long-term expectations of interest rates but also to bond risk premia as recently discussed by Bernanke (2020). Our postcrisis testing results are also in line with the findings of Gurkaynak, Sack, and Swanson (2004); Swanson (2017); Cieslak and Schrimpf (2019), and Cieslak and Pang (2021) for the multifactor structure of asset returns around FOMC announcements.

The present paper is related to several strands of literature. A large literature looks at rank testing in various economic contexts; see Rothenberg (1973), Hsiao (1983), Lewbel (1991), Gill and Lewbel (1992), Cragg and Donald (1997), and Robin and Smith (2000), among others. Our theory is different from that prior literature mainly because of its nonstandard in-fill asymptotic setting, which in particular allows for essentially unrestricted nonstationarity and data heterogeneity in a nonergodic setting. Under the in-fill asymptotic setting, Li et al. (2019) recently studied a rank test for price jumps of multiple assets over a collection of statistically detected jump events. But the econometric analysis here is fundamentally different from that prior work for two reasons. One is that we focus on jumps in the spot covariance matrix, treating the price jumps as a nuisance. The other difference is that we consider prescheduled macro news announcements rather than statistically detected price-jump events.³ These differences also manifest in the methods’ distinct empirical scopes: our test speaks to the dimensionality of announcement-induced policy shocks, whereas the method of Li et al. (2019) mainly concerns the stability of factor loadings over time. Our analysis of the rank of the covariance jump matrix is more generally related to heteroskedasticity-based identification of simultaneous linear equations models (Rigobon, 2003; Rigobon & Sack, 2004), for which the key identification assumption is that the change in the covariance matrix of asset returns has rank one. Our in-fill asymptotic theory also contributes to the high-frequency econometrics literature on nonparametric volatility inference; see, for example, Comte and Renault (1996), Foster and Nelson (1996), Kristensen (2010), and Jacod and Protter (2012). However, unlike the prior work, our focus is not on the spot volatility estimation per se, but rather on the rank of the covariance jump matrix. Finally, our bootstrap inference shares the same “local i.i.d.” resampling scheme as originally proposed by Bollerslev, Li, and Xue (2018). Our new bootstrap result is developed under much weaker conditions (regarding jump activity) in a more general multivariate setting; moreover, Bollerslev et al. (2018) does not consider rank test, which is exactly the focus here.

The remainder of the paper is organized as follows. Section II describes the new inference method and the underlying asymptotic theory. In section III we examine the finite-sample performance of the proposed test in a Monte Carlo experiment. An empirical application on monetary policy shocks is presented in section IV. Section V concludes. The online supplemental appendix contains all proofs and additional numerical results. The following notation will be used. We denote the d -dimensional identity matrix by I_d . The

²See Kuttner (2018) for additional discussions on the conventional event-study approach and the related fragility issue.

³Kandel and Pearson (1995) show both theoretically and empirically that a news announcement does not necessarily cause a price jump, but generally leads to elevated trading activity when investors agree to disagree on the interpretation of the news, which in turn results in a heightened level of asset price volatility after the announcement; see Lucca and Moench (2015) and Bollerslev et al. (2018) for additional empirical evidence for macro news announcements.

Euclidean norm is denoted $\|\cdot\|$. For a matrix A , we use A_{jk} and $A^\top$ to denote its (j, k) element and transpose, respectively. For two real sequences a_n and b_n , we write $a_n \asymp b_n$ if $a_n/C \leq b_n \leq Ca_n$ for some finite constant $C \geq 1$.

II. The Econometric Method

We describe the rank test of spot covariance jump matrix in this section. Section IIA introduces the setting, and section IIB describes the test and establishes its asymptotic validity.

A. The Setting

Suppose that the vector of log price processes X is a d -dimensional Itô semimartingale, $d \geq 2$, defined on a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t \geq 0}, \mathbb{P})$ that can be written as

$$X_t = X_0 + \int_0^t b_s ds + \int_0^t \sigma_s dW_s + J_t, \quad (1)$$

where the d -dimensional drift process b is optional, the stochastic volatility matrix process σ is càdlàg adapted and takes values in $\mathbb{R}^{d \times d}$, and W is a standard d -dimensional Brownian motion. The J process captures the price jumps, defined as the sum of a purely discontinuous local martingale with jump sizes no bigger than 1 and a purejump process with jump sizes bigger than 1, both of which are driven by a homogeneous Poisson random measure on $\mathbb{R}_+ \times \mathbb{R}$. We impose the following regularity conditions on the underlying processes.

Assumption 1. Suppose that X has the form equation (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 $\gamma \in [0, 2)$, $|b_t| + |\sigma_t| + |\sigma_t|^{-1} + \int (|x|^\gamma \wedge 1) F_t(dx) \leq K_m$ 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 and rather standard regularity conditions, allowing for essentially unrestricted price and volatility jumps, leverage effect, and intraday periodicity. Condition (i) imposes local boundedness on various processes, and condition (ii) states that the volatility process σ is locally $(1/2)$ -Hölder continuous under the L_2 norm, which can be readily verified if σ is an Itô semimartingale or a long-memory process driven by a fractional Brownian motion.

The spot covariance matrix process is formally defined as

$$c_t \equiv \sigma_t \sigma_t^\top,$$

which can be interpreted as the instantaneous covariance matrix of the diffusive returns (i.e., $\sigma_t dW_t$). As mentioned in the introduction, we focus not on the level of spot covariance matrix per se, but rather on its jump at the news announce-

ment time, which precisely measures the “abnormal” movement induced by the “lumpy” information embedded in the announcement. More precisely, with $\mathcal{T} \equiv \{\tau_1, \dots, \tau_m\}$ denoting a collection of announcement times, we denote the spot covariance jump matrix at each time $\tau \in \mathcal{T}$ as

$$\Delta c_\tau \equiv c_{\tau+} - c_{\tau-},$$

with $c_{\tau-}$ and $c_{\tau+}$ being the left and right limits of the c process at time τ , respectively.⁴ Note that the spot covariance matrix c_t at a given point in time generally has full rank when the assets under consideration are nonredundant. However, the jump matrix Δc_τ may be of reduced rank if the underlying policy shock has a lower-dimensional structure. For example, if the Federal Reserve surprises the market by *only* altering the short-term interest rate, as is typical under the conventional monetary policy, we may find a one-factor structure in Δc_τ (see, e.g., Rigobon & Sack, 2004). On the other hand, a multidimensional policy shock may arise if the announcement also contains forward guidance regarding the future trajectory of interest rates, which in turn can result in a higher rank in the jump matrix Δc_τ .

The main goal of this paper is to uncover the dimensionality of policy shocks through testing the rank of the spot covariance jump matrix. The method is decidedly nonparametric without imposing any parametric restrictions. That noted, the proposed method also speaks to more specific structural estimation problems concerning policy impact. One case in point is the heteroskedasticity-based identification and estimation of linear simultaneous equation models as considered by Rigobon (2003) and Rigobon and Sack (2004). The key premise of this identification strategy is that there is a single source of policy shock, which implies that the difference between assets' covariance matrices in two subsamples (corresponding to announcement and nonannouncement periods) has rank one. The rank test proposed below may be used to test the underlying identification assumption. Recently, Lewis (2020) proposed an empirical strategy to identify announcement-specific decompositions of asset price changes into monetary policy shocks using high-frequency data. He does not provide formal inference theory for his empirical procedure, as the “parameters” of econometric interest are themselves random quantities, which is typical in nonergodic high-frequency inference problems as studied in the present paper. Our theory may also shed light on further theoretical development in this direction.⁵

These economic considerations motivate us to rigorously test the rank of the spot covariance jump matrix. Formally, the null and alternative hypotheses of interest are

⁴Since the c process is right continuous, we simply have $c_{\tau+} = c_\tau$. Nevertheless, we adopt the $c_{\tau+}$ notation because it matches the same notational convention in equation (4) below and is convenient for later discussions.

⁵It is also worth noting that heteroskedasticity-based inference may encounter weak-identification issues if time variation in some component of the volatility process is moderate. In such a scenario, the econometric method of Cheng et al. (2021) may be applied to achieve robust and efficient inference.

represented, respectively, by the following events:

$$\begin{aligned}\Omega_{0,r} &= \{\text{Rank}(\Delta c_\tau) = r \text{ for all } \tau \in \mathcal{T}\}, \\ \Omega_{a,r} &= \{\text{Rank}(\Delta c_\tau) > r \text{ for some } \tau \in \mathcal{T}\},\end{aligned}\quad (2)$$

where $r \in \{1, \dots, d-1\}$ is the candidate rank to be tested. We note that in the present setting, the spot covariance jump matrix Δc_τ is itself a random matrix, and the null hypothesis $\Omega_{0,r}$ contains the sample paths on which this random matrix has rank r . Specifying hypotheses as random events is common in the high-frequency econometrics literature, because the “population quantity” is the realized sample paths of processes of interest, instead of some constant parameter; see Aït-Sahalia and Jacod (2014) for a comprehensive review.

Our inference is developed under an in-fill asymptotic framework that is now standard in the high-frequency econometrics literature (see, e.g., Jacod & Protter, 2012 and Aït-Sahalia & Jacod, 2014). Suppose that the price vector process X is observed at discrete times $i\Delta_n$ for $i = 0, 1, \dots$, within the fixed time interval $[0, T]$. We denote the i th return of X by

$$R_{n,i} \equiv X_{i\Delta_n} - X_{(i-1)\Delta_n}.$$

Since we are interested in the local behavior of the spot covariance matrix process near announcement times, we focus on return observations in local windows before and after each announcement. For each $\tau \in \mathcal{T}$, let i_τ denote the unique integer such that $\tau \in ((i_\tau - 1)\Delta_n, i_\tau\Delta_n]$. We pick a local window sequence k_n of integers satisfying $k_n \rightarrow \infty$ and $k_n\Delta_n \rightarrow 0$, which plays a similar role as the “bandwidth” parameter in conventional kernel-based nonparametric estimation. Our inference relies on observations in the pre-event window $\mathcal{I}_{n,\tau-}$ and the post-event window $\mathcal{I}_{n,\tau+}$, defined, respectively, as

$$\begin{aligned}\mathcal{I}_{n,\tau-} &\equiv \{i_\tau - k_n, \dots, i_\tau - 1\}, \\ \mathcal{I}_{n,\tau+} &\equiv \{i_\tau + 1, \dots, i_\tau + k_n\}.\end{aligned}\quad (3)$$

Each of these local windows consists of k_n returns. It is instructive to note that the return at the announcement time (indexed by i_τ) is excluded from these windows, as it is likely to be “contaminated” by the price jump at the announcement time, which may bias the estimation of the spot covariance matrix.

The spot covariance matrix estimators before and after each announcement are constructed essentially as the second sample moments of returns from the corresponding local windows. Specifically, we estimate the pre-event and post-event spot covariance matrices using

$$\hat{c}_{n,\tau\pm} \equiv \frac{1}{k_n\Delta_n} \sum_{i \in \mathcal{I}_{n,\tau\pm}} R_{n,i} R_{n,i}^\top 1_{\{\|R_{n,i}\| \leq u_n\}}, \quad (4)$$

where the $1_{\{\|R_{n,i}\| \leq u_n\}}$ indicator is a standard device used to eliminate price jumps in the spot covariance estimation (as originally proposed by Mancini, 2001), with u_n being

a (shrinking) sequence of truncation threshold. Using well-known results in the literature (see, e.g., Chapter 9 of Jacod & Protter, 2012), the spot estimators $\hat{c}_{n,\tau-}$ and $\hat{c}_{n,\tau+}$ consistently estimate the pre-event and post-event spot covariance matrices $c_{\tau-}$ and $c_{\tau+}$, respectively.⁶

An important advantage of the in-fill asymptotic setting considered here is that it permits theoretically valid inference even if one has only a small number of events. This is reflected by the fact that the sample span $[0, T]$ is explicitly fixed under our in-fill asymptotic framework, so the collection $\mathcal{T}$ of event times is treated as a finite set. Indeed, our test can be applied even when $\mathcal{T}$ is a singleton $\{\tau\}$, corresponding to a single-case study. This feature is empirically desirable because important macro news announcements are infrequent (e.g., FOMC meetings are scheduled only eight times per year), and their effects can be highly heterogeneous depending on the prevailing macroeconomic and policy environment. Our approach overcomes this “small-sample” issue by exploiting the “large sample” of high-frequency price observations in the neighborhood of announcement times. We next proceed with the details.

B. The Rank Test

A natural way to carry out the rank test is to examine the number of zero eigenvalues of Δc_τ . Although the c process takes values as positive semidefinite matrices, its jump Δc_τ is only a symmetric matrix and may have negative eigenvalues. It is thus more convenient to consider the squared jump matrix

$$Q_\tau \equiv (\Delta c_\tau)^2,$$

which is positive semidefinite by construction. The rank test can then be implemented by examining whether the smallest $d - r$ eigenvalues of Q_τ are all identically zero. More precisely, let $(\lambda_{j,\tau})_{1 \leq j \leq d}$ denote the eigenvalues of Δc_τ , so each $\lambda_{j,\tau}^2$ is an eigenvalue of Q_τ , and order them as $\lambda_{1,\tau}^2 \geq \lambda_{2,\tau}^2 \geq \dots \geq \lambda_{d,\tau}^2$. We then set

$$S_r \equiv \sum_{\tau \in \mathcal{T}} S_{r,\tau}, \quad \text{where} \quad S_{r,\tau} \equiv \sum_{j=r+1}^d \lambda_{j,\tau}^2. \quad (5)$$

In restriction to $\Omega_{0,r}$ (resp. $\Omega_{a,r}$), we have $S_r = 0$ (resp. $S_r > 0$). Therefore, we can rely on the S_r variable to discriminate the null and alternative hypotheses.

⁶The consistent spot estimation relies on the length of the local windows, $k_n\Delta_n$, shrinking to zero asymptotically, as is standard in kernel-based nonparametric statistics. From their constructions, it is clear that the pre- and post-event spot estimators are in fact consistent estimators of the local averages $(k_n\Delta_n)^{-1} \int_{\tau-k_n\Delta_n}^{\tau} c_s ds$ and $(k_n\Delta_n)^{-1} \int_{\tau}^{\tau+k_n\Delta_n} c_s ds$, respectively. This latter interpretation is somewhat more robust in a finite-sample sense because it holds true even if one treats the window length $k_n\Delta_n$ as fixed, and our results may be interpreted accordingly. That being said, we note that, under commonly used term-structure models (see, e.g., Singleton, 2009), the average volatility is typically close to the spot values over short windows. Hence, the aforementioned distinction is often immaterial, and we adopt the standard “spot” interpretation in our subsequent discussion.

Our test statistic is simply the sample analog of S_r . Based on the spot estimators in equation (4), we estimate Q_τ using

$$\widehat{Q}_{n,\tau} \equiv (\hat{c}_{n,\tau+} - \hat{c}_{n,\tau-})^2.$$

Let $\hat{\lambda}_{n,1,\tau}^2 \geq \dots \geq \hat{\lambda}_{n,d,\tau}^2$ be the (ordered) eigenvalues of $\widehat{Q}_{n,\tau}$. The test statistic is then defined analogously to equation (5) as

$$\widehat{S}_{n,r} \equiv \sum_{\tau \in \mathcal{T}} \widehat{S}_{n,r,\tau}, \quad \text{where} \quad \widehat{S}_{n,r,\tau} \equiv \sum_{j=r+1}^d \hat{\lambda}_{n,j,\tau}^2. \quad (6)$$

To obtain the critical value, we need to characterize the asymptotic distribution of $\widehat{S}_{n,r}$ under the null hypothesis. Some additional notation is needed to describe the asymptotic distribution. Clearly, the sampling variability of the test statistic is solely driven by that of the spot covariance matrix estimators. To represent the latter, we consider $d \times d$ random matrices $(\zeta_{\tau-}, \zeta_{\tau+})_{\tau \in \mathcal{T}}$ that are $\mathcal{F}$ -conditionally independent, centered Gaussian, with the covariances between their components characterized by

$$\mathbb{E}[\zeta_{jk,\tau\pm} \zeta_{lm,\tau\pm} | \mathcal{F}] = c_{jl,\tau\pm} c_{km,\tau\pm} + c_{jm,\tau\pm} c_{kl,\tau\pm},$$

for $1 \leq j, k, l, m \leq d$.

It can be shown that

$$(k_n^{1/2}(\hat{c}_{n,\tau-} - c_{\tau-}), k_n^{1/2}(\hat{c}_{n,\tau+} - c_{\tau+}))_{\tau \in \mathcal{T}} \xrightarrow{\mathcal{L}-s} (\zeta_{\tau-}, \zeta_{\tau+})_{\tau \in \mathcal{T}}, \quad (7)$$

where $\xrightarrow{\mathcal{L}-s}$ denotes stable convergence in law. We also need to consider an eigenvalue decomposition of Δc_τ in the form

$$\Delta c_\tau = U_\tau \Lambda_\tau U_\tau^\top, \quad (8)$$

where Λ_τ is a diagonal matrix collecting the ordered (in magnitude) eigenvalues $(\lambda_{j,\tau})_{1 \leq j \leq d}$, and U_τ is an orthogonal matrix consisting of the corresponding eigenvectors. Finally, we partition $U_\tau = [\Gamma_\tau; V_\tau]$ such that Γ_τ and V_τ contain r and $d - r$ columns, respectively. Theorem 1 establishes the asymptotic distribution of the test statistic $\widehat{S}_{n,r}$ under the null hypothesis.

Theorem 1. Suppose that (i) assumption 1 holds and (ii) $k_n \asymp \Delta_n^{-\rho}$ and $u_n \asymp \Delta_n^\varpi$ such that

$$0 < \rho < \left(\frac{2}{\gamma} - 1\right) \wedge \frac{1}{2}, \quad \frac{\rho}{2(2-\gamma)} < \varpi < \frac{1}{2}. \quad (9)$$

Then, in restriction to the null hypothesis $\Omega_{0,r}$,

$$k_n \widehat{S}_{n,r} \xrightarrow{\mathcal{L}-s} \xi_r \equiv \sum_{\tau \in \mathcal{T}} \|V_\tau^\top (\zeta_{\tau+} - \zeta_{\tau-}) V_\tau\|^2. \quad (10)$$

COMMENTS. (i) The limiting variable ξ_r described in equation (10) has a nonstandard distribution. It is instructive to elaborate on the structure of this variable. Since each $\zeta_{\tau+} - \zeta_{\tau-}$ matrix is an $\mathcal{F}$ -conditional centered Gaussian,

$V_\tau^\top (\zeta_{\tau+} - \zeta_{\tau-}) V_\tau$ is also a mixed centered Gaussian matrix. Consequently, $\|V_\tau^\top (\zeta_{\tau+} - \zeta_{\tau-}) V_\tau\|^2$ is the sum of (conditionally correlated) scaled chi-squared variables. These variables generally have different $\mathcal{F}$ -conditional distributions resulting from the heterogeneity of different events. They are $\mathcal{F}$ -conditionally independent across different τ 's, and ξ_r is their sum.

(ii) Although we are mainly interested in the rank of the jump matrix Δc_τ , our method can be easily adapted to study the rank of $c_{\tau+}$ and $c_{\tau-}$ as well. For example, if $c_{\tau+}$ is of interest, then one can simply replace the returns in the pre-event window with zeros, so that $\Delta c_\tau = c_{\tau+}$. Our theoretical results remain valid in this partially degenerate case.

The nonstandard limiting distribution described in Theorem 1 does not appear to be pivotalizable. We instead propose a bootstrap algorithm for estimating the limiting distribution, particularly its quantiles. Analogous to equation (8), we perform the following eigenvalue decomposition for each estimated jump:

$$\hat{c}_{n,\tau+} - \hat{c}_{n,\tau-} = \widehat{U}_{n,\tau} \widehat{\Lambda}_{n,\tau} \widehat{U}_{n,\tau}^\top,$$

where the eigenvalues and eigenvectors are ordered in magnitude. We then set $\widehat{V}_{n,\tau}$ to be the $d \times (d - r)$ matrix consisting of the last $d - r$ columns of $\widehat{U}_{n,\tau}$, which consistently estimates V_τ up to rotation. The bootstrap algorithm for computing the critical value at significance level $\alpha \in (0, 1)$ is detailed below.

Algorithm 1 (bootstrap critical value for rank test)

Step 1: For each $\tau \in \mathcal{T}$, generate i.i.d. draws $(R_{n,i}^*)_{i \in \mathcal{I}_{n,\tau-}}$ and $(R_{n,i}^*)_{i \in \mathcal{I}_{n,\tau+}}$ from $(R_{n,i})_{i \in \mathcal{I}_{n,\tau-}}$ and $(R_{n,i})_{i \in \mathcal{I}_{n,\tau+}}$, respectively.
 Step 2: Compute $\hat{c}_{n,\tau\pm}^*$ in the same way as $\hat{c}_{n,\tau\pm}$, except that the original data $R_{n,i}$ is replaced with $R_{n,i}^*$. Set $\zeta_{n,\tau\pm}^* = k_n^{1/2}(\hat{c}_{n,\tau\pm}^* - \hat{c}_{n,\tau\pm})$.
 Step 3: Repeat steps 1 and 2 many times. Set the critical value $cv_{n,\alpha}$ as the $1 - \alpha$ quantile of $\sum_{\tau \in \mathcal{T}} \|\widehat{V}_{n,\tau}^\top (\zeta_{n,\tau+}^* - \zeta_{n,\tau-}^*) \widehat{V}_{n,\tau}\|^2$ at significance level α . $\square$

The intuition underlying the bootstrap algorithm is as follows. Step 1 of the algorithm implements i.i.d. resampling of the return vectors in each local window. The “localized” resampling is needed to address the data heterogeneity across different estimation windows (i.e., $\mathcal{I}_{n,\tau\pm}$). Note that even within each local window, the return observations are *not* assumed to be actually i.i.d. In fact, since the returns depend on the stochastic volatility, they are generally heterogeneous and serially highly dependent. Nevertheless, return observations within each local window are approximately conditionally i.i.d., which is the intuition why the local i.i.d. resampling is valid. The $\zeta_{n,\tau\pm}^*$ variables defined in step 2 are exactly the bootstrap analogs of $k_n^{1/2}(\hat{c}_{n,\tau\pm} - c_{\tau\pm})$, and their conditional distribution estimates that of the $\zeta_{\tau\pm}$ limiting variables. The $\|\widehat{V}_{n,\tau}^\top (\zeta_{n,\tau+}^* - \zeta_{n,\tau-}^*) \widehat{V}_{n,\tau}\|^2$ variable in step 3

clearly mirrors the limiting variable $\|V_\tau^\top (\zeta_{\tau+} - \zeta_{\tau-}) V_\tau\|^2$ in Theorem 1. The validity of the bootstrap algorithm and the asymptotic properties of our rank test are described in Theorem 2 below.

Theorem 2. *Suppose that the conditions in Theorem 1 hold and $\varpi > 1/(4 - \gamma)$. Then (a) the conditional distribution function of $\sum_{\tau \in \mathcal{T}} \|\widehat{V}_{n,\tau}^\top (\zeta_{n,\tau+}^* - \zeta_{n,\tau-}^*) \widehat{V}_{n,\tau}\|^2$ given data converges in probability to the $\mathcal{F}$ -conditional distribution of ξ_r under the uniform metric and (b) the test associated with the critical region $\{k_n \widehat{S}_{n,r} \geq cv_{n,\alpha}\}$ has asymptotic level α under the null hypothesis and asymptotic power 1 under the alternative hypothesis, that is,*

$$\mathbb{P}(k_n \widehat{S}_{n,r} \geq cv_{n,\alpha} | \Omega_{0,r}) \rightarrow \alpha, \quad \mathbb{P}(k_n \widehat{S}_{n,r} \geq cv_{n,\alpha} | \Omega_{a,r}) \rightarrow 1.$$

We close this section with two practical remarks about the purposed test. The first concerns the threshold u_n , which is a statistical device for eliminating price jumps in the volatility estimation. In practice, recognizing that the price jump is very likely to occur either at or very close to the announcement time, one may exploit this prior knowledge and eliminate price jumps by simply removing a few returns near the announcement time, without performing any additional “statistical” truncation (i.e., use $u_n = \infty$). This is our recommended method that is adopted in the subsequent numerical work.⁷

Second, we note that null hypotheses with different r values are ordered and the alternative hypotheses are nested. As illustrated in Robin and Smith (2000) (see section 5), this structure naturally suggests that tests with different r values may be interpreted as a sequential testing procedure: we can implement the test with increasing r values and stop at the first nonrejection. Since the proposed test is consistent, this procedure will rule out all r values strictly less than $\text{Rank}(\Delta c_\tau)$ with probability approaching 1. In particular, if Δc_τ is of full rank, the sequential procedure will reject all reduced-rank null hypotheses and, hence, provide a consistent estimator for the true rank. On the other hand, when Δc_τ is of reduced rank, the true null hypothesis at stage $r = \text{Rank}(\Delta c_\tau)$ will be rejected with asymptotic probability α , reflecting the type I error of the test.

III. Monte Carlo Simulations

In this section we examine the performance of the proposed test in a Monte Carlo experiment. The unit of time is one day. Let $(W_{j,t})_{1 \leq j \leq 3}$ and $(B_{j,t})_{1 \leq j \leq 3}$ be independent standard Brownian motions. We consider three assets and simulate their log returns according to

$$dX_t = \begin{pmatrix} 1 & 0.5 & 0.5 \\ 0.5 & 1 & 0.5 \\ 0.5 & 0.5 & 1 \end{pmatrix} df_t,$$

where f is a three-dimensional factor process with its j th component satisfying

$$df_{j,t} = \sigma_{j,t} dW_{j,t}, \quad 1 \leq j \leq 3.$$

The volatility processes are simulated according to the following stochastic differential equations:

$$\begin{aligned} d\sigma_{1,t}^2 &= 0.6930(0.4068 - \sigma_{1,t}^2)dt \\ &\quad + 0.7023\sqrt{\sigma_{1,t}^2}(\rho dW_{1,t} + \sqrt{1 - \rho^2}dB_{1,t}) \\ &\quad + J_\sigma \sigma_{1,t}^2 - 1_{\{t=\tau\}}, \\ d\sigma_{2,t}^2 &= 0.0128(0.4068 - \sigma_{2,t}^2)dt \\ &\quad + 0.0954\sqrt{\sigma_{2,t}^2}(\rho dW_{2,t} + \sqrt{1 - \rho^2}dB_{2,t}) \\ &\quad + \phi J_\sigma \sigma_{2,t}^2 - 1_{\{t=\tau\}}, \\ d\sigma_{3,t}^2 &= 0.0128(0.4068 - \sigma_{3,t}^2)dt \\ &\quad + 0.0954\sqrt{\sigma_{3,t}^2}(\rho dW_{3,t} + \sqrt{1 - \rho^2}dB_{3,t}), \end{aligned}$$

where the parameter values are calibrated according to Bollerslev and Todorov (2011), and we set $\rho = -0.7$ to capture the well-documented negative correlation between price and volatility shocks (i.e., the “leverage” effect). The first volatility factor $\sigma_{1,t}^2$ is quickly mean-reverting with a half-life of one day, and it jumps at the announcement time τ with relative jump size $J_\sigma \sim \text{Exp}(7)$, where the mean of the exponential distribution is calibrated to the empirical estimates of Bollerslev et al. (2018). The other two volatility factors, $\sigma_{2,t}^2$ and $\sigma_{3,t}^2$, are highly persistent with a half-life of 2.5 months. The $\sigma_{2,t}^2$ process also jumps at time τ when $\phi \neq 0$. The ϕ parameter conveniently controls the relative magnitude of the jump in $\sigma_{2,t}^2$ with respect to that in $\sigma_{1,t}^2$. These continuous-time processes are simulated using an Euler scheme on a one-second mesh. The observed returns actually used in the calculations are resampled at $\Delta_n = 1, 3$, and 5 minute intervals.

This data-generating process permits a simple characterization of the factor structure of spot covariance matrix jumps. Let Δc_τ and $\Delta c_{f,\tau}$ denote the jumps of the spot covariance matrices of X and f , respectively. It is easy to see that Δc_τ and $\Delta c_{f,\tau}$ share the same rank, and the latter is diagonal with elements $J_\sigma \sigma_{1,\tau-}^2$, $\phi J_\sigma \sigma_{2,\tau-}^2$, and 0. Therefore,

$$\text{Rank}(\Delta c_\tau) = \begin{cases} 1 & \text{when } \phi = 0, \\ 2 & \text{when } \phi > 0. \end{cases}$$

We can then impose the null hypothesis in two ways: when $\phi = 0$, the null hypothesis corresponds to $r = 1$, and when $\phi > 0$, the null hypothesis corresponds to $r = 2$. Moreover,

⁷More generally, one can use the u_n -based truncation to further guard against price jumps occurring within the local estimation windows. For example, when jumps have finite variation, a theoretically valid adaptive choice is $u_n = 3 \times \widehat{\sigma} \times \Delta_n^{1/3}$, with $\widehat{\sigma}$ being some preliminary estimate of volatility. That said, for frequencies of several minutes, this choice very rarely leads to any additional truncation after the returns near the announcement time are removed.

TABLE 1.—MONTE CARLO REJECTION RATES UNDER NULL HYPOTHESIS

	60-Minute Window			90-Minute Window		
	$\Delta_n = 1$	$\Delta_n = 3$	$\Delta_n = 5$	$\Delta_n = 1$	$\Delta_n = 3$	$\Delta_n = 5$
Panel A: Case $r = 1$						
$\phi = 0$	0.050	0.035	0.036	0.053	0.036	0.034
Panel B: Case $r = 2$						
$\phi = 0.25$	0.037	0.037	0.043	0.040	0.034	0.038
$\phi = 0.50$	0.044	0.043	0.048	0.047	0.040	0.043
$\phi = 0.75$	0.048	0.045	0.051	0.050	0.044	0.046
$\phi = 1.00$	0.051	0.047	0.053	0.051	0.047	0.048

This table reports the finite-sample rejection rates of the 5% level rank test under the null hypothesis for various data generating processes. Panel A reports results for the null hypothesis with $r = 1$ that is imposed by setting $\phi = 0$. Panel B reports results for the null hypothesis with $r = 2$ that is imposed by setting $\phi > 0$. The sampling interval Δ_n ranges from 1 minute to 5 minutes. The length of the local estimation window is fixed in calendar time to be 60 or 90 minutes, corresponding to $k_n \in \{60, 20, 12\}$ and $k_n \in \{90, 30, 18\}$, respectively, for the three different sampling frequencies. The volatility of microstructure noise is set to be $\eta = 0.0156$.

specifications with $\phi > 0$ also provide a range of alternative hypotheses versus the $r = 1$ null. To trace out the power function, we consider a range of ϕ values in $\{0, 0.05, 0.1, \dots, 1\}$ and expect to see higher finite-sample power associated with larger ϕ .

Finally, we allow for the presence of microstructure noise in the data. That is, instead of the “efficient price” X_t , the econometrician observes noise-contaminated price Y_t given by

$$Y_{i\Delta_n}^{(j)} = X_{i\Delta_n}^{(j)} + \eta \times \epsilon_i^{(j)}, \quad j = 1, 2, 3, \quad i \geq 0, \quad (11)$$

where the constant η determines the noise scale and $\{\epsilon_i^{(j)}\}_{i \geq 0}$ are i.i.d. standard normal error terms that are independent across assets. Note that our asymptotic theory is designed for the no-noise case with $\eta = 0$. That being said, in the more realistic setting with microstructure noise, the common practice in the high-frequency econometrics literature is to mitigate the effect of noise by using sparsely sampled data, which we adopt throughout our numerical work. In this regard we use the noisy setting as a robustness check. We calibrate $\eta = 0.0156$ according to the noise-to-signal ratio in our empirical data, determined using the two-scale method of Zhang, Mykland, and Ait-Sahalia (2005). Consistent with aforementioned conventional wisdom, we find that the proposed test is indeed robust to the presence of noise when applied to sparsely sampled data. Therefore, for brevity, we present results only in the (more challenging) case with noisy data in the subsequent discussion.⁸

Below we report rejection frequencies of tests based on two local window specifications. Specifically, we fix the width of each local window to be either 60 minutes or 90 minutes, so k_n takes different values for different sampling frequencies. For example, we have $k_n = 60$ or 90 when $\Delta_n = 1$ minute, and $k_n = 12$ or 18 when $\Delta_n = 5$ minutes. This setup makes numerical results for different sampling frequencies more comparable. The bootstrap algorithm is implemented using 1,000 repetitions. For all tests below, we fix the significance level at 5% and compute finite-

sample rejection frequencies based on 100,000 Monte Carlo replications.

Table 1 reports the test’s rejection rates under the null hypothesis. Panel A reports results for the $r = 1$ null hypothesis imposed by setting $\phi = 0$. We see that the test controls size well across different sampling frequencies and window sizes. The rejection rates are slightly lower than the 5% nominal level, suggesting that the test is somewhat conservative in finite samples. Panel B reports the rejection rates for the $r = 2$ null hypothesis for a range of positive ϕ values. Again, we see that the test controls size quite well across the board.

We next turn to the power analysis. As mentioned above, alternative hypothesis with respect to the $r = 1$ null hypothesis can be imposed by setting $\phi > 0$. The value of ϕ measures the “distance” between the null and the alternative hypotheses. Figure 1 plots the test’s rejection rates as functions of ϕ separately for three sampling frequencies and two local window sizes. Consistent with table 1, we see that the test controls size well in all settings under the null hypothesis (i.e., $\phi = 0$). As ϕ increases, the power of the test clearly increases as predicted by the asymptotic theory. Other things being equal, the test rejects more often for more frequently sampled data (i.e., smaller Δ_n). In addition, the rejection rates for the 90-minute window appear to be higher than those for the 60-minute window, which can be explained by the former’s larger sample size. As a caveat, we note that this finding does not automatically imply one should always use larger local windows, which would be subject to more non-parametric estimation bias stemming from the time-variation in stochastic volatility.


Overall, these simulation results show that the proposed test has excellent size control and adequate power for a range of commonly used sampling frequencies. The results also appear to be stable for different choices of local window sizes. These findings support our asymptotic theory developed above and suggest that the test can be reliably used in empirical work, to which we now turn.

IV. FOMC Announcements and Monetary Policy Shocks

To demonstrate the usefulness of the proposed rank test, we conduct a formal econometric analysis to study the

⁸In the no-noise case, we find that the test has similar size properties and, as expected, its power is higher than that in the noisy case.

FIGURE 1.—MONTE CARLO REJECTION RATE OF RANK TEST



The figure plots the rejection frequencies of the $r = 1$ null hypothesis at 5% nominal level for sampling frequencies at 1, 3, and 5 minutes (solid, dash-dot, and dashed). The test is implemented using 60-minute and 90-minute local windows. The shaded area signifies the 5% nominal level. The null and alternative hypotheses correspond to $\phi = 0$ and $\phi > 0$, respectively. The critical value for each test is computed using 1,000 bootstrap repetitions. The rejection rates are computed for $\phi \in \{0, 0.05, \dots, 1\}$ based on 100,000 Monte Carlo trials.

dimensionality of monetary policy shocks triggered by scheduled FOMC announcements. A large literature in macroeconomics studies the effect of monetary policy shocks induced by FOMC announcements. Early notable contributions include Kuttner (2001), Cochrane and Piazzesi (2002), Rigobon and Sack (2004), and Bernanke and Kuttner (2005). More recently, Nakamura and Steinsson (2018) rely on intraday high-frequency data on short-term interest rate futures to measure policy shocks and then use these measures to test for monetary nonneutrality. Although the target federal funds rate is conventionally the main policy tool of the central bank, the Federal Reserve has relied heavily on unconventional monetary policy tools such as forward guidance and quantitative easing (QE) during the Great Recession and its aftermath (Kuttner, 2018; Bernanke, 2020). The plurality of policy tools naturally suggests that the notion of a “monetary policy shock” is a multidimensional concept in the era of unconventional monetary policy. We apply the rank test to formally test hypotheses regarding the shock’s dimensionality.

Our data consist of intraday transaction prices for four interest rate futures contracts including 12-month Eurodollars, 2-year and 10-year Treasury notes, and 30-year Treasury bonds. The sample period is from October 2, 2001, to August 13, 2018. The data are obtained from Tick Data. We rely on Bloomberg’s Economic Calendar to pinpoint the exact announcement times for each of the 135 scheduled FOMC announcements that occurred during regular trading hours over our sample period.

The rank test is implemented as follows. To help mitigate the effect of market microstructure “noise,” we follow standard practice in the literature to sparsely sample the data at a 3-minute sampling frequency (see, e.g., the discussion in


Zhang et al., 2005).⁹ We take the block size $k_n = 30$, corresponding to a 90-minute window for spot estimation as we have done in the simulation study above. As a conservative way to remove announcement-induced price jumps, we exclude the returns from the five minutes immediately before and the five minutes immediately after each announcement time, so that there is a 10-minute gap between pre-event and post-event windows.¹⁰ Critical values are computed using Algorithm 1 based on 100,000 bootstrap resamples to minimize the effect of random resampling on the testing results.

The most important episode in our sample period arguably is the global financial crisis and the Great Recession of 2007–2009, which marked the beginning of the recent era of unconventional monetary policy. From September 16, 2007, to December 16, 2008, the Federal Reserve gradually lowered the target federal funds rate from the peak level of 5.25% to the $[0, 0.25\%]$ range, namely, the zero lower bound. The rate stayed at the lower bound till the end of 2015, and gradually increased to the $[1.75\%, 2\%]$ range by the end of our sample. During the crisis and the subsequent recovery period, the Federal Reserve was impelled to rely on unconventional monetary policy tools such as

⁹Our choice of the 3-minute sampling frequency is justified by a preliminary analysis based on the volatility signature plot (see figure S.1 in the online supplemental appendix). We find that realized volatility estimates are insensitive to the change of sampling frequency when $\Delta_n \geq 3$ minutes, but they appear to carry large positive bias at higher sampling frequencies. Also recall that with the level of noise calibrated to the data, our Monte Carlo experiment indicates that the 3-minute sampling frequency is sparse enough to sufficiently mitigate the effect of noise on the proposed test.

¹⁰Our subsequent empirical results are not sensitive to the choice of the 10-minute gap between the pre- and post-event windows. In results not presented here, we also conducted the empirical analysis with 8-minute and 12-minute gaps as robustness checks. Our main empirical findings are qualitatively unaltered with respect to these changes.

FIGURE 2.—EIGENVALUES OF SPOT COVARIANCE JUMP MATRIX AT SELECTED ANNOUNCEMENT TIMES



The figure plots the relative magnitudes of the eigenvalues of spot covariance jump matrices for four FOMC announcements. The relative magnitude is computed as the absolute value of each eigenvalue normalized by the sum of the absolute values of all eigenvalues. The two announcements displayed on the left column are examples for forward guidance (January 25, 2012) and quantitative easing (March 18, 2009) as discussed in Rudebusch (2018), respectively, and the right column corresponds to the two subsequent announcements used for comparison.

forward guidance and QE to steer the macroeconomy. It is thus economically important to examine whether, and to which extent, the multifaceted policy tools are associated with multidimensional policy shocks, which should manifest in the factor structure of the spot covariance jump matrix of bond futures.

To have a clear narrative relating unconventional monetary policy and announcement-induced shocks, we start with a case study based on two specific FOMC announcements that involve forward guidance and QE. The first announcement, which occurred on January 25, 2012, provided an explicit forward guidance that “economic conditions . . . are likely to warrant exceptionally low levels for the federal funds rate at least through late 2014.” The second example pertains to the announcement on March 18, 2009, when the Federal Reserve stated that it would be purchasing an additional \$300 billion of Treasury bonds as part of its QE1 program. These two examples are chosen solely based on the recent review article by Rudebusch (2018) on unconventional monetary policy, which also provides additional policy background.

We test the ranks of the spot covariance jump matrices for these two announcements separately. For the announcement on January 25, 2012, we reject the null hypothesis of $r = 1$ at the 5% significance level but do not reject $r = 2$ or $r = 3$. This result thus suggests that the announcement was associated with two distinct sources of policy shocks. According to Rudebusch (2018), during the early recovery period (2010–2011) after the Great Recession, market investors expected a quick policy rate liftoff. Contrary to this anticipation, the


Federal Reserve ruled out the possibility of raising the rate not only in short term, but also throughout a three-year horizon, which may explain the finding of a two-dimensional policy shock. The QE announcement on March 18, 2009, tells a similar story.¹¹ We reject the null hypothesis of $r = 1$ and $r = 2$ both at the 1% significance level, and reject the null hypothesis of $r = 3$ at the 5% significance level, indicating the presence of an even higher dimensional policy shock that had a distinct impact on different sections of the yield curve.

The two examples discussed above are intentionally presented here to highlight how a multidimensional policy shock, backed with a clear economic narrative, may be revealed by the proposed rank test. Because of their unique nature, these announcements (particularly the one on March 18, 2009) are not meant to be thought of as “typical” policy announcements. It is thus instructive to contrast them with others. To prevent confounding factors from complicating the comparison, we compare these two events with the two FOMC announcements that occurred immediately after them (and hence under similar macroeconomic conditions), on March 13, 2012, and April 29, 2009, respectively. For these comparison events, we do not reject the $r = 1$ null hypothesis at any conventional significance levels and, hence, find evidence for a one-dimensional policy shock.

To gain further insight into these testing results, in figure 2 we plot the relative magnitude of the covariance jump

¹¹This announcement also contains a forward guidance; see the discussion in Kuttner (2018), p. 130.

FIGURE 3.—EIGENVALUES OF SPOT COVARIANCE MATRIX AT SELECTED ANNOUNCEMENT TIMES



The figure plots the relative magnitudes of the eigenvalues of spot covariance matrices before and after two FOMC announcements, along with those of the covariance jump matrices. The relative magnitude is computed as the absolute value of each eigenvalue normalized by the sum of the absolute values of all eigenvalues.

matrix's eigenvalues (i.e., $|\hat{\lambda}_j| / \sum_{k=1}^4 |\hat{\lambda}_k|$) for these events. The top panel corresponds to the January 25, 2012, announcement and its comparison event on March 13, 2012. For the former, the contribution of the second largest (in magnitude) eigenvalue is nontrivial, which contrasts sharply with the notable one-factor pattern for the March 13, 2012, event shown on the top-right panel. These plots are consistent with the formal testing results that suggest a two-factor structure for the former, and a one-factor structure for the latter. The bottom panel of figure 2 tells a similar story for the announcements on March 18, 2009, and April 29, 2009.

The sharp one-factor structure in the covariance jump matrix for the March 13, 2012, and April 29, 2009, events seen in figure 2 is remarkable, because it suggests a single source of policy shock to the four interest rate futures contracts. It is important to note that this finding speaks to the covariance jump matrix at the announcement time, rather than the assets' spot covariance matrix itself. Indeed, a reduced rank in the covariance matrix of the asset returns would imply the redundancy of certain assets, which is highly unlikely for the four bond futures studied here. To make this point more concrete, in figure 3 we compare the normalized eigenvalues of $c_{\tau-}$, $c_{\tau+}$, and Δc_{τ} for the announcements on March 13, 2012, and April 29, 2009. From the figure, we see that the low-rank structure is indeed much more pronounced in the jump matrix than the pre-event and post-event covariance matrices. Formally, the $r = 3$ null hypotheses are strongly rejected for both $c_{\tau-}$ and $c_{\tau+}$, suggesting that they have full ranks as expected.

TABLE 2.—EIGENVECTORS OF SPOT COVARIANCE JUMP MATRIX

Maturity	First Eigenvector		Second Eigenvector	
	January 25, 2012	March 18, 2009	January 25, 2012	March 18, 2009
1 year	0.0108	0.0188	−0.0215	−0.0648
2 year	−0.0035	0.0271	−0.0852	−0.0761
10 years	0.3888	0.3920	−0.9178	−0.9140
30 years	0.9212	0.9194	0.3873	0.3932

The table reports the first two eigenvectors of spot covariance jump matrices for FOMC announcements on January 25, 2012, and March 18, 2009.

It is interesting to note that the FOMC meetings on January 25, 2012, and March 18, 2009, for which we find strong evidence for a multifactor structure of the generated policy shock, are both followed by ones for which our test suggests the policy shocks were one-dimensional. This shows that there can be significant time variation in the nature of the policy shocks triggered by the FOMC announcements and further illustrates the benefits of the developed testing procedures based on event-by-event analysis.

To gain further insight about the change in the yield curve triggered by the policy shocks, we present the first two eigenvectors of Δc_{τ} in table 2 for the announcements on January 25, 2012, and March 18, 2009. These eigenvectors have remarkably similar structure across the two dates.¹² The first

¹²The value of the first eigenvector is in fact quite stable across all announcements in our sample. The second eigenvector is more difficult to interpret across the board because it is not uniquely identified for announcements whose policy shock is only one-dimensional.

TABLE 3.—REJECTION DECISION OF RANK TESTS FOR FOMC ANNOUNCEMENTS

Null	2001–2018	Subsample Periods					
		2001–2003	2004–2006	2007–2009	2010–2012	2013–2015	2016–2018
$r = 1$	Yes	No	No	Yes	No	Yes	Yes
$r = 2$	Yes	No	No	Yes	No	No	No
$r = 3$	No	No	No	No	No	No	No

We report the rejection decisions of the rank tests for the spot covariance jump matrix for the full sample (column 1) and 3-year subsamples (columns 2–7). The rows correspond to null hypotheses with rank $r = 1, 2$, and 3, respectively. The test is implemented for $\Delta_n = 3$ minutes and $k_n = 30$. The significance level is fixed at 1%.

eigenvector appears to collect loadings of a “level” factor that moves all yields in the same direction. Not surprisingly, the weights assigned to the 12-month Eurodollar and the 2-year Treasury note are close to zero, whereas most of the weight is assigned to the 30-year Treasury bond. The ranking of these weights are in line with the volatilities of the four assets, with the 12-month Eurodollar contract and the 30-year Treasury bond being the least and most volatile, respectively. The second eigenvector of Δc_τ reveals a “slope” factor as it moves the yields in different directions. Again, we see that the weights assigned to the 12-month Eurodollar contract and the 2-year Treasury note are small in magnitude. In contrast to the first eigenvector of Δc_τ , the second eigenvector puts most weight on the 10-year Treasury note, suggesting that the corresponding policy shock is particularly relevant for the medium term.¹³

To draw more general conclusions regarding lower frequency variation in the nature of the policy shocks, we next turn to aggregated testing results by pooling information from all announcements in our sample. In view of the evolution of monetary policy during this sample period, we also divide the 18-year sample period between 2001 and 2018 into six 3-year subsamples. The two precrisis subsamples spanning 2001–2003 and 2004–2006 are mainly subject to conventional monetary policy. The 2007–2009 period witnessed the financial crisis and the Great Recession, during which the Federal Reserve employed both conventional and unconventional policy tools. The next two subsamples cover the 2010–2015 period, when the target federal funds rate was maintained at the zero lower bound, and the rate was gradually raised during the last subsample between 2016 and 2018. We implement the rank test for the six subsamples separately. To guard against concerns pertaining to “multiple

testing,” we use a relatively stringent 1% significance level for all tests. Formally, because of the asymptotic conditional independence among test statistics formed on the nonoverlapping subsamples, we may equivalently interpret the testing results jointly across all six subsamples at significance level $1 - (99\%)^6 \approx 5.85\%$.

Table 3 reports whether the null hypotheses with different ranks (i.e., r) are rejected.¹⁴ From the full-sample results displayed in column 1, we see that both $r = 1$ and $r = 2$ null hypotheses are rejected at the 1% significance level, whereas the $r = 3$ hypothesis is not rejected. This echoes our earlier finding for the QE announcement on March 18, 2009, further confirming that some announcements within this sample period triggered multidimensional policy shocks.

The subsample analysis reveals richer and economically more interesting information. From columns 2 and 3, we see that none of the null hypotheses (particularly including $r = 1$) is rejected for the first couple of 3-year subsamples. This suggests that during the precrisis 2001–2006 period, when the conventional monetary policy was largely in force, monetary policy shocks delivered by FOMC announcements is one-dimensional, which is consistent with the notion that the Federal Reserve achieves its macroeconomic objectives mainly by altering the short-term interest rate. However, the story is drastically different during the 2007–2009 crisis period. As shown in column 4 of the table, both the $r = 1$ and $r = 2$ null hypotheses are rejected at the 1% significance level, suggesting formally that some FOMC announcements had triggered multi-dimensional policy shocks. The evidence is in line with the policy environment during that period when the Federal Reserve deployed a complex mix of policy tools including the reduction of the target federal funds rate gradually to the zero lower bound, forward guidance about future rate policy, large-scale purchases of both mortgage-backed securities and US Treasury securities, and policies aimed at stabilizing dysfunctional financial markets (Bernanke, 2020).

In the postcrisis period, we again find a low-dimensional structure for policy shocks. Specifically, during the 2010–2012 subsample, we do not reject the $r = 1$ null hypothesis at the 1% significance level. This likely reflects the limited policy tools at the central bank’s disposal during the early recovery phase after the crisis. Interestingly, for the

¹³As mentioned above, the two announcements discussed here were picked on the basis of the review article of Rudebusch (2018). We also implemented a similar analysis on the FOMC announcement days during the QE1 (December 16, 2008, January 28, 2009, and March 18, 2009) and QE2 (August 10, 2010, and September 21, 2010) episodes. These additional results are qualitatively similar to those presented in this paper. Specifically, for all of the FOMC announcements in QE1 and QE2, we find evidence for multidimensional policy shocks. Moreover, the decomposition of the eigenvectors of Δc_τ on these dates is similar to the ones reported in table 2. In particular, we do not see a qualitative difference in the structure of Δc_τ for the events in QE1 and QE2. Vissing Jorgensen and Krishnamurthy (2011) argue for the existence of different economic channels of the QE1 and QE2 programs: QE1 contains liquidity and risk premia channels that are not operative during QE2. These different economic channels may be further analyzed (identified) by the inclusion of additional fixed-income securities, such as mortgage-backed securities and fixed-income derivatives, to the ones considered in our analysis.

¹⁴The null hypothesis with $r = 0$ is overwhelmingly rejected in the data, suggesting the presence of volatility jumps; this is consistent with prior findings of Bollerslev et al. (2018).

TABLE 4.—REJECTION DECISION OF RANK TESTS FOR NON-FOMC ANNOUNCEMENTS

Null	2001–2018	Subsample Periods					
		2001–2003	2004–2006	2007–2009	2010–2012	2013–2015	2016–2018
$r = 1$	No	No	No	No	No	No	No
$r = 2$	No	No	No	No	No	No	No
$r = 3$	No	No	No	No	No	No	No

We report the rejection decisions of the rank tests for the spot covariance jump matrix for the full sample (column 1) and 3-year subsamples (columns 2–7) for the inflation, initial jobless claims, and the nonfarm payroll announcements. The rows correspond to null hypotheses with rank $r = 1, 2$, and 3 , respectively. The test is implemented for $\Delta_n = 3$ minutes and $k_n = 30$. The significance level is fixed at 1%.

2013–2015 and 2016–2018 subsamples, we reject the $r = 1$ null hypothesis and, hence, find evidence for two-dimensional policy shocks. Relative to the 2010–2012 period, the higher rank in these two later periods may be attributed to two important policies. The first pertains to a sequence of decisions on the QE program. After Bernanke’s congressional appearance in May 2013, hints that asset purchases might begin to slow led to a “taper tantrum” in bond markets. In its September 18, 2013, announcement, the FOMC clarified that it would not immediately slow down the pace of asset purchasing. The Federal Reserve then announced its decision to taper QE on December 18, 2013, and ended the asset purchasing program in October 2014. The other policy concerns raising the target federal funds rate. Following the ending of the QE program, the FOMC carefully indicated the possibility of a rate increase in its announcements throughout 2015, and eventually raised the rate by 0.25% on December 16, 2015. The target rate gradually rose to the [1.75%, 2%] range by the end of our sample period (August 13, 2018).

The evidence above shows that the jump in spot covariance matrix of bond prices triggered by the FOMC announcement is generally of reduced rank, even in the more recent period associated with unconventional monetary policy. It is thus possible to carry out heteroskedasticity-based structural identification in the spirit of Rigobon (2003) using changes in high-frequency volatility estimates. Note that this empirical regularity is mainly driven by the elevated trading activity in the (short) post-announcement window, when the new information is incorporated into the asset prices, rather than because of any “abnormal” volatility dynamics in the pre-announcement window. To illustrate this point concretely and further buttress the findings in table 3, we replace the pre-announcement “control” $c_{\tau-}$ with a post-announcement version $c_{(\tau+1)+}$ (i.e., the next-day spot covariance matrix at the same time of day as $c_{\tau+}$) and test the rank of $c_{\tau+} - c_{(\tau+1)+}$. The resulting rejection decisions are exactly the same as those reported in table 3, confirming that the findings are robust to the selection of “control,” or nonannouncement, time windows.

The evidence for increased dimension of the policy shock triggered by FOMC announcements following the start of the unconventional monetary policy by the Federal Reserve can be contrasted with the reaction of the yield curve following the other types of prescheduled macroeconomic announcements. More specifically, we look at the CPI, the initial jobless claims (IC), and the nonfarm payroll (NFP)

announcements. All of them are issued at 8:30 EST, and we find altogether 1,248 of them during our sample period. These events are typically associated with nontrivial change in the volatility matrix of the four fixed-income contracts that we analyze here. We conduct our rank test for Δc_{τ} triggered by these announcements by pooling the data in the same subperiods as those reported in table 3. The results from the test are reported in table 4 and are in sharp contrast to those for the FOMC announcements presented in table 3. Mainly, throughout our sample period, we find no evidence that the jump Δc_{τ} triggered by the CPI, IC, and NFP announcements is multidimensional. This reinforces the observation that the evidence for higher rank of the generated policy shock triggered by FOMC announcements is to do with the change of monetary policy around the period of the financial crisis rather than because of a change in the overall economic environment.

In summary, the multidimensionality of policy shocks uncovered by our rank test is closely in line with the underlying economic narrative and demonstrates the empirical usefulness of the proposed econometric method. The finding that FOMC announcements during the precrisis period triggered one-dimensional monetary shocks confirms the view that conventional monetary policy has its direct impact mainly on the short end of the yield curve. Meanwhile, the rank test also formally reveals the multifaceted nature of the Federal Reserve’s policy shocks during the Great Recession and its aftermath. Our findings also have useful implications for studies on the identification and estimation of the effect of monetary shocks on asset prices and macroeconomic quantities (see, e.g., Rigobon & Sack, 2004; Bernanke & Kuttner, 2005; Nakamura & Steinsson, 2018).

V. Conclusion

Motivated by the recent literature on the high-frequency identification of policy shocks, we propose a test for the rank of spot covariance jump matrix of asset prices at macro news announcement times. The test statistic is formed using the eigenvalues of the covariance jump matrix estimated nonparametrically from high-frequency asset returns in local windows around announcement times. The test statistic has a nonstandard limiting distribution. We propose an easy-to-implement bootstrap algorithm to compute the critical value, and justify its asymptotic validity. Empirically, we apply the method to test the dimensionality of monetary policy shocks triggered by FOMC announcements using

intraday transaction data for interest rate futures contracts in a sample from 2001 to 2018. We document empirical evidence for one-dimensional policy shocks before the Great Recession, when conventional monetary policy largely prevailed, and find support for multidimensional policy shocks in the era of unconventional monetary policy, especially during the 2007–2009 financial crisis.

REFERENCES

- Ai, Hengjie, and Ravi Bansal, "Risk Preferences and the Macroeconomic Announcement Premium," *Econometrica* 86:4 (2018), 1383–1430. 10.3982/ECTA14607
- Aït-Sahalia, Yacine, and Jean Jacod, *High-Frequency Financial Econometrics* (Princeton, NJ: Princeton University Press, 2014).
- Beechey, Meredith J., and Jonathan H. Wright, "The High-Frequency Impact of News on Long-Term Yields and Forward Rates: Is It Real?" *Journal of Monetary Economics* 56:4 (2009), 535–544. 10.1016/j.jmoneco.2009.03.011
- Bernanke, Ben S., "The New Tools of Monetary Policy," *American Economic Review* 110:4 (2020), 943–983. 10.1257/aer.110.4.943
- Bernanke, Ben S., and Kenneth N. Kuttner, "What Explains the Stock Market's Reaction to Federal Reserve Policy?" *Journal of Finance* 60:3 (2005), 1221–1257. 10.1111/j.1540-6261.2005.00760.x
- Bollerslev, Tim, Jia Li, and Yuan Xue, "Volume, Volatility, and Public News Announcements," *Review of Economic Studies* 85:4 (2018), 2005–2041. 10.1093/restud/rdy003
- Bollerslev, Tim, and Viktor Todorov, "Estimation of Jump Tails," *Econometrica* 79:6 (2011), 1727–1783. 10.3982/ECTA9240
- Cheng, Xu, Winston Wei Dou, and Zhipeng Liao, "Macro-Finance Decoupling: Robust Evaluations of Macro Asset Pricing Models," *Econometrica* 90 (2021), 685–713. 10.3982/ECTA18506
- Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans, "Monetary Policy Shocks: What Have We Learned and to What End?" (pp. 65–148), in J. B. Taylor and M. Woodford (eds.), *Handbook of Macroeconomics*, Vol. 1 (Elsevier, 1999). 10.1016/S1574-0048(99)01005-8
- Cieslak, Anna, Adair Morse, and Annette Vissing-Jorgensen, "Stock Returns over the FOMC Cycle," *Journal of Finance* 74:5 (2019), 2201–2248. 10.1111/jofi.12818
- Cieslak, Anna, and Hao Pang, "Common Shocks in Stocks and Bonds," *Journal of Financial Economics* 142 (2021), 880–904. 10.1016/j.jfineco.2021.06.008
- Cieslak, Anna, and Andreas Schrimpf, "Non-Monetary News in Central Bank Communication," *Journal of International Economics* 118 (2019), 293–315.
- Cochrane, John H., and Monika Piazzesi, "The Fed and Interest Rates—A High-Frequency Identification," *American Economic Review* 92:2 (May 2002), 90–95. 10.1257/000282802320189069
- Comte, Fabienne, and Eric Renault, "Long Memory Continuous Time Models," *Journal of Econometrics* 73:1 (1996), 101–149. 10.1016/0304-4076(95)01735-6
- Cragg, John G., and Stephen G. Donald, "Inferring the Rank of a Matrix," *Journal of Econometrics* 76:1 (1997), 223–250. 10.1016/0304-4076(95)01790-9
- Foster, Dean P., and Dan B. Nelson, "Continuous Record Asymptotics for Rolling Sample Variance Estimators," *Econometrica* 64:1 (1996), 139–174. 10.2307/2171927
- Gill, Len, and Arthur Lewbel, "Testing the Rank and Definiteness of Estimated Matrices with Applications to Factor, State-Space and ARMA Models," *Journal of the American Statistical Association* 87:419 (1992), 766–776. 10.1080/01621459.1992.10475278
- Gurkaynak, Refet S., Brian P. Sack, and Eric T. Swanson, "Do Actions Speak Louder than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," FEDS working paper 2004-66 (2004).
- Hsiao, Cheng, "Identification" (pp. 223–283), in Zvi Griliches and Michael D. Intriligator (eds.), *Handbook of Econometrics*, Vol. 1 (Amsterdam: North-Holland/Elsevier, 1983). 10.1016/S1573-4412(83)01008-9
- Jacod, Jean, and Philip Protter, *Discretization of Processes* (Berlin: Springer, 2012).
- Kandel, Eugene, and Neil D. Pearson, "Differential Interpretation of Public Signals and Trade in Speculative Markets," *Journal of Political Economy* 103:4 (1995), 831–872.
- Kristensen, Dennis, "Nonparametric Filtering of the Realized Spot Volatility: A Kernel-Based Approach," *Econometric Theory* 26:1 (2010), 60–93. 10.1017/S0266466609090616
- Kuttner, Kenneth N., "Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market," *Journal of Monetary Economics* 47:3 (2001), 523–544. 10.1016/S0304-3932(01)00055-1
- , "Outside the Box: Unconventional Monetary Policy in the Great Recession and Beyond," *Journal of Economic Perspectives* 32:4 (November 2018), 121–146.
- Lewbel, Arthur, "The Rank of Demand Systems: Theory and Nonparametric Estimation," *Econometrica* 59:3 (1991), 711–730. 10.2307/2938225
- Lewis, Daniel J., "Announcement-Specific Decompositions of Unconventional Monetary Policy Shocks and Their Macroeconomic Effects," Federal Reserve Bank of New York Staff Report 891 (2020).
- Li, Jia, Viktor Todorov, George Tauchen, and Huidi Lin, "Rank Tests at Jump Events," *Journal of Business & Economic Statistics* 37:2 (2019), 312–321.
- Lucca, David O., and Emanuel Moench, "The Pre-FOMC Announcement Drift," *Journal of Finance* 70:1 (2015), 329–371. 10.1111/jofi.12196
- Mancini, Cecilia, "Disentangling the Jumps of the Diffusion in a Geometric Jumping Brownian Motion," *Giornale dell'Istituto Italiano degli Attuari* 64 (2001), 19–47.
- Nakamura, Emi, and Jón Steinsson, "High-Frequency Identification of Monetary Non-Neutrality: The Information Effect," *Quarterly Journal of Economics* 133:3 (2018), 1283–1330. 10.1093/qjle/qjy004
- Rigobon, Roberto, "Identification through Heteroskedasticity," this REVIEW 85:4 (2003), 777–792.
- Rigobon, Roberto, and Brian Sack, "Measuring the Reaction of Monetary Policy to the Stock Market," *Quarterly Journal of Economics* 118:2 (2003), 639–669. 10.1162/003355303321675473
- , "The Impact of Monetary Policy on Asset Prices," *Journal of Monetary Economics* 51:8 (2004), 1553–1575. 10.1016/j.jmoneco.2004.02.004
- Robin, Jean-Marc, and Richard J. Smith, "Tests of Rank," *Econometric Theory* 16 (2000), 151–175. 10.1017/S0266466600162012
- Rothenberg, Thomas J., *Efficient Estimation with A Priori Information* (New Haven, CT: Yale University Press, 1973).
- Rudebusch, Glenn D., "Do Measures of Monetary Policy in a VAR Make Sense?" *International Economic Review* 39:4 (1998), 907–931. 10.2307/2527344
- , "A Review of the Fed's Unconventional Monetary Policy," *FRBSF Economic Letter* 27 (2018), 1–5.
- Savor, Pavel, and Mungo Wilson, "Asset Pricing: A Tale of Two Days," *Journal of Financial Economics* 113 (2014), 171–201. 10.1016/j.jfineco.2014.04.005
- Singleton, Kenneth J., *Empirical Dynamic Asset Pricing: Model Specification and Econometric Assessment* (Princeton, NJ: Princeton University Press, 2009).
- Swanson, Eric T., "Measuring the Effects of Federal Reserve Forward Guidance and Asset Purchases on Financial Markets," NBER working paper w23311 (2017).
- Vissing Jorgensen, Annette, and Arvind Krishnamurthy, "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy," *Brookings Papers on Economic Activity* 43:2 (2011), 215–287.
- Zhang, Lan, Per A. Mykland, and Yacine Aït-Sahalia, "A Tale of Two Time Scales: Determining Integrated Volatility with Noisy High-Frequency Data," *Journal of the American Statistical Association* 100:472 (2005), 1394–1411. 10.1198/016214505000000169