Review of Economic Studies (2018) **85**, 2005–2041 doi:10.1093/restud/rdy003

Advance access publication 22 January 2018

Volume, Volatility, and Public News Announcements

TIM BOLLERSLEV
Duke University
JIA LI
Duke University
and
YUAN XUE
Duke University

First version received June 2016*; Editorial decision December* 2017*; Accepted January* 2018 (*Eds.*)

We provide new empirical evidence for the way in which financial markets process information. Our results rely critically on high-frequency intraday price and volume data for the S&P 500 equity portfolio and U.S. Treasury bonds, along with new econometric techniques, for making inference on the relationship between trading intensity and spot volatility around public news announcements. Consistent with the predictions derived from a theoretical model in which investors agree to disagree, our estimates for the intraday volume-volatility elasticity around important news announcements are systematically below unity. Our elasticity estimates also decrease significantly with measures of disagreements in beliefs, economic uncertainty, and textual-based sentiment, further highlighting the key role played by differences-of-opinion.

Key words: Differences-of-opinion, High-frequency data, Jumps, Macroeconomic news announcements, Trading volume, Stochastic volatility, Economic uncertainty, Textual sentiment

JEL Codes: C51, C52, G12

1. INTRODUCTION

Trading volume and return volatility in financial markets typically, but not always, move in tandem. By studying the strength of this relationship around important public news announcements, we shed new light on the way in which financial markets function and process new information. Our empirical investigations rely critically on the use of high-frequency intraday price and volume data for the S&P 500 equity portfolio and U.S. Treasury bonds, together with new econometric inference procedures explicitly designed to deal with the unique complications that arise in the high-frequency data setting. Consistent with the implications derived from a stylized theoretical model in which investors agree to disagree, our estimates for the intraday volume-volatility elasticity around important news announcements are systematically below unity, the

The editor in charge of this paper was Aureo de Paula.

benchmark case that is obtained in the absence of any disagreement among investors. In line with the theoretical predictions from the same model, our estimates for the elasticity also decrease significantly with proxies for disagreements in beliefs, economic uncertainty, and textual-based sentiment, further corroborating the key role played by differences-of-opinion.

An extensive empirical literature has documented the existence of an on-average strong contemporaneous relation between trading volume and volatility; see Karpoff (1987) for a survey of some of the earliest empirical evidence. The mixture-of-distributions hypothesis (MDH) (see, *e.g.*, Clark, 1973; Tauchen and Pitts, 1983; Andersen, 1996) provides a possible statistical explanation for the positive volume–volatility relationship based on the idea of a common news arrival process driving both the magnitude of returns and trading volume.[1] The MDH, however, remains silent about the underlying economic mechanisms that link the actual trades and price adjustments to the news.

Meanwhile, a variety of equilibrium-based economic models have been developed to help better understand how prices and volume respond to new information. This includes the rational-expectations type models of Kyle (1985) and Kim and Verrecchia (1991, 1994) among many others, in which investors agree on the interpretation of the news, but their information sets differ. In this class of models, the trading volume is mainly determined by liquidity trading and portfolio rebalancing needs. Although this is able to explain the on-average positive correlation between volume and volatility, the underlying trading motives would seem too small to account for the large trading volume observed empirically, especially when the price changes are close to zero (see, *e.g.*, Hong and Stein, 2007, for additional arguments along these lines). Instead, models that feature differences-of-opinion, including those by Harrison and Kreps (1978), Harris and Raviv (1993), Kandel and Pearson (1995), Scheinkman and Xiong (2003) and Banerjee and Kremer (2010) among others, in which investors agree to disagree, may help explain these oft observed empirical phenomena. In the differences-of-opinion class of models, the investors' interpretation of the news and their updated valuations of the assets do not necessarily coincide, thereby providing an additional trading motive that is not directly tied to changes in the equilibrium price.

Much of empirical evidence presented in the literature in regards to the economic models discussed above, and the volume–volatility relationship in particular, have been based on daily or coarser frequency data (see, *e.g.*, Tauchen and Pitts, 1983; Andersen, 1996).[2] Meanwhile, another large and growing strand of literature has emphasized the advantages of the use of high-frequency intraday data for analysing the way in which financial markets respond to new information and more accurately identifying price jumps. Harvey and Huang (1991) and Ederington and Lee (1993), in particular, both find that the strong intraday volatility patterns observed in most financial markets may in part be attributed to regularly scheduled macroeconomic news announcements. Correspondingly large price jumps are often readily associated with specific news announcements; see, *e.g.*, Andersen *et al.* (2007), Lee and Mykland (2008) and Lee (2012).[3] This naturally

1. The MDH formally posits that price changes and trading volume are both subordinated to the same latent information, or news, arrival process. In its simplest form, the hypothesis implies that conditional on the number of information arrivals observed over a non-trivial time interval, the corresponding price change and volume will follow a bivariate normal distribution.

2. One notable exception is Chaboud *et al.* (2008), who document large trading volume in the foreign exchange market in the minutes immediately before macroeconomic announcements, even when the announcements are in line with market expectations and the actual price changes are small. In a recent work, Crego (2017) also uses high-frequency data to study the effect of the Weekly Petroleum Status Reports on bid-ask spreads, trading volume, and returns for "oil firms" and "non-oil firms".

3. Related to this, Savor and Wilson (2013) also document higher average excess market returns on days with important macroeconomic news releases compared to non-announcement days.

suggests that by "zooming in" and analysing how not only prices but also trading volume and volatility evolve around important news announcements, a deeper understanding of the economic mechanisms at work and the functioning of markets may be forthcoming.

Set against this background, we provide new empirical evidence on the volume–volatility relationship around various macroeconomic news announcements. Our leading empirical investigations are based on high-frequency one-minute data for the aggregate S&P 500 equity market portfolio, but we also provide complimentary results for U.S. interest rates using one-minute Treasury bond futures data. We begin by documenting the occurrence of large increases in trading volume intensity around Federal Open Market Committee (FOMC) meetings without accompanying large price jumps. As noted above, this presents a challenge for models in which investors rationally update their beliefs based on the same interpretation of the news, and instead points to the importance of models allowing for disagreements or, differences-of-opinion, among investors.

To help further explore this thesis and guide our more in-depth empirical investigations, we derive an explicit expression for the elasticity of expected trading volume with respect to price volatility within the Kandel and Pearson (1995) differences-of-opinion model. We purposely focus our analysis on the elasticity as it may be conveniently estimated with high-frequency data and, importantly, has a clear economic interpretation in terms of model primitives. In particular, we show theoretically that the volume–volatility elasticity is monotonically decreasing in a well-defined measure of relative disagreement. Moreover, the elasticity is generally below one and reaches its upper bound of unity only in the benchmark case without disagreement.

The theoretical model underlying these predictions is inevitably stylized, focusing exclusively on the impact of public news announcements. As such, the theory mainly speaks to the "abnormal" movements in volume and volatility observed around these news events. To identify the abnormal movements, and thus help mitigate the effects of other confounding forces, we rely on the "jumps" in the volume intensity and volatility around the news announcements. Our estimation of the jumps is based on the differences between the post- and pre-event levels of the instantaneous volume intensity and volatility, which we recover non-parametrically using high-frequency data. Even though the differencing step used in identifying the jumps effectively removes low-frequency dynamics in the volatility and volume series (including daily and lower frequency trending behaviour) that might otherwise confound the estimates, the jump estimates are still affected by the well-documented strong intraday periodic patterns that exist in both volume and volatility (for some of the earliest empirical evidence, see Wood *et al.*, 1985; Jain and Joh, 1988). In an effort to remove this additional confounding effect, we apply a second difference with respect to a control group of non-announcement days. The resulting "doubly-differenced" jump estimates in turn serve as our empirical analogues of the abnormal volume and volatility movements that we use in our regression-based analysis of the theoretical predictions.

Our empirical strategy for estimating the jumps may be viewed as a Difference-in-Difference (DID) type approach, as commonly used in empirical microeconomic studies (see, *e.g.*, Ashenfelter and Card, 1985). The subsequent regression involving the jumps is similar in spirit to the "jump regressions" studied by Li *et al.* (2017), which in turn resembles the non-parametric estimation in (fuzzy) Regression Discontinuity Designs (RDD) (see, *e.g.*, Lee and Lemieux, 2010).[4] We hence refer to our new econometric method as a DID jump regression. However, our setup is distinctly different from conventional econometric settings, and the usual justification

4. In parallel to Li *et al.* (2017), our "jump regressions" are designed to estimate equilibrium relationships between economic variables manifest in unusually large moves, or jumps, in those variables, rather than using the jumps (*i.e.*, the discontinuities) for identifying causal treatment effects. Hence, in spite of the similarity between the two procedures, our goal is distinctly different from that of RDD.

for the use of DID or RDD does not apply in the high-frequency data setting. Correspondingly, our new econometric procedures and the justification thereof entail two important distinctions. First, to accommodate the strong dynamic dependencies in the volatility and volume intensity, we provide a rigorous theoretical justification based on a continuous-time infill asymptotic framework allowing for essentially unrestricted non-stationarity. Secondly, we provide an easy-to-implement local bootstrap method for conducting valid inference. By randomly resampling only locally in time (separately before and after each announcement), the method provides a simple solution to the issue of data heterogeneity, which otherwise presents a formidable challenge in the high-frequency data setting (see, *e.g.*, Gonçalves and Meddahi, 2009).

Our actual empirical findings are closely in line with the theoretical predictions derived from the Kandel and Pearson (1995) model and the differences-of-opinion class of models more generally. In particular, we first document that the estimated volume–volatility elasticity around FOMC announcements is significantly below unity. This finding carries over to other important intraday public news announcements closely monitored by market participants. Interestingly, the volume–volatility elasticity estimates are lower for announcements that are released earlier in the monthly news cycle (see, *e.g.*, Andersen *et al*., 2003), such as the ISM Manufacturing Index and the Consumer Confidence Index, reflecting the importance of the timing across the different announcements and the effect of learning.

Going one step further, we show that the intraday volume–volatility elasticity around news announcements decreases significantly in response to increases in measures of dispersions-in-beliefs (based on the survey of professional forecasters as in, *e.g.*, Van Nieuwerburgh and Veldkamp, 2006; Pasquariello and Vega, 2007) and economic uncertainty (based on the economic policy uncertainty index of Baker *et al*., 2015). This holds true for the S&P 500 aggregate equity portfolio as well as the U.S. Treasury bond market, and again corroborates our theoretical predictions and the key role played by differences-of-opinion. Our more detailed analysis of FOMC announcements, in which we employ an additional textual-based measure for the negative sentiment in the accompanying FOMC statements (based on the methodology of Loughran and McDonald, 2011), further underscores the time-varying nature of the high-frequency volume–volatility relationship and the way in which the market processes new information: when the textual sentiment in the FOMC statement is more negative, the relative disagreement among investors also tends to be higher, pushing down the volume–volatility elasticity.

In contrast to prior empirical studies related to the volume–volatility relationship (see, *e.g.*, Tauchen and Pitts, 1983; Andersen, 1996), our analysis is much more closely guided by an economic model. In particular, based on the implications derived from the Kandel–Pearson model, we provide a new perspective on the volume–volatility relationship by directly linking the strength of the relationship to notions of investors' disagreement. Since the Kandel–Pearson model explicitly concerns "abnormal" variations in volume and volatility around news announcements, we also derive the new DID jump regression framework to rigorously analyse the joint behaviour of volume and volatility jumps. The use of high-frequency data are crucial in this regard, as jumps are invariably short-lived in nature and would be difficult/impossible to accurately identify using data sampled at coarser, say daily, frequencies.

Putting the empirical results in the article into a broader perspective, there is a large literature in market microstructure finance on the price impact of trades, and correspondingly the development of "optimal" trade execution strategies; empirical work along these lines include Hasbrouck (1991), Madhavan *et al*. (1997), and Chordia *et al*. (2002). In contrast to the volume–volatility relationship analysed here, the price impact literature is explicitly concerned with directional price changes predicted by "signed" trading volume, or order flow. Green (2004), in particular, documents a significant increase in the impact of changes in order flow in the U.S. Treasury bond market on intraday bond prices immediately following macroeconomic news announcements, and goes on to suggest that this heightened price impact of trades may be attributed to increased informational asymmetry at the time of the announcements.[5] Further corroborating this idea, Pasquariello and Vega (2007) find that the regression-based estimate for the impact of unanticipated daily order flow in the U.S. Treasury bond market is higher when the dispersion in beliefs among market participants, as measured by the standard deviation of the forecasts from professional forecasters, is high and when the public news announcement is more "noisy". Our main empirical findings based on high-frequency intraday data for the S&P 500 aggregate market portfolio are generally in line with these existing results about the price impact of order flow in the U.S. Treasury bond market. At the same time, however, our focus on trading volume as opposed to order flow presents an important distinction from the aforementioned studies.[6] In line with the insights of Hong and Stein (2007), by focusing on trading volume our empirical findings map more closely into the theoretical predictions from differences-of-opinion class of models in which the disagreement among traders provides an important trading motive beyond conventional rational-expectation type models (as, *e.g.*, the classical model by Kyle, 1985). In addition, our empirical strategy relies critically on the use of high-frequency data coupled with new econometric techniques for non-parametrically uncovering both the instantaneous trading intensity and volatility under minimal statistical assumptions.

The rest of the article is organized as follows. Section 2 presents the basic economic arguments and theoretical model that guide our empirical investigations. Section 3 describes the high-frequency intraday data and news announcements used in our empirical analysis. To help set the stage for our more detailed subsequent empirical investigations, Section 4 discusses some preliminary findings specifically related to the behavior of the aggregate stock market around the FOMC announcements. Section 5 describes the new inference procedures necessitated by our more in-depth high-frequency empirical analysis. Section 6 presents our main empirical findings based on the full set of news announcements, followed by our more detailed analysis of FOMC announcements. In addition to our main results based on data for the S&P 500 aggregate equity portfolio, we also present complementary results for the U.S. Treasury bond market. Section 7 concludes. Technical details concerning the new econometric inference procedure are provided in Appendix A. Appendix B contains further data descriptions. Additional empirical results and robustness checks are relegated to a (not-for-publication) Supplementary Appendix.

2. THEORETICAL MOTIVATION

We rely on the theoretical volume–volatility relations derived from the differences-of-opinion model of Kandel and Pearson (1995) to help guide our empirical investigations. We purposely focus on a simplified version of the model designed to highlight the specific features that we are after, and the volume–volatility elasticity around news arrivals in particular. We begin by discussing the basic setup and assumptions.

2.1. *New information and differences-of-opinion*

Following Kandel and Pearson (1995), henceforth KP, we assume that a continuum of traders trade a risky asset and a risk-free asset in a competitive market. The random payoff of the risky

5. This is consistent with the predictions from the theoretical model in Kim and Verrecchia (1994), in which earnings announcements may increase informational asymmetry due to the superior information processing skills of certain traders.

6. As such, our empirical findings also do not speak directly to prior work that explicitly rely on net order flow data. Instead, we offer a new perspective on the widely studied volume–volatility relationship.

asset, denoted $\tilde{u}$, is unknown to the traders. The risk-free rate is normalized to be zero. The traders' utility functions have constant absolute risk aversion with risk tolerance r. There are only two types of traders, $i \in \{1,2\}$, with the proportion of type 1 traders denoted α. The traders have different prior beliefs about the payoff before the announcement, and they also *disagree* about the interpretation of the public signal at the time of the announcement. Type i trader's prior is given by a normal distribution with precision s_i.

After the announcement, the traders observe the same public signal $\tilde{u}+\tilde{\varepsilon}$, where the noise term $\tilde{\varepsilon}$ is normally distributed. The traders then update their beliefs about $\tilde{u}$ and optimally re-balance their positions. The key feature of the KP model is that the two types of traders agree to disagree on how to interpret the public signal when updating their beliefs about the asset value: type i traders believe that $\tilde{\varepsilon}$ is drawn from the $N(\mu_i, h^{-1})$ distribution. Differences-of-opinion regarding the public signal among the traders thus corresponds to $\mu_1 \neq \mu_2$.

Following KP it is possible to show that in equilibrium[7]

$$\text{Volume} = |\beta_0 + \beta_1 \cdot \text{Price Change}|, \tag{2.1}$$

where

$$\beta_0 = r\alpha(1-\alpha)h(\mu_1-\mu_2), \quad \beta_1 = r\alpha(1-\alpha)(s_1-s_2). \tag{2.2}$$

The coefficient β_0 is directly associated with the degree of differences-of-opinion concerning the interpretation of the public signal (*i.e.* $\mu_1-\mu_2$), while β_1 depends on the dispersion in the precisions of the prior beliefs about the payoff (*i.e.* s_1-s_2). Both of the coefficients are increasing in the degree of risk tolerance r, and the degree of heterogeneity among the traders as measured by $\alpha(1-\alpha)$.

Looking at the equilibrium relationship in equation (2.1), the first β_0 term represents the "disagreement component". This term becomes increasingly more important for higher levels of disagreement (*i.e.* $|\mu_1-\mu_2|$ is large) and/or when the traders are more confident about their interpretation of the public signal (*i.e.*, h is large). Hence, other things being equal, the higher the level of disagreement, the weaker the relationship between trading activity and price changes. In the extreme case, when the equilibrium price does not change as a result of the announcement and the second term on the right-hand side of (2.1) equals zero, there can still be large trading volume arising from disagreement among the traders because of the β_0 term. In this sense, disagreement generates an additional endogenous trading motive that is effectively "orthogonal" to any revision in the equilibrium price.

2.2. *Expected volume and volatility*

The implication of the KP model for the relationship between price adjustment and trading volume in response to new information is succinctly summarized by equations (2.1) and (2.2). These equations, however, depict an exact functional relationship between (observed) random quantities. A weaker, but empirically more realistic, implication can be obtained by thinking of this equilibrium relationship as only holding "on average". Moment conditions corresponding to the stochastic version (2.1) formally capture this idea.

Specifically, let $m(\sigma)$ denote the expected volume as a function of the volatility σ (*i.e.* the standard deviation of price change). We assume that the price changes are normally distributed with mean zero and standard deviation σ. In particular, we note that the zero-mean assumption

7. See equation (5) in Kandel and Pearson (1995).

TABLE 1

Summary statistics of high-frequency price returns and volume data

	Mean	Min	1%	25%	50%	75%	99%	Max
Return (percent)	0.000	−2.014	−0.149	−0.021	0.000	0.021	0.149	2.542
Volume (1,000 shares)	289.5	0.0	2.4	55.4	147.6	350.7	2005	31153

Notes: The table reports summary statistics of the one-minute returns and one-minute trading volumes for the SPY ETF during regular trading hours from 10, April 2001 to 30, September 2014.

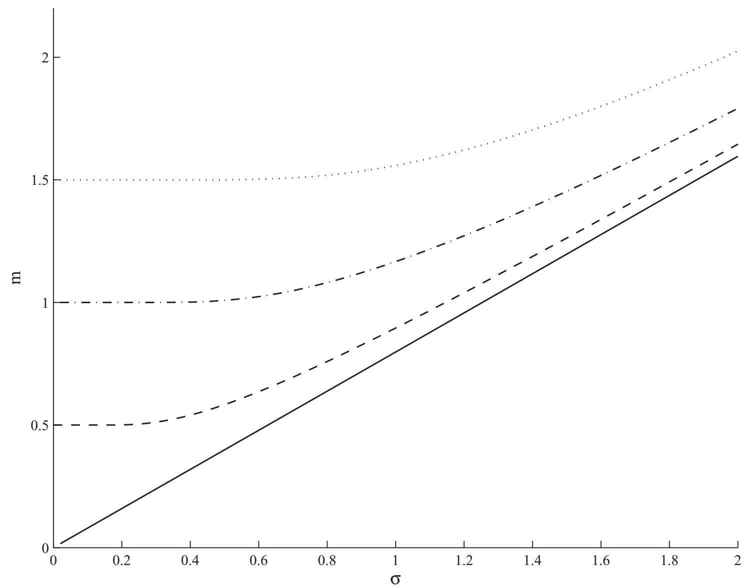

FIGURE 1

Equilibrium volume–volatility relations

Notes: The figure shows the equilibrium relationship between expected trading volume m and price volatility σ in the Kandel–Pearson model for various levels of disagreement, ranging from $\beta_0=0$ (bottom) to 0.5, 1 and 1.5 (top). β_1 is fixed at one in all of the graphs.

is empirically sound for high-frequency returns (Table 1). It follows by direct integration of (2.1) that

$$m(\sigma)=\sqrt{\frac{2}{\pi}}|\beta_1|\sigma\exp\left(-\frac{\beta_0^2}{2\beta_1^2\sigma^2}\right)+|\beta_0|\left(2\Phi\left(\frac{|\beta_0|}{|\beta_1|\sigma}\right)-1\right), \tag{2.3}$$

where Φ denotes the cumulative distribution function of the standard normal distribution. The expected volume $m(\sigma)$ depends on σ and the (β_0,β_1) coefficients in a somewhat complicated fashion. However, it is straightforward to show that $m(\sigma)$ is increasing in σ.

To gain further insight regarding this (expected) volume–volatility relationship, Figure 1 illustrates how the $m(\sigma)$ function varies with the "disagreement component" β_0. Locally, when the volatility σ is close to zero, the expected volume is positive if and only if the opinions of the traders differ (*i.e.* $\mu_1\neq\mu_2$). Globally, as β_0 increases (from the bottom to the top curves in the figure), the equilibrium relationship between the expected volume and volatility "flattens out". This pattern is consistent with the aforementioned intuition that disagreement among traders provide an additional trading motive which loosens the relationship between volume and volatility.

The exact non-linear expression in equation (2.3) inevitably depends on the specific assumptions and setup underlying the KP model. Although this expression helps formalize the intuition about the way in which disagreement affects the volume–volatility relationship, the KP model is obviously too stylized to allow for a direct structural estimation. In our empirical analysis below, we therefore rely on a reduced-form approach. In so doing, we aim to test the basic economic intuition and empirical implications stemming from the differences-of-opinion class of models more generally. From an empirical perspective this is naturally accomplished by focusing on the volume–volatility relationship expressed in terms of the elasticity of $m(\sigma)$ with respect to σ.

Not only does this volume–volatility elasticity provide a convenient "reduced-form" summary statistic, it also admits a clear economic interpretation within the KP model. Let $\mathcal{E}$ denote the volume–volatility elasticity. A straightforward calculation then yields

$$\mathcal{E} \equiv \frac{\partial m(\sigma)/m(\sigma)}{\partial \sigma/\sigma} = \frac{1}{1+\psi(\gamma/\sigma)}, \tag{2.4}$$

where

$$\gamma \equiv \frac{|\beta_0|}{|\beta_1|} = \frac{h|\mu_1-\mu_2|}{|s_1-s_2|}, \tag{2.5}$$

and the function ψ is defined by $\psi(x) \equiv x(\Phi(x)-1/2)/\phi(x)$, with ϕ being the density function of the standard normal distribution. The function ψ is strictly increasing on $[0,\infty)$, with $\psi(0)=0$ and $\lim_{x\to\infty}\psi(x)=\infty$.

The expressions in (2.4) and (2.5) embody two key features in regards to the volume–volatility elasticity. First, $\mathcal{E} \leq 1$ with the equality and an elasticity of unity obtaining if and only if $\gamma=0$. Secondly, $\mathcal{E}$ only depends on and is decreasing in γ/σ. This second feature provides a clear economic interpretation of the volume–volatility elasticity $\mathcal{E}$: it is low when differences-in-opinion is *relatively* high, and vice versa, with γ/σ serving as the *relative* measure of the differences-of-opinion. This relative measure is higher when traders disagree more on how to interpret the public signal (*i.e.* larger $|\mu_1-\mu_2|$) and with more confidence (*i.e.* larger h), relative to the degree of asymmetric private information (*i.e.* $|s_1-s_2|$), and the overall price volatility (*i.e.* σ).

These two features in turn translate into directly testable implications that we use to guide our empirical analysis. In particular, taking the KP model and the expressions in (2.4) and (2.5) at face value, it is possible to test for the presence of differences-of-opinion among traders by testing whether the volume–volatility elasticity around important news arrivals is less than or equal to unity. This strict implication, however, hinges on a number of specific parametric and distributional modeling assumptions.[8] We therefore also investigate the second more qualitative implication arising from the model that the volume–volatility elasticity is decreasing with the overall level of disagreement. This implication reflects the more general economic intuition that disagreement among traders provides an extra trading motive, and as such this implication should hold true more broadly. To allow for a focused estimation of the elasticities, we base our empirical investigations on intraday high-frequency transaction data around well-defined public macroeconomic news announcements, along with various proxies for the heterogeneity in beliefs and economic uncertainty prevailing at the exact time of the announcements. We turn next to a discussion of the data that we use in doing so.

8. Among others, the assumption of CARA and normally distributed prior beliefs.

3. DATA DESCRIPTION AND SUMMARY STATISTICS

Our empirical investigations are based on high-frequency intraday transaction prices and trading volume, together with precisely timed macroeconomic news announcements. We describe our data sets in turn.

3.1. *High-frequency market prices and trading volume*

Our primary data are composed of intraday transaction prices and trading volume for the S&P 500 index ETF (ticker: SPY). All of the data are obtained from the TAQ database. The sample covers all regular trading days from April 10th, 2001 through September 30th, 2014. The raw data are cleaned following the procedures detailed in Brownlees and Gallo (2006) and Barndorff-Nielsen *et al.* (2009). Further, to mitigate the effect of market microstructure noise, we follow standard practice in the literature to sparsely sample the data at a one-minute sampling interval, resulting in a total of 1,315,470 one-minute return and volume observations.

Summary statistics for the SPY returns and trading volumes (number of shares) are reported in Table 1. Consistent with prior empirical evidence (see, *e.g.* Bollerslev and Todorov, 2011), the high-frequency one-minute returns appear close to be symmetrically distributed. The one-minute volume series, on the other hand, is highly skewed to the right, with occasionally very large values.

To highlight the general dynamic dependencies inherent in the data, Figure 2 plots the daily logarithmic trading volume (constructed by summing the one-minute trading volumes over each of the different days) and the logarithmic daily realized volatilities (constructed as the sum of squared one-minute returns over each of the days in the sample). As the figure shows, both of the daily series vary in a highly predictable fashion. The volume series, in particular, seem to exhibit an upward trend over the first half of the sample, but then levels off over the second half. Meanwhile, consistent with the extensive prior empirical evidence discussed above, there are strong dynamic commonalities evident in the two series.

The volume and volatility series also both exhibit strong intraday patterns. To illustrate this, Figure 3 plots the square-root of the one-minute squared returns averaged across each minute-of-the-day (as an estimate for the volatility over that particular minute) and the average trading volume over each corresponding minute. To prevent abnormally large returns and volumes from distorting the picture, we only include non-announcement days that are discussed in Section 3.2. Consistent with the evidence in the extant literature, there is a clear U-shaped pattern in the average volatility and trading activity over the active part of the trading day.[9]

In addition to our main empirical results based on the high-frequency intraday data for the S&P 500 aggregate equity portfolio discussed above, we also report complementary empirical evidence for the U.S. Treasury bond market. Our intraday price and volume data for the ten-year U.S. T-bond futures are obtained from TickData, and spans the slightly shorter sample period from 1, July 2003 to 30, September 2014. The intraday T-bond volume and volatility series naturally exhibit their own distinct dynamic dependencies and intraday patterns. However, as discussed further below, the key features pertaining to the volume and volatility "jumps" observed around important macroeconomic news announcements closely mirror those for the aggregate equity market.

9. See Wood *et al.* (1985), Harris (1986), Jain and Joh (1988), Baillie and Bollerslev (1990), and Andersen and Bollerslev (1997) for some of the earliest empirical evidence on the intraday patterns in volatility and volume.

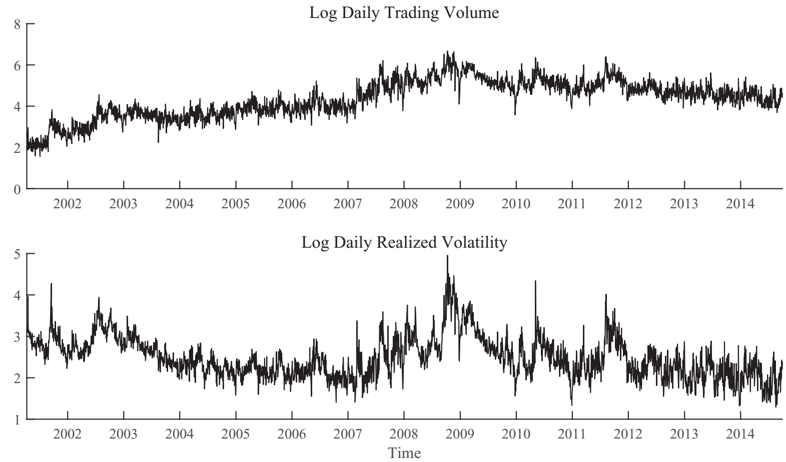

FIGURE 2

Time series of volume and volatility

Notes: The figure shows the daily logarithmic trading volume (top panel) and logarithmic realized volatility (bottom panel) for the SPY ETF. The daily volume (in millions) is constructed by accumulating the intraday volume. The daily realized volatility (annualized in percentage) is constructed as the sum of one-minute squared returns over the day.

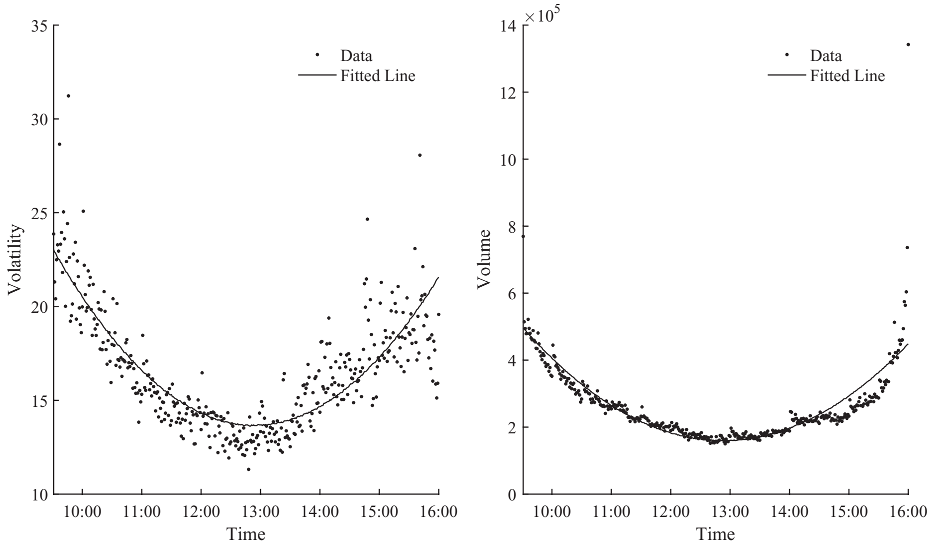

FIGURE 3

Intraday patterns of volatility and volume

Notes: The figure shows the intraday volatility (annualized in percentage) for the SPY ETF (left panel) constructed as the square root of the one-minute squared returns averaged across all non-announcement days, along with the best quadratic fit. The intraday trading volume (in shares) for the SPY ETF (right panel) is similarly averaged across all non-announcement days.

3.2. *Macroeconomic news announcements*

The Economic Calendar Economic Release section in Bloomberg includes the date and exact within day release time for over one hundred regularly scheduled macroeconomic news

TABLE 2
Macroeconomic news announcements.

	No. of Obs.	Time	Source
FOMC	109	Varies	Federal Reserve Board
ISM Manufacturing (ISMM)	160	10:00	Institute of Supply Management
ISM Non-Manufacturing (ISMNM)	158	10:00	Institute of Supply Management
Consumer Confidence (CC)	160	10:00	Conference Board
Other Indicators (Others)	1682	Varies	See Appendix B.1

Notes: The table reports the total number of observations, release time, and data source for each of the news announcements over the 10, April 2001 to 30, September 2014 sample.

announcements. Most of these announcements occur before the market opens or after it closes. We purposely focus on announcements that occur during regular trading hours only.[10] While this leaves out some key announcements (most notably the monthly employment report), importantly it allows us to accurately estimate volume and volatility "jumps" by harnessing the rich information inherent in intraday high-frequency data about the way in which markets process new information. In addition to the intraday announcements pertaining to specific macroeconomic variables and indicators, we also consider FOMC rate decisions. FOMC announcements have been analysed extensively in the existing literature, and we also pay special attention to those announcements in our empirical analysis discussed below.

Based on prior empirical evidence (see, *e.g.* Andersen *et al.*, 2003; Boudt and Petitjean, 2014; Jiang *et al.*, 2011; Lee, 2012), we identify four announcements as being the most important overall in the sense of having the on-average largest price impact among all of the regularly scheduled intraday news announcements. In addition to the aforementioned FOMC rate decisions, this includes announcements pertaining to the ISM Manufacturing Index (ISMM), the ISM Non-Manufacturing Index (ISMNM), and the Consumer Confidence (CC) Index. Table 2 provides the typical release times and the number of releases over the April 10, 2001 to September 30, 2014 sample period for each of these important announcements.[11] The remaining announcements are categorized as "Others", a full list of which is provided in Table B.1 in Appendix B. Some of the announcement times for these other indicators invariably coincide, so that all-in-all our sample is composed of a total of 2,130 unique intraday public news announcement times.

4. A PRELIMINARY ANALYSIS OF FOMC ANNOUNCEMENTS

To set the stage for our more in-depth subsequent empirical investigations, we begin by presenting a set of simple summary statistics and illustrative figures related to the volume–volatility relationship around FOMC announcements. We focus our preliminary analysis on FOMC announcements, because these are arguably among the most important public news announcements that occur during regular trading hours.[12]

10. To ensure that there is a 30-minute pre-event (resp. post-event) window before (resp. after) each announcement, we also exclude announcements that are released during the first and the last 30 minutes of the trading day.

11. The Economic Release section in Bloomberg reports 108 FOMC meetings during our sample period. However, there are also additional FOMC intermeeting announcements. All of these occurred outside the intraday trading hours, except the one at 10:55am on 18, April 2001. Including this one additional intermeeting announcement, thus leaves us with a total 109 FOMC announcements. Almost identical empirical results to the ones reported below are obtained by excluding this one FOMC intermeeting observation.

12. The reaction of market prices to FOMC announcements has been extensively studied in the recent literature; see, for example, Johnson and Paye (2015) and the many references therein.

For each announcement, let τ denote the pre-scheduled announcement time. The time τ is naturally associated with the integer $i(\tau)$ such that $\tau=(i(\tau)-1)\Delta_n$, where $\Delta_n=1$ minute is the sampling interval of our intraday data. We define the event window as $((i(\tau)-1)\Delta_n, i(\tau)\Delta_n]$. Further, we define the pre-event (resp. post-event) window to be the k_n-minute period immediately before (resp. after) the event. We denote the return and trading volume over the jth intraday time-interval $((j-1)\Delta_n, j\Delta_n]$ by r_j and $V_{j\Delta_n}$, respectively. The volume intensity m (*i.e.*, the instantaneous mean volume) and the spot volatility σ before and after the announcement, denoted by $m_{\tau-}$, m_τ, $\sigma_{\tau-}$ and σ_τ, respectively, can then be estimated by

$$\begin{aligned} &\hat{m}_{\tau-} \equiv \frac{1}{k_n}\sum_{j=1}^{k_n} V_{(i(\tau)-j)\Delta_n}, \qquad \hat{m}_{\tau} \equiv \frac{1}{k_n}\sum_{j=1}^{k_n} V_{(i(\tau)+j)\Delta_n}, \\ &\hat{\sigma}_{\tau-} \equiv \sqrt{\frac{1}{k_n\Delta_n}\sum_{j=1}^{k_n} r^2_{i(\tau)-j}}, \qquad \hat{\sigma}_{\tau} \equiv \sqrt{\frac{1}{k_n\Delta_n}\sum_{j=1}^{k_n} r^2_{i(\tau)+j}}. \end{aligned} \tag{4.1}$$

These estimators are entirely non-parametric, in the sense that they only rely on simple averages of the data in local windows around the event time. The window size k_n plays the same role as the bandwidth parameter in conventional non-parametric analysis and regressions; the standard technical assumptions needed to formally justify the estimators are further detailed in Appendix A.[13] In the empirical analysis reported below we set $k_n=30$, corresponding to a 30-minute pre-event (resp. post-event) window.[14]

As an initial illustration of the volume and volatility jumps observed around FOMC announcements, Figure 4 plots the average (across all FOMC announcements in the sample) estimated volume intensity (top panel) and spot volatility (bottom panel) processes for the 15 minutes before and after the announcements. As the figure clearly shows, both volume and volatility sharply increase at the time of the announcement. Consistent with the recent work of Bernile *et al.* (2016) and Kurov *et al.* (2016) and the finding that some price-adjustment seemingly occurs in anticipation of the actual news release, there is also a slight increase in both series leading up to the news announcement time. However, this increase is clearly small compared to the "jumps" that manifest at the time of the announcement, and as such will not materially affect any of our subsequent empirical analysis.[15]

To examine whether the "average jumps" evident in Figure 4 are representative, Figure 5 plots the time series of individually estimated logarithmic volume intensities (top panel) and logarithmic spot volatilities (bottom panel) before and after each of the FOMC announcements.[16] Consistent with the on-average estimates depicted in Figure 4, the figure shows marked bursts in the trading volume following each of the FOMC announcements, accompanied by positive jumps

13. Technically, the consistent estimation of the volume intensity and spot volatility before and after the announcements only requires the processes to be right-continuous with left-limits. Correspondingly, the local window parameter k_n is assumed to satisfy $k_n \to \infty$ and $k_n\Delta_n \to 0$, which directly mirrors the usual regularity conditions on the bandwidth parameter in nonparametric regressions.

14. Our main empirical findings are not sensitive to the choice of k_n; see the Supplementary Appendix for some robustness checks.

15. As a simple robustness check, we also estimated the pre-event quantities excluding the one-minute observation immediately before the announcement, resulting in virtually the same estimates as the ones reported below; further details are available in the Supplementary Appendix.

16. The logarithmic transform is naturally motivated by our interest in the volume–volatility elasticity derived in Section 2.2. The log-transformation also helps reduce the heteroskedasticity in both series. The heteroskedasticity in the spot volatility estimates is directly attributable to estimation errors. Formally, the standard error of the $\hat{\sigma}_\tau$ estimate equals $\hat{\sigma}_\tau/\sqrt{2k_n}$, so that by the delta-method the standard error of $\log(\hat{\sigma}_\tau)$ equals the constant $1/\sqrt{2k_n}$. In addition, the logarithm effectively transforms the salient multiplicative trend in the volume series over the first half of the sample to an additive trend, which as seen in Figure 2 is close to linear in time.

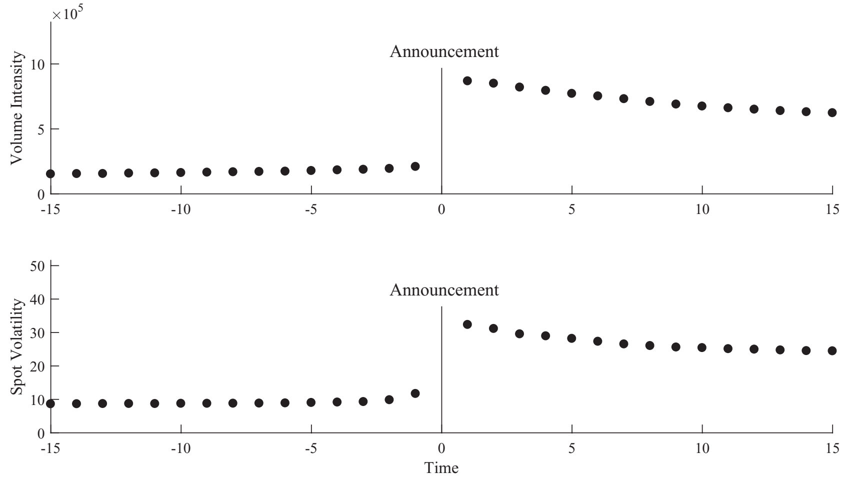

FIGURE 4

Average volume intensity and volatility jumps around FOMC announcements

Notes: The figure plots the non-parametric estimates of the volume intensity (in shares) and volatility (annualized in percentage) before and after an "average" FOMC announcement. We compute $\hat{m}_{\tau\pm s}$ (resp. $\hat{\sigma}_{\tau\pm s}$) for s ranging from 15 minutes before to 15 minutes after each FOMC announcement, and plot the averages across all of the announcements in the top (resp. bottom) panel. The local window is set to $k_n = 30$.

in the volatility. These jumps in the volume intensity and volatility are economically large, with average jump sizes (in log) of 1.41 and 1.09, respectively. As discussed in Section 6, they are also highly statistically significant. This therefore suggests that traders revise their beliefs about the stock market differently upon seeing the FOMC announcement. It is, of course, possible that traders have asymmetric private information about the overall market. However, in line with the reasoning of Hong and Stein (2007) regarding the large observed burst in trading volume, it seems much more likely that the differences are attributed to the traders' disagreement regarding the news.[17]

To further buttress the importance of differences-of-opinion among investors, we consider two additional empirical approaches. First, as discussed in Section 2.1, if all investors agreed on the interpretation of the FOMC announcements, the trading volumes observed around the news releases should be approximately proportional to the price changes. Consequently, if there were no disagreement we would expect to see large trading volume accompanied by large changes in prices and vice versa. To investigate this hypothesis, we sort all of the FOMC announcements by the normalized one-minute event returns $r_i/\hat{\sigma}_{\tau-}\sqrt{\Delta_n}$, and plot the resulting time series of pre- and post-event volume intensity estimates in Figure 6.[18] Consistent with the findings of KP (based on daily data), Figure 6 shows no systematic association between trading volumes and returns. Instead, we observe many sizable jumps in the volume intensity for events with absolute returns "close" to zero, that is, when they are less than one instantaneous standard deviation (highlighted by the shaded area).

17. It also appears highly unlikely that there is any "insider information" pertaining to the actual FOMC announcements.

18. The normalization with respect to the spot volatility serves as a scale adjustment to make the returns across announcements more comparable. A similar figure based on the five-minute returns is included in the Supplementary Appendix.

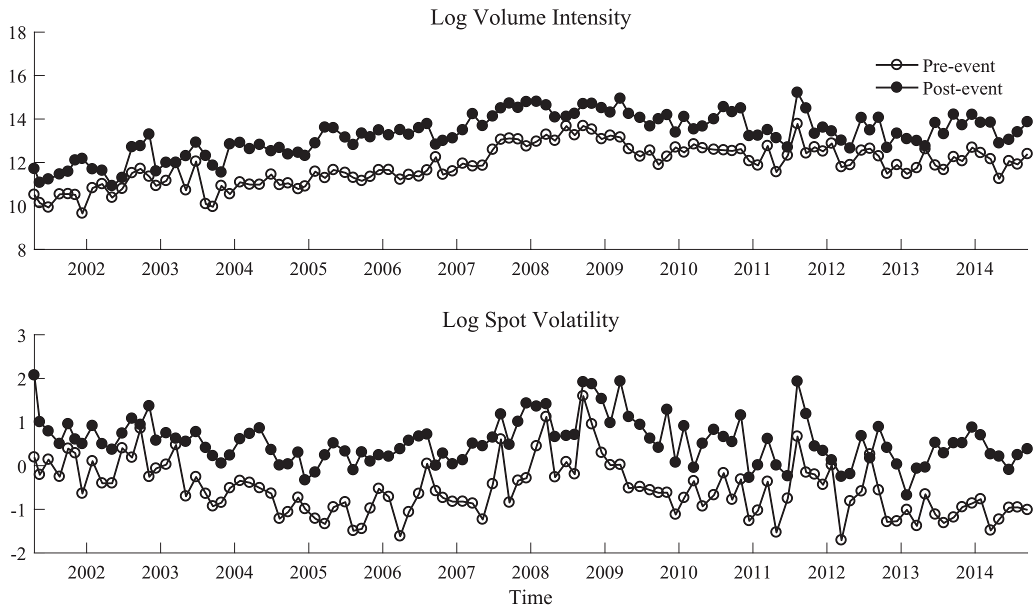

FIGURE 5

Volume and volatility around FOMC announcements

Notes: The figure plots the log-volume intensity (top panel, in shares) and log-volatility (bottom panel, in percentage) around FOMC announcements. The volume intensity and volatility are calculated using equation (4.1) with $k_n = 30$.

The empirical approach underlying Figure 6 mainly focuses on events with price changes close to zero and, hence, is *local* in nature. Our second empirical approach seeks to exploit a more *global* feature of the differences-of-opinion type models, namely that the volume–volatility elasticity should be below unity. While this prediction was derived from the explicit solution for the KP model in equation (2.3), the underlying economic intuition holds more generally: differences-of-opinion provides an additional trading motive that is not tied to the traders' average valuation of the asset and, hence, serves to "loosen" the relationship between trading volume and volatility.

To robustly examine this prediction for the volume–volatility elasticity, without relying on the specific functional form in (2.4), we adopt a less restrictive reduced-form estimation strategy. Further along those lines, the theoretical models discussed in Section 2 are inevitably stylized in nature, abstracting from other factors that might affect actual trading activity (*e.g.*, liquidity or life-cycle trading, reduction in trading costs, advances in trading technology, to name but a few). As such, the theoretical predictions are more appropriately thought of as predictions about "abnormal" variations in the volume intensity and volatility. In the high-frequency data setting, abnormal movements conceptually translate into "jumps". Below, we denote the corresponding log-volatility and log-volume intensity jumps by $\Delta\log(\sigma_\tau) \equiv \log(\sigma_\tau) - \log(\sigma_{\tau-})$ and $\Delta\log(m_\tau) \equiv \log(m_\tau) - \log(m_{\tau-})$, respectively.

Figure 7 shows a scatter plot of the estimated $\Delta\log(\sigma_\tau)$ and $\Delta\log(m_\tau)$ jumps around FOMC announcement times. As expected, there is a clear positive association between the two series, with a correlation coefficient of 0.57. Moreover, consistent with the theoretical predictions and the idea that traders interpret FOMC announcements differently, the estimate for the volume–volatility elasticity implied by the slope coefficient from a simple linear fit equals 0.66, which is much less than unity.

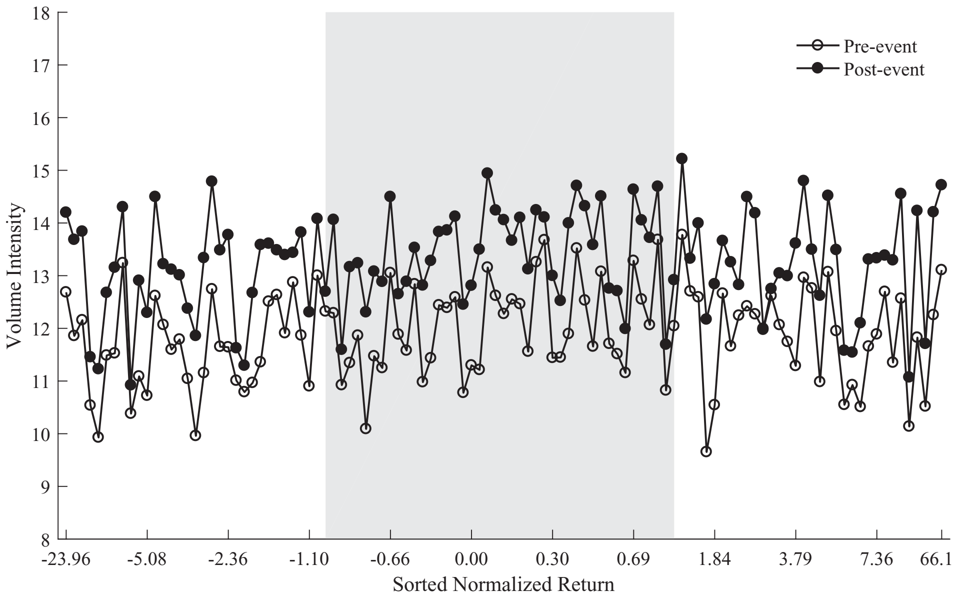

FIGURE 6

Sorted volume around FOMC announcements

Notes: The figure shows the pre- and post-event log volume intensities (in shares) sorted on the basis of the one-minute normalized returns $r_{i(\tau)}/\hat{\sigma}_{\tau-}\sqrt{\Delta_n}$ around FOMC announcements (dots). The normalized return increases from left to right. Announcements with normalized returns less than 1 are highlighted by the shaded area.

The summary statistics and figures discussed above all corroborate the conjecture that differences-of-opinion among investors play an important role in the way in which the market responds to FOMC announcements. To proceed with a more formal empirical analysis involving other announcements and explanatory variables, we need econometric tools for conducting valid inference, to which we now turn.

5. HIGH-FREQUENCY ECONOMETRIC PROCEDURES

The econometrics in the high-frequency setting is notably different from more conventional settings, necessitating the development of new econometric tools properly tailored to the data and the questions of interest.[19] To streamline the discussion, we focus on the practical implementation and heuristics of the underlying econometric theory, deferring the technical details to Appendix A.

Following the discussion and the theoretical implications developed in Section 2.2, our primary interest centers on estimation and inference concerning the volume–volatility elasticity, and in particular, whether the elasticity decreases with the level of disagreement. To this end, we estimate the elasticity between announcement-induced "abnormal" volume and volatility variations by regressing jumps in the log-volume intensity on jumps in the log-volatility. Further, to help address how the elasticity is affected by disagreement, we parameterize the regression coefficient as a function of various measures of disagreement prevailing at the time of the announcement

19. Aït-Sahalia and Jacod (2014) provide a comprehensive review of recent development on the econometrics of high-frequency data.

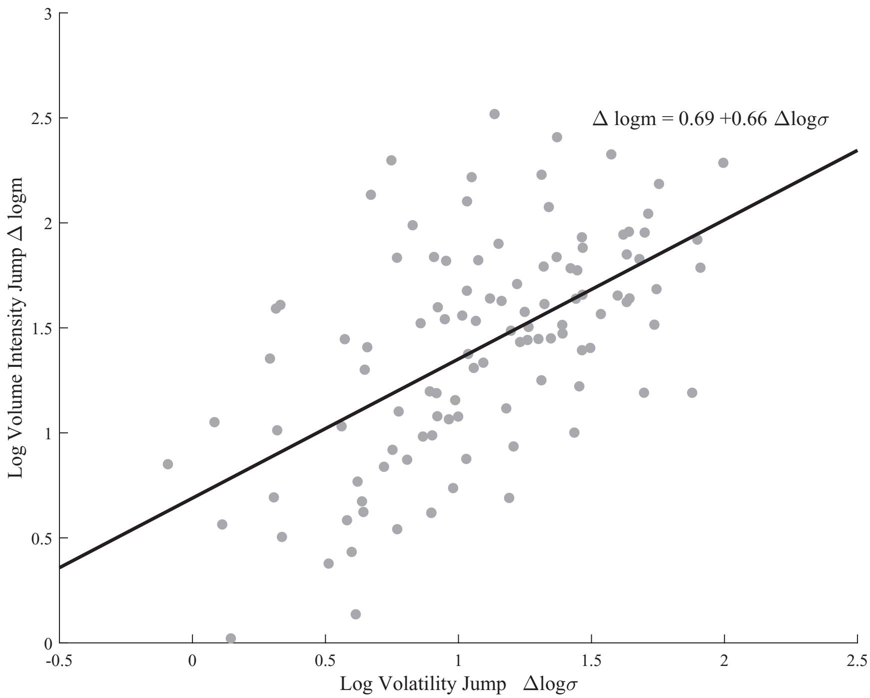

FIGURE 7

Volume and volatility jumps around FOMC announcements

Notes: The figure shows the scatter of the jumps in the log-volume intensity (in shares) versus the jumps in the log-volatility (in percentage) around FOMC announcements. The line represents the least-square fit.

(*e.g.*, the dispersion among professional forecasters). The resulting econometric model may be succinctly expressed as

$$\Delta\log(m_\tau) = (a_0 + b_0^\top X_{0,\tau}) + (a_1 + b_1^\top X_{1,\tau}) \cdot \Delta\log(\sigma_\tau), \tag{5.1}$$

where $X_\tau \equiv (X_{0,\tau}, X_{1,\tau})$ is composed of the different explanatory/control variables employed in the estimation. In particular, by including measures of disagreement in $X_{1,\tau}$ it is possible to directly assess the aforementioned theoretical predictions based on the statistical significance of the estimated b_1 coefficients.

From an econometric perspective, equation (5.1) is best understood as an instantaneous moment condition, in which the volume intensity process m_t (resp. the spot volatility process σ_t) represents the latent local first (resp. second) moment of the volume (resp. price return) process.[20] The sample analogue of (5.1) therefore takes the form

$$\Delta\widehat{\log(m_\tau)} = (a_0 + b_0^\top X_{0,\tau}) + (a_1 + b_1^\top X_{1,\tau}) \cdot \Delta\widehat{\log(\sigma_\tau)} + e_\tau, \tag{5.2}$$

20. This setting resembles the fuzzy RDD. To illustrate, consider the special case of (5.1) where a_0, b_0 and b_1 equal zero. In this situation, (5.1) becomes $\Delta\log(m_\tau) = a_1\Delta\log(\sigma_\tau)$, so a_1 is exactly equal to the ratio between the jump of $\log(m_\tau)$ and that of $\log(\sigma_\tau)$, where m and σ are local moments of the volume and return processes, respectively. Analogously, in the fuzzy RDD setting, the treatment effect is identified exactly as the ratio between the jump in the conditional mean of the outcome variable and that in the treatment propensity; see, *e.g.*, Lee and Lemieux (2010). In the more general setting with multiple discontinuities (induced by announcements), we aggregate the information among the announcements assuming that heterogeneous effects can be controlled by the observed covariates in X_τ.

where the error term e_τ arises from the estimation errors associated with the local moments (*i.e.*, m and σ). Our goal is to conduct valid inference about the parameter vector $\theta \equiv (a_0, b_0, a_1, b_1)$, especially the components a_1 and b_1 that determine the volume–volatility elasticity.

Consider the group $\mathcal{A}$ composed of a total of M announcement times.[21] Further, define $S_\tau \equiv (m_{\tau-}, m_\tau, \sigma_{\tau-}, \sigma_\tau, X_\tau)$ and $\mathbf{S} \equiv (S_\tau)_{\tau\in\mathcal{A}}$, where the latter collects the information on all announcements. Our estimator of $\mathbf{S}$ may then be expressed as $\widehat{\mathbf{S}}_n \equiv (\widehat{S}_\tau)_{\tau\in\mathcal{A}}$, where $\widehat{S}_\tau \equiv (\hat{m}_{\tau-}, \hat{m}_\tau, \hat{\sigma}_{\tau-}, \hat{\sigma}_\tau, X_\tau)$ is formed using the non-parametric pre- and post-event volume intensity and volatility estimators previously defined in (4.1). Correspondingly, summary statistics pertaining to the jumps in the volume intensity and volatility for the group of announcement times $\mathcal{A}$ may be succinctly expressed as $f(\widehat{\mathbf{S}})$ for some smooth function $f(\cdot)$.[22]

Moreover, we may estimate the parameter vector $\theta \equiv (a_0, b_0, a_1, b_1)$ in (5.1) for the group $\mathcal{A}$ using the following least-square estimator

$$\hat{\theta} \equiv \underset{\theta}{\operatorname{argmin}} \sum_{\tau\in\mathcal{A}} \left(\Delta\widehat{\log(m_\tau)} - (a_0 + b_0^\top X_{0,\tau}) - (a_1 + b_1^\top X_{1,\tau}) \cdot \Delta\widehat{\log(\sigma_\tau)} \right)^2. \tag{5.3}$$

This estimator may similarly be expressed as $\hat{\theta} = f(\widehat{\mathbf{S}})$, albeit for a more complicated transform $f(\cdot)$. It can be shown that $\widehat{\mathbf{S}}$ is a consistent estimator of $\mathbf{S}$, which in turn implies that $f(\widehat{\mathbf{S}})$ consistently estimates $f(\mathbf{S})$, provided $f(\cdot)$ is a smooth function of the estimated quantities. The estimates that we reported in our preliminary analysis in Section 4 may be formally justified this way.[23]

The "raw" estimator defined above is asymptotically valid under general regularity conditions (again, we refer to Appendix A for the specific details). However, the non-parametric estimators $\Delta\widehat{\log(\sigma_\tau)}$ and $\Delta\widehat{\log(m_\tau)}$ underlying the simple estimator in (5.3) do not take into account the strong intraday U-shaped patterns in trading volume and volatility documented in Figure 3. While the influence of the intraday patterns vanishes asymptotically, they invariably contaminate our estimates of the jumps in finite samples, and thus affect our use of the jump estimates as measures of "abnormal" volume and volatility movements, to which the economic theory speaks. A failure to adjust for this may therefore result in a mismatch between the empirical strategy and the economic theory.[24]

To remedy this, we correct for the influence of the intraday pattern by differencing it out with respect to a control group. Since this differencing step is applied to the jumps, which are

21. Importantly, unlike conventional econometric settings, our asymptotic inference does not rely on an increasingly large number of announcements. Indeed, we assume that the sample span and the number of announcements within the span (*i.e.*, M) are fixed. Our econometric theory exploits the fact that the high-frequency data are sampled at (asymptotically increasingly) short intervals. Our econometric setting allows for essentially arbitrary heterogeneity across the announcements and empirically realistic strong persistence in the volume intensity and volatility processes.

22. For example, the average jump sizes in the logarithmic volatility and volume intensity around the announcements are naturally measured by

$$f_1(\widehat{\mathbf{S}}) = \frac{1}{M}\sum_{\tau\in\mathcal{A}} \Delta\widehat{\log(\sigma_\tau)} \quad \text{and} \quad f_2(\widehat{\mathbf{S}}) = \frac{1}{M}\sum_{\tau\in\mathcal{A}} \Delta\widehat{\log(m_\tau)},$$

respectively.

23. This econometric framework also readily accommodates more general functional forms than the log-linear specification in (5.3), corresponding to more complicated transforms $f(\cdot)$. That said, we focus our empirical investigations on the log-linear specification, as the scatter plot in Figure 7 does not suggest any obvious non-linear dependencies between the volume intensity and volatility jumps (in logs).

24. For instance, the ISM indices and the Consumer Confidence Index are all released at 10:00 when volume and volatility both tend to be decreasing even on non-announcement days, while FOMC announcements mostly occur at 14:15 when volume and volatility are generally increasing (though to a much smaller extent).

themselves differences between the post- and pre-event quantities, our empirical strategy for jump estimation may be naturally thought of as a high-frequency DID type approach, in which we consider the event-control difference of the jump estimates as our measure for the abnormal movements in the volume intensity and volatility. We then regress these DID jump estimates to obtain the volume–volatility elasticity. Following the terminology of Li *et al*. (2017), who considered "jump regressions" involving discontinuous price increments, we refer to our new procedure as a *DID jump regression estimator*.[25]

Formally, with each announcement time τ, we associate a control group $\mathcal{C}(\tau)$ of non-announcement times. Based on this control group, we then correct for the intraday patterns in the "raw" jump estimators by differencing out the corresponding estimates averaged within the control group, resulting in the adjusted jump estimators

$$\begin{aligned} \Delta\widetilde{\log(m_\tau)} &\equiv \Delta\widehat{\log(m_\tau)} - \frac{1}{N_C}\sum_{\tau'\in C(\tau)} \Delta\widehat{\log(m_{\tau'})}, \\ \Delta\widetilde{\log(\sigma_\tau)} &\equiv \Delta\widehat{\log(\sigma_\tau)} - \frac{1}{N_C}\sum_{\tau'\in C(\tau)} \Delta\widehat{\log(\sigma_{\tau'})}, \end{aligned} \tag{5.4}$$

where N_C refers to the number of times in the control group.[26] Analogously to (5.3), we then estimate the parameters of interest by regressing the DID jump estimates as

$$\tilde{\theta} \equiv \underset{\theta}{\operatorname{argmin}} \sum_{\tau\in\mathcal{A}} \left(\Delta\widetilde{\log(m_\tau)} - (a_0 + b_0^\top X_{0,\tau}) - (a_1 + b_1^\top X_{1,\tau}) \cdot \Delta\widetilde{\log(\sigma_\tau)} \right)^2. \tag{5.5}$$

Note that $\tilde{\theta}$ depends not only on $(\widehat{S}_\tau)_{\tau\in\mathcal{A}}$ but also on $(\widehat{S}_{\tau'})_{\tau'\in\mathcal{C}}$, where $\mathcal{C}\equiv\cup_{\tau\in\mathcal{A}}\mathcal{C}(\tau)$ contains the times of all control groups. This estimator can be expressed as $\tilde{\theta} = f(\widetilde{\mathbf{S}})$ where $\widetilde{\mathbf{S}} \equiv (\widehat{S}_\tau)_{\tau\in\mathcal{T}}$ for $\mathcal{T}\equiv\mathcal{A}\cup\mathcal{C}$.

In summary, our estimation procedure consists of two steps. The first step is to estimate the jumps in the volume intensity and volatility processes via DID. The second step consists in estimating the parameters that describe the relationship between the two via a least-square regression. Since jumps are discontinuities, the resulting DID jump regression estimator formally bears some resemblance to that in RDD commonly used in empirical microeconomics for the estimation of treatment effects (see, *e.g.*, Lee and Lemieux, 2010). Importantly, however, our new econometric inference procedures (including the computation of standard errors) is very different and non-standard. Explicitly allowing for heterogeneity and dependence in the return and volume data, the sampling variability in $\tilde{\theta}$ arises exclusively from the non-parametric estimation errors in the pre- and post-event high-frequency-based volume intensity and volatility estimators, $\hat{m}_{\tau\pm}$ and $\hat{\sigma}_{\pm}$, respectively. While in theory we can characterize the resulting asymptotic covariance matrix and it would be possible to use it in the design of "plug-in" type standard errors, the control groups $\mathcal{C}(\tau)$ used for the different announcement times often partially overlap, which would severely complicate the formal derivation and implementation of the requisite formulas.

To facilitate the practical implementation, we instead propose a novel easy-to-implement local *i.i.d.* bootstrap procedure for computing the standard errors, which is a localized version

25. Unlike the price, the volume intensity and volatility are both latent processes and, as a result, our new estimator converges at a slower non-parametric rate, compared to the "jump regressions" studied by Li *et al*. (2017).

26. In our empirical analysis below, $\mathcal{C}(\tau)$ consists of the same time-of-day as τ over the previous $N_C = 22$ non-announcement days (roughly corresponding to the length of one trading month). We also experimented with the use of other control periods, including periods comprised of future non-event days, resulting in the same general conclusions as the DID results reported.

of the *i.i.d.* bootstrap of Gonçalves and Meddahi (2009). This procedure does not require the exact dependence of $\tilde{\theta}$ on $\widetilde{\mathbf{S}}$ to be fully specified. Instead, it merely requires repeated estimation over a large number of locally *i.i.d.* bootstrap samples for the pre-event and post-event windows around each of the announcement and control times. The "localization" is important, as it allows us to treat the conditional distributions as (nearly) constant, in turn permitting the use of an *i.i.d.* re-sampling scheme. The actual procedure is summarized by the following algorithm, for which the formal theoretical justification is given in the technical Appendix A.

5.1. *Bootstrap Algorithm*

Step 1: For each $\tau \in \mathcal{T}$, generate *i.i.d.* draws $(V^*_{i(\tau)-j}, r^*_{i(\tau)-j})_{1\le j\le k_n}$ and $(V^*_{i(\tau)+j}, r^*_{i(\tau)+j})_{1\le j\le k_n}$ from $(V_{i(\tau)-j}, r_{i(\tau)-j})_{1\le j\le k_n}$ and $(V_{i(\tau)+j}, r_{i(\tau)+j})_{1\le j\le k_n}$, respectively.

Step 2: Compute $\widetilde{\Delta\log(m_\tau)}^*$ and $\widetilde{\Delta\log(\sigma_\tau)}^*$ the same way as $\widetilde{\Delta\log(m_\tau)}$ and $\widetilde{\Delta\log(\sigma_\tau)}$, respectively, except that the original data $(V_{i(\tau)-j}, r_{i(\tau)-j})_{1\le |j|\le k_n}$ is replaced with $(V^*_{i(\tau)-j}, r^*_{i(\tau)-j})_{1\le |j|\le k_n}$. Similarly, compute $\tilde{\theta}^*$ according to (5.5) using re-sampled data.

Step 3: Repeat steps 1 and 2 a large number of times. Report the empirical standard errors of (the components of) $\tilde{\theta}^* - \tilde{\theta}$ as the standard errors of the original estimator $\tilde{\theta}$. □

It is important to stress that even though the re-sampling scheme is locally *i.i.d.*, the validity of this bootstrap procedure does *not* require the volume and return data to be actually *i.i.d.* Rather, the procedure allows for very general dynamic dependencies, and potentially highly persistent and non-stationary volatility and volume intensity processes.

Equipped with the new high-frequency DID jump regression estimator defined in equation (5.5) and the accompanying bootstrap procedure outlined above for calculating standard errors and conducting valid inference, we now turn to our main empirical findings.

6. VOLUME–VOLATILITY RELATIONSHIP AROUND PUBLIC ANNOUNCEMENTS

We begin our empirical investigations by verifying the occurrence of (on average) positive jumps in both trading volume intensity and return volatility around scheduled macroeconomic announcements. We document how these jumps, and the volume–volatility elasticity in particular, vary across different types of announcements. We then show how the variation in the elasticities observed across different announcements may be related to explanatory variables that serve as proxies for differences-of-opinion and, relatedly, notion of economic uncertainty. A more detailed analysis of FOMC announcements further highlights the important role played by the sentiment embedded in the FOMC statements accompanying each of the rate decisions. In addition to our main results based on the S&P 500 aggregate equity portfolio, we also report complementary results based on U.S. Treasury bond futures.

6.1. *Jumps and announcements*

Consistent with the basic tenet of information-based trading around public news announcements, the preliminary analysis underlying Figure 5 clearly suggests an increase in both volatility and trading intensity from 30 minutes before an FOMC announcement to the 30 minutes after the announcement. In order to more formally corroborate these empirical observations and extend them to a broader set of announcements, we report in Table 3 the average magnitudes of the logarithmic volatility and volume intensity jumps observed around our full combined set of news

TABLE 3

Volume and volatility jumps around public news announcements

Events	All	FOMC	ISMM	ISMNM	CC	Others
No. of Obs.	2130	109	160	158	160	1682
			No DID			
Log Volatility	0.090**	1.088**	0.162**	0.072**	0.114**	0.023**
	(0.004)	(0.027)	(0.017)	(0.018)	(0.015)	(0.005)
Log Volume	0.034**	1.410**	0.056**	−0.045**	0.049**	−0.045**
	(0.004)	(0.023)	(0.015)	(0.013)	(0.015)	(0.005)
			DID			
Log Volatility	0.152**	1.037**	0.256**	0.165**	0.209**	0.087**
	(0.006)	(0.029)	(0.017)	(0.020)	(0.017)	(0.005)
Log Volume	0.204**	1.329**	0.335**	0.233**	0.328**	0.118**
	(0.005)	(0.025)	(0.015)	(0.015)	(0.015)	(0.006)

Notes: The table reports the average logarithmic volatility jumps and the average logarithmic volume intensity jumps around all announcements (All), which are further categorized into FOMC announcements (FOMC), ISM Manufacturing Index (ISMM), ISM Non-Manufacturing Index (ISMNM), Consumer Confidence Index (CC), and other pre-scheduled macroeconomic announcements (Others). The top panel reports the raw statistics. The bottom panel adjusts for the intraday pattern via the DID method using the past 22 non-announcement days as the control group. The sample period spans 10, April 2001 to 30, September 2014. Bootstrapped standard errors (1000 repetitions) are reported in parentheses. **Indicates significance at the 1% level.

announcements, as well as the five specific news categories explicitly singled out in Table 2. To highlight the importance of properly correcting for the intraday patterns seen in Figure 3, we report and contrast both the "raw" $\Delta\widehat{\log(m_\tau)}$ and $\Delta\widehat{\log(\sigma_\tau)}$ jump estimates obtained without the DID correction and the DID corrected estimates formally based on equation (5.4), together with the new bootstrapped standard errors described in Section 5.

The top panel presents the "raw" jump statistics without the DID correction. The volatility jumps are always estimated to be positive and highly statistically significant. This is true for all of the announcements combined, as well as within each of the five separate categories. The volume jumps averaged across all news announcements are also significantly positive. However, the jumps in the volume intensities are estimated to be *negative* for two of the news categories: ISM Non-Manufacturing and Others. This is difficult to reconcile with any of the economic mechanisms and theoretical models discussed in Section 2. Instead, these negative estimates may be directly attributed to the strong diurnal pattern evident in Figure 3. The ISM indices and most of the economic news included in the Others category are announced at 10:00am, when both volatility and trading volume tend to be falling, thus inducing a downward bias in the jump estimation.

To remedy this, the bottom panel of Table 3 reports the corresponding DID estimates based on equation (5.4) in which we rely on the previous 22 non-announcement days as the control group. As these estimates show, applying the DID correction results in significantly *positive* jumps for the spot volatility and trading intensity across *all* of the different news categories, ISMNM and Others included. This contrast directly underscores the importance of properly controlling for the intradaily features outside the stylized theoretical models when studying volume and volatility at the high-frequency intraday level. At the same time, the magnitude of the jump estimates associated with FOMC announcements, which mostly occur between 14:00 and 14:15 when volatility and trading volume both tend to be rising, is actually reduced by the DID correction. Nevertheless, FOMC clearly stands out among all of the different news categories, as having the largest (by a wide margin) average jump sizes in both volume and volatility.

TABLE 4
Volume–volatility elasticities around public news announcements

	FOMC	ISMM	ISMNM	CC	Others
Constant (a_0)	0.586**	0.199**	0.119**	0.218**	0.050**
	(0.078)	(0.022)	(0.020)	(0.023)	(0.006)
Elasticity (a_1)	0.716**	0.529**	0.688**	0.521**	0.787**
	(0.067)	(0.064)	(0.069)	(0.075)	(0.022)
R^2	0.330	0.155	0.220	0.109	0.287

Notes: The table reports the results from the DID jump regression in equation (5.5) for the specification $\Delta\widetilde{\log(m_\tau)} = a_0 + a_1 \cdot \Delta\widetilde{\log(\sigma_\tau)}$, using the past 22 non-announcement days as the control group. The sample spans the period from 10, April 2001 to 30, September 2014. Bootstrapped standard errors (1000 repetitions) are reported in parentheses. **indicates significance at the 1% level.

Having documented the existence of highly significant positive jumps in both volume and volatility around public announcements, we next turn to the joint relationship between these jumps, focusing on the volume–volatility elasticity and the implications of the theoretical models discussed in Section 2.

6.2. *Volume–volatility elasticities around public news announcements*

The theoretical models that guide our empirical investigations are explicitly designed to highlight how trading volume and return volatility respond to well-defined public news announcements. As such, the models are inevitably stylized, with other influences (such as those underlying the intraday patterns and long-term trends evident in Figures 3 and 5, respectively) deliberately abstracted away. As discussed above, the DID method provides a way to guard against the non-trivial influence of the systematic intraday patterns. It also conveniently differences out other unmodelled nuisances, like trends, which would otherwise contaminate the estimates. Consequently, we rely on the DID jump regression approach throughout.

To begin, consider a basic specification of equation (5.5) without any explanatory variables (*i.e.*, $X_{0,\tau}$ and $X_{1,\tau}$ are both absent). Table 4 reports the resulting estimates for each of the different news categories. All of the estimated intercepts (*i.e.* a_0) are positive and highly statistically significant, indicative of higher trading intensities following public news announcements, even in the absence of heightened return volatility.[27] Put differently, abnormal bursts in trading volume around announcements are not always associated with abnormal price changes. This, of course, is directly in line with the key idea underlying the KP model that differences-of-opinion provides an additional trading motive over explicit shifts in investors' average opinion. Further corroborating the rank of FOMC as the most important news category released during regular trading hours, the estimated intercept is the largest for FOMC announcements.[28]

27. In contrast, the estimates obtained for a_0 without the DID correction, reported in the Supplementary Appendix, are significantly negative for ISMM, ISMNM and the Others news categories, underscoring the importance of properly controlling for the strong intraday patterns in the volume intensity and volatility.

28. In our reduced-form specification (5.1), the intercept term captures the abnormal trading volume in the "limiting" case in which the price change is zero. In the KP model this corresponds to the β_0 term in equation (2.1), which from equation (2.2) is directly related to the level of disagreement, but is also confounded by the degree of risk tolerance. In addition to these influences, the large estimated intercepts for the FOMC announcements may simply reflect the fact that the FOMC announcements are more closely watched by a larger pool of investors than the other announcements. In contrast, the volume–volatility elasticity and the slope coefficient being a relative measure, as formalized within the context of the KP model in equations (2.4) and (2.5), is free from these confounding influences. We therefore focus our attention on the easier to interpret elasticity estimates.

Turning to the volume–volatility elasticities (*i.e.* a_1), all of the estimates are below unity, and significantly so.[29] The theoretical derivations in (2.4) and (2.5) based on the KP model also predict that in the presence of differences-of-opinion the volume–volatility elasticity should be below unity. Our empirical findings are therefore directly in line with this theoretical prediction, and further support the idea that disagreements among investors often provide an important motive for trading.

6.3. *Volume–volatility elasticities and disagreement measures*

In addition to the prediction that the volume–volatility elasticity should be less than unity, our theoretical derivations in Section 2 also predict that the elasticity should be decreasing with the overall level of disagreement among investors. To examine this more refined theoretical prediction, we include a set of additional explanatory variables (in the form of the $X_{1,\tau}$ variable in the specification in equation (5.5)) that serve as proxies for disagreement. To account for the category-specific heterogeneity in the volume–volatility elasticity estimates reported in Table 4, we also include a full set of category dummy variables (*i.e.* one for each of the FOMC, ISMM, ISMNM and CC news categories).

We consider two proxies for the overall level of investors' disagreement that prevails at the time of the announcement. The first is the forecast dispersion of the one-quarter-ahead unemployment rate from the Survey of Professional Forecasters (SPF).[30] For aggregated macroeconomic time series, like the unemployment rate, the dispersion among forecasters arises more likely from differences in their economic/statistical models, rather than their "private" information regarding the aggregate macroeconomy. Hence, the SPF dispersion measure speaks directly to the notion of disagreement in the Kandel–Pearson model, where investors' disagreement arises from differences in their models used for interpreting the public signal. This same measure has also previously been used in other studies to gauge the degree of disagreement; see, *e.g.*, Van Nieuwerburgh and Veldkamp (2006) and Ilut and Schneider (2014) among others.

Secondly, as an indirect proxy for differences-of-opinion, we employ a weekly moving average of the economic policy uncertainty index developed by Baker *et al.* (2015).[31] There is a voluminous literature that addresses the relation between disagreement and uncertainty, generally supporting the notion of a positive relation between the two; see, for example, Acemoglu *et al.* (2006) and Patton and Timmermann (2010). Below, we refer to these two proxies as Dispersion and Weekly Policy, respectively. To facilitate comparisons, we scale both measures with their own sample standard deviations.

The estimation results for different specifications including these additional explanatory variables in the volume–volatility elasticity are reported in Table 5.[32] As a reference, the first column reports the results from a basic specification without any explanatory variables. The

29. The robust DID estimate for the elasticity around FOMC announcements reported in Table 4 is slightly larger than the preliminary raw non-DID estimate discussed in Section 4.

30. The SPF is a quarterly survey. It is released and collected in the second month of each quarter. To prevent any look-ahead bias (which may arise from using future information in the estimation of the current elasticity), we use the value from the previous quarter. Additional results for other forecast horizons and dispersion measures pertaining to other economic variables are reported in the Supplementary Appendix.

31. The economic policy uncertainty index of Baker *et al.* (2015) is based on newspaper coverage frequency. We use the weekly moving average so as to reduce the noise in the daily index. The averaging also naturally addresses the weekly cycle in the media. Comparable results based on the monthly index are available in the Supplementary Appendix.

32. We also include news-category dummies in the intercept $X_{0,\tau}$ in all of the different specifications, so as to control for the heterogeneity in the a_0 estimates in Table 4. Since our main focus centers on the volume–volatility elasticity, to conserve space we do not report these estimated b_0 dummy coefficients.

TABLE 5
Volume–volatility elasticity estimates and disagreement measures

Baseline estimates					
Constant (a_0)	0.044**	0.041**	0.040**	0.041**	0.041**
	(0.006)	(0.007)	(0.007)	(0.007)	(0.007)
Elasticity (a_1)	0.733**	0.776**	0.906**	0.921**	0.984**
	(0.020)	(0.024)	(0.043)	(0.036)	(0.045)
Estimates for explanatory variables in elasticity (b_1)					
News-category dummy variables					
FOMC		−0.060	−0.058	−0.048	−0.049
		(0.072)	(0.072)	(0.072)	(0.072)
ISMM		−0.238**	−0.238**	−0.220**	−0.222**
		(0.070)	(0.070)	(0.070)	(0.070)
ISMNM		−0.090	−0.092	−0.082	−0.085
		(0.077)	(0.077)	(0.077)	(0.077)
CC		−0.244**	−0.244**	−0.202**	−0.207**
		(0.073)	(0.073)	(0.073)	(0.073)
Disagreement measures					
Dispersion			−0.051**		−0.031*
			(0.013)		(0.014)
Weekly Policy				−0.079**	−0.070**
				(0.013)	(0.014)
R^2	0.481	0.482	0.483	0.486	0.486

Notes: The table reports the results from the DID jump regression in equation (5.5) for the specification $\Delta\widetilde{\log(m_\tau)} = a_0 + b_0^\top X_{0,\tau} + (a_1 + b_1^\top X_{1,\tau}) \cdot \Delta\widetilde{\log(\sigma_\tau)}$ based on all of the public announcements, using the past 22 non-announcement days as the control group. In all specifications, $X_{0,\tau}$ includes category dummy variables for FOMC rate decision (FOMC), ISM Manufacturing Index (ISMM), ISM Non-Manufacturing Index (ISMNM) and Consumer Confidence Index (CC); the estimates of these dummies (*i.e.* b_0) are not reported for brevity. The Dispersion variable is constructed as the latest forecast dispersion of the one-quarter-ahead unemployment rate from the Survey of Professional Forecasters before the announcement. The Weekly Policy variable is constructed as the weekly moving average before the announcement of the economic policy uncertainty index developed by Baker *et al.* (2015). Both variables are scaled by their own sample standard deviations. The sample spans 10, April 2001 to 30, September 2014. Bootstrapped standard errors (1000 repetitions) are reported in parentheses. * and ** indicate significance at the 5% and 1% level, respectively.

common elasticity is estimated to be 0.733 which, not surprisingly, is close to the average value of the category-specific estimates reported in Table 4. Underscoring the importance of disagreement more generally, the estimate is also significantly below one.

The specification in the second column includes the full set of news category dummies in the elasticity, with the baseline category being Others.[33] The elasticity for the Others category, which includes by far the largest number of announcements, is estimated to be 0.776 and close to the value of 0.733 from the specification without any dummies. The estimates for FOMC and ISM Non-Manufacturing announcements are also both statistically indistinguishable from this value of 0.776. On the other hand, the volume–volatility elasticities estimated around ISM Manufacturing and Consumer Confidence announcements are both significantly lower, indicating that the levels of disagreement among investors are higher for these events. To help understand this latter finding, we note that ISM Manufacturing and Consumer Confidence announcements are both released early in the macroeconomic news cycle, as described in Andersen *et al.* (2003). As documented in Andersen *et al.* (2003, 2007), the first announcements in a given news cycle tend to have larger price impacts than later related news announcements, as much of the information contained in the later releases may have already been gleaned from the earlier news announcements. Our

33. Although the full set of news-category dummy variables are included in both the intercept and the elasticity specifications, the estimates in the second column in Table 5 are not exactly identical to those in Table 4, because some of the announcements across the different news categories occur concurrently.

findings are in line with this logic. The estimated elasticities indicate relatively high levels of disagreement around the two early ISMM and CC announcements, and indirectly suggest that the release of these help resolve some of the economic uncertainty and decrease the overall level of disagreement, as manifest in the closer-to-one elasticity estimate for the later (in the news cycle) ISMNM releases.

The next two columns in the table report the results for specifications that include either Dispersion or Weekly Policy as an additional explanatory variable. Consistent with our theoretical prediction, both of these disagreement proxies significantly *negatively* impact the volume–volatility elasticity, and the estimated effects are also economically non-trivial. Moreover, the estimated elasticities obtained in the absence of any dispersion or economic policy uncertainty (*i.e.* the a_1 coefficients) are much closer to the no-disagreement benchmark of unity implied by the theoretical expressions in (2.4) and (2.5), than the corresponding estimates obtained without controlling for disagreement.

The last column shows that both of the disagreement measures remain statistically significant when included jointly, although less so for Dispersion. Interestingly, the baseline elasticity of 0.984 for Others (*i.e.* a_1) is no longer statistically different from unity, nor are the estimated elasticities for FOMC and ISM Non-Manufacturing.

Taken as a whole, these results strongly corroborate the existence of a negative relationship between the volume–volatility elasticity and the level of disagreement among investors. The results also suggest that for a majority of the public news announcements, our two specific disagreement proxies, involving measures of forecast dispersion and economic policy uncertainty, are able to explain the deviation in the volume–volatility elasticity from the no-disagreement benchmark of unity.

6.4. *Further analysis of FOMC announcements*

The results discussed in the previous section were based on the joint estimation involving all of the macroeconomic news announcements that occur during regular trading hours. Meanwhile, as documented in Table 3, the FOMC rate decisions rank supreme in inducing the on-average largest jumps in both trading activity and return volatility over our sample. These large responses occur in spite of the fact that the federal funds rate was fixed at the effective zero lower bound over much of the later half of the sample. Moreover, economists also routinely disagree about the interpretation of monetary policy. All of these unique features grant FOMC announcements of particular interest for our analysis pertaining to the role of disagreement in financial markets.

Before discussing the results from our more detailed analysis of the volume–volatility elasticity estimated exclusively around the times of FOMC announcements, it is important to stress some crucial differences between our analysis and prior empirical work related to FOMC announcements. In particular, there is already an extensive literature devoted to the study of the impact of FOMC announcements on equity returns. In addition to the earlier influential work by Bernanke and Kuttner (2005), this includes more recent studies specifically related to the behaviour of monetary policy and market reactions when the rate is at or near the zero lower bound (*e.g. Bernanke, 2012; Wright, 2012; van Dijk et al., 2014; Johnson and Paye, 2015*). Other recent studies have also documented that most of the equity risk premium is earned in specific phases of the FOMC news release cycle (*e.g. Savor and Wilson, 2014; Lucca and Moench, 2015; Cieslak et al., 2015*). It is not our intent to add to this burgeoning literature on the determinant of the equity risk premium, and the functioning of monetary policy per se. Instead, we simply recognize the unique position of FOMC announcements as the most important news category in our sample. Motivated by this fact, we further investigate how the variation in the

TABLE 6
Volume–volatility elasticity estimates around FOMC announcements

Baseline estimates						
Constant (a_0)	0.586**	0.579**	0.583**	0.545**	0.550**	0.545**
	(0.078)	(0.079)	(0.079)	(0.078)	(0.079)	(0.079)
Elasticity (a_1)	0.716**	0.996**	0.790**	0.913**	1.089**	1.055**
	(0.067)	(0.085)	(0.082)	(0.082)	(0.092)	(0.093)
Estimates for explanatory variables in elasticity (b_1)						
Dispersion		−0.108**			−0.088**	−0.096**
		(0.019)			(0.020)	(0.021)
Weekly Policy			−0.037			0.046
			(0.021)			(0.028)
FOMC Sentiment				−0.104**	−0.077**	−0.097**
				(0.023)	(0.024)	(0.028)
R^2	0.330	0.382	0.329	0.370	0.400	0.400

Notes: The table reports the results from the DID jump regression in equation (5.5) for the specification $\widetilde{\Delta\log(m_\tau)} = a_0+(a_1+b_1^\top X_{1,\tau})\cdot\widetilde{\Delta\log(\sigma_\tau)}$ based on FOMC announcements, using the past 22 non-announcement days as the control group. The news-category dummies, *i.e.*, $X_{0,\tau}$, are absent from this specification for the FOMC subsample. Dispersion and Weekly Policy are constructed as in Table 5. FOMC Sentiment is a textual measure constructed using financial-negative words in the FOMC press release. These variables are scaled by their own sample standard deviations. The sample spans 10, April 2001 to 30, September 2014. Bootstrapped standard errors (1000 repetitions) are reported in parentheses. * and ** indicate significance at the 5% and 1% level, respectively.

volume–volatility elasticity more generally varies with measures of disagreement prevailing at the exact time of and directly extracted from the FOMC news releases.

For ease of reference, the first column in Table 6 reports the DID estimation results for the FOMC subsample and the benchmark specification that does not include any explanatory variables in the elasticity, as previously reported in Table 4. The second column includes the previously defined Dispersion measure as an explanatory variable. In parallel to the full-sample results in Table 5, the estimates show that higher levels of forecast dispersions are generally associated with lower volume–volatility elasticities. The estimate of -0.108 for the b_1 coefficient is also highly statistically significant. Moreover, after controlling for Dispersion, the baseline elasticity (*i.e.* a_1) is virtually one. This finding thus suggests that Dispersion alone, as a measure of differences-of-opinion, is able to successfully explain much of the deviation from unity in the volume–volatility elasticity at the times of FOMC announcements.

The b_1 estimate for the Weekly Policy variable reported in the third column of Table 6 is also negative. However, it is not significant at conventional levels. This lack of significance of the Weekly Policy variable in the FOMC subsample, stands in sharp contrast with its highly significant effect in Table 5 based on the full sample of all announcements. This therefore suggests that the Weekly Policy variable, which is constructed as a "catch-all" measure of economic uncertainty, is simply too diverse (or noisy) to satisfactorily explain the variation in the volume–volatility relationship observed exclusively around FOMC announcements.

To remedy this, we construct an alternative textual measure based on the actual FOMC press releases. The FOMC statements, in addition to announcing the new target rates, also signal the future path of policy. In recent years, the statements also include brief summaries of the state of the economy, providing additional context underlying the rate decisions.[34] We succinctly summarize this additional information by counting the number of negative words, in accordance with the

34. This is especially important over the later half of our sample period, when the target rate was consistently stuck at the zero lower bound and, hence, offered little new information by itself. Bernanke (2012) also explicitly emphasized the important role of "public communications" as a nontraditional policy tool of the Fed.

financial-negative (Fin-Neg) word list compiled by Loughran and McDonald (2011). We refer to this textual measure as the FOMC Sentiment. A more detailed description of the construction is provided in Appendix B.2.[35] Assuming that the use of more negative words provides additional room for investors to differ in their interpretation of the news, we consider this alternative FOMC Sentiment measure as a more direct proxy for the level of disagreement at the exact times of the FOMC announcements.

From the theoretical relations derived in Section 2, we would therefore expect to see lower volume–volatility elasticities in response to higher FOMC Sentiment measures. The estimation results reported in the fourth column of Table 6 supports this theoretical prediction. The estimated b_1 coefficient for our FOMC Sentiment measure equals -0.104. It is also highly statistically significant. Moreover, controlling for the FOMC Sentiment, the baseline elasticity (*i.e.* a_1) is estimated to be 0.914, and this estimate is statistically indistinguishable from unity at conventional significance levels. Interestingly, the b_1 estimate of -0.104 for the FOMC Sentiment variable is also very close to the -0.108 estimate for the Dispersion measure reported in the second column.

To further gauge the relative merits of the Dispersion and FOMC Sentiment measures, the specification reported in the fifth column includes both as explanatory variables in the elasticity. Both of the estimated coefficients are negative and statistically significant. The coefficient estimates are also similar in magnitude, suggesting that the Dispersion and FOMC Sentiment measures are equally important in terms of capturing the disagreements-in-beliefs that motivate the abnormal trading at the times of FOMC announcements. Again, the estimate for the baseline elasticity a_1 is also not statistically different from the theoretical prediction of unity that should obtain in the absence of differences-of-opinion.

Further augmenting the specification to include the Weekly Policy measure as an additional explanatory variable in the elasticity does not change the key aspects of any of these findings, as shown by the results reported in the last column of the table. Counter to the previous empirical results and theoretical predictions, the estimated coefficient for the Weekly Policy variable in this expanded DID jump regression is actually positive, albeit not significant at conventional levels. This also indirectly supports our earlier conjecture that the textual-based FOMC Sentiment measure affords a much more pointed and accurate characterization of the economic uncertainty and differences-in-opinion at the exact times of the FOMC announcements, compared to the "catch-all" Weekly Policy measure.

6.5. *Bond futures and FOMC announcements*

Our empirical analysis so far has been focused exclusively on the aggregate stock market, as proxied by the S&P 500.[36] The theoretical implications developed in Section 2, however, should hold true more broadly for other markets. Hence, to further buttress the theoretical implications,

35. Loughran and McDonald (2011) originally constructed their Fin-Neg list for the purpose of analysing corporate 10-K reports. Compared with negative words, positive words tend to be less informative due to their more frequent negation. Along those lines, some Fin-Neg words detected in the FOMC statement may not actually have a negative meaning. One example is the word "late" in the context "... are likely to warrant exceptionally low levels for the federal funds rate at least through late 2014". Another example is "unemployment", which routinely appears in the first two paragraphs of the statements over the later half of our sample. To examine the severity of this issue, we manually checked every word that was classified as negative using Loughran and McDonald's list, and then refined this selection by only keeping words with an unambiguous negative meaning within their context. The regression results based on this refined measure are very similar to those based on the FOMC Sentiment measure reported here, and hence are omitted for brevity.

36. Additional robustness checks based on an ETF tracking the Dow Jones Industrial Average are provided in the Supplementary Appendix.

TABLE 7
Volume–volatility elasticity estimates around FOMC announcements for Treasury bond futures

Baseline estimates						
Constant(a_0)	1.245**	1.197**	1.227**	1.241**	1.203**	1.218**
	(0.076)	(0.075)	(0.078)	(0.076)	(0.076)	(0.077)
Elasticity(a_1)	0.531**	0.813**	0.627**	0.649**	0.832**	0.786**
	(0.066)	(0.086)	(0.089)	(0.074)	(0.087)	(0.094)
Estimates for explanatory variables in elasticity (b_1)						
Dispersion		−0.094**			−0.079**	−0.092**
		(0.018)			(0.019)	(0.021)
Weekly Policy			−0.039*			0.052
			(0.020)			(0.027)
FOMC Sentiment				−0.076**	−0.041	−0.067*
				(0.021)	(0.022)	(0.026)
R^2	0.314	0.364	0.315	0.337	0.364	0.364

Notes: The table reports the results for the ten-year Treasury bond futures data from the DID jump regression in equation (5.5) for the specification $\widetilde{\Delta\log(m_\tau)} = a_0 + (a_1 + b_1^\top X_{1,\tau}) \cdot \widetilde{\Delta\log(\sigma_\tau)}$ based on FOMC announcements, using the past 22 non-announcement days as the control group. The news-category dummies, *i.e.*, $X_{0,\tau}$, are absent from this specification for the FOMC subsample. Dispersion and Weekly Policy are constructed as in Table 5. FOMC Sentiment is a textual measure constructed using financial-negative words in the FOMC press release. These variables are scaled by their own sample standard deviations. The sample spans 1, July 2003 to 30, September 2014. Bootstrapped standard errors (1000 repetitions) are reported in parentheses. * and ** indicate significance at the 5% and 1% level, respectively.

we briefly analyse the behaviour of the volume–volatility elasticity in the fixed income market.[37] We rely on high-frequency price and volume data for the ten-year U.S. Treasury bond futures sampled at the one-minute frequency. As discussed in Section 3, our sample for the T-bond futures spans the slightly shorter period from 1, July 2003 to 30, September 2014.

In parallel to the preliminary analysis for the stock market in Section 4 and the more detailed analysis described in the previous subsection, we focus our analysis on the volume–volatility elasticity estimated around the FOMC announcements.[38] Table 7 reports the resulting estimates. The patterns are generally consistent with the comparable findings for the stock market reported in Table 6. Specifically, looking at the first column without any controls, the elasticity is estimated to be 0.531, and significantly lower than unity. Moreover, the coefficients associated with the various disagreement proxies included in columns two through five are all estimated to be negative (*i.e.* $b_1 < 0$), and controlling for these proxies pushes the baseline elasticity (*i.e.* a_1) closer towards one. In the most general specification reported in the last column, the Dispersion and FOMC Sentiment measures both exert a statistically significant influence on the elasticity, while the Weekly Policy measure is rendered insignificant by the inclusion of the former two measures. This compares directly to the results for the stock market in Table 6. Meanwhile, the baseline elasticity estimates for the bond market are slightly lower than for the stock market, suggesting that bond traders tend to disagree more (in a relative sense as formalized by equation (2.5) for the KP model) about the interpretation of the monetary policy compared to stock market traders.

37. As previously noted, prior work, notably Green (2004) and Pasquariello and Vega (2007), have analysed the price impact of order flow in the bond market around the time of macroeconomic news announcements. Our analysis specifically pertaining to the volume–volatility elasticity complements this work by allowing for a more direct interpretation of the estimates vis-a-vis the KP and differences-of-opinion class of models.

38. The graph for the log-volume intensity and log-spot volatility in the T-bond market immediately before and after all of the FOMC announcements included in the Supplementary Appendix closely mirrors the corresponding Figure 5 for the aggregate stock market, clearly suggesting the presence of jumps in both series at the FOMC news announcement times.

All in all, however, our findings for the bond futures data are very similar to those for the aggregate stock market portfolio, supporting the theoretical implications of the KP model and the importance of differences-of-opinion more generally.

7. CONCLUSION

We provide new empirical evidence concerning the behaviour of financial market volatility and trading activity in response to public news announcements. Our results are based on intraday prices and trading volume for the aggregate stock market portfolio and Treasury bond futures, along with new econometric procedures specifically designed to deal with the unique complications that arise in the high-frequency data setting. Explicitly zooming in on the volume–volatility changes right around the exact times of the announcements allows us to cast new light on the way in which financial markets process new information and function more generally.

Consistent with the implications from theoretical models involving economic agents who agree-to-disagree, we find that the sensitivity of abnormal volume changes with respect to those of volatility estimated around the times of the most important public news announcements, as embedded within the volume–volatility elasticity, is systematically below unity. Further corroborating the important role played by differences-of-opinion among market participants, the elasticity tends to be low during times of high economic policy uncertainty and high dispersion among professional economic forecasters. A direct textual-based measure of the negative sentiment in the FOMC statements accompanying the actual rate decisions also negatively impacts the elasticity estimated at FOMC announcement times, lending additional empirical support to our key theoretical predictions.

APPENDIX

A. TECHNICAL BACKGROUND FOR ECONOMETRIC PROCEDURES

This appendix presents the formal econometric theory behind our high-frequency econometric estimation and inference procedures discussed in Section 5. Appendix A.1 describes the continuous-time setup for modelling the high-frequency price and volume data used in our empirical analysis. Appendix A.2 presents the main theoretical results, which we prove in Appendix A.3.

A.1. *Continuous-time setup and definitions*

Throughout, we fix a filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$. Let $(P_t)_{t\geq 0}$ denote the logarithmic price process of an asset. As is standard in the continuous-time finance literature (see, *e.g.*, Merton (1992) and Duffie (2001)), we assume that P is a jump-diffusion process of the form

$$dP_t = b_t dt + \sigma_t dW_t + dJ_t, \tag{A.1}$$

where b is an instantaneous drift process, σ is a stochastic *spot volatility* process, W is a Brownian motion, and J is a pure jump process. The price is sampled at discrete times $\{i\Delta_n : 0 \leq i \leq [T/\Delta_n]\}$, where T denotes the sample span and Δ_n denotes the sampling interval of the high-frequency data. We denote the corresponding high-frequency asset returns by $r_i \equiv P_{i\Delta_n} - P_{(i-1)\Delta_n}$.

Our empirical analysis is justified using an infill econometric theory with $\Delta_n \to 0$ and T fixed. This setting is standard for analysing high-frequency data (see, *e.g.*, Aït-Sahalia and Jacod (2014) and Jacod and Protter (2012)) and it allows us to non-parametrically identify processes of interest in a general setting with essentially unrestricted nonstationarity and persistence.

We denote the trading volume within the high-frequency interval $((i-1)\Delta_n, i\Delta_n]$ by $V_{i\Delta_n}$. Unlike the price, the high-frequency volume data cannot be realistically modeled using the jump-diffusion model. Following Li and Xiu (2016), we consider a general state-space model

$$V_{i\Delta_n} = \mathcal{V}(\zeta_{i\Delta_n}, \epsilon_{i\Delta_n}), \tag{A.2}$$

where ζ is a latent state process, $(\epsilon_{i\Delta_n})$ are i.i.d. transitory shocks with distribution F_ϵ, and $\mathcal{V}(\cdot)$ is a possibly unknown transform. The latent state process ζ captures time-varying conditioning information such as the intensity of order arrival and the shape of the order size distribution. Technically, this state-space model can be formally defined on an *extension* of the space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$; see Li and Xiu (2016) for a formal construction.

This state-space model suits our empirical modelling of the trading volume well in two ways. First, it conveniently allows the observed volume data to be discretely valued, while still being completely flexible on the state space of the latent process ζ. This allows us to model ζ under minimal statistical restrictions without introducing unintended model inconsistency. Secondly, the setup does not restrict the dynamic persistence of the state process ζ. This is important as volume data often exhibit non-stationary behaviour.

Finally, to complete the notation, let

$$m_t = \int \mathcal{V}(\zeta_t, \epsilon) F_\epsilon(d\epsilon) \tag{A.3}$$

denote the instantaneous conditional mean process of V, that is, the *volume intensity* process. We use this process as the instantaneous empirical analogue to the expected volume in the theoretical models discussed in Section 2. Correspondingly, we use the spot volatility processes σ as the instantaneous analogue of the return standard deviation. Empirically, we should and do allow for general stochastic behaviour in these instantaneous moments so as to accommodate essentially arbitrary stochastic behaviours of time-varying conditioning information.

A.2. *Estimation and inference*

We now formally justify our inference procedures used in the main text. Below, we write $\tilde{\theta}_n$ in place of $\tilde{\theta}$ and similarly for other estimators, so as to emphasize the asymptotic stage. We focus on the DID jump regression estimator $\tilde{\theta}_n$ defined by equation (5.5), which includes the "raw" estimator $\hat{\theta}_n$ as a special case with the control group $\mathcal{C}$ set to be empty. We write $\tilde{\theta}_n = f(\widetilde{\mathbf{S}}_n)$, where $f(\cdot)$ is defined implicitly by the definitions (5.4) and (5.5).

We assume the following regularity conditions for the underlying processes.

Assumption 1. *(i) The price process P is given by (A.1) for $J_t = \int_0^t \xi_s dN_s + \int_0^t \int_{\mathbb{R}} \delta(s,z)\mu(ds,dz)$, where the processes b and σ are càdlàg (i.e. right continuous with left limit) and adapted; σ_t is positive for $t \in [0,T]$ almost surely; the process ξ is predictable and locally bounded; N is a counting process that jumps at the scheduled announcement times which are specified by the set $\mathcal{A}$; δ is a predictable function; μ is a Poisson random measure with compensator $\nu(ds,dz) = ds \otimes \lambda(dz)$ for some finite measure λ.*

(ii) The volume process V satisfies (A.2). The process ζ is càdlàg and adapted. The error terms (ϵ_i) take values in some Polish space, are defined on an extension of $(\Omega, \mathcal{F})$, i.i.d. and independent of $\mathcal{F}$.

(iii) For a sequence of stopping times $(T_m)_{m\geq 1}$ increasing to infinity and constants $(K_m)_{m\geq 1}$, we have $\mathbb{E}\left|\sigma_{t\wedge T_m} - \sigma_{s\wedge T_m}\right|^2 + \mathbb{E}\left|\zeta_{t\wedge T_m} - \zeta_{s\wedge T_m}\right|^2 \leq K_m |t-s|$ for all t,s such that $[s,t] \cap \mathcal{A} = \emptyset$.

Assumption 1 is fairly standard in the study of high-frequency data. Condition (i) allows the price process to contain jumps at both scheduled times and random times. Condition (ii) separates the conditional i.i.d. shocks $(\epsilon_{i\Delta_n})$ at observation times from the latent continuous-time state process (ζ_t). This condition only mildly restricts the volume series, which can still exhibit essentially unrestricted conditional and unconditional heterogeneity through the (typically highly persistent) time-varying state process (ζ_t). Condition (iii) imposes a mild smoothness condition on σ and ζ only in expectation, while allowing for general forms of jumps in their sample paths. This condition is satisfied for any semimartingales with absolutely continuous predictable characteristics (possibly with discontinuity points in $\mathcal{A}$) and for long-memory type processes driven by the fractional Brownian motion.

In addition, we need the following conditions for the non-parametric analysis, where we denote $M_p(\cdot) \equiv \int \mathcal{V}(\cdot, \epsilon)^p F_\epsilon(d\varepsilon)$ for $p \geq 1$.

Assumption 2. *$k_n \to \infty$ and $k_n^2 \Delta_n \to 0$.*

Assumption 3. *(i) The function $M_1(\cdot)$ is Lipschitz on compact sets and the functions $M_2(\cdot)$ and $M_4(\cdot)$ are continuous.*

(ii) Almost surely, the function f is well-defined and continuously differentiable in a neighborhood of $\widetilde{\mathbf{S}}$.

Assumption 2 specifies the growth rate of the local window size k_n. As typical in nonparametric analysis, this condition features a type of undersmoothing, so as to permit feasible inference. Assumption 3 imposes some smoothness conditions that are very mild.

We need some notations for stating the asymptotic results. For notational simplicity, we denote $v_t = M_2(\zeta_t) - M_1^2(\zeta_t)$. Consider variables $(\eta_{\tau-}, \eta_\tau, \eta'_{\tau-}, \eta'_\tau)_{\tau\in\mathcal{T}}$ which, conditionally on $\mathcal{F}$, are mutually independent, centered Gaussian with variances $(v_{\tau-}, v_\tau, \sigma_{\tau-}^2/2, \sigma_\tau^2/2)_{\tau\in\mathcal{T}}$. We denote the first differential of f at $\widetilde{\mathbf{S}}$ with increment $d\widetilde{\mathbf{S}}$ by $F(\widetilde{\mathbf{S}}; d\widetilde{\mathbf{S}})$. For a sequence Y_n of random variables, we write $Y_n \xrightarrow{\mathcal{L}\text{-}s} Y$ if Y_n converges stably in law towards Y, meaning that (Y_n, U) converges in distribution to (Y, U) for any bounded $\mathcal{F}$-measurable random variable U.

Theorem 1 characterizes the asymptotic distribution of the estimator $\tilde{\theta}_n$, and further shows that the asymptotic distribution can be consistently approximated using the Bootstrap Algorithm described in Section 5.

Theorem 1. *(a) Under Assumptions 1, 2 and 3*

$$\sqrt{k_n}(\tilde{\theta}_n - \theta) \xrightarrow{\mathcal{L}\text{-}s} F(\widetilde{\mathbf{S}}; (\eta_{\tau-}, \eta_\tau, \eta'_{\tau-}, \eta'_\tau, 0)_{\tau\in\mathcal{T}}). \tag{A.4}$$

(b) Moreover, the conditional distribution function of $\sqrt{k_n}(\tilde{\theta}^*_n - \tilde{\theta}_n)$ given the original data converges in probability to that of $F(\widetilde{\mathbf{S}}; (\eta_{\tau-}, \eta_\tau, \eta'_{\tau-}, \eta'_\tau, 0)_{\tau\in\mathcal{T}})$ under the uniform metric.

Comment. Since the variables $(\eta_{\tau-}, \eta_\tau, \eta'_{\tau-}, \eta'_\tau)_{\tau\in\mathcal{T}}$ are jointly $\mathcal{F}$-conditionally centered Gaussian, so is the limiting distribution given by (A.4).

A.3. *Proofs*

Throughout the proofs, we use K to denote a generic constant which may change from line to line. For a generic random sequence Y_n, we write $Y_n \xrightarrow{\mathcal{L}|\mathcal{F}} Y$ if the $\mathcal{F}$-conditional distribution function of Y_n converges in probability to that of Y under the uniform metric. By a classical localization procedure (see, *e.g.*, Section 4.4.1 in Jacod and Protter (2012)), we can assume that the processes σ and ζ are bounded, and piecewise $(1/2)$-Hölder continuous under the L^2-norm (with possible discontinuity points given by $\mathcal{A}$) without loss of generality. We consider a sequence Ω_n of events given by

$$\Omega_n \equiv \{\text{the intervals } [\tau - 2k_n\Delta_n, \tau + 2k_n\Delta_n],\ \tau \in \mathcal{T},\ \text{are mutually disjoint}\}.$$

Since $\mathcal{T}$ is finite, $\mathbb{P}(\Omega_n) \to 1$. Therefore, we can focus attention on Ω_n, again without loss of generality.

Proof of Theorem 1(a). We first show that

$$\sqrt{k_n}(\hat{m}_{n,\tau-} - m_{\tau-}, \hat{m}_{n,\tau} - m_\tau)_{\tau\in\mathcal{T}} \xrightarrow{\mathcal{L}|\mathcal{F}} (\eta_{\tau-}, \eta_\tau)_{\tau\in\mathcal{T}}. \tag{A.5}$$

In restriction to Ω_n, the estimators on the left-hand side of (A.5) are $\mathcal{F}$-conditionally independent. Therefore, it suffices to establish the convergence for the marginal distributions. We thus focus on $\sqrt{k_n}(\hat{m}_{n,\tau} - m_\tau)$, noting that the proof concerning $\sqrt{k_n}(\hat{m}_{n,\tau-} - m_{\tau-})$ is similar.

We decompose $\sqrt{k_n}(\hat{m}_{n,\tau} - m_\tau) = A_n + R_n$, where

$$A_n \equiv \frac{1}{\sqrt{k_n}} \sum_{j=1}^{k_n} \left(V_{(i(\tau)+j)\Delta_n} - M_1(\zeta_{(i(\tau)+j)\Delta_n})\right),$$

$$R_n \equiv \frac{1}{\sqrt{k_n}} \sum_{j=1}^{k_n} \left(M_1(\zeta_{(i(\tau)+j)\Delta_n}) - M_1(\zeta_\tau)\right).$$

By Assumption 3, $\mathbb{E}|R_n| \le K k_n \Delta_n^{1/2}$, which goes to zero by Assumption 2.

It remains to consider the convergence of A_n. Since $\left(V_{i\Delta_n}\right)_{i\ge 0}$ are $\mathcal{F}$-conditionally independent,

$$\mathbb{E}\left[A_n^2|\mathcal{F}\right] = \frac{1}{k_n}\sum_{j=1}^{k_n} v_{(i(\tau)+j)\Delta_n} \to v_\tau, \tag{A.6}$$

where the convergence holds because the process v is càdlàg. Moreover, by the continuity of $M_4(\cdot)$ and the boundedness of ζ, we can verify a Lyapunov-type condition:

$$\mathbb{E}\left[\left|V_{(i(\tau)+j)\Delta_n} - M_1(\zeta_{(i(\tau)+j)\Delta_n})\right|^4 |\mathcal{F}\right] \le K M_4\left(\zeta_{(i(\tau)+j)\Delta_n}\right) \le K. \tag{A.7}$$

By (A.6) and (A.7), we can apply the Lindeberg central limit theorem under the $\mathcal{F}$-conditional probability, resulting in $\sqrt{k_n}(\hat{m}_\tau - m_\tau) \xrightarrow{\mathcal{L}|\mathcal{F}} \eta_\tau$ as claimed.

Since the jumps of P are of finite activity, the returns involved in $\hat{\sigma}_{n,\tau-}$ and $\hat{\sigma}_{n,\tau}$ do not contain jumps, with probability approaching one. By Theorem 13.3.3(c) of Jacod and Protter (2012), we have

$$\sqrt{k_n}(\hat{\sigma}_{n,\tau-} - \sigma_{n,\tau-}, \hat{\sigma}_{n,\tau} - \sigma_{n,\tau})_{\tau\in\mathcal{T}} \xrightarrow{\mathcal{L}\text{-}s} \left(\eta'_{\tau-}, \eta'_\tau\right)_{\tau\in\mathcal{T}}. \tag{A.8}$$

By Proposition 5 of Barndorff-Nielsen *et al*. (2008) and the property of stable convergence in law, we can combine (A.5) and (A.8), yielding

$$\sqrt{k_n}(\tilde{\mathbf{S}}_n-\tilde{\mathbf{S}}) \xrightarrow{\mathcal{L}\text{-}s} \left(\eta_{\tau-},\eta_\tau,\eta'_{\tau-},\eta'_\tau,0\right)_{\tau\in\mathcal{T}}. \tag{A.9}$$

The assertion (A.4) then follows from (A.9) and the delta method. *Q.E.D.*

Proof of Theorem 1(b). Step 1. We divide the proof into several steps. Denote $\mathcal{G}\equiv\mathcal{F}\vee\sigma\{\epsilon_i : i\geq 0\}$. In this step, we show that, for each $\tau\in\mathcal{T}$,

$$\sqrt{k_n}(\hat{m}^*_{n,\tau-}-\hat{m}_{\tau-},\hat{\sigma}^*_{n,\tau-}-\hat{\sigma}_{\tau-}) \xrightarrow{\mathcal{L}|\mathcal{G}} \left(\eta_{\tau-},\eta'_{\tau-}\right). \tag{A.10}$$

Observe that

$$\sqrt{k_n}\begin{pmatrix}\hat{m}^*_{\tau-}-\hat{m}_{\tau-}\\ \hat{\sigma}^{*2}_{\tau-}-\hat{\sigma}^2_{\tau-}\end{pmatrix}=\frac{1}{\sqrt{k_n}}\sum_{j=1}^{k_n}\left(\begin{pmatrix}V^*_{(i(\tau)-j)\Delta_n}\\ r^{*2}_{i(\tau)-j}/\Delta_n\end{pmatrix}-\frac{1}{k_n}\sum_{j=1}^{k_n}\begin{pmatrix}V_{(i(\tau)-j)\Delta_n}\\ r^2_{i(\tau)-j}/\Delta_n\end{pmatrix}\right). \tag{A.11}$$

By the construction of the bootstrap sample, the summands in the right-hand side of (A.11) are i.i.d. with zero mean conditional on $\mathcal{G}$. We denote the $\mathcal{G}$-conditional covariance matrix of $\sqrt{k_n}(\hat{m}^*_{\tau-}-\hat{m}_{\tau-},\hat{\sigma}^{*2}_{\tau-}-\hat{\sigma}^2_{\tau-})$ by

$$\Sigma_{n,\tau-}=\begin{pmatrix}\Sigma^{(11)}_{n,\tau-} & \Sigma^{(12)}_{n,\tau-}\\ \Sigma^{(12)}_{n,\tau-} & \Sigma^{(22)}_{n,\tau-}\end{pmatrix},$$

where

$$\begin{cases}\Sigma^{(11)}_{n,\tau-}=\frac{1}{k_n}\sum_{j=1}^{k_n}V^2_{(i(\tau)-j)\Delta_n}-\left(\frac{1}{k_n}\sum_{j=1}^{k_n}V_{(i(\tau)-j)\Delta_n}\right)^2,\\ \Sigma^{(12)}_{n,\tau-}=\frac{1}{k_n\Delta_n}\sum_{j=1}^{k_n}V_{(i(\tau)-j)\Delta_n}r^2_{i(\tau)-j}-\left(\frac{1}{k_n}\sum_{j=1}^{k_n}V_{(i(\tau)-j)\Delta_n}\right)\left(\frac{1}{k_n\Delta_n}\sum_{j=1}^{k_n}r^2_{i(\tau)-j}\right),\\ \Sigma^{(22)}_{n,\tau-}=\frac{1}{k_n\Delta_n^2}\sum_{j=1}^{k_n}r^4_{i(\tau)-j}-\left(\frac{1}{k_n\Delta_n}\sum_{j=1}^{k_n}r^2_{i(\tau)-j}\right)^2.\end{cases}$$

In step 2, we shall show that

$$\left(\Sigma^{(11)}_{n,\tau-},\Sigma^{(12)}_{n,\tau-},\Sigma^{(22)}_{n,\tau-}\right)\xrightarrow{\mathbb{P}}\left(v_{\tau-},0,2\sigma^4_{\tau-}\right). \tag{A.12}$$

For each subsequence $\mathbb{N}_1\subseteq\mathbb{N}$, we can find a further subsequence $\mathbb{N}_2\subseteq\mathbb{N}_1$ such that the convergence in (A.12) holds almost surely along $\mathbb{N}_2$. Moreover, on the paths for which this convergence holds, we apply a central limit theorem under the $\mathcal{G}$-conditional probability to deduce the convergence of the conditional law of $\sqrt{k_n}(\hat{m}^*_{\tau-}-\hat{m}_{\tau-},\hat{\sigma}^*_{\tau-}-\hat{\sigma}_{\tau-})$ towards that of $\left(\eta_{\tau-},\eta'_{\tau-}\right)$ along the subsequence $\mathbb{N}_2$. From here, we deduce (A.10) for the original sequence by reversing the subsequence argument.

Step 2. We show (A.12) in this step. We start with $\Sigma^{(11)}_{n,\tau-}$. Consider the decomposition

$$\frac{1}{k_n}\sum_{j=1}^{k_n}V^2_{(i(\tau)-j)\Delta_n}=A_{1,n}+A_{2,n},$$

where

$$A_{1,n}\equiv\frac{1}{k_n}\sum_{j=1}^{k_n}\left(V^2_{(i(\tau)-j)\Delta_n}-M_2(\zeta_{(i(\tau)-j)\Delta_n})\right),$$

$$A_{2,n}\equiv\frac{1}{k_n}\sum_{j=1}^{k_n}M_2(\zeta_{(i(\tau)-j)\Delta_n}).$$

Note that the summands in $A_{1,n}$ are $\mathcal{F}$-conditionally independent with zero mean. Therefore,

$$\mathbb{E}\left[A^2_{1,n}|\mathcal{F}\right]\leq Kk_n^{-2}\sum_{j=1}^{k_n}M_4(\zeta_{(i(\tau)-j)\Delta_n})\leq Kk_n^{-1}\to 0,$$

which further implies that $A_{1,n}=o_p(1)$. In addition, since $M_2(\cdot)$ is continuous and ζ is càdlàg, $A_{2,n}\to M_2(\zeta_{\tau-})$. From here and (A.9), it follows that

$$\Sigma^{(11)}_{n,\tau-}\xrightarrow{\mathbb{P}}M_2(\zeta_{\tau-})-M_1(\zeta_{\tau-})^2=v_{\tau-}.$$

Next, we consider the limiting behaviour of $\Sigma^{(12)}_{n,\tau-}$. We decompose

$$\frac{1}{k_n\Delta_n}\sum_{j=1}^{k_n}V_{(i(\tau)-j)\Delta_n}r^2_{i(\tau)-j}=A_{3,n}+A_{4,n}+A_{5,n}+A_{6,n},$$

where

$$A_{3,n} \equiv \frac{1}{k_n \Delta_n} \sum_{j=1}^{k_n} \left(V_{(i(\tau)-j)\Delta_n} - M_1(\zeta_{(i(\tau)-j)\Delta_n}) \right) r_{i(\tau)-j}^2,$$

$$A_{4,n} \equiv \frac{1}{k_n \Delta_n} \sum_{j=1}^{k_n} \left(M_1(\zeta_{(i(\tau)-j)\Delta_n}) - M_1(\zeta_{(i(\tau)-j-1)\Delta_n}) \right) r_{i(\tau)-j}^2,$$

$$A_{5,n} \equiv \frac{1}{k_n \Delta_n} \sum_{j=1}^{k_n} M_1(\zeta_{(i(\tau)-j-1)\Delta_n}) \left(r_{i(\tau)-j}^2 - \mathbb{E}\left[r_{i(\tau)-j}^2 | \mathcal{F}_{i(\tau)-j-1} \right] \right),$$

$$A_{6,n} \equiv \frac{1}{k_n \Delta_n} \sum_{j=1}^{k_n} M_1(\zeta_{(i(\tau)-j-1)\Delta_n}) \mathbb{E}\left[r_{i(\tau)-j}^2 | \mathcal{F}_{i(\tau)-j-1} \right].$$

We now show that

$$A_{j,n} = o_p(1), \quad \text{for} \quad j = 3,4,5. \tag{A.13}$$

We note that $A_{3,n}$ is an average of variables that are, conditionally on $\mathcal{F}$, independent with zero mean. Therefore, $\mathbb{E}\left[A_{3,n}^2 | \mathcal{F}\right] \leq k_n^{-2} \Delta_n^{-2} \sum_{j=1}^{k_n} M_2(\zeta_{(i(\tau)-j)\Delta_n}) r_{i(\tau)-j}^4$, which further implies $\mathbb{E}\left[A_{3,n}^2\right] \leq K k_n^{-1} \to 0$. Turning to $A_{4,n}$, we observe

$$\begin{aligned}
\mathbb{E}|A_{4,n}| &\leq \frac{K}{k_n \Delta_n} \sum_{j=1}^{k_n} \mathbb{E}\left[\left| \zeta_{(i(\tau)-j)\Delta_n} - \zeta_{(i(\tau)-j-1)\Delta_n} \right| \left| r_{i(\tau)-j} \right|^2 \right] \\
&\leq \frac{K}{k_n \Delta_n} \sum_{j=1}^{k_n} \left(\mathbb{E}\left[\left| \zeta_{(i(\tau)-j)\Delta_n} - \zeta_{(i(\tau)-j-1)\Delta_n} \right|^2 \right] \mathbb{E}\left[r_{i(\tau)-j}^4 \right] \right)^{1/2} \\
&\leq K \Delta_n^{1/2} \to 0,
\end{aligned}$$

where the first inequality holds because $M_1(\cdot)$ is locally Lipschitz and ζ is bounded; the second inequality follows from the Cauchy–Schwarz inequality; the third inequality is due to the (1/2)-Hölder continuity of ζ and $\mathbb{E}[r_{i(\tau)-j}^4] \leq K \Delta_n^2$. Finally, we note that $A_{5,n}$ is an average of k_n martingale difference terms with bounded second moments. Therefore, $\mathbb{E}[A_{5,n}^2] \leq K k_n^{-1} \to 0$. The proof for (A.13) is now finished.

We now consider $A_{6,n}$. By Itô's formula and some standard estimates for continuous Itô semimartingales (noting that $r_{i(\tau)-j}$ does not contain jumps in restriction to Ω_n),

$$\begin{aligned}
&\mathbb{E}\left[r_{i(\tau)-j}^2 | \mathcal{F}_{i(\tau)-j-1} \right] - \sigma_{(i(\tau)-j-1)\Delta_n}^2 \Delta_n \\
&\quad = \int_{(i(\tau)-j-1)\Delta_n}^{(i(\tau)-j)\Delta_n} \mathbb{E}\left[\sigma_s^2 - \sigma_{(i(\tau)-j-1)\Delta_n}^2 | \mathcal{F}_{(i(\tau)-j-1)\Delta_n} \right] ds + o_{pu}(\Delta_n),
\end{aligned}$$

where $o_{pu}(\Delta_n)$ denotes a term that is $o_p(\Delta_n)$ uniformly in j. Hence,

$$\begin{aligned}
&\left| A_{6,n} - \frac{1}{k_n} \sum_{j=1}^{k_n} M_1(\zeta_{(i(\tau)-j-1)\Delta_n}) \sigma_{(i(\tau)-j-1)\Delta_n}^2 \right| \\
&\quad \leq \frac{K}{k_n} \sum_{j=1}^{k_n} \frac{1}{\Delta_n} \int_{(i(\tau)-j-1)\Delta_n}^{(i(\tau)-j)\Delta_n} \mathbb{E}\left[\left| \sigma_s^2 - \sigma_{(i(\tau)-j-1)\Delta_n}^2 \right| \middle| \mathcal{F}_{(i(\tau)-j-1)\Delta_n} \right] ds + o_p(1).
\end{aligned}$$

We further note that the expectation of the first term on the majorant side of the above display is bounded by $K\mathbb{E}\left[\sup_{s \in [\tau - 2k_n\Delta_n, \tau)} |\sigma_s^2 - \sigma_{\tau-}^2|\right]$, which goes to zero since σ is càdlàg and bounded. Therefore,

$$A_{6,n} = \frac{1}{k_n} \sum_{j=1}^{k_n} M_1(\zeta_{(i(\tau)-j-1)\Delta_n}) \sigma_{(i(\tau)-j-1)\Delta_n}^2 + o_p(1) \xrightarrow{\mathbb{P}} M_1(\zeta_{\tau-}) \sigma_{\tau-}^2.$$

From here, it readily follows that

$$\Sigma^{(12)}_{n,\tau-} \xrightarrow{\mathbb{P}} 0.$$

We now turn to the limiting behavior of $\Sigma^{(22)}_{n,\tau-}$. For each $i \in \{i(\tau)-j : 1 \le j \le k_n\}$, we can decompose $r_i = r_i' + r_i''$ in restriction to Ω_n where

$$\begin{aligned} r_i' &= \sigma_{(i-1)\Delta_n}(W_{i\Delta_n} - W_{(i-1)\Delta_n}), \\ r_i'' &= \int_{(i-1)\Delta_n}^{i\Delta_n} b_s ds + \int_{(i-1)\Delta_n}^{i\Delta_n} \left(\sigma_s - \sigma_{(i-1)\Delta_n}\right) dW_s. \end{aligned}$$

We can then decompose

$$\frac{1}{k_n \Delta_n^2} \sum_{j=1}^{k_n} r^4_{i(\tau)-j} = A_{7,n} + A_{8,n}, \tag{A.14}$$

where $A_{7,n} \equiv \frac{1}{k_n\Delta_n^2}\sum_{j=1}^{k_n} r'^4_{i(\tau)-j}$ and $A_{8,n}$ is defined implicitly by (A.14). Note that $\mathbb{E}[r_i'^4/\Delta_n^2|\mathcal{F}_{(i-1)\Delta_n}] = 3\sigma^4_{(i-1)\Delta_n}$ and the variance of the martingale difference sequence $r_i'^4/\Delta_n^2 - \mathbb{E}\left[r_i'^4/\Delta_n^2|\mathcal{F}_{(i-1)\Delta_n}\right]$ is bounded. It is then easy to see that $A_{7,n} \xrightarrow{\mathbb{P}} 3\sigma^4_{\tau-}$. Turning to the term $A_{8,n}$, we first observe

$$\left|A_{8,n}\right| \le \frac{K}{k_n \Delta_n^2} \sum_{j=1}^{k_n} \left|r'_{i(\tau)-j}\right|^3 \left|r''_{i(\tau)-j}\right| + \frac{K}{k_n \Delta_n^2} \sum_{j=1}^{k_n} \left|r''_{i(\tau)-j}\right|^4. \tag{A.15}$$

By the Burkholder–Davis–Gundy inequality and Hölder's inequality,

$$\mathbb{E}\left[\left|r''_{i(\tau)-j}\right|^4\right] \le K\Delta_n^4 + K\Delta_n \mathbb{E}\left[\int_{(i(\tau)-j-1)\Delta_n}^{(i(\tau)-j)\Delta_n} \left(\sigma_s - \sigma_{(i(\tau)-j-1)\Delta_n}\right)^4 ds\right]. \tag{A.16}$$

Hence,

$$\begin{aligned} \mathbb{E}\left[\frac{1}{k_n\Delta_n^2}\sum_{j=1}^{k_n}\left|r''_{i(\tau)-j}\right|^4\right] &\le K\Delta_n^2 + K\frac{1}{k_n\Delta_n}\mathbb{E}\left[\int_{(i(\tau)-k_n-1)\Delta_n}^{(i(\tau)-1)\Delta_n}\left(\sigma_s - \sigma_{(i(\tau)-j-1)\Delta_n}\right)^4 ds\right] \\ &\le K\Delta_n^2 + K\mathbb{E}\left[\sup_{t\in[\tau-2k_n\Delta_n,\tau)} |\sigma_t - \sigma_{\tau-}|^4\right] \to 0, \end{aligned}$$

where the convergence follows from the bounded convergence theorem, because σ is càdlàg and bounded. Therefore, the second term on the majorant side of (A.15) is $o_p(1)$. Since $A_{7,n} = O_p(1)$, we can use Hölder's inequality to show that the first term on majorant side of (A.15) is also $o_p(1)$. Hence, $A_{8,n} = o_p(1)$ and $k_n^{-1}\Delta_n^{-2}\sum_{j=1}^{k_n} r^4_{i(\tau)-j} \xrightarrow{\mathbb{P}} 3\sigma^4_{\tau-}$. Recall that $\frac{1}{k_n\Delta_n}\sum_{j=1}^{k_n} r^2_{i(\tau)-j} \xrightarrow{\mathbb{P}} \sigma^2_{\tau-}$. Hence, $\Sigma^{(22)}_{n,\tau-} \xrightarrow{\mathbb{P}} 2\sigma^4_{\tau-}$ as asserted. The proof for (A.12) is now complete.

Step 3. We finish the proof of part (b) of Theorem 1 in this step. By essentially the same argument for (A.10), we can show that

$$\sqrt{k_n}(\hat{m}^*_{n,\tau} - \hat{m}_{n,\tau}, \hat{\sigma}^*_{n,\tau} - \hat{\sigma}_{n,\tau}) \xrightarrow{\mathcal{L}|\mathcal{G}} \left(\eta_\tau, \eta'_\tau\right). \tag{A.17}$$

Moreover, since the variables in (A.10) and (A.17) are $\mathcal{G}$-conditionally independent between the pre- and post-event windows and across $\tau \in \mathcal{T}$ (by the construction of the local i.i.d. resampling scheme), the joint convergence holds as well. Therefore,

$$\sqrt{k_n}(\widetilde{\mathbf{S}}^*_n - \widetilde{\mathbf{S}}_n) \xrightarrow{\mathcal{L}|\mathcal{G}} \left(\eta_{\tau-}, \eta_\tau, \eta'_{\tau-}, \eta'_\tau, 0\right)_{\tau\in\mathcal{T}}. \tag{A.18}$$

Note that, with probability approaching one, $\widetilde{\mathbf{S}}_n$ falls in the neighbourhood of $\widetilde{\mathbf{S}}$ on which f is continuously differentiable. The assertion of the theorem then follows from the delta method. *Q.E.D.*

B. ADDITIONAL DATA DESCRIPTION

B.1. *Macroeconomic news announcements*

TABLE B.1
Other macroeconomic announcements

Category by Bloomberg	Index	No.Obs.	Time[a]
Labor Market	JOLTS Job Openings	47	10:00
Retail and Wholesale	Wards Total Vehicle Sales	10	14:41
Retail and Wholesale	Wholesale Inventories	160	10:00
Cyclical Indicators	Chicago Purchasing Manager	69	10:00
Cyclical Indicators	Dallas Fed Manf Activity	68	10:30
Cyclical Indicators	ISM Milwaukee	58	10:00
Cyclical Indicators	Kansas City Fed Manf Activity	37	11:00
Cyclical Indicators	Leading Index	161	10:00
Cyclical Indicators	Philadelphia Fed Business Outlook	162	10:00
Cyclical Indicators	Richmond Fed Manufact Index	108	10:00
Cyclical Indicators	U of Mich Sentiment	78	10:00
Personal Household Sector	Consumer Credit	159	15:00
Housing and Real Estate	Existing Home Sales	116	10:00
Housing and Real Estate	FHFA House Price Index	56	10:00
Housing and Real Estate	House Price Purchase Index	16	10:00
Housing and Real Estate	NAHB Housing Market Index	138	13:00
Housing and Real Estate	New Home Sales	160	10:00
Housing and Real Estate	One Family Home Resales	47	10:00
Housing and Real Estate	Pending Home Sales	114	10:00
Government Finance And Debt	Monthly Budget Statement	155	14:00

Notes: The table lists all of the intraday news announcements included in the "Others" category. [a]Some of these release times vary over the sample. To ensure that the pre- and post-event windows both have thirty observations, we only keep announcements between 10:00 and 15:30. The times indicated in the table refer to the most common release times.

B.2. *A detailed description for the FOMC sentiment measure*

Our detailed analysis of the volume–volatility relationship around FOMC announcements in Section 6.4 uses a textual measure of the negative sentiment in the actual written statements accompanying each of the FOMC rate decisions. Our construction of this sentiment measure closely follows Loughran and McDonald (2011), henceforth LM.

Specifically, we begin by extracting the actual text for each of the FOMC statements in our sample, excluding the paragraph detailing the voting decisions and the paragraph stating the new target funds rate. The individual words in each of the statements are then compared with the Fin-Neg list of LM, resulting in a total of 79 different negative words. Table B.2 provides a count of ten such words that occur most frequently over the full sample. The count of words that are of different forms of the same word (*i.e.* inflections) are summed up. For example, the 65 count for the word "slow" also includes occurrences of "slow" "slowed" "slower" "slowing" and "slowly".

Following LM, we then construct a weighted measure for the *ith* negative word in the *jth* statement according to the following formula:

$$w_{i,j} \equiv \begin{cases} \dfrac{1+\log(tf_{i,j})}{1+\log(a_j)} \log\left(\dfrac{N}{df_i}\right), & tf_{i,j} \geq 1, \\ 0, & \text{otherwise}, \end{cases} \tag{B.1}$$

where $tf_{i,j}$ refers to the raw count of the *ith* word in the *jth* statement, a_j denotes the total number of words in the *jth* statement, and df_i is the number of statements out of the total number of statements N containing at least one occurrence of the *ith* word. Importantly, this non-linear weighting scheme implies that negative words that seldom occur will receive a higher weight than more commonly used negative words. The overall negative sentiment for the *jth* statement is then simply obtained by summing these individual weights over all of the 79 negative words that occurred over the full sample, that is,

$$\text{FOMC Sentiment}_j = \sum_{i=1}^{79} w_{i,j}. \tag{B.2}$$

The words included in the Fin-Neg list are, of course, somewhat subjective. Also, certain words do not necessarily have a negative meaning in the context of FOMC announcements. For example, the most frequently used word "unemployment",

TABLE B.2
Word counts in FOMC sentiment

Word	Unemployment	Slow	Weak	Decline	Against	Depressed	Downward	Diminish	Concern	Persist
Count	73	65	53	39	23	23	21	18	16	16

Notes: The table reports the counts of the ten most frequently occurring financial-negative words in the FOMC statements over the April 2001 to September 2014 sample period.

which may sound negative in the 10-K filings analysed by LM, may be a neutral word in the FOMC statements, simply used to summarize the economic conditions, whether good or bad. Similarly, the word "late" does not have a negative meaning when used in a sentence like: "... are likely to warrant exceptionally low levels for the federal funds rate at least through late 2014". To guard against such ambiguities, we explicitly checked the usage of each of the words on the Fin-Neg list in the FOMC statements, and then select only those occurrences whose meaning were unambiguously negative. The resulting "selective" measure based on a total of sixty-seven different words is highly correlated with FOMC Sentiment (correlation coefficient of 0.73). All of the results based on this more selective sentiment measure are also very close to those based on the original FOMC Sentiment measure reported in Section 6.4 of the main text. These additional results are available upon request.

Acknowledgments. We would like to thank the Editor (Aureo de Paula) and three anonymous referees for their many insightful comments and suggestions. We would also like to thank Yacine Aït-Sahalia, Julio Crego, Joel Hasbrouck, Nour Meddahi, Ulrich Müller, Enrique Sentana, Neil Shephard, George Tauchen, Mark Watson, Brian Weller, Wei Xiong, Dacheng Xiu, and seminar participants at Duke University, Princeton University, Toulouse School of Economics and CEMFI for their helpful comments. Bollerslev and Li gratefully acknowledge support from CREATES funded by the Danish National Research Foundation (DNRF78) and the NSF (SES-1326819), respectively.

Supplementary Data

Supplementary data are available at Review of Economic Studies online.

REFERENCES

ACEMOGLU, D., CHERNOZHUKOV, V. and YILDIZ, M. (2006), "Learning and disagreement in an uncertain world" (*NBER Working Paper No.19814*).

AÏT-SAHALIA, Y. and JACOD, J. (2014), *High-Frequency Financial Econometrics* (Princeton, NJ: Princeton University Press).

ANDERSEN, T. G. (1996), "Return Volatility and Trading Volume: An Information Flow Interpretation of Stochastic Volatility", *Journal of Finance*, **51**, 169–204.

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

ANDERSEN, T. G., BOLLERSLEV, T., DIEBOLD, F. X., *et al.* (2003), "Micro Effects of Macro Announcements: Real-Time Price Discovery in Foreign Exchange", *American Economic Review*, **93**, 38–62.

—— (2007), "Real-Time Price Discovery in Global Stock, Bond and Foreign Exchange Markets", *Journal of International Economics*, **73**, 251–277.

ASHENFELTER, O. and CARD, D. (1985), "Using the Longitudinal Structure of Earnings to Estimate the Effect of Training Programs", *Review of Economics and Statistics*, **67**, 648–660.

BAILLIE, R. T. and BOLLERSLEV, T. (1990), "Intra-Day and Inter-Market Volatility in Foreign Exchange Rates", *Review of Economic Studies*, **58**, 565–585.

BAKER, S. R., BLOOM, N. and DAVIS, S. J. (2015), "Measuring Economic Policy Uncertainty" (*NBER Working Paper No.21633*).

BANERJEE, S. and KREMER, I. (2010), "Disagreement and Learning: Dynamic Patterns of Trade", *Journal of Finance*, **65**, 1269–1302.

BARNDORFF-NIELSEN, O. E., HANSEN, P. R., LUNDE, A., *et al*. (2008), "Designing Realized Kernels to Measure the Ex Post Variation of Equity Prices in the Presence of Noise", *Econometrica*, **76**, 1481–1536.

—— (2009), "Realized Kernels in Practice: Trades and Quotes", *Econometrics Journal*, **12**, C1–C32.

BERNANKE, B. (2012), "Monetary Policy Since the Onset of the Crisis", In *Federal Reserve Bank of Kansas City Economic Symposium, Jackson Hole, Wyoming*, vol. 31. https://www.kansascityfed.org/publications/research/escp/symposiums/escp-2012

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

BERNILE, G., HU, J. and TANG, Y. (2016), "Can Information be Locked up? iNformed Trading Ahead of Macro-News Announcements", *Journal of Financial Economics*, **121**, 496–520.

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

BOUDT, K. and PETITJEAN, M. (2014), "Intraday Liquidity Dynamics and News Releases Around Price Jumps: Evidence from the DJIA Stocks", *Journal of Financial Markets*, **17**, 121–149.

BROWNLEES, C. T. and GALLO, G. M. (2006), "Financial Econometric Analysis at Ultra-High Frequency: Data Handling Concerns", *Computational Statistics & Data Analysis*, **51**, 2232–2245.

CHABOUD, A. P., CHERNENKO, S. V. and WRIGHT, J. H. (2008), "Trading Activity and Macroeconomic Announcements in High-Frequency Exchange Rates", *Journal of the European Economic Association*, **6**, 589–596.

CHORDIA, T., ROLL, R. and SUBRAHMANYAM, A. (2002), "Order Imbalance, Liquidity, and Market Returns", *Journal of Financial Economics*, **65**, 111–130.

CIESLAK, A., MORSE, A. and VISSING-JORGENSEN, A. (2015), Stock Returns Over the FOMC Cycle (Working Paper, Duke University and University of California Berkeley).

CLARK, P. K. (1973), "A Subordinated Stochastic Process Model with Finite Variance for Speculative Prices", *Econometrica*, **41**, 135–155.

CREGO, J. (2017), Does Public News Decrease Information Asymmetries? Evidence from the Weekly Petroleum Status Report (Technical report, CEMFI).

DUFFIE, D. (2001), *Dynamic Asset Pricing Theory*, 3rd edn. (Princeton, NJ: Princeton University Press).

EDERINGTON, L. H. and LEE, J. H. (1993), "How Markets Process Information: News Releases and Volatility", *Journal of Finance*, **48**, 1161–1191.

GONÇALVES, S. and MEDDAHI, N. (2009), "Bootstrapping Realized Volatility", *Econometrica*, **77**, 283–306.

GREEN, C. T. (2004), "Economic News and the Impact of Trading on Bond Prices", *Journal of Finance*, **59**, 1201–1233.

HARRIS, L. (1986), "A Transaction Data Study of Weekly and Intradaily Patterns in Stock Returns", *Journal of Financial Economics*, **16**, 99–117.

HARRIS, M. and RAVIV, A. (1993), "Differences of Opinion Make a Horse Race", *Review of Financial Studies*, **6**, 473–506.

HARRISON, J. M. and KREPS, D. M. (1978), "Speculative Investor Behavior in a Stock Market with Heterogeneous Expectations", *Quarterly Journal of Economics*, **92**, 323–336.

HARVEY, C. R. and HUANG, R. D. (1991), "Volatility in the Foreign Currency Futures Market", *Review of Financial Studies*, **4**, 543–569.

HASBROUCK, J. (1991), "Measuring the Information Content of Stock Trades", *Journal of Finance*, **46**, 178–208.

HONG, H. and STEIN, J. C. (2007), "Disagreement and the Stock Market", *Journal of Economic Perspectives*, **21**, 109–128.

ILUT, C. and SCHNEIDER, M. (2014), "Ambiguous Business Cycles", *American Economic Review*, **104**, 2368–2399.

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

JAIN, P. C. and JOH, G. (1988), "The Dependence Between Hourly Prices and Trading Volume", *Journal of Financial and Quantitative Analysis*, **23**, 269–283.

JIANG, G. J., LO, I. and VERDELHAN, A. (2011), "Information Shocks, Liquidity Shocks, Jumps, and Price Discovery: Evidence from the U.S. Treasury Market", *Journal of Financial and Quantitative Analysis*, **46**, 527–551.

JOHNSON, J. A. and PAYE, B. S. (2015), "Stock Market Reactions to Unconventional Monetary Policy Announcements" (*Working Paper, University of Georgia*).

KANDEL, E. and PEARSON, N. D. (1995), "Differential interpretation of public signals and trade in speculative markets", *Journal of Political Economy*, **103**, 831–872.

KARPOFF, J. M. (1987), "The Relation Between Price Changes and Trading Volume: A Survey", *Journal of Financial and Quantitative Analysis*, **22**, 109–126.

KIM, O. and VERRECCHIA, R. E. (1991), "Trading Volume and Price Reactions to Public Announcements", *Journal of Accounting Research*, **29**, 302–321.

KIM, O. and VERRECCHIA, R. E. (1994), "Market Liquidity and Volume Around Earnings Announcements", *Journal of Accounting and Economics*, **17**, 41–67.

KUROV, A., SANCETTA, A., STRASSER, G., *et al*. (2016), Price Drift Before U.S. Macroeconomic News: Private Information About Public Announcements? (*Working Paper, European Central Bank*).

KYLE, A. S. (1985), "Continuous Auctions and Insider Trading", *Econometrica*, **53**, 1315–1335.

LEE, D. S. and LEMIEUX, T. (2010), "Regression Discontinuity Designs in Economics", *Journal of Economic Literature*, **48**, 281–355.

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

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

LI, J., TODOROV, V. and TAUCHEN, G. (2017), "Jump Regressions", *Econometrica*, **85**, 173–195.

LI, J. and XIU, D. (2016), "Generalized Methods of Integrated Moments for High-Frequency Data", *Econometrica, Forthcoming*.

LOUGHRAN, T. and McDONALD, B. (2011), "When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-ks", *Journal of Finance*, **66**, 35–65.

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

MADHAVAN, A., RICHARDSON, M. and ROOMANS, M. (1997), "Why do Security Prices Change? a Transaction-Level Analysis of NYSE Stocks", *Review of Financial Studies*, **10**, 1035–1064.

MERTON, R. C. (1992), *Continuous Time Finance* (New York, NY: Basil Blackwell).

PASQUARIELLO, P. and VEGA, C. (2007), "Informed and strategic order flow in the bond markets", *Review of Financial Studies*, **20**, 1975–2019.

PATTON, A. J. and TIMMERMANN, A. (2010), "Why do forecasters disagree? Lessons from the term structure of cross-sectional dispersion", *Journal of Monetary Economics*, **57**, 803–820.

SAVOR, P. and WILSON, M. (2013), "How Much do Investors Care About Macroeconomic Risk? Evidence from Scheduled Economic Announcements", *Journal of Financial and Quantitative Analysis*, **48**, 343–375.

—— (2014), "Asset Pricing: A Tale of Two Days", *Journal of Financial Economics*, **113**, 171–201.

SCHEINKMAN, J. A. and XIONG, W. (2003), "Overconfidence and Speculative Bubbles", *Journal of Political Economy*, **111**, 1183–1220.

TAUCHEN, G. E. and PITTS, M. (1983), "The Price Variability-Volume Relationship on Speculative Markets", *Econometrica*, **51**, 485–505.

VAN DIJK, D., LUMSDAINE, R. L. and VAN DER WEI, M. (2014), "Market Set-Up in Advance of Federal Reserve Policy Decisions" (NBER Working Paper No.19814).

VAN NIEUWERBURGH, S. and VELDKAMP, L. (2006), "Learning asymmetries in real business cycles", *Journal of Monetary Economics*, **53**, 753–772.

WOOD, R. A., MCINISH, T. H. and ORD, J. K. (1985), "An Investigation of Transactions Data for NYSE Stocks", *Journal of Finance*, **40**, 723–739.

WRIGHT, J. H. (2012), "What does monetary policy do at the zero lower bound?", *Economic Journal*, **122**, 447–466.