Macro Announcement Disagreement with Jump Regressions

by

Leonardo Salim Saker Chaves

Department of Economics
Duke University

Date: _____

Approved:

_____
Tim Bollerslev, Advisor

_____
Jia Li, Co-Advisor

_____
Andrew Patton

_____
Federico Bugni

Dissertation submitted in partial fulfillment of the requirements for the degree of
Doctor of Philosophy in the Department of Economics
in the Graduate School of Duke University
2021

Abstract

Macro Announcement Disagreement with Jump Regressions

by

Leonardo Salim Saker Chaves

Department of Economics
Duke University

Date: _____

Approved:

_____
Tim Bollerslev, Advisor

_____
Jia Li, Co-Advisor

_____
Andrew Patton

_____
Federico Bugni

An abstract of a dissertation submitted in partial fulfillment of the requirements for the degree of Doctor of Philosophy in the Department of Economics in the Graduate School of Duke University
2021

Abstract

This dissertation consists of two main essays in which it extends our knowledge on how stock market investors process the information from macro announcements.

In the first essay, we extend the existing econometric theory to study the relation between jumps in multiple processes at a high-frequency. More specifically, we develop new high-frequency-based inference procedures for analyzing the relationship between jumps in instantaneous moments of stochastic processes. The estimation consists of two steps: the nonparametric determination of the jumps as differences in local averages, followed by a minimum-distance type estimation of the parameters of interest under general loss functions that include both least-square and more robust quantile regressions as special cases. The resulting asymptotic distribution of the estimator, derived under an infill asymptotic setting, is highly nonstandard and generally not mixed normal. In addition, we establish the validity of a novel bootstrap algorithm for making feasible inference including bias-correction.

In the second essay, the new methods are applied to determine whether investors disagree when they process relevant macro-news announcements. If investors do disagree, we investigate the systematic components that drive disagreement. The high frequency data on stocks price and trade enable us to precisely isolate the news impact, and we use the volume-volatility elasticity framework to interpret our estimation. We

consider a set of stock characteristics that might contribute to investor disagreement: idiosyncratic volatility, market size, value, and institutional ownership. Our findings suggest that investors do disagree whenever there is more uncertainty about future payoffs. Furthermore, the different stock characteristics explain, to a large extent, the deviation from the case of no disagreement. For last, we explore how the direction of stock misprice affects the elasticity and verify that the overall investor disagreement may not be entirely observed due to arbitrage constraints.

Acknowledgements

I would like to start by thanking my wife, family, and friends who are all part of this accomplishment. I would not be able to complete this dissertation without their love and support. My wife, my true love, a true hero and caring person, has been on my side from day one and thanks to her I could make through the doctorate program! I am very grateful for my mother and father, who always provided endless love and care, and I aim to be like them in the future. Also, my grandmother, sister, and brothers have also been extremely supportive and excited with this opportunity that I enjoyed for the last 6 years. The journey was not always bright but they kept me in the right direction which made me continue persevering towards this degree.

Turning to my academic inspirations and examples, I would like to thank Tim Bollerslev for his guidance, insightful comments, and for the continuous support during my PhD, specially in the last year when we needed to change to virtual meetings. Knowing that I could count on him has made me push it through those challenging times. Additionally, I want to thank Jia Li for his guidance and knowledge over many meetings we had to discuss different research projects. I am very thankful for the rest of the committee members, Prof. Andrew Patton and Federico Bugni, for their thoughtful contributions to the advance of my research. For last but not less important, I want to thank everyone that participated in the financial econometric

lunch group. They helped me sharpen my arguments and pushed me forward in my projects with their questions and suggestions.

Also, I am very indebted to Duke University and its Economics department which have provided great support, in particular last year when they went beyond expectations and helped me overcome the challenges imposed by the pandemic. I am also very proud to have met brilliant minds and developed great friendships at Duke which made the PhD a pleasant experience.

Contents

List of Figures

List of Tables

1

Introduction

A major function of financial markets is the ability of aggregating information, public and private, into prices through trades. By doing this, they are able to funnel investments towards the most profitable options and make the economy more efficient. An important set of events are the macroeconomic announcements such as the Federal Open Market Committee (FOMC) decision on interest rates. Additionally, it is reasonable to assume that investors do not have the same interpretation of the news and it lead to market inefficiencies. Therefore, it is important that we have a measure of the existing disagreement among market participants.

However, we do not directly observe each investor actions and neither his belief about the value of a stock, only the price changes and the traded volume. This poses a challenge for our understanding of how financial markets process information and in this dissertation we propose a method to overcome this adversity and be able to explore investor disagreement by looking at aggregate market variables.

In Chapter 2, we extend the theory and develop methods to better estimate and make inference on the relationship between abnormal changes (jumps) in volume intensity and spot volatility for multiple stocks. This procedure relies on the high-frequency data for nonparametrically estimating the jumps with a minimum distance type estimator and accompanying bootstrap procedure for making robust inference about the parameters.

This proposed solution is able to overcome two major econometric challenges. First, the jumps are latent processes and the data is sampled discretely. Thus, we need to estimate them as a first step before considering the relation of interest. This creates additional source of uncertainty adding a finite sample bias that we are able to address. Second, the asymptotic distribution of our proposed method is usually not mixed Gaussian and, consequently, we cannot rely on the standard inference procedures. As an alternative, we propose a bootstrap algorithm for conducting inference.

In Chapter 3, we provide an application of the method developed in the previous chapter. We use the stocks in the $S\&P500$ to investigate the systematic components driving investor disagreement and a parsimonious model is able to explain most of the deviation from the no disagreement case. This model contains four stock characteristics: the idiosyncratic volatility, market capitalization, book-to-market ratio, and the percentage of shares owned by institutional investors. These variables capture two important characteristics of information processing, how complex it is to understand the implications of the event for the asset and how diverse is the pool of investors of that stock. Additionally, we interact the misprice direction of a stock with its degree of arbitrage constraint and notice that not all disagreement is reflected into

market variables and, as a result, this method captures a lower bound for investor disagreement.

2

Generalized Jump Regressions for Local Moments

2.1 Introduction

Many stochastic processes of practical empirical interest exhibit jump-like behavior. We propose a new statistical framework for analyzing the relationship between such jumps and other explanatory variables, as well as the relationship between simultaneously occurring jumps in multiple stochastic processes. Our approach relies crucially on the availability of high-frequency data for nonparametrically estimating the jumps together with a general minimum distance type estimator and accompanying bootstrap procedure for making robust inference about the parameters describing the relationship of interest.

Our new procedure is broadly applicable for studying the relationship between jumps of instantaneous moment processes associated with semimartingales. In financial applications, arguably the most important example of these instantaneous

moments is the spot variance of asset prices, formally defined as the local second moment of the return process. However, the local moment processes of other market variables such as trading volume, the time between trades, and quoted spreads, to name a few, are also of empirical interest as measures of trading activity and market liquidity. Jumps in these local moments are often triggered by macroeconomic news announcements occurring at specific times.

To illustrate, Figure 2.1 plots the price and trading volume of the S&P 500 E-mini futures contract on September 18, 2013, when the Federal Open Market Committee (FOMC) announced its decision *not* to taper the quantitative easing in place at the time. As the figure clearly shows, following the 2pm announcement there was a sharp increase in the volatility of the price (i.e., a positive volatility jump). This increase in the volatility was accompanied by an equally abrupt increase in trading activity (i.e., a positive volume jump). These types of jumps associated with clearly identifiable news events provide an ideal framework for studying the economic mechanisms at work, as exemplified by the economic theory of Kandel and Pearson (1995) and the recent empirical study of Bollerslev et al. (2018) concerning the relationship between jumps in the spot volatility and volume intensity at FOMC announcement times. This same "identification-by-discontinuity" empirical strategy using jumps has also been used in many other settings; see, for example, Jacod and Todorov (2010), Li et al. (2017a), Alexeev et al. (2017), Bibinger et al. (2017), among others.

The key statistical challenge in analyzing these types of jump relations stems from the fact that the jumps are latent processes. Only if the full continuous-time sample path of the underlying processes were available would the jumps be exactly identified. In practice, however, empirical researchers are almost always limited to discretely,

FIGURE 2.1: Price and Volume around an FOMC Announcement

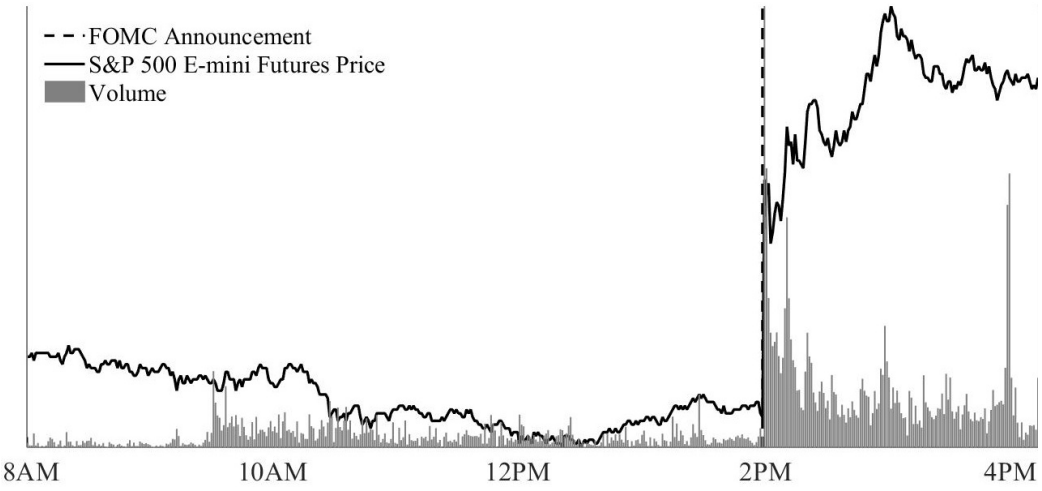

Note: The figure shows the price and volume of the S&P 500 E-mini futures on September 18, 2013. On that day, the FOMC announced its decision not to taper the quantitative easing in effect at the time.

albeit sometimes very finely, sampled data. As such, the jumps are invariably latent quantities that need to be estimated. Moreover, in our application, the local moments (such as the spot volatility of an asset) are themselves latent, creating an additional source of estimation error uncertainty. Our new two-step estimation procedure for addressing these issues builds on, and importantly extends, the least-squares approach of Bollerslev et al. (2018) to allow for the use of general convex loss functions and corresponding minimum-distance type estimators to assess the relationship between the first-stage jump estimates. Notably, this includes lin-lin loss, in which case the second-stage may be implemented via quantile regressions, as a special case.

Our motivation for considering more general loss functions is twofold. Firstly, compared to the quadratic loss employed by Bollerslev et al. (2018), the lin-lin loss is known to be more robust against influential observations in the sense of Huber

and Ronchetti (2009) (see, e.g., Koenker and Bassett (1978), Koenker (2005)). This type of robustness is especially relevant in the high-frequency data setting to help guard against overly influential "observations" associated with "noisy" data and potentially imprecise first-stage nonparametric jump estimates. Secondly, in parallel to standard quantile regressions, estimators based on different lin-lin losses have the potential to reveal heterogeneous responses across quantiles (see, e.g., Koenker and Bassett (1982)). As such, the different estimates may be used as a diagnostic tool for examining the assumption of a homogeneous response that is routinely, but implicitly, imposed in most empirical work. In our leading empirical example, discussed further below, we find that this is indeed a relevant concern.

The generalization to accommodate more general, possibly non-smooth loss functions like lin-lin, also requires the use of a distinctly different asymptotic theory and method of proof from that of existing work. The strategy typically adopted to address the complications stemming from the use of non-smooth loss functions relies on a quadratic expansion of an appropriately defined limiting criterion function, as the latter will be smooth in conventional settings (see, e.g., Huber (1967), Pollard (1985), Koenker (2005)). However, this approach does not work in the present setting, as the aggregation in the second-step estimation is based on only a fixed number of jumps. Hence, the non-smoothness of the loss function cannot simply be "averaged away." Instead, we derive the asymptotic distribution of our new estimator (in terms of stable convergence in law) using a novel convexity argument (as in, e.g., Knight (1989), Knight (1998)), in which the distribution is characterized as the argmin of a localized version of the limiting objective function.

Our theoretical results include those of Bollerslev et al. (2018) based on the

quadratic loss, for which the asymptotic distribution is mixed normal, as a special case. However, the mixed normality property that obtains under the quadratic loss does *not* hold true more generally with non-smooth loss functions, like lin-lin. Our theoretical arguments are also related to those underlying the so-called jump regressions recently analyzed by Li et al. (2017a,b). In contrast to that setting, however, which involves the jumps inferred from discretely observable processes, our setting entails an "extra layer" of latency associated with the nonparametric estimation of the local moment processes, in turn resulting in an overall slower rate of convergence.

The non-standard asymptotic distribution of the proposed estimator also renders routine "studentization" infeasible. Instead, we propose an easy-to-implement bootstrap algorithm as a natural alternative for conducting feasible inference (see, e.g., Efron and Tibshirani (1994), Hall (1997), Davison and Hinkley (1997)). The bootstrap consists of two steps: resampling the data in an *i.i.d.* fashion within local windows around the jump times, followed by repeating the estimation using the resampled data (after proper re-centering). The use of a local resampling scheme conveniently addresses the issue of data heterogeneity, which constitutes one of the key complications for bootstrapping in the high-frequency data setting (see Gonçalves and Meddahi (2009)). We prove the asymptotic validity of the bootstrap in this non-standard statistical setting under general conditions that permit both data heterogeneity and strong persistence. In particular, we do not need the data to be actually *i.i.d.* for the bootstrap to work. As such, our approach is distinctly different from the bootstrap used in conventional quantile regressions (see, e.g., Angelis et al. (1993), Hahn (1995)). It also differs from the block-bootstrap sometimes used for capturing time-series dependency (see, e.g., Carlstein (1986), Kunsch (1989)).

Going one step further, we demonstrate how the resampled bootstrap estimates readily allow for the implementation of finite-sample bias-correction (as in Efron and Tibshirani (1994), Horowitz (2001)). An empirically realistically calibrated Monte Carlo experiment in Section 2.4 further shows that the resulting bootstrap confidence intervals have good coverage properties, and that the bias-correction is indeed useful in reducing any finite-sample biases.

We apply the new method to study the relationship between the volume and volatility jumps, and how that relationship is affected by investors' disagreement. Consistent with the implications from an extensive theoretical literature in economics (see, e.g., Kandel and Pearson (1995)), we find that the volume-volatility elasticity is generally below unity and decreasing in the level of investors' disagreement. These findings confirm the recent results of Bollerslev et al. (2018). Importantly, however, we also find that these relations not only hold true "on average," but across a broad range of different quantiles. At the same time, we also uncover notable systematic heterogeneity in the elasticity estimates for certain types of announcements, thus directly highlighting the empirical relevance of using the more general loss functions and corresponding inference procedures developed here. In the Supplemental Appendix, we further illustrate the method with another empirical application, in which we study how the estimated jumps in the spot volatility and volume intensity around macroeconomic news announcements are related to the magnitude of the announcement surprises.

The rest of the chapter is organized as follows. Section 2.2 introduces the statistical setting and describes a few motivating examples. Section 2.3 presents the statistical inference methods. In Section 2.4, we examine the finite-sample performance of the

proposed method in a Monte Carlo experiment. Section 2.5 details our main empirical findings. Section 2.6 concludes. For last, Section 2.7 contains all the technical proofs.

2.2 The setting

2.2.1 Underlying stochastic processes

We begin by introducing the general statistical setting. We assume that the data are observed at discrete times $i\Delta_n$, $0 \leq i \leq [T/\Delta_n]$, and that the sampling interval $\Delta_n \to 0$ asymptotically over the fixed sample span $[0, T]$. This hypothetical setting of ever finer sampled data over a fixed time-interval is now standard in the analysis of high-frequency intraday financial data (see, e.g., Aït-Sahalia and Jacod (2014) and Jacod and Protter (2012)). Our statistical analysis concerns two types of high-frequency data: asset prices, which following standard practice we model as a semimartingale, and other possibly discrete-valued market variables, for which we rely on a more general state-space representation.

Let P denote the (log) price of an asset. We will assume that P is defined on some probability space $(\Omega, \mathcal{F}, \mathbb{P})$ and can be described by a continuous-time Itô semimartingale of the form,

$$dP_t = \alpha_t dt + \sigma_t dW_t + dJ_t, \tag{2.1}$$

where α_t denotes the drift process, σ_t is the stochastic volatility process, W_t is a standard Brownian motion, and J_t collects the jumps in the price process. We denote the spot variance process by $c_t \equiv \sigma_t^2$, which is the instantaneous variance of the diffusive price moves, that is,

$$c_t = \mathbb{E}_t \left[(\sigma_t dW_t)^2 \right] / \, dt. \tag{2.2}$$

We relegate the specifics of the regularity conditions concerning the σ_t volatility process to the supplemental appendix. However, the assumptions are extremely general, allowing for intraday periodicity, stochastic volatility-of-volatility, volatility jumps, leverage effects, and long-memory type dynamic dependencies.

In contrast to the (log) price process, other types of market data have only limited support. For instance, trading volume or quote spreads are typically integer multiples of a given lot or tick size. This in turn necessitates a different modeling framework from the Itô semimartingale in (2.1). Hence, following Li and Xiu (2016) and Bollerslev et al. (2018), we consider a univariate process V generated by the state-space model on the same discrete-time sampling grid,

$$V_{i\Delta_n} = \mathcal{V}\left(\zeta_{i\Delta_n}, \epsilon_i\right), \quad 0 \leq i \leq [T/\Delta_n], \tag{2.3}$$

where $\zeta_{i\Delta_n}$ is a latent state process, ϵ_i is a random shock, and the function $\mathcal{V}(\cdot,\cdot)$ transforms these two variables into the observed time series $(V_{i\Delta_n})_{i\geq 0}$. By integrating out the random shock ϵ_i with respect to its distribution $F_\epsilon(\cdot)$, one naturally obtains the instantaneous mean process of V, i.e.,

$$m_{i\Delta_n} \equiv \int \mathcal{V}\left(\zeta_{i\Delta_n}, \epsilon\right) F_\epsilon(d\epsilon). \tag{2.4}$$

This type of state-space representation embodies two useful features that we exploit in our statistical inference. Firstly, by assuming that the shocks (ϵ_i) are *i.i.d.*, the observations $(V_{i\Delta_n})$ become *conditionally* (given the state process ζ) independent. However, unconditionally, V is still allowed to be highly serially dependent (and heterogeneous) through the state process ζ. Secondly, we do not need to impose any specific assumptions on the transformation function $\mathcal{V}(\cdot,\cdot)$. Instead, we merely re-

quire some rather mild smoothness conditions on the m and ζ processes to allow for the construction of valid nonparametric inference procedures (see Assumptions A1 and A3 in the supplemental appendix for further technical details).

Our analysis focuses on the *jumps* in the local instantaneous moments, that is, the c and m processes. Formally, for a generic process Z, its jump at time τ is defined by $\Delta Z_\tau \equiv Z_\tau - Z_{\tau-}$, where $Z_{\tau-} = \lim_{s \nearrow \tau} Z_s$ is the left limit. We are primarily concerned with jumps that occur at known (announcement) times, corresponding to the setup commonly used in event-type studies. However, the proposed statistical methods remain valid with a finite set of unobserved jump times, provided that the jump times may be recovered with probability approaching one up to the sampling precision Δ_n. As a case in point, in the setting of Li et al. (2017a), the times of "large" price jumps may be consistently recovered using the thresholding technique of Mancini (2001).

2.2.2 Motivating examples

Intuitively, the jumps in economic variables may be seen as capturing "abnormal" moves induced by the arrival of new "lumpy" information, a prime example being regularly scheduled macroeconomic news announcements. Unlike "everyday" trading environments, in which it is difficult to clearly pinpoint specific shocks that drive the market, important macroeconomic announcements provide a convenient "laboratory" for isolating well-defined news from other confounding factors (see, e.g., the discussion in Andersen et al. (2003)). Correspondingly, insights as to what drives the jumps and the relationship among the jumps in different variables can help shed new light on the underlying economic mechanisms at work.

To fix ideas, we discuss two motivating examples. We will later return to these examples in our empirical analysis. Both examples concern the price volatility σ and the volume intensity m, defined as the (square root of) the local second moment of returns and the local mean of the observed trading volume, respectively. For each announcement time τ, we denote the jumps in the log levels of these local moment processes as,

$$\Delta \log(\sigma_\tau) \equiv \log(\sigma_\tau) - \log(\sigma_{\tau-}), \quad \Delta \log(m_\tau) \equiv \log(m_\tau) - \log(m_{\tau-}). \tag{2.5}$$

Empirically, as illustrated in Figure 2.1 above, $\Delta \log(\sigma_\tau)$ and $\Delta \log(m_\tau)$ are both generally positive at the time of important macroeconomic news announcements.

In a recent paper, Law et al. (2018) study how the surprise component of an announcement determine price jumps. Taking this analysis one step further, it is possible to examine more broadly the relationship between surprises and jumps in local moments, such as the volatility and the volume intensity. This empirical question in turn motivates the following specification,

$$\Delta \log(Y_\tau) = \boldsymbol{\theta}^\top \boldsymbol{X}_\tau, \qquad Y \in \{\sigma, m\}, \tag{2.6}$$

where the explanatory variable $\boldsymbol{X}_\tau$ would include proxies for the announcement surprises, and possibly other control variables. The expression in equation (2.6) is naturally interpreted as an *instantaneous moment condition*, necessitating the use of specialized inference procedures.

The study of volume and volatility jumps is also related to the large existing literature on volume-volatility relations more generally (see, e.g., Clark (1973) and Tauchen and Pitts (1983)). In particular, following the analysis of Bollerslev et al. (2018), the

oft-cited Kandel–Pearson equilibrium model (Kandel and Pearson (1995)) predicts that the volume-volatility elasticity should be below unity, and a decreasing function of the level of investor disagreement. Meanwhile, since the Kandel–Pearson theory concerns "abnormal moves" of market variables induced by news announcements, this naturally suggests identifying the elasticity as the slope coefficient θ_2 in the following log-linear specification (this is also the specification adopted by Bollerslev et al. (2018)),

$$\Delta \log (m_\tau) = \theta_1 + \theta_2 \Delta \log (\sigma_\tau) . \tag{2.7}$$

In parallel to equation (2.6) above, this baseline specification may also be extended to include covariates. Specifically, one may investigate the hypothesis that the elasticity is indeed a decreasing function of the level of investors' disagreement, by parameterizing the elasticity (and the intercept) as a linear function of other explanatory variables $(\boldsymbol{X}_{1,\tau}, \boldsymbol{X}_{2,\tau})$, that is,

$$\Delta \log (m_\tau) = \boldsymbol{\theta}_1^\top \boldsymbol{X}_{1,\tau} + (\boldsymbol{\theta}_2^\top \boldsymbol{X}_{2,\tau}) \, \Delta \log (\sigma_\tau) . \tag{2.8}$$

A test of the aforementioned hypothesis thus amounts to testing whether the component of the $\boldsymbol{\theta}_2$ parameter associated with the investor disagreement proxy is negative.

The instantaneous moment conditions in (2.6) and (2.8) may both be seen as specific examples of the following more general form,

$$G (m_{\tau-}, m_\tau, c_{\tau-}, c_\tau) = \sum_{k=1}^{K} \boldsymbol{\theta}_k^\top \boldsymbol{X}_{k,\tau} H_k (m_{\tau-}, m_\tau, c_{\tau-}, c_\tau) , \tag{2.9}$$

where $G(\cdot)$ and $H_k(\cdot)$ are continuously differentiable functions, and $\boldsymbol{\theta} = (\boldsymbol{\theta}_1, \ldots, \boldsymbol{\theta}_K)$ denotes the parameter vector of interest. We turn next to the development of the new statistical methods designed to allow for robust inference in this general setting.

2.3 Statistical methods

2.3.1 Estimation procedure

The practical estimation of $\boldsymbol{\theta}$ is complicated by the fact that the local moments σ and m (and hence their jumps) are not directly observable. In response to this, we rely on a two-step estimation procedure in which we first recover the jumps nonparametrically through the use of properly designed "spot" estimators, followed by a minimum-distance type estimation of $\boldsymbol{\theta}$.

Specifically, for each announcement time τ associated with the jumps, let $i(\tau) = \tau/\Delta_n + 1$ denote the corresponding observation count. The volume intensity and spot volatility after/before time τ (denoted by $+/-$) are then estimated by,

$$\hat{m}_{\tau\pm} \equiv \frac{1}{k_n} \sum_{j=1}^{k_n} V_{(i(\tau)\pm j)\Delta_n}, \quad \hat{c}_{\tau\pm} \equiv \frac{1}{k_n \Delta_n} \sum_{j=1}^{k_n} r^2_{(i(\tau)\pm j)}, \tag{2.10}$$

where $r_i \equiv P_{i\Delta_n} - P_{(i-1)\Delta_n}$ denotes the ith return, and the integer sequence k_n used in determining the size of the local window formally satisfies $k_n \to \infty$ and $k_n \Delta_n \to 0$. In general, one could also apply the thresholding technique of Mancini (2001) to construct a jump-robust estimator for the spot variance, although this is not formally needed under our maintained assumption of finitely active jumps.

Armed with these spot estimators, the sample analogue of (2.8) may be expressed as,

$$\Delta \widehat{\log(m_\tau)} = \boldsymbol{\theta}_1^\top \boldsymbol{X}_{1,\tau} + (\boldsymbol{\theta}_2^\top \boldsymbol{X}_{2,\tau})\, \Delta \widehat{\log(\sigma_\tau)} + e_\tau, \tag{2.11}$$

with the corresponding jump estimates defined by,

$$\Delta \widehat{\log(m_\tau)} \equiv \log(\hat{m}_{\tau+}) - \log(\hat{m}_{\tau-}), \quad \Delta \widehat{\log(\sigma_\tau)} \equiv (\log(\hat{c}_{\tau+}) - \log(\hat{c}_{\tau-}))/2. \tag{2.12}$$

Note that $c_t = \sigma_t^2$ implies $\log(\sigma_t) = \log(c_t)/2$. The error term e_τ in (2.11) arises from the estimation errors associated with the local moments c and m (e.g., $\hat{c}_{\tau-} - c_{\tau-}$). Similarly, the sample analogue for the more general possibly non-linear functional form in (2.9) may be expressed as,

$$G\left(\hat{m}_{\tau-}, \hat{m}_{\tau+}, \hat{c}_{\tau-}, \hat{c}_{\tau+}\right) = \sum_{k=1}^{K} \boldsymbol{\theta}_k^\top \boldsymbol{X}_{k,\tau} H_k\left(\hat{m}_{\tau-}, \hat{m}_{\tau+}, \hat{c}_{\tau-}, \hat{c}_{\tau+}\right) + e_\tau. \tag{2.13}$$

In view of equations (2.11) and (2.13), the $\boldsymbol{\theta}$ parameter could in principle be estimated by linear least squares. However, as is well-known in the literature on robust statistics, the implicit use of a quadratic loss function is potentially problematic for at least two reasons. Firstly, the estimates may be driven by a few highly influential "extreme" observations that manifest in the high frequency data. Secondly, it rules out the possibility that the strength of the relationship is not necessarily the same across all announcements included in the estimation (i.e., heterogeneous responses). Hence, we adopt a more general minimum-distance type estimation framework,

$$\begin{aligned}
&\widehat{\boldsymbol{\theta}}_n \equiv \underset{\boldsymbol{\theta}}{\operatorname{argmin}}\ Q_n(\boldsymbol{\theta}), \\
&Q_n\left(\boldsymbol{\theta}\right) \equiv \sum_{\tau \in \mathcal{T}} L\left(G\left(\hat{m}_{\tau-}, \hat{m}_{\tau+}, \hat{c}_{\tau-}, \hat{c}_{\tau+}\right) - \sum_{k=1}^{K} \boldsymbol{\theta}_k^\top \boldsymbol{X}_{k,\tau} H_k\left(\hat{m}_{\tau-}, \hat{m}_{\tau+}, \hat{c}_{\tau-}, \hat{c}_{\tau+}\right) \right),
\end{aligned} \tag{2.14}$$

where the set $\mathcal{T}$ identifies the specific announcements (as given by the announcement times) included in the estimation.

In the formal analysis below, we will further assume that the loss function $L(\cdot)$ satisfies the following very general set of assumptions.

The loss function $L(\cdot)$ is convex, and for some constant $p > 0$, $L(cx) = |c|^p L(x)$

for all $c, x \in \mathbb{R}$.

This setup differs from the setting commonly studied in the literature on M-estimation with possibly non-smooth objective functions (see, e.g., Huber (1967), Pollard (1985), Koenker (2005)). In that extant literature, the distribution of the estimator is typically characterized through the use of a quadratic approximation to a smooth limiting objective function, even if the sample objective function is non-smooth. By contrast, in the present setting with high-frequency data sampled over a fixed time span, the aggregation in (2.14) is invariably over finitely many announcement times $\mathcal{T}$, thereby rendering the use of a quadratic approximation to a possibly non-smooth $L(\cdot)$ loss function inappropriate, and in turn complicating the characterization of the $\widehat{\boldsymbol{\theta}}_n$ estimator by conventional methods.

The setup also differs from that of more conventional robust quantile regressions. In particular, even though the lin-lin loss function (i.e., $L(x) = x(q - 1_{\{x<0\}})$ for $q \in (0, 1)$) satisfies Assumption 1 and directly mirrors the loss function used in standard quantile regressions (see, e.g., Koenker and Bassett (1978), Koenker and Bassett (1982), Koenker (2005)), the $\widehat{\boldsymbol{\theta}}_n$ estimator is distinctly different as it involves non-parametrically estimated (latent) jumps, as opposed to directly observed data.

Assumption 1 pertaining to the form of the loss function obviously also includes quadratic loss (i.e., $L(x) = x^2$) as a special case. Further assuming the linear functional form in (2.11), $\widehat{\boldsymbol{\theta}}_n$ may be expressed in closed form as a function of the non-parametric jump estimates. In this situation, it is also relatively straightforward to show that the asymptotic distribution of $\widehat{\boldsymbol{\theta}}_n$ is centered at the true value with a mixed Gaussian distribution (that is indeed the method of proof adopted in Bollerslev et al.

(2018)). However, that same method of proof is not applicable for more general possibly non-smooth loss functions. Correspondingly, the asymptotic distribution of $\widehat{\boldsymbol{\theta}}_n$ is generally *not* mixed Gaussian either.

For the empirical results reported below, we will primarily rely on the non-smooth lin-lin loss function. As noted above, our motivation for doing so is twofold. Firstly, since the lin-lin loss is less sensitive to outliers than the quadratic loss, the resulting estimators will be more robust against data imperfections in the sense of Koenker and Bassett (1978) and Huber and Ronchetti (2009). This feature is particularly desirable in our study of (major) news announcements, as the market tends to be especially turbulent during such times. Secondly, estimators associated with different quantiles may reveal heterogeneous responses across announcements, with their own distinct economic interpretations. This feature of the lin-lin loss function has also previously been emphasized by Koenker and Bassett (1982) as providing a useful tool for detecting heteroskedasticity and evaluating the validity of a given specification more generally.

To derive the limit distribution of $\widehat{\boldsymbol{\theta}}_n$, it is helpful to reparametrize the sample objective function via a change of variable $\boldsymbol{\theta} \to \boldsymbol{\theta}_0 + k_n^{-1/2}\boldsymbol{h}$, where $\boldsymbol{\theta}_0$ denotes the true parameter, and the local parameter $\boldsymbol{h} = (\boldsymbol{h}_1, \ldots, \boldsymbol{h}_K)$ quantifies the deviation of $\boldsymbol{\theta}$ from the true parameter in a $k_n^{-1/2}$-neighborhood (this also corresponds to the convergence rate of the nonparametric jump estimates in (2.12) that enter the objective function in (2.14)). Correspondingly, we define the reparametrized objective function as,

$$M_n(\boldsymbol{h}) \equiv k_n^{p/2} Q_n(\boldsymbol{\theta}_0 + k_n^{-1/2}\boldsymbol{h}), \tag{2.15}$$

where the scaling factor $k_n^{p/2}$ is included to ensure that $M_n(\cdot)$ is well behaved asymptotically. It follows readily that since $\widehat{\boldsymbol{\theta}}_n$ minimizes $Q_n(\boldsymbol{\theta})$, the normalized estimator $\widehat{\boldsymbol{h}}_n = k_n^{1/2}(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0)$ minimizes $M_n(\boldsymbol{h})$, that is,

$$\widehat{\boldsymbol{h}}_n = \underset{\boldsymbol{h}}{\operatorname{argmin}}\ M_n(\boldsymbol{h}). \tag{2.16}$$

Moreover, under mild regularity conditions, the localized objective function $M_n(\cdot)$ converges stably in law (i.e., joint with any bounded random variables that are measurable to the underlying σ-field) to a limiting process, say $M(\cdot)$, thereby providing a framework for deriving the distribution of $\widehat{\boldsymbol{\theta}}_n$ through that of $\widehat{\boldsymbol{h}}_n$.

Some additional notation is required for characterizing the process $M(\cdot)$. For each t, we set $v_t \equiv \int \mathcal{V}(\zeta_t, \epsilon)^2 F_\epsilon(d\epsilon) - m_t^2$. Further, let $\partial G(x; dx)$ and $\partial H_k(x; dx)$ denote the first differential of $G(\cdot)$ and $H_k(\cdot)$, respectively. In order to represent the asymptotic distribution, we consider the random variables $(\eta_{m,\tau-}, \eta_{m,\tau+}, \eta_{c,\tau-}, \eta_{c,\tau+})_{\tau\in\mathcal{T}}$ that are, conditionally on $\mathcal{F}$, mutually independent, centered Gaussian with conditional variances $\mathbb{E}\left[\eta_{m,\tau\pm}^2|\mathcal{F}\right] = v_{\tau\pm}$ and $\mathbb{E}[\eta_{c,\tau\pm}^2|\mathcal{F}] = 2c_{\tau\pm}^2$. These η variables capture the sampling variability of the spot estimators. Finally, we set,

$$\begin{aligned} \xi_\tau &\equiv \partial G\left(m_{\tau-}, m_\tau, c_{\tau-}, c_\tau; \eta_{m,\tau-}, \eta_{m,\tau+}, \eta_{c,\tau-}, \eta_{c,\tau+}\right), \\ \xi'_{k,\tau} &\equiv \partial H_k\left(m_{\tau-}, m_\tau, c_{\tau-}, c_\tau; \eta_{m,\tau-}, \eta_{m,\tau+}, \eta_{c,\tau-}, \eta_{c,\tau+}\right), \end{aligned} \tag{2.17}$$

and define the limiting process $M(\cdot)$ as

$$M(\boldsymbol{h}) = \sum_{\tau\in\mathcal{T}} L\left(\xi_\tau - \sum_{k=1}^{K} \boldsymbol{\theta}_{0,k}^\top \boldsymbol{X}_{k,\tau}\xi'_{k,\tau} - \sum_{k=1}^{K} \boldsymbol{h}_k^\top \boldsymbol{X}_{k,\tau} H_k\left(m_{\tau-}, m_\tau, c_{\tau-}, c_\tau\right)\right). \tag{2.18}$$

Since the objective function $M_n(\cdot)$ converges stably in law to $M(\cdot)$ in finite dimensions, we can appeal to a convexity argument (as in Knight (1989), Knight (1998)) to deduce

that $\widehat{\boldsymbol{h}}_n$ converges stably in law to the argmin of the $M(\cdot)$ limiting process, that is,

$$\widehat{\boldsymbol{h}} \equiv \underset{\boldsymbol{h}}{\operatorname{argmin}}\ M(\boldsymbol{h}). \tag{2.19}$$

In the special case when the loss function $L(\cdot)$ is quadratic, the limit minimization problem in (2.19) may be solved analytically. In that situation, it is also relatively straightforward to show that the distribution of $\widehat{\boldsymbol{h}}$ is centered mixed Gaussian. In general, however, with non-quadratic loss, even though $\widehat{\boldsymbol{h}}$ is symmetrically distributed, the estimator will *not* be mixed Gaussian. For example, with absolute deviation loss (i.e., $L(x) = |x|$), the distribution of $\widehat{\boldsymbol{h}}$ is given by that of the regression coefficient in a median regression for the mixed Gaussian variables $\xi_\tau - \sum_{k=1}^{K} \boldsymbol{\theta}_{0,k}^{\top} \boldsymbol{X}_{k,\tau} \xi'_{k,\tau}$ against $\boldsymbol{X}_{k,\tau} H_k\left(m_{\tau-}, m_\tau, c_{\tau-}, c_\tau\right)$, $1 \leq k \leq K$, for τ in the finite set $\mathcal{T}$ (see Section 3.1 of Koenker (2005) for details on the finite-sample behavior of regression quantiles).

The following theorem summarizes the asymptotic behavior of $M_n(\cdot)$ and $\widehat{\boldsymbol{h}}_n$ in terms of $M(\cdot)$ and $\widehat{\boldsymbol{h}}$, and in turn the distribution of $\widehat{\boldsymbol{\theta}}_n$, for general loss function $L(\cdot)$.

Under Assumption 2.3.1 and Assumptions A1–A3 in the supplemental appendix, the sequence $M_n(\cdot)$ of processes converges stably in law to $M(\cdot)$ in finite dimensions. Moreover, if $\widehat{\boldsymbol{h}}$ uniquely minimizes $M(\cdot)$ almost surely, then $\widehat{\boldsymbol{h}}_n \equiv k_n^{1/2}(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0)$ converges stably in law to $\widehat{\boldsymbol{h}}$.

Proof: See Section 2.7.

Theorem 1 establishes that $\widehat{\boldsymbol{\theta}}_n$ is indeed a $k_n^{1/2}$-consistent estimator of the true $\boldsymbol{\theta}_0$ parameter. Since the number of (announcement-induced) jumps is finite within the fixed sample span in our infill asymptotic setting, $\widehat{\boldsymbol{\theta}}_n$ inherits the nonparametric $k_n^{1/2}$-rate of the corresponding finite collection of spot estimators. We note that

the uniqueness condition of $\widehat{\boldsymbol{h}}$ corresponds to the identification condition, which typically amounts to ruling out multi-collinearity among the regressors in specific settings. Moreover, it characterizes the limiting distribution of the normalized estimator $k_n^{1/2}(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0)$ in terms of the argmin (i.e., $\widehat{\boldsymbol{h}}$) of the $M(\cdot)$ limiting process. However, as the discussion above makes clear, the resulting asymptotic distribution of $\widehat{\boldsymbol{\theta}}_n$ can be highly non-standard, and it is fundamentally different from those in conventional M-estimation theory.

2.3.2 Feasible inference via bootstrap

The non-standard distribution of $\widehat{\boldsymbol{\theta}}_n$ that obtains under general non-smooth loss does not allow for the use of standard Gaussian-based inference procedures. Instead, we propose an easy-to-implement bootstrap approach for computing confidence intervals for the true parameter $\boldsymbol{\theta}_0$. The bootstrap has two distinct advantages in the current setting. First, since the asymptotic distribution of $\widehat{\boldsymbol{\theta}}_n$ is generally not (mixed) Gaussian, there is no clear way to render the estimator pivotal via "studentization." By contrast, the bootstrap readily approximates the non-standard asymptotic distribution. Second, the same bootstrap resampling scheme may be used for multiple competing estimators associated with different loss functions, thereby facilitating any formal statistical comparisons of the different estimators. The bootstrap algorithm is defined by

Algorithm 1.

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

Step 2: Compute $(\hat{m}^*_{\tau-}, \hat{m}^*_{\tau+}, \hat{c}^*_{\tau-}, \hat{c}^*_{\tau+})_{\tau\in\mathcal{T}}$ the same way as $(\hat{m}_{\tau-}, \hat{m}_{\tau+}, \hat{c}_{\tau-}, \hat{c}_{\tau+})_{\tau\in\mathcal{T}}$ except that the original data $(V_{i(\tau)+j}, r_{i(\tau)+j})_{1\leq|j|\leq k_n}$ are replaced with $(V^*_{i(\tau)+j}, r^*_{i(\tau)+j})_{1\leq|j|\leq k_n}$.

Step 3: Estimate $\widehat{\boldsymbol{\theta}}^*_n = \text{argmin}_{\boldsymbol{\theta}} Q^*_n(\boldsymbol{\theta})$, where

$$\begin{aligned} Q^*_n(\boldsymbol{\theta}) &\equiv \sum_{\tau\in\mathcal{T}} L\left(G\left(\hat{m}^*_{\tau-}, \hat{m}^*_{\tau+}, \hat{c}^*_{\tau-}, \hat{c}^*_{\tau+}\right) - \hat{\varepsilon}_\tau - \sum_{k=1}^{K} \boldsymbol{\theta}_k^\top \boldsymbol{X}_{k,\tau} H_k\left(\hat{m}^*_{\tau-}, \hat{m}^*_{\tau+}, \hat{c}^*_{\tau-}, \hat{c}^*_{\tau+}\right)\right), \\ \hat{\varepsilon}_\tau &\equiv G\left(\hat{m}_{\tau-}, \hat{m}_{\tau+}, \hat{c}_{\tau-}, \hat{c}_{\tau+}\right) - \sum_{k=1}^{K} \widehat{\boldsymbol{\theta}}_k^\top \boldsymbol{X}_{k,\tau} H_k\left(\hat{m}_{\tau-}, \hat{m}_{\tau+}, \hat{c}_{\tau-}, \hat{c}_{\tau+}\right). \end{aligned}$$

Step 4: Repeat steps 1–3 a large number of times. Use the Monte Carlo distribution of $k_n^{1/2}(\widehat{\boldsymbol{\theta}}^*_n - \widehat{\boldsymbol{\theta}}_n)$ to approximate that of $k_n^{1/2}(\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0)$. In particular, a symmetric two-sided confidence interval for $\theta_{0,j}$ (i.e., the jth element of $\boldsymbol{\theta}_0$) is given by $CI_n = [\hat{\theta}_{n,j} - z_{n,1-\alpha/2}, \hat{\theta}_{n,j} + z_{n,1-\alpha/2}]$, where $z_{n,1-\alpha/2}$ is the $(1-\alpha/2)$-quantile of $|\hat{\theta}^*_{n,j} - \hat{\theta}_{n,j}|$ in the Monte Carlo sample.

The bootstrap described in Algorithm 1 relies on an *i.i.d.* resampling scheme within local windows before and after the announcement times $\tau \in \mathcal{T}$ to account for temporal heterogeneity in the data. Intuitively, within each of these local windows, the state processes σ and ζ are approximately constant, thereby permitting the use of an *i.i.d.* scheme. However, it is important to stress that the validity of this bootstrap does *not* require the data to actually be *i.i.d.* We only require the observations of V to be *conditionally* independent, which allows for both heterogeneity and persistence in the underlying processes. As such, the bootstrap theory is also very different from the type of bootstrap traditionally used in quantile regressions (see, e.g., Angelis et al. (1993), Hahn (1995)). It is possible that the wild bootstrap (see Wu (1986))

may similarly be used to address the issue of heterogeneity in the present context. However, we leave this question for future research.

Theorem 2.3.2 formally establishes the asymptotic validity of this new bootstrap procedure proposed here.

Under the same conditions as in Theorem 2.3.1, the conditional distribution function of $k_n^{1/2}(\widehat{\boldsymbol{\theta}}_n^* - \widehat{\boldsymbol{\theta}}_n)$ given data converges in probability to the $\mathcal{F}$-conditional distribution of $\widehat{\boldsymbol{h}}$ under the uniform metric. Consequently, the confidence interval CI_n described in the bootstrap Algorithm 1 has asymptotic level $1-\alpha$.

Proof: See Section 2.7.

In addition to constructing confidence intervals, the same bootstrap algorithm may also be used in correcting finite-sample biases in the $\widehat{\boldsymbol{\theta}}_n$ estimator and the bootstrap confidence intervals (see also the discussion in Horowitz (2001)). In particular, the bias in $\widehat{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0$ is naturally approximated by,

$$\widehat{\boldsymbol{\beta}}_n \equiv \text{Med}^* \left[\widehat{\boldsymbol{\theta}}_n^* - \widehat{\boldsymbol{\theta}}_n \right],$$

where Med* denotes the median in the bootstrap sample. This in turn suggests the bias-corrected estimator,

$$\widehat{\boldsymbol{\theta}}_n^c \equiv \widehat{\boldsymbol{\theta}}_n - \widehat{\boldsymbol{\beta}}_n.$$

Similarly, let $z_{n,1-\alpha/2}^c$ denote the $(1-\alpha/2)$-quantile of $|\hat{\theta}_{n,j}^* - \hat{\theta}_{n,j} - \hat{\beta}_{n,j}|$ in the Monte Carlo sample. A bias-corrected confidence interval may then be constructed as,

$$CI_n^c \equiv [\hat{\theta}_{n,j}^c - z_{n,1-\alpha/2}^c, \hat{\theta}_{n,j}^c + z_{n,1-\alpha/2}^c].$$

Since $k_n^{1/2}\widehat{\boldsymbol{\beta}}_n$ is $o_p(1)$ (by Theorem 2.3.2), it follows readily that $\widehat{\boldsymbol{\theta}}_n^c$ (resp. CI_n^c) will have the same asymptotic properties as $\widehat{\boldsymbol{\theta}}_n$ (resp. CI_n) described in Theorem 2.3.1

(resp. Theorem 2). However, as shown by the Monte Carlo simulations presented in the supplemental appendix, the bias-corrected estimator and confidence intervals tend to be better behaved in finite samples.

2.3.3 Intraday patterns and difference-in-difference estimation

A further complication, and a potential source of finite-sample bias, that arise in the analysis of high-frequency financial data stems from the marked intraday periodic patterns that exist in such data. In particular, volatility, trading activity, bid-ask spreads and many other financial variables all tend to be higher around the time of market opening and closing (see, e.g., Wood et al. (1985) for some of the earliest empirical evidence). To further complicate matters, these intraday patterns also tend to vary somewhat both over time and across assets. A failure to account for this may result in systematically biased parameter estimates if the jumps underlying the estimation occur at specific times-of-day. To remedy this, Bollerslev et al. (2018) proposed a simple difference-in-difference (DID) type approach based on an appropriate control group. This same DID strategy can be applied in the current more general setting.

Formally, for each announcement time τ, define the control group $\mathcal{C}(\tau)$ of N_C non-announcement times, the implicit assumption being that the processes of interest do not jump at time τ in the control group. The intraday patterns in the "raw" jump estimators defined in (2.12) may then be controlled for by "differencing out" the

corresponding estimates averaged within the control group,

$$\begin{aligned}\widetilde{\Delta \log(m_\tau)} &\equiv \Delta \widehat{\log(m_\tau)} - \frac{1}{N_C}\sum_{\eta \in C(\tau)} \Delta \widehat{\log(m_\eta)},\\ \widetilde{\Delta \log(\sigma_\tau)} &\equiv \Delta \widehat{\log(\sigma_\tau)} - \frac{1}{N_C}\sum_{\eta \in C(\tau)} \Delta \widehat{\log(\sigma_\eta)}.\end{aligned} \tag{2.20}$$

In our empirical analysis below, we take $\mathcal{C}(\tau)$ to be 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).

These DID jump estimators can be incorporated in the estimation straightforwardly by allowing the $G(\cdot)$ and $H_k(\cdot)$ transformations in the instantaneous moment condition (2.9) to also depend on the spot estimators in the control group. To simplify the notation, define $\widehat{\boldsymbol{S}}_\tau \equiv (\hat{m}_{\tau-}, \hat{m}_{\tau+}, \hat{c}_{\tau-}, \hat{c}_{\tau+})$ and $\widetilde{\boldsymbol{S}}_\tau \equiv (\widehat{\boldsymbol{S}}_t)_{t\in\{\tau\}\cup\mathcal{C}(\tau)}$. The DID estimator for $\boldsymbol{\theta}$ is then given by,

$$\begin{aligned}\widetilde{\boldsymbol{\theta}}_n &\equiv \underset{\boldsymbol{\theta}}{\operatorname{argmin}}\ \widetilde{Q}_n(\boldsymbol{\theta}),\\ \widetilde{Q}_n(\boldsymbol{\theta}) &\equiv \sum_{\tau\in\mathcal{T}} L\left(G(\widetilde{\boldsymbol{S}}_\tau) - \sum_{k=1}^{K} \boldsymbol{\theta}_k^\top \boldsymbol{X}_{k,\tau} H_k(\widetilde{\boldsymbol{S}}_\tau)\right).\end{aligned} \tag{2.21}$$

Compared to the no-DID objective function Q_n, the DID counterpart involves the additional spot estimators in the control groups. Since the different control groups may overlap with each other, possibly in a highly irregular fashion, the asymptotic distribution of $\widetilde{\boldsymbol{\theta}}_n$ becomes much more cumbersome to characterize analytically than that of $\widehat{\boldsymbol{\theta}}_n$. However, the bootstrap Algorithm 1 is readily adapted to accommodate this additional complication. Algorithm 2 spells out the necessary adjustments.

Algorithm 2.

Step 1: For each $\tau \in \mathcal{T} \cup (\cup_{\tau' \in \mathcal{T}} \mathcal{C}(\tau'))$, 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 $\widetilde{\boldsymbol{S}}^*_\tau$ the same way as $\widetilde{\boldsymbol{S}}_\tau$, except that the original data $(V_{i(\tau)+j}, r_{i(\tau)+j})_{1 \le |j| \le k_n}$ are replaced with $(V^*_{i(\tau)+j}, r^*_{i(\tau)+j})_{1 \le |j| \le k_n}$.

Step 3: Estimate $\widetilde{\boldsymbol{\theta}}^*_n = \text{argmin}_{\boldsymbol{\theta}} \widetilde{Q}^*_n(\boldsymbol{\theta})$, where

$$\begin{aligned} \widetilde{Q}^*_n(\boldsymbol{\theta}) &\equiv \sum_{\tau \in \mathcal{T}} L\left(G(\widetilde{\boldsymbol{S}}^*_\tau) - \tilde{\varepsilon}_\tau - \sum_{k=1}^K \boldsymbol{\theta}_k^\top \boldsymbol{X}_{k,\tau} H_k(\widetilde{\boldsymbol{S}}^*_\tau) \right), \\ \tilde{\varepsilon}_\tau &\equiv G(\widetilde{\boldsymbol{S}}_\tau) - \sum_{k=1}^K \widetilde{\boldsymbol{\theta}}_k^\top \boldsymbol{X}_{k,\tau} H_k(\widetilde{\boldsymbol{S}}_\tau). \end{aligned}$$

Step 4: Repeat steps 1–3 a large number of times. Use the Monte Carlo distribution of $k_n^{1/2}(\widetilde{\boldsymbol{\theta}}^*_n - \widetilde{\boldsymbol{\theta}}_n)$ to approximate that of $k_n^{1/2}(\widetilde{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0)$. In particular, a symmetric two-sided confidence interval for $\theta_{0,j}$ (i.e., the jth element of $\boldsymbol{\theta}_0$) is given by $\widetilde{CI}_n = [\tilde{\theta}_{n,j} - \tilde{z}_{n,1-\alpha/2}, \tilde{\theta}_{n,j} + \tilde{z}_{n,1-\alpha/2}]$, where $\tilde{z}_{n,1-\alpha/2}$ is the $(1-\alpha/2)$-quantile of $|\tilde{\theta}^*_{n,j} - \tilde{\theta}_{n,j}|$ in the Monte Carlo sample. □

The theoretical justification for the DID estimator and Algorithm 2 essentially mirrors the theory described in the previous subsection. To proceed with the details, define the modified limiting variables corresponding to (2.17) as,

$$\begin{aligned} &\tilde{\xi}_\tau \equiv \partial G((m_{t-}, m_t, c_{t-}, c_t)_{t \in \{\tau\} \cup \mathcal{C}(\tau)}; (\eta_{m,t-}, \eta_{m,t+}, \eta_{c,t-}, \eta_{c,t+})_{t \in \{\tau\} \cup \mathcal{C}(\tau)}), \\ &\tilde{\xi}'_{k,\tau} \equiv \partial H_k((m_{t-}, m_t, c_{t-}, c_t)_{t \in \{\tau\} \cup \mathcal{C}(\tau)}; (\eta_{m,t-}, \eta_{m,t+}, \eta_{c,t-}, \eta_{c,t+})_{t \in \{\tau\} \cup \mathcal{C}(\tau)}), \end{aligned} \tag{2.22}$$

and, correspondingly, modify the definition in (2.18) as,

$$\widetilde{M}(\boldsymbol{h}) = \sum_{\tau \in \mathcal{T}} L \left(\tilde{\xi}_\tau - \sum_{k=1}^{K} \boldsymbol{\theta}_{0,k}^\top \boldsymbol{X}_{k,\tau} \tilde{\xi}'_{k,\tau} - \sum_{k=1}^{K} \boldsymbol{h}_k^\top \boldsymbol{X}_{k,\tau} H_k \left((m_{t-}, m_t, c_{t-}, c_t)_{t \in \{\tau\} \cup \mathcal{C}(\tau)} \right) \right). \tag{2.23}$$

Theorem 2.3.3, below, characterizes the asymptotic distribution of the DID estimator $\widetilde{\boldsymbol{\theta}}_n$ and justifies the asymptotic validity of Algorithm 2.

Under the same conditions as Theorem 2.3.1, the following statements hold:

(a) The sequence $\widetilde{M}_n(\boldsymbol{h}) = k_n^{p/2} \widetilde{Q}_n(\boldsymbol{\theta}_0 + k_n^{-1/2} \boldsymbol{h})$ of processes converges stably in law to $\widetilde{M}(\boldsymbol{h})$ in finite dimensions. Moreover, if $\widetilde{\boldsymbol{h}}$ uniquely minimizes $\widetilde{M}(\cdot)$ almost surely, then $k_n^{1/2}(\widetilde{\boldsymbol{\theta}}_n - \boldsymbol{\theta}_0)$ converges stably in law to $\widetilde{\boldsymbol{h}}$.

(b) The conditional distribution function of $k_n^{1/2}(\widetilde{\boldsymbol{\theta}}_n^* - \widetilde{\boldsymbol{\theta}}_n)$ given data converges in probability to the $\mathcal{F}$-conditional distribution of $\widetilde{\boldsymbol{h}}$ under the uniform metric. Consequently, the confidence interval $\widetilde{CI}_n$ described in the bootstrap Algorithm 2 has asymptotic level $1 - \alpha$.

Proof: See Section 2.7.

Note that the same bootstrap-based bias correction used in adjusting the non-DID estimates described in the previous subsection may similarly be used in bias correcting the DID estimates. The requisite modifications to the expressions for $\widehat{\boldsymbol{\theta}}_n^c$ and CI_n^c are obvious, albeit notationally cumbersome, and we omit the details for brevity.

2.3.4 Discussion

The proposed new methods are related to several studies on regression-type analysis of jumps. In particular, Li et al. (2017a) first introduced the notion of least-squares jump regressions for analyzing the relationship among price jumps, while Li et al. (2017b) extended that framework to allow for the use of general loss functions. Unlike these prior studies, however, the present analysis pertains to the jumps in *local moments*, such as price volatility and volume intensity, rather than the jumps in the price process itself. The estimation and inference for these types of jumps are notably more complicated. For one, jumps in the local moments are estimated at a nonparametric rate, whereas the price jumps can be recovered at a parametric rate. The much more pronounced intraday diurnal patterns that exist in both volatility and trading volume, and the DID estimation strategy based on the inclusion of irregularly spaced control groups developed here to address this issue, also results in additional sampling errors that are quite cumbersome to characterize analytically. Our new bootstrap-based inference procedure conveniently solves this problem.

The current study is also closely related to the recent work of Bollerslev et al. (2018) and the analysis therein pertaining to regressions involving jumps in volume intensity and spot volatility. However, our method generalizes this prior work by allowing for non-linear functional forms and general possibly non-smooth loss functions, like the lin-lin loss function. All of this in turn necessitates a different strategy for developing the asymptotic distribution of the estimators. Thus, even though our *i.i.d.* bootstrap resampling scheme bears close resemblance to that of Bollerslev et al. (2018), the validity of the bootstrap inference for our new estimator demands its own

(new) and very different method of proof.

The present research also extends the scope of the possible empirical investigations from the univariate volume-volatility relations analyzed in Bollerslev et al. (2018) to more general event type analysis involving the jumps in other instantaneous moments, and our theory can be readily extended to a multivariate setting (which is done in a working paper version of this paper). Further along these lines, we also explicitly recognize a nontrivial finite-sample bias in this type of analysis that could severely distort any empirical conclusions. Our new bootstrap provides a simple, yet effective, way of correcting this bias.

It would be interesting to extend the new theory developed here to explicitly allow for the presence of microstructure noise in the spot volatility estimation. The same proof strategy underlying Theorem 2.3.1 could in principle be used to characterize the asymptotic distribution, provided that the joint asymptotic distribution of the spot volatility estimator and the $\hat{m}_{\tau\pm}$ local mean estimator is known. Results on noise-robust spot volatility estimation (see, e.g., Bibinger and Winkelmann (2015)) could possibly be extended to verify this "high-level" condition.

2.4 Monte Carlo study

This section discusses the results from a Monte Carlo simulation study designed to assess the finite sample behavior of the new estimators and bootstrap inference procedures in an empirically realistic setting that closely mimic our actual empirical analysis. For concreteness, we focus on the estimation of the volume-volatility elasticity as in Section 2.5.2. We begin by describing the data generating process.

2.4.1 Data generating process

We normalize the unit of time to be one day, and set the total span of the sample to $2,500$ days. The log price and stochastic volatility processes are then simulated according to the following stochastic differential equations,

$$\begin{aligned} dP_t &= \sigma_t dW_t + \varphi_{P,t} dN_t, \\ d\log(\sigma_t) &= -0.03\log(\sigma_t)dt + 0.1\left(\rho dW_t + \sqrt{1-\rho^2}dB_t\right) + \varphi_{\sigma,t}dN_t, \end{aligned}$$

where W_t and B_t are independent standard Brownian motions, $\rho = -0.6$ accounts for the widely documented "leverage effect" (see, e.g., Wang and Mykland (2015), Kalina and Xiu (2017), Aït-Sahalia et al. (2017)), and the jump sizes when a jump occurs (i.e., $dN_t = 1$) are drawn according to $\varphi_{P,t} \sim$Uniform$[-1, 1]$ and $\varphi_{\sigma,t} \sim$Uniform$[0.25, 2]$, respectively. The occurrences of jumps in turn are determined by simulating 100 equally-spaced announcement times over the sample, with N_t denoting the counting process for the total number of jumps within the $[0, t]$ time interval. The continuous-time model is approximated using an Euler scheme on a 5-second grid, and then aggregated to a $\Delta_n = 1$ minute sampling interval, paralleling the discrete-time sampling in our actual empirical applications.

Our analysis and new estimation procedures only involve data in local windows before and after the announcement times, or $\tau \in \mathcal{T}$, corresponding to $dN_t = 1$ in the above notation. We simulate the log volume intensity in the relevant local windows

according to,

$$\log(m_{i\Delta_n}) = 7 + (0.7 - 0.06X_\tau) \cdot \Delta_n^{-1} \int_{(i-1)\Delta_n}^{i\Delta_n} \log(\sigma_s) ds, \qquad i\Delta_n \in [\tau - k_n\Delta_n, \tau),$$

$$\log(m_{i\Delta_n}) = 7.8 + (0.7 - 0.06X_\tau) \cdot \Delta_n^{-1} \int_{(i-1)\Delta_n}^{i\Delta_n} \log(\sigma_s) ds, \qquad i\Delta_n \in [\tau, \tau + k_n\Delta_n],$$

where the parameters are calibrated using data from our empirical analysis. We rely on the same X_τ disagreement measure used in our empirical analysis, defined as the dispersion in the survey of professional forecasters for the one-quarter-ahead unemployment rate. Finally, the volume data that actually enters the estimation is simulated as,

$$V_{i\Delta_n} = \mathcal{V}\left(m_{i\Delta_n}, \epsilon_i\right) = m_{i\Delta_n}\epsilon_i,$$

where $\epsilon_i = \tilde{\epsilon}_i/df$, and $\tilde{\epsilon}_i$ are *i.i.d.* draws from a chi-square distribution with $df \in \{10, 30\}$ degrees-of-freedom. Note that even though the simulated $V_{i\Delta_n}$ series is conditionally independent given m, it is unconditionally serially correlated. In fact, the specific choices of $df = 10$ and 30 are purposely calibrated so that the range of the autocorrelations for the simulated volume series bracket those observed in the actual volume data.

Altogether, this simulation setup implies that the log volume intensity and log volatility jumps around announcement times satisfy the following relation,

$$\Delta \log(m_\tau) = \theta_1 + (\theta_2 + \theta_3 X_\tau)\, \Delta \log(\sigma_\tau),$$

with the true value of $\boldsymbol{\theta} = (\theta_1, \theta_2, \theta_3)$ being $(0.8, 0.7, -0.06)$. Since the specific value of θ_1 is of little economic interest, we focus on the performance of the new statistical

procedures for making valid inference about the θ_2 and θ_3 parameters that characterize the volume-volatility elasticity.

2.4.2 Simulation results

We report the results for both least-square regression estimates (corresponding to a quadratic loss function) and q-quantile regression estimates with $q \in \{0.1, 0.25, 0.5, 0.75, 0.9\}$ (corresponding to a lin-lin loss function, $L(x) = x(q - 1_{\{x<0\}})$). In addition to the θ_2 and θ_3 parameter estimates, we also compute the 90% and 95% level two-sided symmetric confidence intervals (CI) based on the bootstrap Algorithm 1, along with the bias-corrected versions thereof. In line with Bollerslev et al. (2018), and the actual empirical application discussed below, we fix the local window parameter $k_n = 30$ for all of the estimates, corresponding to half an hour before and after each announcement. In results not presented here, we find that varying the window k_n between 25 to 35 has little impact on the simulation results. Table 2.1 (resp. Table 2.2) reports the results where the ϵ_i shocks are drawn from a scaled chi-square distribution with $df = 10$ (resp. $df = 30$) degrees-of-freedom. Both of the tables are based on a total of 1,000 Monte Carlo replications.

We begin our discussion with Table 2.1, pertaining to $df = 10$ and the more weakly autocorrelated trading volume process. Looking first at the results in Panel A for the estimates of θ_2, we find that the coverage rates of the CIs associated with the uncorrected estimators (reported in the left part of the table) are all close to, albeit mostly slightly below, the corresponding nominal levels. At the same time, the relative biases in the uncorrected estimates are nontrivial. By contrast, the bootstrap bias-correction (reported in the right part of the table) substantially reduces the

Table 2.1: Monte Carlo Simulation Results: Weakly Autocorrelated Trading Volume

	Uncorrected Estimator				Corrected Estimator			
	Bias	RMSE	90% CI	95% CI	Bias	RMSE	90% CI	95% CI
				Panel A: θ_2				
Least-square	-10%	0.079	83%	88%	-2%	0.048	87%	91%
$q = 0.10$	-12%	0.108	86%	91%	-3%	0.083	88%	93%
$q = 0.25$	-10%	0.091	85%	90%	-2%	0.063	87%	92%
$q = 0.50$	-9%	0.083	85%	90%	-2%	0.058	88%	92%
$q = 0.75$	-8%	0.083	85%	90%	-2%	0.062	87%	92%
$q = 0.90$	-7%	0.090	88%	93%	-2%	0.076	91%	94%
				Panel B: θ_3				
Least-square	-2%	0.012	94%	96%	0%	0.012	92%	95%
$q = 0.10$	-12%	0.021	94%	97%	-2%	0.022	91%	96%
$q = 0.25$	-6%	0.016	94%	97%	0%	0.017	91%	96%
$q = 0.50$	-1%	0.015	93%	96%	1%	0.015	91%	96%
$q = 0.75$	3%	0.017	92%	96%	0%	0.016	92%	96%
$q = 0.90$	8%	0.022	93%	96%	0%	0.020	93%	97%

Note: This table reports the relative bias (Bias), root mean squared error (RMSE) and the coverage rates for 90% and 95% confidence intervals (CI) for the elasticity slope parameters θ_2 (Panel A) and θ_3 (Panel B) in the specification $\Delta\widehat{\log(m_\tau)} = \theta_1 + (\theta_2 + \theta_3 X_\tau)\Delta\widehat{\log(\sigma_\tau)} + e_\tau$. The results are based on a total of 1,000 Monte Carlo replications. The local window used in the estimation is fixed at $k_n = 30$. The degrees-of-freedom in the chi-square distribution for $\tilde{\epsilon}_i$ is set to $df = 10$. The number of bootstrap resampling is 1,000. Results for the uncorrected and bias-corrected procedures are reported in the left and right set of columns, respectively. The rows labeled least-square report the least-squares regression estimates. The rows labeled $q = 0.10, ..., 0.90$ report the corresponding quantile-regression estimates.

relative bias from roughly -10% to just -2%. (Examining the higher order bias in the bias-corrected estimator remains to be an open question.) The bias correction generally also improves the size of the CIs.

Turning to Panel B and the estimates for the θ_3 parameter, the results again indicate quite good coverage properties of the CIs, although the intervals now appear to be somewhat conservative. The relative biases for the θ_3 parameter are generally smaller than for the θ_2 parameter, but sill quite large for some of the more extreme

Table 2.2: Monte Carlo Simulation Results: Strongly Autocorrelated Trading Volume

	Uncorrected Estimator				Corrected Estimator			
	Bias	RMSE	90% CI	95% CI	Bias	RMSE	90% CI	95% CI
				Panel A: θ_2				
Least-square	-10%	0.075	77%	82%	-2%	0.040	81%	87%
$q = 0.10$	-12%	0.103	82%	87%	-3%	0.067	85%	90%
$q = 0.25$	-11%	0.089	78%	85%	-3%	0.053	82%	88%
$q = 0.50$	-10%	0.078	80%	85%	-2%	0.046	83%	89%
$q = 0.75$	-8%	0.072	83%	88%	-2%	0.049	86%	91%
$q = 0.90$	-7%	0.074	87%	91%	-2%	0.058	89%	94%
				Panel B: θ_3				
Least-square	-2%	0.009	92%	95%	0%	0.010	89%	94%
$q = 0.10$	-16%	0.019	90%	94%	-4%	0.019	87%	92%
$q = 0.25$	-9%	0.014	90%	94%	-2%	0.014	87%	93%
$q = 0.50$	-2%	0.011	93%	96%	-1%	0.012	92%	94%
$q = 0.75$	5%	0.013	92%	96%	1%	0.012	92%	96%
$q = 0.90$	11%	0.017	94%	97%	0%	0.014	95%	98%

Note: This table reports the relative bias (Bias), root mean squared error (RMSE) and the coverage rates for 90% and 95% confidence intervals (CI) for the elasticity slope parameters θ_2 (Panel A) and θ_3 (Panel B) in the specification $\Delta\widehat{\log(m_\tau)} = \theta_1 + (\theta_2 + \theta_3 X_\tau)\Delta\widehat{\log(\sigma_\tau)} + e_\tau$. The results are based on a total of 1,000 Monte Carlo replications. The local window used in the estimation is fixed at $k_n = 30$. The degrees-of-freedom in the chi-square distribution for $\tilde{\epsilon}_i$ is set to $df = 30$. The number of bootstrap resampling is 1,000. Results for the uncorrected and bias-corrected procedures are reported in the left and right set of columns, respectively. The rows labeled least-square report the least-squares regression estimates. The rows labeled $q = 0.10, ..., 0.90$ report the corresponding quantile-regression estimates.

uncorrected quantile estimates. Meanwhile, the corrected estimates are effectively all unbiased. Interestingly, the relative biases for the uncorrected least-square and median ($q = 0.5$) regressions are both close to zero, even without any bias-correction, suggesting that finite-sample bias is not a major concern for properly assessing the "central" dependency of the slope coefficient on other covariates.

The results in Table 2.2 for $df = 30$ and the more persistent volume process are fairly similar to those in Table 2.1. Again, we find that the bootstrap method

effectively reduces the finite sample biases, and that the coverage rates of the CIs are close to the corresponding nominal levels. The θ_2 coefficients have slightly larger size distortions in this situation, but the CIs for the θ_3 coefficients still exhibit quite good coverage properties.

All-in-all, the Monte Carlo results clearly underscore the reliability of the new estimation method and accompanying bootstrap inference procedures in a realistically calibrated simulation setting. In particular, the bootstrap-based bias correction is quite effective in reducing the finite-sample bias for all of the estimators. In additional simulations not presented here, we further verify that the bias is indeed a finite-sample issue, as it becomes much smaller for data sampled at higher frequencies.

2.5 Macroeconomic news, volume, and volatility

2.5.1 Data description

Our primary data consists of intraday observations on trading volume and transaction prices for the E-mini futures contract on the S&P 500 index obtained from TickData. We sample the data at every minute to help mitigate the effect of market microstructure noise (see, e.g., the discussion in Zhang et al. (2005)). The sample spans 7:00am to 4:15pm from July 1, 2003 to March 2, 2017. We further removed days with irregular trading hours. In the end, we are left with a total 3,383 trading days, comprising 1,880,948 one-minute return and trading volume observations.

In addition to the price and volume data, we also utilize information about the date and time of two important macroeconomic announcements: namely the Federal Open Market Committee (FOMC) rate decisions and statements about monetary policy, and the nonfarm payroll (NFP) employment report. These particular an-

nouncements are generally considered to be the two most important macroeconomic news announcements (see, e.g., Andersen et al. (2003), Andersen et al. (2007)). The FOMC decision is typically announced every six-week at 2:15pm, while the NFP report is released at 8:30am on the first Friday of each month. We rely on Bloomberg's Economic Calendar to pinpoint the exact time and date. Importantly, our use of futures data spanning several hours before the opening of the "cash" market at 9:30am allows us to study the all-important NFP report (this contrasts with many other studies, including Bollerslev et al. (2018), which rely on data during regular trading hours only). In total our sample contains 110 FOMC and 157 NFP announcements.

2.5.2 Volume-volatility elasticity and investor disagreement

We apply the proposed method to investigate the relationship among volatility and volume jumps, by revisiting the analysis in Bollerslev et al. (2018) pertaining to the volume-volatility elasticity. Our baseline specification, corresponding to equation (2.7) above, takes the form,

$$\widetilde{\Delta \log(m_\tau)} = \theta_1 + \theta_2 \widetilde{\Delta \log(\sigma_\tau)} + e_\tau, \tag{2.24}$$

where again $\widetilde{\Delta \log(m_\tau)}$ and $\widetilde{\Delta \log(\sigma_\tau)}$ denote the DID jump estimates of the volume intensity and spot volatility, respectively. Based on least-square estimation methods, Bollerslev et al. (2018) found the elasticity (i.e., θ_2) to be generally below unity, which according to the economic theory of Kandel and Pearson (1995) is indicative of disagreement among investors in interpreting the macroeconomic news announcements. Furthermore, by parameterizing the elasticity as a function of proxies of investors'

disagreement X_τ,

$$\Delta \widetilde{\log(m_\tau)} = \theta_1 + (\theta_2 + \theta_3 X_\tau) \Delta \widetilde{\log(\sigma_\tau)} + e_\tau, \tag{2.25}$$

Bollerslev et al. (2018) also found the elasticity to be generally lower for higher levels of disagreement (i.e., θ_3 is negative). This again accords with the theoretical implications derived from Kandel and Pearson (1995).

Following the analysis in Bollerslev et al. (2018), we will consider two different disagreement proxies: (i) the dispersion in the forecasts in the Survey of Professional Forecasters (SPF) for the one-quarter-ahead unemployment rate (the unemployment rate serves a natural gauge for the state of the macro economy, but the dispersion in the forecasts for other macroeconomic variables, like GDP growth, leads to very similar results), and (ii) the Economic Policy Uncertainty (EPU) index of Baker et al. (2016) (a more detailed rationale for the use of this specific disagreement proxy is provided in Bollerslev et al. (2018)). As in Bollerslev et al. (2018), we also set the local window $k_n = 30$ as a reasonable rule-of-thumb in view of the simulation results shown in the supplemental appendix. Recall that the theory features undersmoothing (i.e., k_n being small) in order to reduce the bias in the spot estimation resulting from time-varying volatility and volume intensity, which suggests using a small k_n. However, at the same time, the inference is based on central limit theorems under $k_n \to \infty$, so the limit theory would not "kick in" when k_n is too small. A data-driven choice of k_n in the present context is a challenging open question, which may be an interesting topic for future research.

Our analysis advances Bollerslev et al. (2018) in three important ways. First, our use of futures data, which is available before the regular trading hours for the

SPY ETF used by Bollerslev et al. (2018), allows us to study the all-important NFP announcement. Second, we complement the least-square estimation strategy used by Bollerslev et al. (2018) with the new quantile-regression type estimators formally developed here, so as to uncover (potentially) heterogeneous responses in the volume-volatility relationship across quantiles. Third, since the regressors in (2.24) and (2.25) are estimated with error, the findings reported in Bollerslev et al. (2018) could be affected by finite-sample "attenuation" biases, which our new bootstrap bias-correction technique conveniently circumvents.

To begin, Figure 2.2 plots the least-square and quantile-regression estimates for the θ_2 volume-volatility elasticity parameter based on the baseline specification in (2.24), along with their 90% two-sided CIs. For the FOMC (resp. NFP) announcements reported in the left (resp. right) panel, the bias-corrected least-square elasticity estimate equals 0.714 (resp. 0.733). Although both of these estimates exceed their uncorrected counterparts (equal to 0.697 and 0.687, respectively), they are still significantly below unity, consistent with the implications from the economic theory of Kandel and Pearson (1995) and the presence of disagreement among investors.

The median regression estimate ($q = 0.5$) of 0.703 for the FOMC announcements is also very close to the least-square estimate of 0.714. Hence, the least-square estimate appears robust, in the sense that it is not driven by a few influential outliers. Importantly, all of the elasticity estimates for the FOMC announcements are also below unity and generally statistically significantly so. As such , this further buttresses the idea that investor disagreement plays an important role in the functioning of markets.

The quantile regression estimates for θ_2 for the NFP announcements are also

FIGURE 2.2: Baseline Volume-Volatility Elasticity Estimates

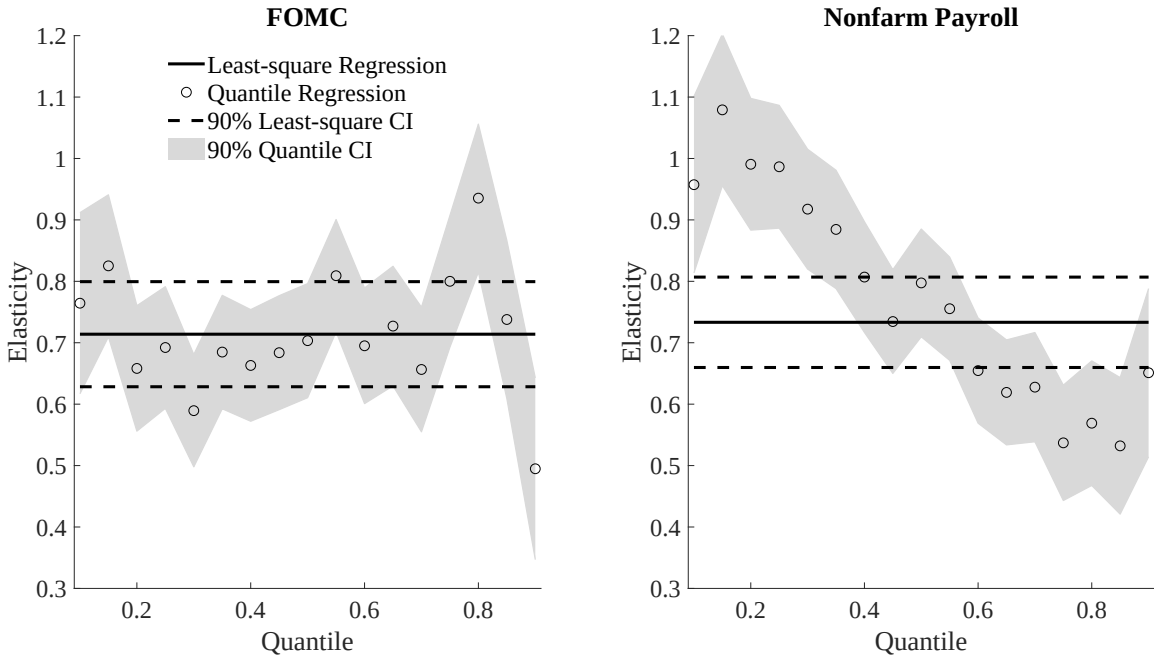

Note: This figure reports the least-square (solid line) and quantile-regression (circles) estimates of the θ_2 elasticity coefficient, along with their confidence intervals (CI), for the baseline specification without any covariates, $\widetilde{\Delta \log(m_\tau)} = \theta_1 + \theta_2 \widetilde{\Delta \log(\sigma_\tau)} + e_\tau$. The left (resp. right) panel gives the estimates around FOMC (resp. NFP) announcements.

mostly below unity. However, in contrast to the fairly homogeneous FOMC quantile estimates, there is a clear downward pattern in the quantile elasticity estimates for the NFP announcements. In particular, the estimates for the lower quantiles are all close to, and from a statistical perspective equivalent to, unity. This therefore suggests that for the NFP announcements a rational-expectation type interpretation, in which most investors agree, is sometimes operative.

A central tenet of all economic disagreement models, the Kandel–Pearson model Kandel and Pearson (1995) included, is that higher levels of disagreement among in-

FIGURE 2.3: Volume-Volatility Elasticity and Disagreement

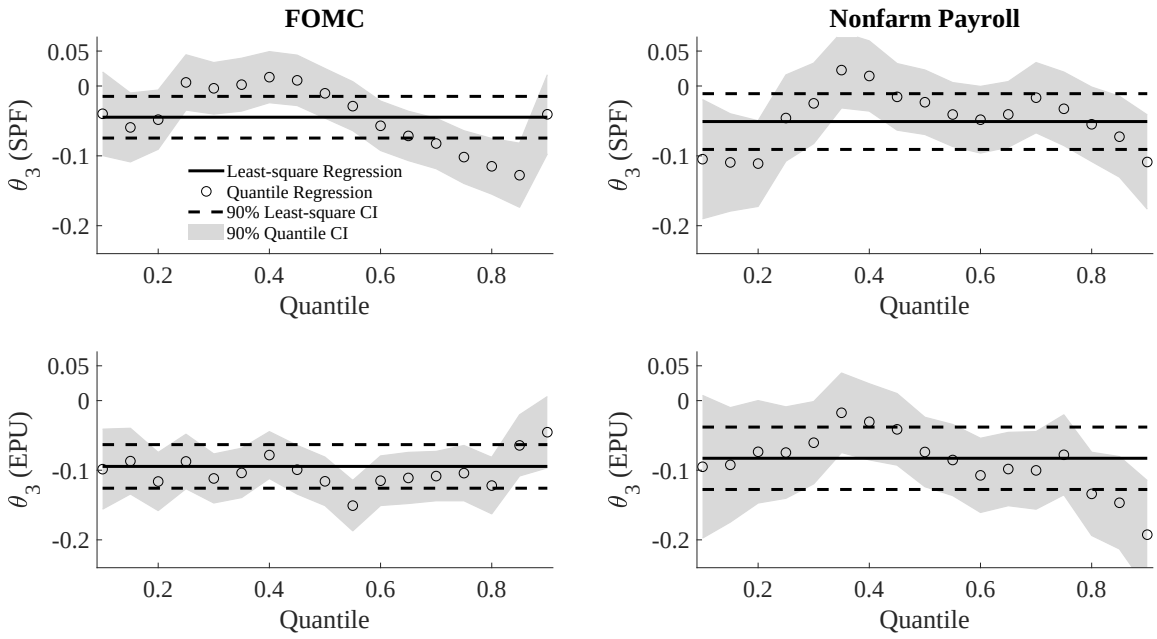

Note: This figure reports the least-square (solid line) and quantile-regression (circles) estimates of the θ_3 coefficient, along with their confidence intervals (CI), for the specification $\widetilde{\Delta \log(m_\tau)} = \theta_1 + (\theta_2 + \theta_3 X_\tau)\widetilde{\Delta \log(\sigma_\tau)} + e_\tau$, where X_τ denotes the disagreement proxy. The top (resp. bottom) row reports the results based on the dispersion measure among professional forecasters (SPF) (resp. the Economic Policy Uncertainty (EPU) index). The left (resp. right) panel gives the estimates around FOMC (resp. NFP) announcements.

vestors should "loosen" the relationship between trading volume and volatility. More specifically, following the analysis of Bollerslev et al. (2018) this should manifest in the volume-volatility elasticity being a decreasing function of the level of disagreement. To examine this hypothesis, Figure 2.3 plots the least-square and quantile-regression estimates for the θ_3 parameter from the specification in (2.25), along with their 90% two-sided CIs. The left (resp. right) two panels report the estimates for FOMC (resp. NFP).

Looking first at the results in the top row based on the use of the forecast dispersion among professional forecasters (SPF) as a measure of disagreement, the θ_3 estimates for both the FOMC and NFP announcements are generally below zero across all quantiles, and often significantly so. This finding is quite remarkable, as it suggests that the negative relationship between the volume-volatility elasticity and disagreement predicted by the economic theory, holds not only on average (consistent with the least-square estimates previously reported in Bollerslev et al. (2018)), but across all quantiles. In other words, this negative relation is a robust feature that does not seem to depend on a particular set of announcements. The results reported in the bottom row based on the Economic Policy Uncertainty (EPU) index further reinforces this same conclusion. In fact, if anything these results are even stronger, with all of the estimates below zero.

In sum, our new bias-corrected estimators confirm prior (potentially biased) empirical evidence that the volume-volatility elasticity around important news announcements is generally below unity. Moreover, this holds true not only on average, but across all quantiles. It also holds true not only for FOMC announcements, but also for the nonfarm payroll employment report, often referred to as the "king" of announcements by market participants. Finally, further corroborating the underlying economic theory and the import of investor disagreement, the new methods reveal the elasticity to be a robustly decreasing function of aggregate levels of disagreement.

2.6 Conclusion

We propose a general minimum-distance type estimator for estimating the relationship between jumps in instantaneous moments of stochastic processes. The asymp-

totic distribution of the proposed estimator, derived under an in-fill asymptotic setting, is generally non-standard. We propose an easy-to-implement bootstrap algorithm for conducting feasible inference and bias-correction. Using high-frequency intraday data for the S&P 500 E-mini futures contract, we apply the new methods to study the behavior of trading intensity and spot volatility at the time of important macroeconomic news announcement. Consistent with the implications from economic theory and a model in which investors agree-to-disagree, we find that the estimated volume-volatility elasticities are below unity and negatively related to the level of investor disagreement, not only "on average," but across all quantiles.

2.7 Technical assumptions and proofs

This appendix presents the technical details underlying our statistical inference procedures discussed in Section 2.3. Section 2.7.1 collects the regularity conditions. Section 2.7.2 provides the proof of Theorem 2.3.3, which include Theorems 2.3.1 and 2.3.2 as special cases.

2.7.1 Assumptions

(i) The price process P is defined by (2.1) on some filtered probability space $(\Omega, \mathcal{F}, (\mathcal{F}_t)_{t\geq 0}, \mathbb{P})$ for

$$J_t = \int_0^t \varphi_s dN_s + \int_0^t \int_{\mathbb{R}} \delta(s,z)\mu(ds,dz),$$

where the processes α and σ are càdlàg (i.e., right continuous with left limit) and adapted; the process φ is predictable and locally bounded; N is a counting process that jumps at the scheduled announcement times which are specified by the set $\mathcal{T}$; 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 process V satisfies (2.3). The state 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 $[s,t] \cap \mathcal{T} = \emptyset$.

(iv) The process X is adapted.

Assumption 2.7.1 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 (ϵ_i) at observation times from the latent continuous-time state process (ζ_t). This condition only mildly restricts the V 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 $\mathcal{T}$) and for long-memory type processes driven by the fractional Brownian motion.

In addition, we need the following conditions for the nonparametric analysis, where we denote $M_q(\cdot) \equiv \int \mathcal{V}(\cdot, \epsilon)^q F_\epsilon(d\varepsilon)$ for $q \geq 1$.

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

The function $M_1(\cdot)$ is Lipschitz on compact sets and the functions $M_2(\cdot)$ and $M_4(\cdot)$

are continuous.

Assumption 2.7.1 specifies the growth rate of the local window size k_n. As typical in nonparametric analysis, this condition features a type of undersmoothing (i.e., $k_n \ll \Delta_n^{-1/2}$), so as to permit feasible inference. Assumption 2.7.1 imposes some smoothness conditions that are very mild.

2.7.2 Proofs of main results

We note that Theorems 2.3.1 and 2.3.2 are special cases of Theorem 2.3.3 with the control group being empty. Hence, it suffices to prove Theorem 2.3.3. Below, we denote $\mathcal{A} \equiv \mathcal{T} \cup (\cup_{\tau \in \mathcal{T}} \mathcal{C}(\tau))$, which collects all times for announcements and associated control groups. In the setting of Theorems 2.3.1 and 2.3.2, $\mathcal{C}(\tau) = \emptyset$ and $\mathcal{A} = \mathcal{T}$.

Proof of Theorem 2.3.3(a). Denote

$$\begin{aligned}
\hat{\xi}_\tau &\equiv k_n^{1/2} \left[G\left((\hat{m}_{t-}, \hat{m}_t, \hat{c}_{t-}, \hat{c}_t)_{t \in \{\tau\} \cup \mathcal{C}(\tau)} \right) - G\left((m_{t-}, m_t, c_{t-}, c_t)_{t \in \{\tau\} \cup \mathcal{C}(\tau)} \right) \right], \\
\hat{\xi}'_{k,\tau} &\equiv k_n^{1/2} \left[H_k\left((\hat{m}_{t-}, \hat{m}_t, \hat{c}_{t-}, \hat{c}_t)_{t \in \{\tau\} \cup \mathcal{C}(\tau)} \right) - H_k\left((m_{t-}, m_t, c_{t-}, c_t)_{t \in \{\tau\} \cup \mathcal{C}(\tau)} \right) \right].
\end{aligned}$$

By Theorem 1 in Bollerslev et al. (2018),

$$\begin{aligned}
&k_n^{1/2}(\hat{m}_{\tau-} - m_{\tau-}, \hat{m}_{\tau+} - m_\tau, \hat{c}_{\tau-} - c_{\tau-}, \hat{c}_{\tau+} - c_\tau)_{\tau \in \mathcal{A}} \\
&\xrightarrow{\mathcal{L}\text{-}s} (\eta_\tau)_{\tau \in \mathcal{A}} = (\eta_{m,\tau-}, \eta_{m,\tau-}, \eta_{c,\tau-}, \eta_{c,\tau+})_{\tau \in \mathcal{A}},
\end{aligned}$$

where $\xrightarrow{\mathcal{L}\text{-}s}$ denotes stable convergence in law. By the delta method, we further have

$$(\hat{\xi}_\tau, (\hat{\xi}'_{k,\tau})_{1 \le k \le K})_{\tau \in \mathcal{A}} \xrightarrow{\mathcal{L}\text{-}s} \left(\tilde{\xi}_\tau, (\tilde{\xi}'_{k,\tau})_{1 \le k \le K} \right)_{\tau \in \mathcal{A}}.$$

Recall that, at the true parameter value,

$$G\left((m_{t-}, m_t, c_{t-}, c_t)_{t \in \{\tau\} \cup \mathcal{C}(\tau)} \right) = \sum_{k=1}^{K} \boldsymbol{\theta}_{0,k}^\top \boldsymbol{X}_{k,\tau} H_k\left((m_{t-}, m_t, c_{t-}, c_t)_{t \in \{\tau\} \cup \mathcal{C}(\tau)} \right).$$

Hence,

$$
\begin{aligned}
&G\left((\hat{m}_{t-},\hat{m}_t,\hat{c}_{t-},\hat{c}_t)_{t\in\{\tau\}\cup\mathcal{C}(\tau)}\right)\\
&\qquad-\sum_{k=1}^{K}(\boldsymbol{\theta}_{0,k}+k_n^{-1/2}\boldsymbol{h}_k)^\top \boldsymbol{X}_{k,\tau}H_k\left((\hat{m}_{t-},\hat{m}_t,\hat{c}_{t-},\hat{c}_t)_{t\in\{\tau\}\cup\mathcal{C}(\tau)}\right)\\
&=G\left((m_{t-},m_t,c_{t-},c_t)_{t\in\{\tau\}\cup\mathcal{C}(\tau)}\right)+k_n^{-1/2}\hat{\xi}_\tau\\
&\qquad-\sum_{k=1}^{K}(\boldsymbol{\theta}_{0,k}+k_n^{-1/2}\boldsymbol{h}_k)^\top \boldsymbol{X}_{k,\tau}H_k\left((m_{t-},m_t,c_{t-},c_t)_{t\in\{\tau\}\cup\mathcal{C}(\tau)}\right)\\
&\qquad-\sum_{k=1}^{K}(\boldsymbol{\theta}_{0,k}+k_n^{-1/2}\boldsymbol{h}_k)^\top \boldsymbol{X}_{k,\tau}k_n^{-1/2}\hat{\xi}'_{k,\tau}\\
&=k_n^{-1/2}\left(\hat{\xi}_\tau-\sum_{k=1}^{K}\boldsymbol{\theta}_{0,k}\boldsymbol{X}_{k,\tau}\hat{\xi}'_{k,\tau}-\sum_{k=1}^{K}\boldsymbol{h}_k^\top\boldsymbol{X}_{k,\tau}H_k\left((m_{t-},m_t,c_{t-},c_t)_{t\in\{\tau\}\cup\mathcal{C}(\tau)}\right)+o_p(1)\right).
\end{aligned}
$$

In view of the property that $L(cx)=|c|^p\,L(x)$, we further deduce

$$
\widetilde{M}_n(\boldsymbol{h})=\sum_{\tau\in\mathcal{T}}L\left(\hat{\xi}_\tau-\sum_{k=1}^{K}\boldsymbol{\theta}_{0,k}^\top\boldsymbol{X}_{k,\tau}\hat{\xi}'_{k,\tau}-\sum_{k=1}^{K}\boldsymbol{h}_k^\top\boldsymbol{X}_{k,\tau}H_k\left((m_{t-},m_t,c_{t-},c_t)_{t\in\{\tau\}\cup\mathcal{C}(\tau)}\right)+o_p(1)\right).
$$

Since $L(\cdot)$ is convex, it is necessarily continuous. Therefore, by the continuous mapping theorem, we deduce that for any $m\geq 1$ and $\boldsymbol{h}^{(1)},\ldots,\boldsymbol{h}^{(m)}$, the variables $(\widetilde{M}_n(\boldsymbol{h}^{(1)}),\ldots,\widetilde{M}_n(\boldsymbol{h}^{(m)}))$ converges stably in law to $(\widetilde{M}(\boldsymbol{h}^{(1)}),\ldots,\widetilde{M}(\boldsymbol{h}^{(m)}))$.

Fix any bounded $\mathcal{F}$-measurable random variable U. Given the finite-dimensional convergence above, there exists a probability space, on which processes $\bar{M}_n(\cdot)$ and $\bar{M}(\cdot)$ and variable $\bar{U}$ are defined, such that $(\bar{M}_n(\cdot),\bar{M}(\cdot),\bar{U})$ has the same finite-dimensional distributions as $(\widetilde{M}_n(\cdot),\widetilde{M}(\cdot),U)$ and $\bar{M}_n(\boldsymbol{h})\to\bar{M}(\boldsymbol{h})$ almost surely. Let

$$
\bar{\boldsymbol{h}}_n=\operatorname*{argmin}_{\boldsymbol{h}}\bar{M}_n(\boldsymbol{h}),\quad \bar{\boldsymbol{h}}=\operatorname*{argmin}_{\boldsymbol{h}}\bar{M}(\boldsymbol{h}).
$$

Since $\widetilde{M}_n(\cdot)$ is convex, we can use the same argument as in Lemma A of Knight

(1989) to deduce that $(\bar{\boldsymbol{h}}_n, \bar{U}) \to (\bar{\boldsymbol{h}}, \bar{U})$ almost surely. This further implies that $(\tilde{\boldsymbol{h}}_n, U)$ converges to $(\tilde{\boldsymbol{h}}, U)$ in law. Since U is arbitrary, we deduce that $\tilde{\boldsymbol{h}}_n$ converges stably in law to $\tilde{\boldsymbol{h}}$. □

Proof of Theorem 2.3.3(b). Denote $\mathcal{G} \equiv \mathcal{F} \vee \sigma(\epsilon_i : i \geq 0)$. Below, we denote

$$\begin{aligned}
\hat{\xi}^*_\tau &\equiv k_n^{1/2}\left[G\left(\left(\hat{m}^*_{t-}, \hat{m}^*_{t+}, \hat{c}^*_{t-}, \hat{c}^*_{t+}\right)_{t\in\{\tau\}\cup\mathcal{C}(\tau)}\right) - G\left(\left(\hat{m}_{t-}, \hat{m}_{t+}, \hat{c}_{t-}, \hat{c}_{t+}\right)_{t\in\{\tau\}\cup\mathcal{C}(\tau)}\right)\right], \\
\hat{\xi}'^*_{k,\tau} &\equiv k_n^{1/2}\left[H_k\left(\left(\hat{m}^*_{t-}, \hat{m}^*_{t+}, \hat{c}^*_{t-}, \hat{c}^*_{t+}\right)_{t\in\{\tau\}\cup\mathcal{C}(\tau)}\right) - H_k\left(\left(\hat{m}_{t-}, \hat{m}_{t+}, \hat{c}_{t-}, \hat{c}_{t+}\right)_{t\in\{\tau\}\cup\mathcal{C}(\tau)}\right)\right].
\end{aligned}$$

By Theorem 1(b) in Bollerslev et al. (2018),

$$k_n^{1/2}\left(\hat{m}^*_{\tau-} - \hat{m}_{\tau-}, \hat{m}^*_{\tau+} - \hat{m}_{\tau+}, \hat{c}^*_{\tau-} - \hat{c}_{\tau-}, \hat{c}^*_{\tau+} - \hat{c}_{\tau+}\right)_{\tau\in\mathcal{A}} \xrightarrow{\mathcal{L}|\mathcal{G}} (\eta_{m,\tau-}, \eta_{m,\tau+}, \eta_{c,\tau-}, \eta_{c,\tau+})_{\tau\in\mathcal{A}},$$

where $\xrightarrow{\mathcal{L}|\mathcal{G}}$ denotes the convergence in probability of the $\mathcal{G}$-conditional distribution functions under the uniform metric. Consequently, by the delta method,

$$(\hat{\xi}^*_\tau, (\hat{\xi}'^*_{k,\tau})_{1\leq k\leq K})_{\tau\in\mathcal{A}} \xrightarrow{\mathcal{L}|\mathcal{G}} (\tilde{\xi}_\tau, (\tilde{\xi}'_{k,\tau})_{1\leq k\leq K})_{\tau\in\mathcal{A}}.$$

We now note that

$$\begin{aligned}
&G(\widetilde{\boldsymbol{S}}^*_\tau) - \tilde{\varepsilon}_\tau - \sum_{k=1}^K \boldsymbol{\theta}_k^\top \boldsymbol{X}_{k,\tau} H_k(\widetilde{\boldsymbol{S}}^*_\tau) \\
&\quad = G(\widetilde{\boldsymbol{S}}^*_\tau) - G(\widetilde{\boldsymbol{S}}_\tau) + \sum_{k=1}^K \widetilde{\boldsymbol{\theta}}_k^\top \boldsymbol{X}_{k,\tau}\left(H_k(\widetilde{\boldsymbol{S}}^*_\tau) - k_n^{-1/2}\hat{\xi}'^*_{k,\tau}\right) - \sum_{k=1}^K \boldsymbol{\theta}_k^\top \boldsymbol{X}_{k,\tau} H_k(\widetilde{\boldsymbol{S}}^*_\tau) \\
&\quad = k_n^{-1/2}\hat{\xi}^*_\tau - k_n^{-1/2}\sum_{k=1}^K \widetilde{\boldsymbol{\theta}}_k^\top \boldsymbol{X}_{k,\tau}\hat{\xi}'^*_{k,\tau} - \sum_{k=1}^K (\boldsymbol{\theta}_k - \widetilde{\boldsymbol{\theta}}_k)^\top \boldsymbol{X}_{k,\tau} H_k(\widetilde{\boldsymbol{S}}^*_\tau).
\end{aligned}$$

Furthermore, with the reparameterization $\boldsymbol{\theta}_k = \widetilde{\boldsymbol{\theta}}_k + k_n^{-1/2}\boldsymbol{h}_k$, we can rewrite the

above as

$$k_n^{-1/2}\left(\hat{\xi}_\tau^* - \sum_{k=1}^{K} \widetilde{\boldsymbol{\theta}}_k^\top \boldsymbol{X}_{k,\tau}\hat{\xi}_{k,\tau}'^* - \sum_{k=1}^{K} \boldsymbol{h}_k^\top \boldsymbol{X}_{k,\tau} H_k(\widetilde{\boldsymbol{S}}_\tau^*)\right).$$

Consider the reparameterized objective function $\widetilde{M}_n^*(\boldsymbol{h}) = k_n^{p/2}\widetilde{Q}_n^*(\widetilde{\boldsymbol{\theta}}_n + k_n^{-1/2}\boldsymbol{h})$. From the derivation above, we see that $\widetilde{M}_n^*(\boldsymbol{h})$ can be rewritten as

$$\widetilde{M}_n^*(\boldsymbol{h}) = \sum_{\tau\in\mathcal{T}} L\left(\hat{\xi}_\tau^* - \sum_{k=1}^{K} \widetilde{\boldsymbol{\theta}}_k^\top \boldsymbol{X}_{k,\tau}\hat{\xi}_{k,\tau}'^* - \sum_{k=1}^{K} \boldsymbol{h}_k^\top \boldsymbol{X}_{k,\tau} H_k(\widetilde{\boldsymbol{S}}_\tau^*)\right).$$

We now show that $\widetilde{M}_n^*(\boldsymbol{h}) \xrightarrow{\mathcal{L}|\mathcal{G}} \widetilde{M}(\boldsymbol{h})$ for any fixed $\boldsymbol{h}$. Consider any subsequence $\mathbb{N}_1 \subseteq \mathbb{N}$. Since $\tilde{\boldsymbol{\theta}}_n \xrightarrow{\mathbb{P}} \boldsymbol{\theta}_0$ and $(\hat{\xi}_\tau^*, (\hat{\xi}_{k,\tau}'^*)_{1\le k\le K})_{\tau\in\mathcal{A}} \xrightarrow{\mathcal{L}|\mathcal{G}} (\tilde{\xi}_\tau, (\tilde{\xi}_{k,\tau}')_{1\le k\le K})_{\tau\in\mathcal{A}}$, there exists a further subsequence $\mathbb{N}_2 \subseteq \mathbb{N}_1$ such that, along $\mathbb{N}_2$, $\widetilde{\boldsymbol{\theta}}_n \to \boldsymbol{\theta}_0$ and the $\mathcal{G}$-conditional distribution of $(\hat{\xi}_\tau^*, (\hat{\xi}_{k,\tau}'^*)_{1\le k\le K})_{\tau\in\mathcal{A}}$ converges to that of $(\tilde{\xi}_\tau, (\tilde{\xi}_{k,\tau}')_{1\le k\le K})_{\tau\in\mathcal{A}}$ almost surely. On each path with these convergences, we also see that $H_k(\widetilde{\boldsymbol{S}}_\tau^*) \xrightarrow{\mathbb{P}} H_k((m_{t-}, m_t, c_{t-}, c_t)_{t\in\{\tau\}\cup\mathcal{C}(\tau)})$ under the transition probability conditionally on $\mathcal{G}$. Hence, by the continuous mapping theorem, we deduce that the $\mathcal{G}$-conditional distribution of $\widetilde{M}_n^*(\boldsymbol{h})$ converges to that of $\widetilde{M}(\boldsymbol{h})$ almost surely along $\mathbb{N}_2$. By another use of the subsequence argument, we see that $\widetilde{M}_n^*(\boldsymbol{h}) \xrightarrow{\mathcal{L}|\mathcal{G}} \widetilde{M}(\boldsymbol{h})$ as wanted.

It is easy to extend $\widetilde{M}_n^*(\boldsymbol{h}) \xrightarrow{\mathcal{L}|\mathcal{G}} \widetilde{M}(\boldsymbol{h})$ to a joint convergence on finite-dimensions. By Lemma A of Knight (1989), we deduce that $k_n^{1/2}(\widetilde{\boldsymbol{\theta}}_n^* - \widetilde{\boldsymbol{\theta}}_n) \xrightarrow{\mathcal{L}|\mathcal{G}} \widetilde{\boldsymbol{h}}$. The assertions concerning the coverage rates of the confidence intervals readily follows from this convergence and the symmetry of the conditional distribution of $\widetilde{\boldsymbol{h}}$. □

3

Macro Announcement Disagreement Observed in the Cross-Section of Stocks

3.1 Introduction

One of the primary functions of financial markets is to aggregate public and private information into prices through investor tradings. An examination of investors forecasts suggests that they have different models to understand the economy (Patton and Timmermann, 2010). In addition, a quick glance at the Survey of Professional Forecasters (SPF) rarely reveals a unanimous view among forecasters. Given this difference on how investors forecast important economic variables, it seems reasonable to consider that investors have different interpretations regarding the impact of news announcements on future payoffs for any stock. Therefore, using high frequency data on stocks prices and trades, we answer the question of whether investors disagree when processing macro-news announcement. In the case of observing disagreement, we identify the systematic components that drive said disagreement. Lastly, we de-

termine whether investors disagreement is fully observable from market variables.

Our results suggest that investors do disagree on the impact of macro announcement news for each individual stock future cash flow. Several stock characteristics explain a large portion of the disagreement: idiosyncratic volatility, market capitalization, book-to-market ratio, and institutional ownership. The need to account for such variables suggests that disagreement has many dimensions and evaluating a stock price is a complex task. In addition, we combine the misprice direction with the cost difference to take a short versus long position to indicate that the observed disagreement might be a lower bound for the actual disagreement.

To infer the disagreement between investors using only information about price and trade, we rely on parsimonious economic models. There are two main types of models that differ by the implication as to what should be the relationship between observed volume and volatility. First, Kim and Verrecchia (1991) proposed a model assuming that investors have different information sets but the same interpretation about new public information. As a result, every volume change is tied to a change in prices due to a shift in market expectations for the asset payoff. On the contrary, Kandel and Pearson (1995) observed that usually stock volatility does not move as much as volume around earnings release, and they suggested a model in which investors have different interpretations about the new public information. This development provided a mechanism where volume change is not related to changes on aggregate expectations for future payoffs. With these competing theories, it is possible to derive different implications for what should be the volume-volatility elasticity.

In the aforesaid economic models, there is no specification of what asset class we should use to study the volume-volatility relationship. In fact, the models are

intended for any particular asset which has many traders, suggesting that individual stock elasticity needs to be explored for a better understanding of the relation between volume and volatility. Because the models providing the economic intuition focus on how investors process new information, we use important macroeconomic announcements and high frequency data to isolate the news from other possible effects.

As an example of macro announcement and its non-zero impact on both volume and volatility, Figure 3.1 shows the intraday volume intensity and price changes of the $S\&P500$ index around the FOMC announcement on June 10, 2020.[1] Clearly, these variables underwent large changes at the announcement time, with the price change in the minute after the announcement being 0.36% and the post-event average traded volume at $284,460$ shares. Both realizations represent an event with more than 2 standard deviations away from the mean. These large changes are what we examine to properly estimate the volume-volatility elasticity. In Figure 3.2, we plot how this same event impacted individual stocks and there is an heterogeneity on how price and volume change after the event. It further stresses the need to understand what might be the reason for this heterogeneity around macro announcements and if investor disagreement is a possible explanation.

To accomplish this goal, we need to properly estimate the volume-volatility elasticity at high-frequency. For this reason we rely on the high frequency econometrics literature to have consistent estimators for the volume intensity and spot volatility (Aït-Sahalia and Jacod (2014), Li and Xiu (2016), and Bollerslev et al. (2018)). Since

[1] In this announcement, the Federal Reserve released its forecast for the economy for the first time in 2020 because it was delayed due to the Covid-19 pandemic which started in March.

FIGURE 3.1: SPY Intraday Price and Volume for FOMC Announcement on June 10, 2020

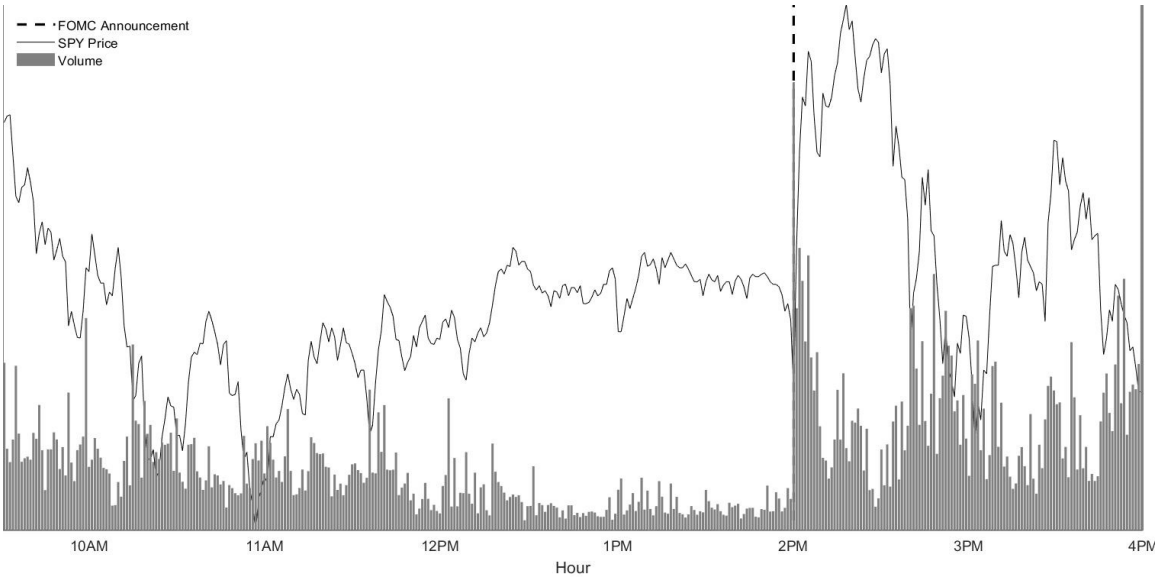

Note: In this figure, we report the intraday price and volume for the SPY and the exact minute of the FOMC announcement, represented by a vertical dashed line, that occurred on June 10, 2020.

our parameter of interest is the volume-volatility elasticity, we must estimate the spot volatility and volume intensity. In addition, following the bootstrap method proposed in (Bollerslev, Li, and Chaves (2020)) we can address the potential attenuation bias.

Once we set the econometric methodology to estimate the elasticity, we explore the literature to understand how to proxy for disagreement in the cross section of stocks. From the information processing literature, it usually there is more room for disagreement whenever the future earnings of a stock are more challenging to forecast and, consequently, its fundamental value (Peng and Xiong (2006), Dávila and Parlatore (2019)). A few characteristics to proxy for this complexity to evaluate future cash-flows and, consequently, investor disagreement, are the idiosyncratic volatility,

FIGURE 3.2: Individual Stocks Intraday Price and Volume for FOMC Announcement on June 10, 2020

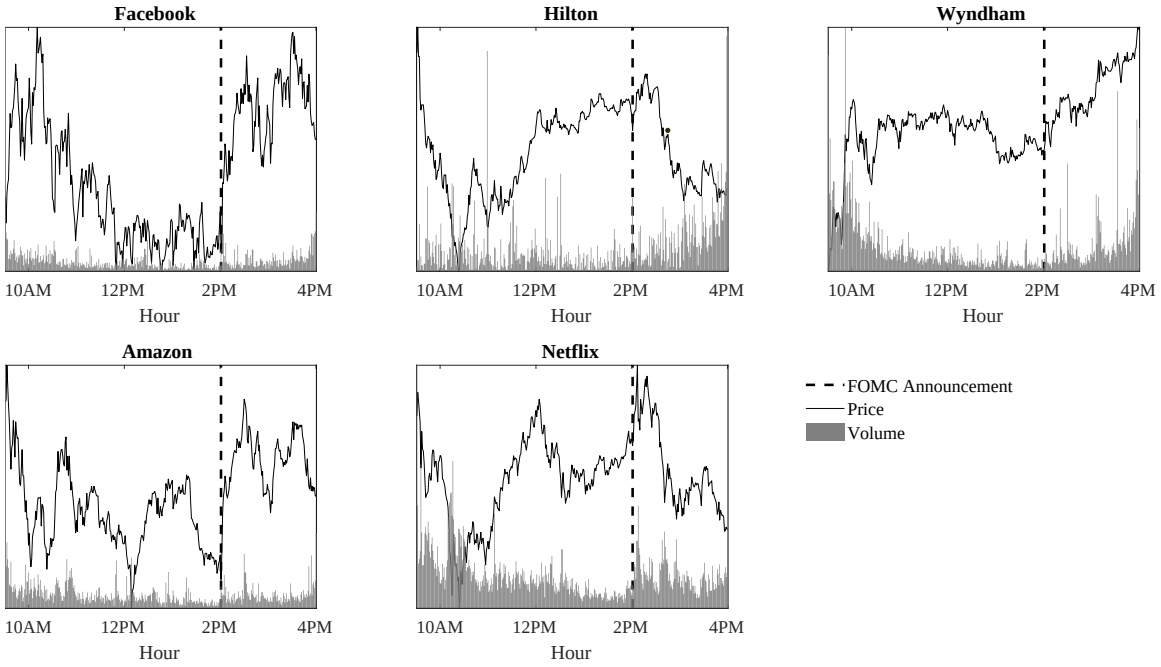

Note: In this figure, we report the intraday price and volume for individual stocks and the exact minute of the FOMC announcement, represented by a vertical dashed line, that occurred on June 10, 2020.

capitalization, and book-to-market ratio (Kumar, 2009).

In addition to characteristics related to information processing, we consider that disagreement can emerge simply from having a diverse pool of investors. According to the arbitrage constraint literature, the riskier it is to take a stock position, the less informed investors will be on the market (De Long et al. (1990), Shleifer and Vishny (1997), and Stambaugh et al. (2015)) and the less heterogeneous will be the observed investors views. To account for the heterogeneity of investors, we use the institutional ownership of a stock because it proxies for the cost of lending a stock and taking a

short position (D'avolio, 2002). Hence, whenever there is lower ownership, the high cost to short should keep investors out of the market and the elasticity would reflect less disagreement.

Motivated by the arbitrage constraint literature, we also assess whether arbitrage constraints keep different opinions outside of the market. That is to say, is the disagreement that we estimate from market variables equal to the total disagreement among investors? We design a setup following the asset pricing literature wherein disagreement can explain overprice and asset bubbles (Miller (1977) and Hong and Stein (2007)). This deviation from the true value can be generated only if there are arbitrage constraints which matter more for investors willing to take short than long positions. Therefore, we explore the direction of misprice, a measure proposed by Stambaugh et al. (2015), and we use the percent of shares owned by institutional investors to proxy for short constraint (D'avolio (2002)). According to our estimate, the volume-volatility elasticity shows an asymmetric pattern when looking at different misprice and different arbitrage constraints. As expected, the short constraint matters for elasticity only when the asset is overpriced. Among these stocks, there is less disagreement whenever it is more costly to take a short position, suggesting that investors who are more pessimistic tend to not take a position on the market when the cost of shorting is high.

This study is related to different strands of the literature. The first connection is to the recently increasing studies of macroeconomic announcements in conjunction with intraday data from financial markets, which allows economists to precisely isolate the impact of the news on the market from other confounding sources. Most of this literature has focused on variables such as investors expectations and asset returns

(Savor and Wilson (2014), Lucca and Moench (2015), Nakamura and Steinsson (2018), Law et al. (2018), and Cieslak and Vissing-Jorgensen (2020)). Our study differs from these papers insofar as it goes beyond the use of asset return and volatility, bringing attention to the volume as another important variable to understand the impact of intraday events.

Another strand of literature is the study of the relationship between asset volume and volatility. This literature started by documenting a positive correlation between volume and volatility and formulating possible hypothesis that could explain how these variables are connected (Tauchen and Pitts (1983) and Andersen (1996)). On the theoretical side, there was the need to depart from representative agent models to generate trade volume. The first attempts assumed different information sets between investors (Kyle (1985) and Kim and Verrecchia (1991)). However, further research showed that sometimes an abnormal change in volume occurred without a corresponding move of similar magnitude in return. To account for this empirical observation, Kandel and Pearson (1995) developed the difference-of-opinion models. In these models, investors not only have different information sets, but also have different interpretations about a public signal/announcement. Recently, as high frequency data has become more available, there has been a series of studies to help distinguish between the two types of models (Bollerslev et al. (2018), and Bollerslev, Li, and Chaves (2020)). With a few additional assumptions, they demonstrate that the volume-volatility elasticity should be bounded between 0 and 1, and it should have a monotonic relationship with disagreement among investors. Using intraday data from the $S\&P500$ index, those recent studies have found evidence supporting models in which investors have different interpretations of the macro announcement,

i.e., volume would move more than the volatility. Given this evidence, researchers also used the volume-volatility elasticity to verify whether shareholder disagreement was reflected in market variables (Li and Schwartz-Ziv (2019)). Our research is most connected with this literature, and we bring another perspective of volume-volatility elasticity; the analysis of a cross-section of stocks combined with high frequency data to identify the systematic components of investor disagreement.

We have organized the rest of this chapter as follows. In Section 3.2, we describe the economic model used to discipline our interpretation of the data. In Section 3.3, we focus on the econometric methodology to properly estimate the volume-volatility elasticity. Section 3.4 provides a discussion of the challenges to obtain the data and a statistical summary of each variable. With the method and data in hand, in Section 3.5 we present and explain the empirical results. Finally, Section 3.6 concludes and provides directions for further research.

3.2 Theoretical motivation

We develop a simple model to help interpret our empirical estimates, and we discuss the literature on possible proxies to capture the disagreement among investors.

3.2.1 Investor Disagreement Model

Consider a 2-period economy in which investors trade two assets: a risk-free asset with payoff 1 and a risky asset with unknown payoff $\tilde{v}$. We assume that investors have the same Constant Absolute Risk Aversion (CARA) utility with risk aversion $\lambda \neq 0$ and they differ according to their prior about $\tilde{v}$ with the proportion of each investor type being α and $(1-\alpha)$. Additionally, we assume their prior is $\tilde{v}_i \sim N(x_i; s_i^{-1})$, where

s_i is the precision each investor has in his own assessment. At time 0 the investors trade based on their priors. Then, they observe a public signal $L = \tilde{v} + \tilde{\varepsilon}$, but they interpret this signal differently. More precisely, each investor considers $\tilde{\varepsilon}$ to come from $N(\mu_i, h^{-1})$. After updating their prior about $\tilde{v}$ they trade again and get the payoff from each asset.

In this setup, the volatility and volume for the risky asset are derived from the price change between time 0 and time 1 and the change in risky asset held by one investor type respectively. The price variation is given by:

$$\Delta P^* = P_1^* - P_0^* = \frac{h(L - P_0^*) - \overline{\mu}}{\overline{s} + h} = \frac{h}{(\overline{s} + h)} \times (L - (P_0^* + \overline{\mu}/h)) \tag{3.1}$$

where $\overline{\mu} = \alpha\mu_1 + (1-\alpha)\mu_2$, $\overline{s} = \alpha s_1 + (1-\alpha)s_2$, ie, they are weighted averages of each parameter in the investors population. The expression for volatility will simply be the absolute value of equation (3.1). To simplify further, from equation (3.1) and considering the decomposition of the public signal into two components, it is clear that the price variation will be proportional to the surprise from the public signal L and the weighted average expected for the signal, ie, $P_0^* + \overline{\mu}/h$.

$$|\Delta P^*| \propto \text{Average Announcement Surprise}$$

For the volume, we just need to consider the changes in risky asset of one investor type. Let $m_{t,i}$ be the fraction held at time t by investor type i,

$$Volume = m_{1,2} - m_{1,1} = |\kappa_0 + \kappa_1 \Delta P^*| \tag{3.2}$$

where $\kappa_0 = \frac{\alpha(1-\alpha)}{\lambda} h(\mu_2 - \mu_1)$ and $\kappa_1 = \frac{\alpha(1-\alpha)}{\lambda}(s_2 - s_1)$.

From equation (3.2), we can conclude the existence of two channels that drive the volume in this economy. The changes in traded volume can be either from the difference in information at time 0, $(s_2 - s_1)$ or from investors having different interpretations of the public signal $(\mu_2 - \mu_1)$.[2]

$$\text{Volume} \propto$$

$$|\text{Announcement Disagreement} + (\text{Difference in Information at time 0}) \times \Delta P^*|$$

Based on this conclusion, if there is no announcement disagreement, the traded volume should move one-to-one with the volatility. As a consequence, in this scenario, the volume-volatility elasticity should not deviate from 1. Moreover, with a few additional assumptions, it is possible to show that the elasticity is bounded by 1. However, we sometimes observe volume moving by more than the absolute value of the stock return. This empirical observation makes it reasonable to allow for disagreement and to test whether we can explain the elasticity deviations from 1 using characteristics that capture investor disagreement.

3.2.2 Stock Characteristics and Investors Disagreement

Given the growing evidence in favor of investors having different priors and our exploration of the cross section (Bollerslev et al. (2018), Bollerslev et al. (2020), and Fedyk (2021)), we sought in the literature the asset characteristics that matter for disagreement. There were essentially two characteristics: how difficult is it to evaluate future payoffs and how risky it is for informed traders to arbitrage the price.

The intuition for the uncertain payoff channel dates to Miller (1977). He intended

[2] We should also consider the role played by investor confidence h, which is multiplicative on the disagreement.

to explain why stocks with more exposure to systemic risk had lower expected returns in the data. Thus, he proposed a static model wherein investors disagree and with limited supply of the stock. Consequently, the price will reflect only the view of the most optimistic investors who will hold the asset driving the price up, and, subsequently, the price should converge to the true value. The effect of uncertain future payoff is to increase the probability of different states in the future. Consequently, if the investors disagree, and we have equally likely states in the future, there should be more differences about what each investor thinks of the asset price. Although it is not plausible to directly apply Miller's model to any asset, the model provided intuition on how to think about stock uncertainty and investor disagreement.

To determine what variables we can use to proxy for the uncertainty of future payoffs, we turn to empirical studies. Among these articles, some connected hard to value stocks with different asset anomalies. Baker and Wurgler (2006) argued that stocks from companies with greater uncertainty tend to be highly affected by investors sentiment. One of their proxies for uncertainty is the high market-to-book ratio. In another paper related to our study, Kumar (2009) asserted that uncertainty about future payoffs is affected by the aggregate uncertainty and the idiosyncratic component. The assets which are more complex to evaluate tend to be influenced by investor bias such as overconfidence. As mentioned in the previous subsection, overconfidence can increase or decrease the impact of any existent disagreement. Kumar used idiosyncratic volatility as a measure for the difficulty in evaluating the stock.

In our study, we will use three measures: idiosyncratic volatility, market capitalization, and growth stocks. If these variables explain disagreement, they should have a negative impact on elasticity because elasticity is a parameter bounded above by 1.

No less important than uncertainty is the arbitrage risk. There is a literature that models (De Long et al. (1990) and Shleifer and Vishny (1997)) and documents (D'avolio (2002) and Stambaugh et al. (2015) possible restrictions for the market to reach the fundamental price of the asset. These papers are more concerned with the ability of informed investors to take short positions in the market when they think the stock is overvalued. In addition, they argue that an arbitrage strategy is not necessarily without risks because some operations involve advanced capital and the price might deviate even further from the true value. Hence, when there are more arbitrage costs, fewer investors are willing to express their opinions in the market. Consequently, it is likely for the researcher to observe less disagreement when looking at those asset transactions.

As possible variables to proxy for arbitrage constraint we have institutional ownership and idiosyncratic volatility. D'avolio (2002) presented institutional ownership to proxy for the cost of lending a stock because institutional traders are the most relevant agents to determine the cost of shorting. Thus, the more institutions that own the stock, the less expensive it is to arbitrage and the more likely it is to observe disagreement. Stambaugh et al. (2015) used idiosyncratic volatility to assess the difference in arbitrage when the price was under or over what it should be. The more idiosyncratic the movement, the more risky it is for investors to take positions in the asset market and the less likely it is to have disagreement.

3.3 Identification and Estimation of Volume-Volatility Elasticity

In this section, we present the econometric methodology to estimate the volume-volatility elasticity using high frequency data. In our approach, we will use different

specifications to explore different dimensions of the data, to verify that the elasticity is identified, and proceed to estimation. For the estimation, we make sure it has desirable properties both asymptotically and in finite samples.

3.3.1 Setup and Notation

Here we present the asymptotic setup and the statistical model in our analysis of volume and return volatility. In this setting, volume and return are sampled from time 0 to T at discrete times $i\Delta_n$, $0 \leq i \leq [T/\Delta_n]$. The asymptotic results are based on the sampling interval going to zero ($\Delta_n \to 0$) while keeping T fixed. Those assumptions are standard for analyzing high frequency data (see, e.g., Aït-Sahalia and Jacod (2014) and Jacod and Protter (2012)) Moreover, asymptotically, each macroeconomic announcement is independent, thereby enabling a different set of stocks for each event. The asymptotic independence of events is helpful because we will track the components of the $S\&P500$ over time.

The trading volume and log price increment over the sampling interval $((i-1)\Delta_n, i\Delta_n]$ are defined, by $V_{i\Delta_n}$ and $r_i = P_{i_n} - P_{(i-1)\Delta_n}$, respectively. We assume the $N \times 1$ log price vector P follows a jump-diffusion process of the form

$$dP_t = \alpha_t dt + c_t dW_t + dJ_t, \tag{3.3}$$

where α_t is the drift vector, W_t is a N-dimensional Brownian motion, $c_t = \sigma_t \sigma_t'$ is the $N \times N$ covariance matrix, and J_t is the $N \times 1$ price jump vector. The stochastic volatility process σ can be interpreted as the instantaneous standard deviation of the diffusive price moves.

Unlike the log price, the observed volume is usually an integer multiple of a given lot size. Given the difference we cannot use the same methodology as the log price.

Hence, we follow Li and Xiu (2016) and we model the $N-$dimension volume process generated by a state-space model on the discrete sampling grid as

$$V_{i\Delta_n} = \mathcal{V}(\xi_{i\Delta_n}, \varepsilon_i), \quad 0 \leq i \leq [T/\Delta_n], \tag{3.4}$$

where ξ is a latent state process, ε is a random shock, and the function $\mathcal{V}(\cdot,\cdot)$ transforms these variables into the observed volume vector. By integrating out the random shock ε_i with respect to its distribution $F_\varepsilon(\cdot)$, we obtain the instantaneous mean process of the trading volume

$$m_{i\Delta_n} \equiv \int \mathcal{V}(\xi_{i\Delta_n}, \varepsilon) F_\varepsilon(d\varepsilon). \tag{3.5}$$

For simplicity, in the remainder of this chapter, we refer to m as the *volume intensity*.

Because the theoretical implications of Kandel and Pearson (1995) concern the incremental variations in return volatility and volume intensity induced by new information we study the "jumps" in spot volatility and volume intensity from observations before and after scheduled macro announcements. For a particular event τ and stock j, the changes are defined as

$$\Delta log(\sigma_{\tau,j}) = log(\sigma_{\tau,j}) - log(\sigma_{\tau-,j}), \quad \Delta log(m_{\tau,j}) = log(m_{\tau,j}) - log(m_{\tau-,j})$$

However, it is not possible to directly observe the jumps in volume and volatility because of the sampling frequency that generates discrete observations. As the "abnormal" changes in the volume intensity and spot volatility from each stock are essential to understand the elasticity, we need to recover them first. We will estimate the changes by using a nonparametric approach with the data before and after each announcement.

To estimate the jumps in volume and volatility, we consider an integer sequence k_n of local windows with $k_n \to \infty$ and $k_n \Delta_n \to 0$. For each announcement τ we associate the integer $i(\tau) = \tau/\Delta_n + 1$ which is the observation within the day that the event occurred. The volume intensity and spot volatility after and before (denoted by + and -) are then estimated according to

$$\hat{m}_{\tau\pm,j} \equiv \frac{1}{k_n} \sum_{k=1}^{k_n} V_{\Delta_n((i(\tau)\pm k),j} \quad \hat{\sigma}_{\tau\pm,j} \equiv \sqrt{\frac{1}{\Delta_n k_n} \sum_{k=1}^{k_n} r^2_{\Delta_n(i(\tau)\pm k),j}} \quad \forall\ j = 1, \ldots, N \tag{3.6}$$

Then, the estimated jumps in log levels are

$$\begin{cases} \widehat{\Delta log(m_{\tau,j})} \equiv log(\hat{m}_{\tau+,j}) - log(\hat{m}_{\tau-,j}), \\ \widehat{\Delta log(\sigma_{\tau,j})} \equiv log(\hat{\sigma}_{\tau+,j}) - log(\hat{\sigma}_{\tau-,j}) \end{cases} \qquad \forall\ j = 1, \ldots, N \tag{3.7}$$

With reasonable assumptions, Bollerslev et al. (2018) showd that the estimate for jump in volume intensity was consistent, and Jacod and Protter (2012) provide the proof for the spot volatility.

Although these jump estimates are consistent, there is evidence of a strong U-shape intraday pattern for both volume intensity (Wood et al. (1985) for early evidence) and spot volatility (Andersen and Bollerslev (1997), Andersen and Bollerslev (1998)). This intraday shape has different consequences depending on the announcement time. If the announcement occurs in the beginning of the day, the jumps are likely to be overestimated; near the closing time, the jumps are likely to be underestimated. Therefore, if we mix different events with different times, the elasticity estimate will be biased.

To avoid the bias in the jump estimates, Bollerslev et al. (2018) proposed a simple DID method, and we use it in our analysis. Particularly, the control group will be the set of observations around the same time as the announcement on the 20 days preceding it. A more formal description follows.

For each announcement time τ, we set a control group $\mathcal{C}(\tau)$ of non-announcement times. For each element of the control group, we estimate the jump in volume and volatility, and later average for the number of elements (N_c). With those averages, we subtract from the initial jump estimation for the announcement. Finally, our adjusted estimates become:

$$\begin{cases} \widetilde{\Delta log(m_{\tau,j})} \equiv \widehat{\Delta log(m_{\tau,j})} - \frac{1}{N_c} \sum_{\tau' \in \mathcal{C}(\tau)} \widehat{\Delta log(m_{\tau',j})}, \\ \widetilde{\Delta log(\sigma_{\tau,j})} \equiv \widehat{\Delta log(\sigma_{\tau,j})} - \frac{1}{N_c} \sum_{\tau' \in \mathcal{C}(\tau)} \widehat{\Delta log(\sigma_{\tau',j})} \end{cases} \quad \forall\ j = 1, \ldots, N \tag{3.8}$$

3.3.2 Identification and Estimation of Elasticity

With a reliable method to recover the changes in volume and volatility around each event, we proceed to define specifications to study the volume-volatility elasticity $\mathcal{E}$ and test for the disagreement proxies from Section 3.2. In all our specifications to estimate the elasticity, we will use the least squares approach. Furthermore, because we are interested in an elasticity, it is reasonable to consider log-log equations between the change in volume intensity and spot volatility as the following

$$\Delta log(m_\tau) = Intercept + \mathcal{E} \Delta log(\sigma_\tau) + \eta_\tau \tag{3.9}$$

where $\Delta log(m_\tau)$ and $\Delta log(\sigma_\tau)$ are vectors containing the changes from each stock, $\mathcal{E}$ is the parameter for the volume volatility elasticity, and η_τ is the $N \times 1$ unobserved

component of this relationship. As for the intercept, based on evidence in the literature, we allow changes depending on the announcement type. The difference between each specification depends on the hypothesis we are testing and the assumptions needed for identification.

Constant Volume volatility elasticity ($\mathcal{E}$)

First, we consider a simple specification to identify and estimate the elasticity ($\mathcal{E}$):

$$\Delta log(m_\tau) = Intercept + \mathcal{E}\Delta log(\sigma_\tau) + \eta_\tau, \tag{3.10}$$

where η_τ is the unobserved $N \times 1$ component of this relationship. As we assume $\mathcal{E}_j = \mathcal{E}\ \forall j = 1, \ldots, N$, for identification of the intercept and slope we require non zero volatility jumps, and $\mathbb{E}(\eta_{\tau,j}) = 0$.

Different elasticity by stock ($\mathcal{E}_j$)

Because Equation (3.10) might be too restrictive, we also allow elasticity to change across stocks, i.e., we consider $\mathcal{E}_j$. We will use covariates to explain this difference in $\mathcal{E}$. The variables used to explain the elasticity for different stocks come from the intuition presented in Section 3.2. Then, the equation for $\mathcal{E}_j$ is

$$\Delta log(m_\tau) = Intercept + \mathcal{E}_j\Delta log(\sigma_\tau) + \eta_\tau = Intercept + (\theta^T X_\tau)\Delta log(\sigma_\tau) + \eta_\tau, \tag{3.11}$$

where θ is a $k \times 1$ vector of parameters and X_τ is a $k \times N$ vector with covariates to explain the differences in elasticity for each stock. For Equation (3.11) to be identified we need $k < N$ and we need to replace the same elasticity condition with a nonsingularity restriction, i.e., $det(X_\tau^T X_\tau) \neq 0$.

Changes in elasticity over time ($\mathcal{E}_\tau$)

In addition the first two specifications, we consider possible changes in the elasticity over time. In contrast to previous work, the new cross-section dimension in our dataset enables us to recover the volume-volatility elasticity for the set of stocks included at each time point, and this estimate might provide a new measure for the dynamics of disagreement and uncertainty.

$$\Delta log(m_\tau) = Intercept + \mathcal{E}_\tau \Delta log(\sigma_\tau) + \eta_\tau, \tag{3.12}$$

To have Equation (3.12) identified, we need again the assumptions required for Equation (3.10).

Section 3.2 sets out a few candidates to explain variations in the volume-volatility elasticity. To properly test the stock characteristics, we parametrize the elasticity as a linear function of the explanatory variables. In addition, in a few designs we parametrize the intercept to capture heterogeneity in attention and possible noise trade across event types. As a result, we have

$$\begin{cases} \mathcal{E} = b_0 + b_1 X_{\tau,j} \\ Intercept = a_0 D_\tau \end{cases} \tag{3.13}$$

where D_τ is a set of dummies for each macro announcement type and $X_{\tau,j}$ is a set of explanatory variables for elasticity.

With these different specifications to estimate the elasticity, our next challenge is to obtain a measure of precision of those estimates. Particularly, because of the nonstandard asymptotic distribution of the estimator, we cannot rely on the usual

standard errors. Hence, Bollerslev, Li, and Chaves (2020) developed a local bootstrap method to account for multiple stocks in the same announcement.

To describe the bootstrap procedure, we present the optimization problem in the most general approach to obtain volume volatility elasticity among our specifications:

$$\hat{\theta}_n = \left(\hat{\mathbf{a}}_n, \hat{\mathbf{b}}_n\right) \equiv \min_{\theta \in \Theta} \sum_{\tau \in \mathcal{T}} \sum_{j \in \mathcal{N}_\tau} \left(\Delta \widetilde{log(m_{\tau,j})} - (a_0 + a_1 D_\tau) - (b_0 + b_1 X_{\tau,j}) \Delta \widetilde{log(\sigma_{\tau,j})}\right)^2 \tag{3.14}$$

where $\mathcal{T}$ denotes the set of announcements in the sample, $\mathcal{N}_\tau$ the set of stocks used for the announcement τ, and the bold letters denote a vector of parameters. Then, the local bootstrap proceeds according to the following algorithm.

Local Bootstrap Algorithm

Step 1: For each $\tau \in \mathcal{T} \cup (\cup_{\tau' \in \mathcal{T}} C(\tau'))$, generate i.i.d. draws $\left(\mathbf{V}^*_{i(\tau)-k}, \mathbf{r}^*_{i(\tau)-k}\right)_{1 \le k \le k_n}$ and $\left(\mathbf{V}^*_{i(\tau)+k}, \mathbf{r}^*_{i(\tau)+k}\right)_{1 \le k \le k_n}$ from $\left(\mathbf{V}_{i(\tau)-k}, \mathbf{r}_{i(\tau)-k}\right)_{1 \le k \le k_n}$ and $\left(\mathbf{V}_{i(\tau)+k}, \mathbf{r}_{i(\tau)+k}\right)_{1 \le k \le k_n}$, respectively.

Step 2: Compute $\Delta \widetilde{log(m^*_{\tau,j})}$ and $\Delta \widetilde{log(\sigma^*_{\tau,j})}$ replacing the original data with the bootstrap version.

Step 3: Estimate $\hat{\theta}^*_n = \left(\hat{\mathbf{a}}^*, \hat{\mathbf{b}}^*\right)$ from minimizing the adjusted objective function

$$\sum_{\tau \in \mathcal{T}} \sum_{j \in \mathcal{N}_\tau} \left(\Delta \widetilde{log(m^*_{\tau,j})} - \hat{u}_\tau - (a_0 + a_1 D_\tau) - (b_0 + b_1 X_{\tau,j}) \Delta \widetilde{log(\sigma^*_{\tau,j})}\right)^2$$

where $\hat{u}_\tau = \Delta \widetilde{log(m_{\tau,j})} - (\hat{a}_0 + \hat{a}_1 D_\tau) - (\hat{b}_0 + \hat{b}_1 X_{\tau,j}) \Delta \widetilde{log(\sigma_{\tau,j})}$

Step 4: Repeat previous steps a large number of times. Report the standard errors from $\left(\hat{\mathbf{a}}^*, \hat{\mathbf{b}}^*\right)$ as an approximation for the standard errors of $\left(\hat{\mathbf{a}}_n, \hat{\mathbf{b}}_n\right)$

□

3.4 Data

In this section, we describe how we collect and clean the data, and provide a description of each series. Our focus of analysis is the set of stocks that comprised the $S\&P500$ index each month from July 2001 until December 2016. We start in July 2001 because the date corresponds to the implementation of decimalization of stock quotes and prices. In addition to the intraday 1-minute observation on price and volume from 9:30am to 4:00pm during regular trading days, we have 670 macro announcements and a few market and accounting characteristics for each stock.

3.4.1 Price and Volume Data for Stocks and the Market Index

The data on price and volume were obtained from the Transactions and Quotes (TAQ) database. This database contains all trades and quotes that took place in every American stock exchange since 1993. For this research, we focus on consolidated trades. To find the information about the stocks and indexes at TAQ we also used the Compustat and Center for Research in Security Prices (CRSP) datasets. To handle this large amount of data, we developed a script in SAS to export the 1-minute observations on price and volume.

Before going into TAQ files, we searched at Compustat for the stocks that comprised the $S\&P500$ for each month in the sample. Compustat uses the Committee on Uniform Security Identification Procedures (cusip) number as the identifier for

each stock. However, this number is not unique, and it changes over time for the same security. A preferable identifier is the Permanent Security Identification Number (permno) because it is invariant. Consequently, with the cusip we can retrieve the permno from CRSP, and, from there, we are ready to look into TAQ daily files[3]. We applied the algorithm outlined below for each weekday in our sample.

TAQ data Algorithm

Step 1: With a list of desired permnos, get the historical cusip from CRSIP which is the cusip associated with the security for that day.

Step 2: In the TAQ Master file, search for securities with the desired historical cusip and get their trade symbols.

Step 3: With the symbols for the desired stocks, search in the TAQ Trades for trades satisfying the guidelines of Brownlees and Gallo (2006) and Barndorff-Nielsen et al. (2009).

Step 4: With the filtered trade observations, consolidate into minute trades, volume, and price (from last trade).

Step 5: Save the table with permno and minute information into a .csv file.

□

After gathering the trade information, we convert the files into Matlab for further cleaning and processing the volume-volatility elasticity analysis. This additional cleaning consists of removing half trading days, constructing the return sequence, and filling in the information on minutes without trades in the data during regular

[3] It comprises a Master file with identifiers for TAQ Trades and TAQ Quotes.

trading minutes.

In the end, we obtain price and volume data for a median of 491 stocks every month[4]. This data represents 712,342,800 observations of approximately 80GB in size.

3.4.2 Idiosyncratic Volatility

We follow the methodology used by Patton and Verardo (2012) to measure the idiosyncratic volatility with high frequency data. Although Patton and Verardo (2012) does not focus on the idiosyncratic component, their study provides a reliable way of measuring at the monthly frequency. In summary, Patton and Verardo (2012) applied the capm model to the high frequency data, and they were attentive to problems that may have arisen from intraday observations. They used 25 minute returns instead of higher frequencies to avoid having the β biased towards zero. For our application, there are a few differences. First, we use a model with six factors to estimate the idiosyncratic volatility. By expanding the number of factors, we expect that the idiosyncratic volatility will provide a reliable estimate of for how difficult it is to understand the asset future cash-flows. The factor model includes the 5 Fama-French factors (market, size, value, investment, profitability) from Fama and French (2015) and the momentum factor (Carhart, 1997). The second difference comes from the return data. We rely on the data obtained from Aït-Sahalia et al. (2020) which takes into account the information from trades and makes intraday portfolios for each factor based on all the stocks traded in NYSE, AMEX, and NASDAQ. A more detailed description of the factor model and how we derive our Ivol measure follows.

[4] We miss some stocks because we consider only one from dual class stocks, and we require the stock to be listed on the index in the preceding month.

For each month t and each asset j, using intraday data we estimate the following equation

$$\begin{aligned} r_{i\Delta_n,j} = \alpha_j + \beta_{mkt,j} r_{smb,i\Delta_n} + \beta_{smb,j} r_{smb,i\Delta_n} + \beta_{hml,j} r_{hml,i\Delta_n} + \cdots \\ \beta_{rmw,j} r_{rmw,i\Delta_n} + \beta_{cma,j} r_{cma,i\Delta_n} + \beta_{mom,j} r_{mom,i\Delta_n} + \varepsilon_{i\Delta_n,j} \end{aligned} \tag{3.15}$$

Once we have estimates for the monthly parameters, we proceed to estimate the monthly idiosyncratic volatility and realized volatility:

$$IdioV_{t,j} = \sum_{i=1}^{n} \widehat{\varepsilon_{i\Delta_n,j}}^2 \quad RV_{t,j} = \sum_{i=1}^{n} r_{i\Delta_n,j}^2$$

Then, we construct the Ivol that we use in our analysis:

$$Ivol_{t,j} = \sqrt{\frac{IdioV_{t,j}}{RV_{t,j}}} \in [0,1] \tag{3.16}$$

We take the ratio because it is bounded and makes this characteristic comparable across different assets. In addition, the ratio provides a straightforward connection between individual stocks and the market return.

In Figure 3.3 we have a detailed evolution of the estimated $\beta_{m,j}$ and Ivol. For the market beta, we have its median fluctuating around 1 which should be the market index, and its distribution might have a uniform dispersion over time. Regarding to Ivol, the median fluctuates around 0.81, i.e., the factor fluctuations explain approximately 19% of the volatility. This high median value, even after taking into account six widely used factors, suggests there are a lot of idiosyncratic movements in the cross section and, by capturing the complexity of each company, they may be useful to explain differences in the elasticity.

FIGURE 3.3: Quantiles for Monthly β_{mkt} and Ivol

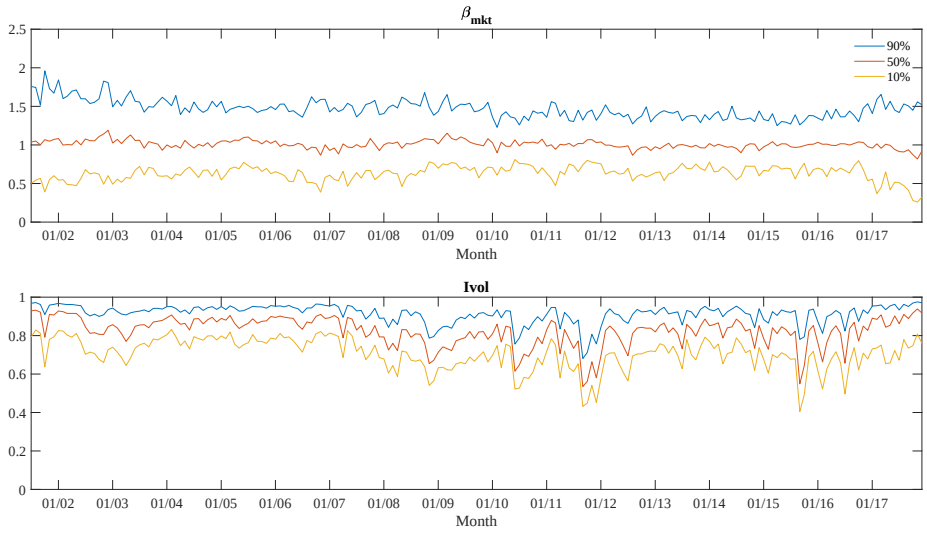

Note: In this figure, we report the specific quantiles from the distribution of the estimated β_{mkt} (upper panel) and Ivol (lower panel) according to Equations 3.15 and 3.16.

3.4.3 Announcement Day and Time

With transaction data at our disposal, the next step was to gather data on the macroeconomic announcements used as news sources in our analysis. The time and day of scheduled announcements is available from Bloomberg. We decided to restrict the analysis to the most important announcements detected by Bloomberg during market hours. These announcements are the FOMC decision on interest rates (FOMC), the ISM Manufacturing (ISMM) and Non-Manufacturing (ISMNM), and the Consumer Confidence (CC). Although the FOMC occurs every 45 days, the other events are monthly with CC being on the last Tuesday, ISMM on the first business day, and ISMNM on the third business day. There were 670 of those events from July 2001 to December 2016. Table 3.1 shows how many each of them happened in the

Table 3.1: Macroeconomic Announcements

	No. Obs.	EST (24hrs)
FOMC	124	14:00
ISM Manufacturing (ISMM)	184	10:00
ISM Non-Manufacturing (ISMNM)	179	10:00
Consumer Confidence (CC)	183	10:00

Note: This table decomposes the events used in the study and describes the time the event occurs during the day. For FOMC, before March 2013 most releases happened at 14:15 and some at 12:30. After March 2013, the time was moved to 14:00.

sample and at what time[5]. In addition, it is important to highlight that FOMC is the only government announcement in the list and the only announcement with empirical evidence of a premium (Lucca and Moench, 2015).

3.4.4 Firm Characteristics

Following studies in the empirical finance literature (e.g. Fama and French (1992) and Baker and Wurgler (2006)), we collect data on financial characteristics for each stock from its annual financial statements. The balance sheet data are available from the Compustat database. Also, following the procedure in the literature and avoiding forward looking, we match the information from fiscal year-ends in "t-1" with the events from July of year "t" until June of year "t+1". For example, the fiscal information that we used for July 2012 through June 2013 came from releases up to June 2012 regarding the 2011 fiscal year for all companies.

The characteristics of our interest are the book-to-market ratio and size. The book equity is defined as total common equity plus balance sheet deferred taxes. Size is obtained from CRSP using price times shares outstanding at June of year t. For these

[5] Although ISMM, ISMNM, and CC are monthly reports, they do not have the same number of occurrence because some of them take place in half-trading days.

variables we winsorise the annual sample at the 0.05% and 99.5% to avoid outliers driving the estimates.

FIGURE 3.4: Quantiles for Stock Characteristics

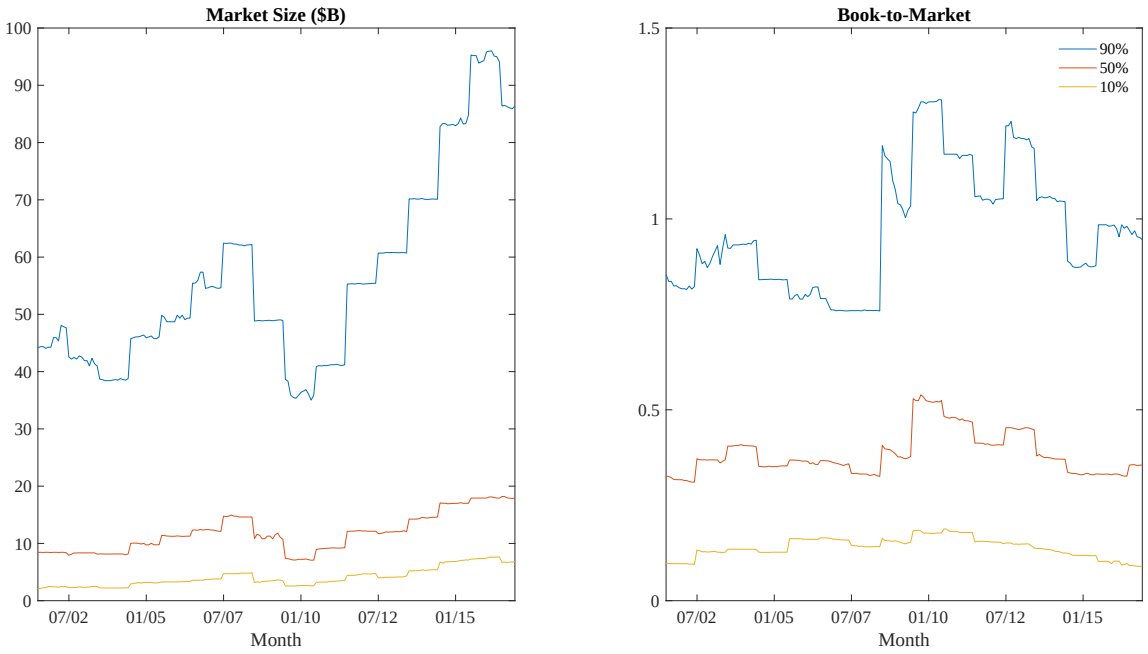

Note: In this figure, we report how specific quantiles of the sampled market size (left panel) and Book-to-Market ratio (right panel) changed over time.

In Figure 3.4 we describe the evolution of these variables. For the market capitalization there is a clear distinction between the largest 10% of stocks and the rest. This gap has only increased after the financial crisis. Moreover, the evolution of each quantile highlights that size is a non stationary variable. Moving to the book-to-market ratio, this variable is stationary as evidenced by a more flat evolution of quantiles and a mean close to 0.5. In addition, not surprisingly, we observe a spike in the ratio during the financial crisis in 2008 when market prices fell.

3.4.5 Institutional Ownership

In addition to characteristics about each stock's company financial performance, we also use a market measure to proxy for the market structure of each stock. This variable is the institutional ownership, i.e., the percent of stocks owned by institutional investors.

The institutional ownership is used widely in the asset pricing literature to proxy for the cost of shorting a stock. D'avolio (2002) was the first to study the market for lending stock and documented the relevance of institutional ownership. This variable is measured quarterly using the 13F form filed by institutional investors as required by the U.S. Securities and Exchange Commisssion (SEC).

Starting in the 1980s, Thompson Reuters has provided a summary of the reports of ownership. However, after June 2013, there are multiple issues in the data provided by Reuters such as omitted institutional reports and excluded securities from coverage (Ben-David et al. (2016) for more details). After detecting these problems, the Wharton Research Data Services (WRDS) decided to collect and make available all the data contained in the 13f reports from the SEC. In addition, they recommend using this dispersed data to complement the Reuters information after June 2013.

The data from WRDS contains both the value of ownership and the number of stocks owned by an institution. To avoid noise from using a different price from the reported value, we focus on shares owned by institutions. For every reported stock we sum the number of shares owned by institutional investors and to make the data comparable across stocks we divide by the shares outstanding provided by CRSP at the end of each quarter.

After cleaning the data after June 2013, we compare them with the information available from Reuters, and we confirm the evidence for omitted data. The difference between Reuters and our data is negative for more than 95% of stocks. On the basis of that difference, we decide to use the institutional ownership data from Reuters up to June 2013 and summarized from 13f reports until December 2016. Figure 3.5 reports a description of the cross-section over time.

FIGURE 3.5: Distribution of Institutional Ownership over Time

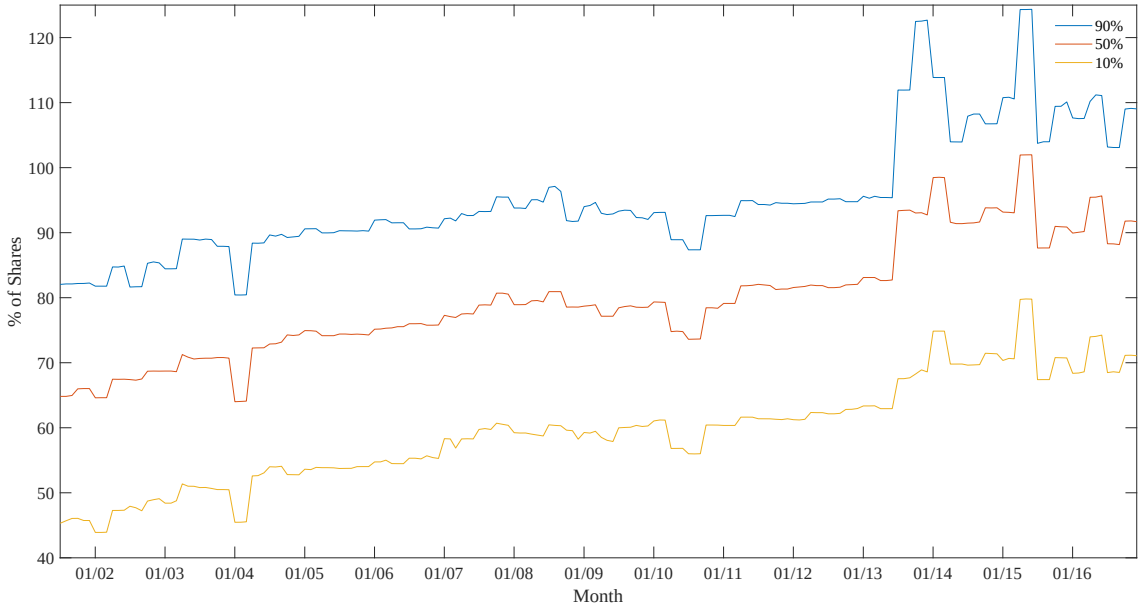

Note: In this figure, we report how certain quantiles of the sampled institutional ownership, measured as the percent of outstanding shares owned by institutional investors, changed over time.

From Figure 3.5, we notice a spike at June 2013 that confirms the WRDS recommendation to use both datasets when constructing institutional ownership series because of possible idleness in updating by Thomas Reuters. This break will not cause problems with our estimation because we will use the cross-section distribu-

tion and not the evolution over time. Therefore, as long as the monthly ranking is preserved in each dataset, we can proceed with our estimation. Another interesting observation from Figure 3.5 is that high values of ownership fluctuate around 100% or more. Observing institutional ownership above 100% occurs because those institutions report with different dates, and sometimes they consider borrowed stock from short operations. Finally, we do see an increase in overall ownership over time since the beginning of the sample.

3.4.6 Merged Dataset

After obtaining all the different data for multiple stocks, we need to merge the data into a useful dataset. We merged the data according to the permno presented in section 3.4.1. However, not all data sources had their stocks identified according to this number. We make the proper conversion to permno using information from CRSP.

There were also some challenges that made the sample window and the number of stocks to be smaller than the price and trade data obtained from TAQ. With respect to the time frame, we did not have information on certain characteristics after December 2016. Regarding the stocks, we removed from analysis those without data on at least one of the cross-sectional explanatory variables.

In the end, the clean and merged data encompasses July 2001 to December 2016, it contains price and volume information for 385 stocks per month, on average, and 670 events (124 FOMC, 184 ISMM, 179 ISMNM, 183 CC). Although 23% of stocks were unavailable for analysis, our sample is still representative of the composition of

$S\&P500$. The most common industries[6] on average are manufacturing (39%), finance and insurance (15%), information (8%), and utilities (7%).

To address any remaining concerns about the data, we present the distribution of return and volume data within a symmetric window of 30 minutes before and after the announcements, excluding the event minute.

Table 3.2: Intraday Return - Before and After Announcements

Industry	Min	1%	10%	25%	50%	75%	90%	99%	Max
Whole Sample	-26.93	-0.44	-0.14	-0.05	0.00	0.05	0.14	0.44	23.64
Manufacturing	-26.93	-0.42	-0.14	-0.05	0.00	0.05	0.14	0.42	14.66
Finance and Insurance	-11.91	-0.48	-0.14	-0.05	0.00	0.05	0.14	0.48	23.64
Information	-14.23	-0.41	-0.13	-0.05	0.00	0.05	0.13	0.41	9.10
Utilities	-22.35	-0.34	-0.11	-0.04	0.00	0.04	0.11	0.34	19.43

Note: The table describes the distribution of 1-minute returns for all the events in the sample using a window of 30min before and after the announcement.

Table 3.3: Intraday Volume (100 shares) - Before and After Announcements

Industry	Min	1%	10%	25%	50%	75%	90%	99%	Max
Whole Sample	0.0	0.0	4.0	15.0	45.0	124.0	313.8	1730.0	143629.7
Manufacturing	0.0	0.0	4.0	14.0	42.8	125.9	344.0	1958.0	110749.6
Finance and Insurance	0.0	0.0	5.0	17.0	49.1	138.1	367.0	2621.9	128373.9
Information	0.0	0.0	6.0	20.0	62.0	194.6	567.6	2649.6	74815.9
Utilities	0.0	0.0	3.0	11.0	28.1	63.3	123.6	410.7	12271.0

Note: The table describes the distribution of observed 1-minute volume for all the events in the sample using a window of 30min before and after every announcement.

As expected, the returns in Table 3.2 are symmetric with few extreme values. Conversely, the volume in Table 3.3 and the turnover ratio in Table 3.4 are left-skewed with some minutes lacking observations from less liquid stocks in the sample.

[6] The industry classification is based on North American Industry Classification System (NAICS).

Table 3.4: Turnover Ratio (%) - Before and After Announcements

Industry	Min	1%	10%	25%	50%	75%	90%	99%	Max
Whole Sample	0.0	0.0	0.1	0.2	0.5	1.2	2.6	10.1	2038.9
Manufacturing	0.0	0.0	0.1	0.2	0.5	1.2	2.7	10.7	855.3
Finance and Insurance	0.0	0.0	0.1	0.2	0.5	1.1	2.3	8.9	437.6
Information	0.0	0.0	0.1	0.3	0.6	1.2	2.6	12.1	640.8
Utilities	0.0	0.0	0.1	0.2	0.4	0.8	1.5	5.2	340.8

Note: The table describes the distribution of observed 1-minute turnover (normalized for day frequency) for all the events in the sample using a window of 30min before and after every announcement.

Nevertheless, out of the total minutes of observations we have for each stock around announcements, only 3.6% of them do not have any trades.

3.5 Empirical Analysis of Volume-Volatility Elasticity

After describing our econometric procedure to estimate the elasticity and obtain a measure for its precision, we are ready to evaluate the results. As explained earlier, we try to determine whether investors disagree and, if they do, how disagreement is represented in the cross section. We will answer these questions by examining the changes in the volume-volatility elasticity estimate.

First, we describe our choice for the parameters from Section 3.5. We set $k_n = 30$ minutes to estimate the volume intensity and spot volatility because some of the announcements in Table 3.1 happened at 10am and we did not want to exclude them. To adjust for intraday patterns, the control group is the set of non-announcement days starting from the previous month until the event date corresponding, on average, to 20 trading days. For the bootstrap standard deviation we set the number of replications at 1,000.

3.5.1 Disagreement in the Cross Section

With these parameters set, we verify the theoretical predictions from Section 3.2. Our initial approach was to generate portfolios by deciles for each stock characteristic. By constructing these portfolios we expect to have a simple characterization of how the elasticity changes as each explanatory variable changes. We generated both the volume and return for each portfolio by using an equal weight of the actual traded volume and return for each stock.

First, we certify that the announcements used in the analysis have non-zero jump for both log volume and log volatility after accounting for the observed intraday pattern. Figures 3.6 and 3.7 report the estimates for volume and volatility respectively sorted by deciles. The average jump size for log volume (volatility) is 0.166 (0.318) with bootstrap standard error 0.004 (0.009). Those estimates confirm the relevance of events to understand how investors process the announcements in the cross-section, and the jump estimates confirm that the elasticity is identified in this setup.

Now, we proceed to use the announcement data to estimate each elasticity. In addition, given the measurement error in the spot volatility estimates, we use the bootstrap approach from Bollerslev, Li, and Chaves (2020) to adjust for the attenuation bias. The specification used to recover the volume-volatility elasticity is like equation 3.10 and we rewrite here:

$$\Delta log(m_\tau) = a_0 + b_0 \Delta log(\sigma_\tau) + \eta_\tau \tag{3.17}$$

In Figure 3.8, we report the unbiased elasticity estimates (b_0) for each portfolio formed by deciles from the smallest to largest, and, in Table 3.5 we report the average

FIGURE 3.6: Log Volume Jump (DID) - All Events

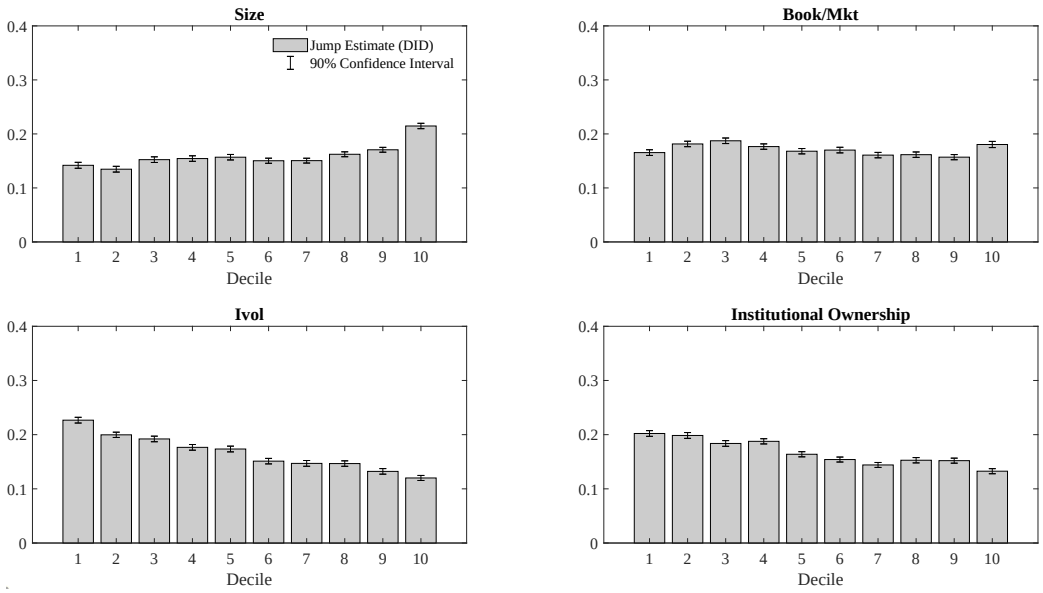

Note: In this figure, we report the estimated log volume jump with intraday adjustment for the equally weighted portfolios based on characteristic deciles. For each estimate, we also describe its 90% *confidence interval based on the bootstrap standard error and by using all the events in the sample.*

R^2 from each regression. For all the characteristics used in Figure 3.8, we observe the expected difference in elasticity between extreme quantiles. The largest difference and stronger monotonicity comes from portfolios based on Ivol. These portfolios have elasticity that starts at 0.58 for the lowest decile and declines until 0.36 at the highest.

Finally, in Table 3.5, we report the average R^2 from each characteristic, which fluctuates around 50%. The R^2 values suggest that the proposed characteristics provide a reasonable explanation for the volume-volatility elasticity and the explanatory power is in line with what other researchers have observed when studying the volume-volatility elasticity in diversified portfolios as the market index.

To confirm this result, we also double sort the stocks to compute their elasticity. In

FIGURE 3.7: Log Volatility Jump (DID) - All Events

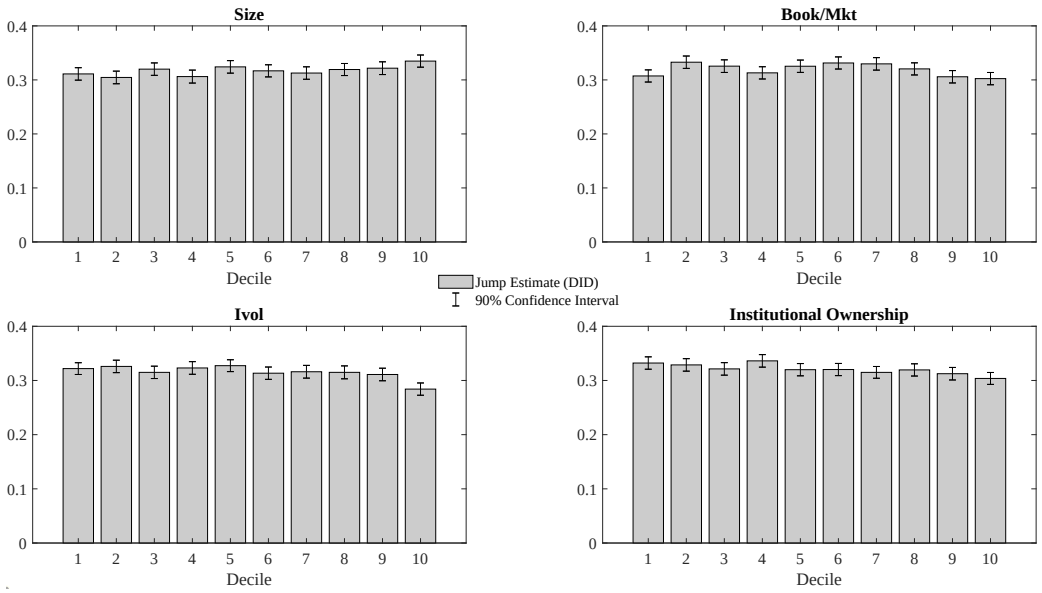

Note: In this figure, we report the estimated log volatility jump with intraday adjustment for the equally weighted portfolios based on characteristic deciles. For each estimate, we also describe its 90% *confidence interval based on the bootstrap standard error and by using all the events in the sample.*

Figure 3.9, we fix Ivol quartiles, and, in each subplot, we consider the quartiles from a different characteristic. The estimates confirm the estimates sorted by a unique characteristic. For high quartiles of Ivol and size, the estimated elasticity is low, whereas for institutional ownership, we observe the opposite effect on elasticity.

Table 3.5: Average R^2 From Elasticity Estimates by Characteristic

	Size	Book/Mkt	Ivol	Inst. Onwership
$\overline{R^2}$	48.03%	50.33%	49.18%	50.15%

Note: This table reports the average R^2 from elasticity estimates of portfolios based on stocks sorted into deciles according to each characteristic.

Although the effect of each characteristic is in accordance with the theory, none of

FIGURE 3.8: Volume-Volatility Elasticity for Sorted Portfolios

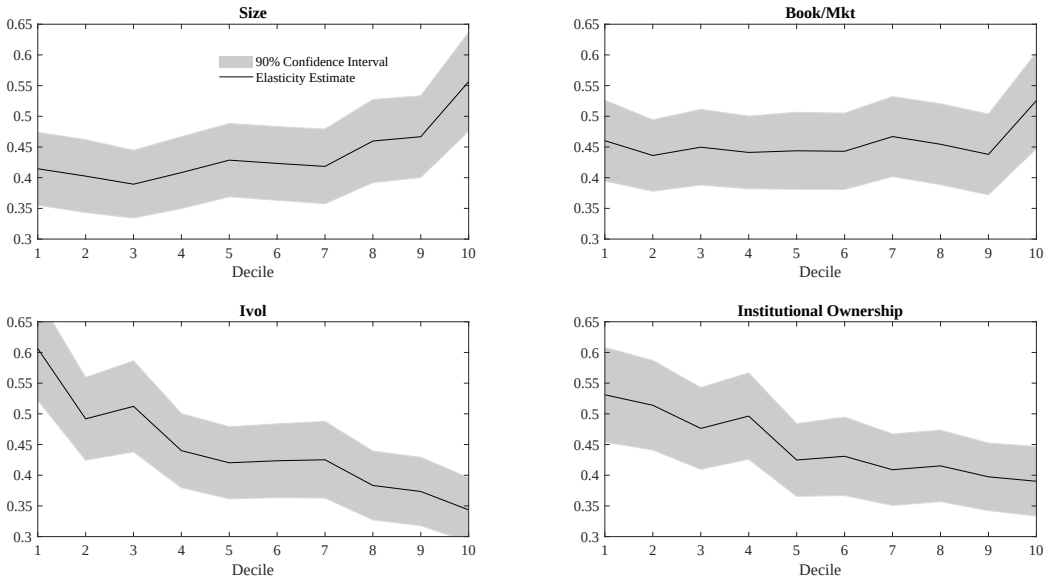

Note: In this figure, we report the estimated unbiased elasticity using specification $\widetilde{\Delta log(m_\tau)} = a + b\widetilde{\Delta log(\sigma_{\tau,j})}$ *with its* 90% *confidence interval for portfolios constructed by sorting the stocks into deciles for each characteristic for each month in the sample from Jul.2001 until Dec.2016. Each plot starts at the bottom decile and goes to the highest.*

the portfolios had an elasticity close to one. This inability of a single variable to explain the deviation from the case with no disagreement might be evidence of how complex it is to evaluate a stock. Thus, our next step to understand the volume-volatility relationship is to assess the information from each stock at each announcement in a panel estimation.

3.5.2 Exploring Panel to Estimate $\mathcal{E}$

In addition to using the jumps in volume and volatility at the individual stock, we now allow for heterogeneity in elasticity across event types. There are at least two reasons for different reactions by event, e.g., the news cycle in which some announce-

FIGURE 3.9: Volume-Volatility Elasticity Sorted for Given Ivol Quartile

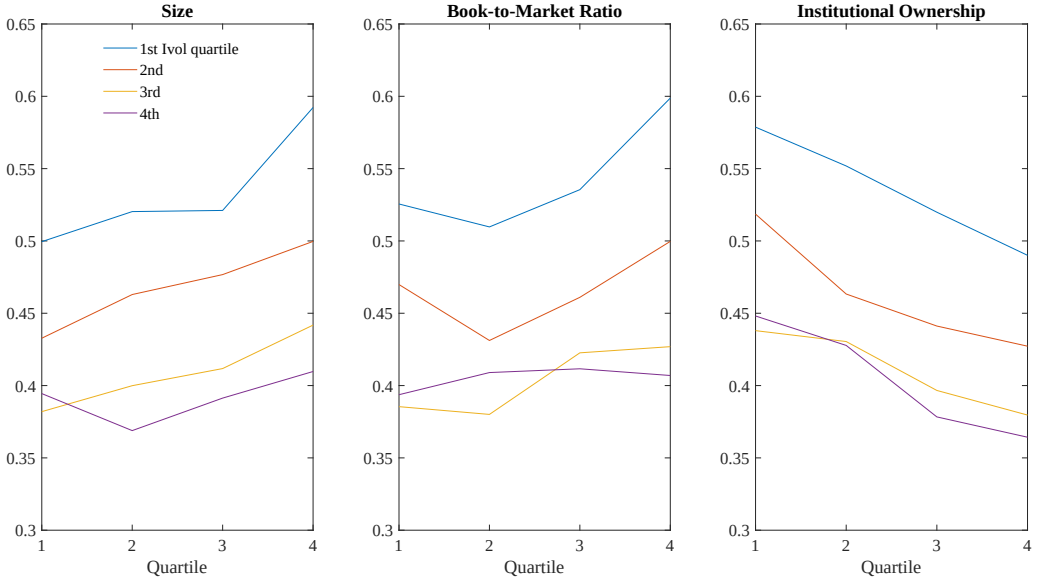

Note: In this figure, we report the estimated unbiased elasticity using specification $\widetilde{\Delta log(m_\tau)} = a + b\widetilde{\Delta log(\sigma_{\tau,j})}$ *for portfolios constructed by double sorting the stocks into quartiles for each characteristic for each month in the sample from Jul.2001 until Dec.2016. Each plot starts at the bottom quartile and goes to the highest.*

ments happen before others, and the possibility of investors devoting different levels of attention for each event type.

Hence, for each stock j and public announcement τ, we estimate according to the following specification

$$\Delta log(m_{\tau,j}) = a_0 D_\tau + (b_0 D_\tau + b_1 X_{\tau,j})\Delta log(\sigma_{\tau,j}) + \eta_{\tau,j} \tag{3.18}$$

where $\eta_{\tau,j}$ is the residual with zero conditional mean, D_τ is the matrix containing the dummies to control for heterogeneity across events, and $X_{\tau,j}$ is the set of individual stock variables used to explain the elasticity.

We consider as explanatory variables different stock characteristics. As docu-

mented in Section 3.4, some of those variables are not stationary. To circumvent this issue and make them comparable not only in the cross-section but also across different announcements, we employ a normalization. The only variable that does not need the transformation is Ivol which, by construction, is bounded and comparable. For the other variables, whenever the theory predicts a negative relationship between elasticity and that a given variable (e.g. institutional ownership), we perform the following transformation for every month t and stock j:

$$X_{t,j} = \frac{\tilde{X}_{t,j}}{\max_{j \in \mathcal{N}} \{\tilde{X}_{t,j}\}} \in [0,1] \tag{3.19}$$

When the theory suggests a positive relationship with elasticity (e.g. size and book to market), we transform the variable as:

$$X_{t,j} = \left| \frac{\tilde{X}_{t,j}}{\max_{j \in \mathcal{N}} \{\tilde{X}_{t,j}\}} - 1 \right| \in [0,1] \tag{3.20}$$

In Table 3.6, we describe the estimates from the unbalanced panel and consider multiple explanatory variables for the volume-volatility elasticity specified in equation (3.18). As explained before, the variables range from 0 to 1, and all but the Ivol represent the deviation from the expected maximum elasticity for that variable in the cross-section.

The variables do have the sign as expected from Section 3.2. The estimates from 3.6 also confirms our initial elasticity estimates from Figure 3.8 where we used artificial portfolios instead of each stock. If we consider a stock at the median values of each characteristic, according to the estimates its elasticity at FOMC will be $0.83 - 0.22 \times$

$(0.82) - 0.06 \times (0.95) - 0.10 \times (0.80) - 0.05 \times (0.62) = 0.48$, which is lower than the estimates in the literature for the $S\&P500$ index. This lower estimate is evidence of additional disagreement in the cross section that we can capture using this set of four explanatory variables.

Notice that Ivol is the variable with the most influence on elasticity, changing from -0.252 to 0 in the estimated elasticity. Possibly Ivol is the best variable to capture the disagreement among investors for each specific stock because it can capture multiple dimensions of a company. Also, the elasticity generated when setting Ivol to zero is similar to the elasticity obtained by Bollerslev et al. (2018) when they used the $S\&P500$ index. The larger market elasticity is intuitive because private information on individual stock should not be meaningful for a diversified portfolio, and it shows that we might be able to proxy for the market elasticity using the cross-section of stocks.

When including multiple explanatory variables, we find that the institutional ownership has a weaker effect, and the effect might be due to its asymmetric relevance for short and long positions. The asymmetric relevance of institutional ownership for mispriced stocks will be explored in Section 3.5.3. Despite this weaker effect, it is clear from Table 3.6 that, by combining these variables, we can explain much of the elasticity deviation from unity. It is interesting that, to capture the disagreement in the cross section, we need to use different dimensions of a stock instead of a single explanatory variable which was the case for the $S\&P500$ index.

To further confirm the results from Table 3.6, we illustrate the implications for elasticity using a couple of examples. In Figure 3.10, we provide two distinct stock comparisons. In the plot at the left, we observe the elasticity for General Electric

(GE) and Newmont Mining (NEM) over time. The estimated elasticity for GE is predominantly larger than the estimate for NEM. The difference can be attributed to GE being a conglomerate while NEM is a mining company that might be more exposed to the business cycle. For the plot on the right, we compare the elasticity for Amazon (AMZN) and Walmart (WMT). It is clear that, over time, the elasticity for AMZN has catch up with WMT and this may be a result from AMZN establishing itself as a major company in online shopping and, consequently, being less volatile as WMT is in its grocery sector.

Table 3.6: Volume-Volatility Elasticity Using Individual Stocks

Baseline Elasticity by Event (b_0):						
CC	**0.298**	**0.505**	**0.379**	**0.378**	**0.347**	**0.639**
	(0.006)	(0.026)	(0.009)	(0.009)	(0.012)	(0.026)
ISMM	**0.267**	**0.474**	**0.348**	**0.348**	**0.316**	**0.608**
	(0.006)	(0.026)	(0.010)	(0.010)	(0.012)	(0.026)
FOMC	**0.496**	**0.696**	**0.576**	**0.576**	**0.544**	**0.827**
	(0.013)	(0.028)	(0.015)	(0.016)	(0.017)	(0.029)
ISMNM	**0.282**	**0.490**	**0.363**	**0.363**	**0.331**	**0.624**
	(0.006)	(0.026)	(0.009)	(0.009)	(0.012)	(0.025)
Explanatory Variables (b_1):						
Ivol		**-0.252**				**-0.222**
		(0.030)				(0.031)
Size (0=Large)			**-0.090**			**-0.059**
			(0.007)			(0.009)
Book/Mkt (0=Value)				**-0.107**		**-0.102**
				(0.010)		(0.010)
Inst.Ownership					**-0.082**	**-0.046**
					(0.017)	(0.018)

Note: This table reports the parameter estimates with its respective bootstrap standard deviation based on the regression $\Delta\widetilde{log(m_{\tau,j})} = aD_\tau + (b_0 D_\tau + b_1 X_{\tau,j})\Delta\widetilde{log(\sigma_{\tau,j})}$. *The estimates are obtained with approximately 385 stocks per month and announcements from Jul.2001 until Dec.2016. The coefficients in bold are significant at* 1% *level.*

FIGURE 3.10: Examples of Elasticity from Panel Estimates

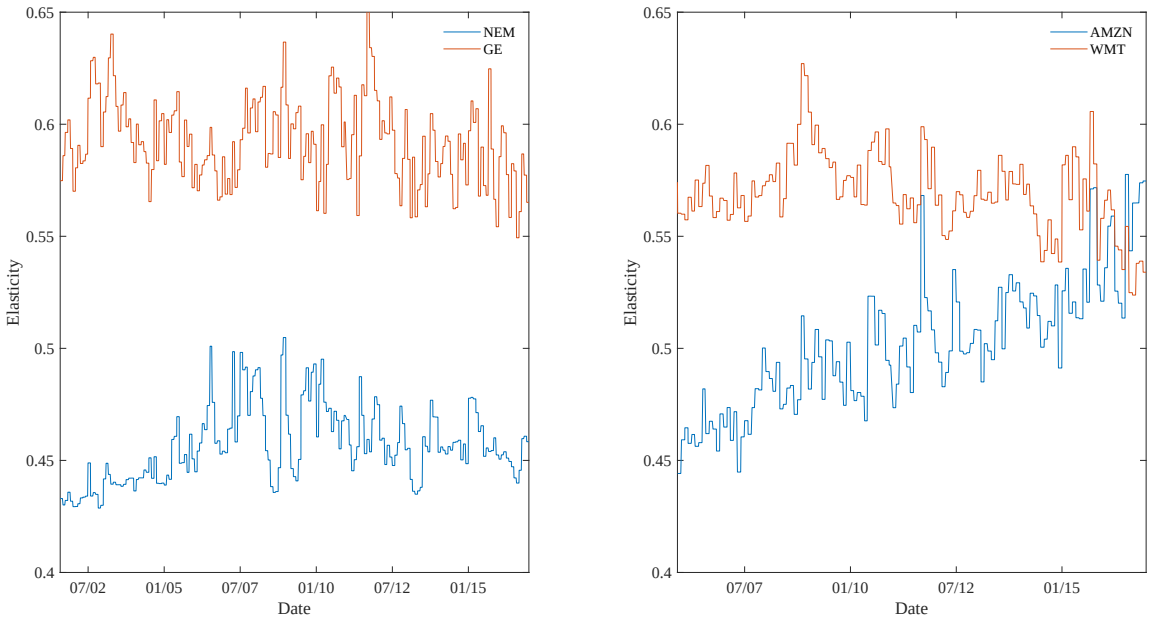

Note: In this figure, we report the implied elasticities for a few selected stocks using the estimates from Table 3.6. In other words, for each selected stock we plot $\hat{b}_{0,FOMC} + \hat{b}_{1,Ivol} \times Ivol_{\tau,j} + \hat{b}_{1,Size} \times Size_{\tau,j} + \hat{b}_{1,BM} \times BM_{\tau,j} + \hat{b}_{1,IO} \times IO_{\tau,j}$*. In panel on the left, we compare Newmon Mining (NEM) and General Electric (GE). In the panel on the right, we compare Amazon (AMZN) and Walmart (WMT).*

Robustness and Time Variation in $\mathcal{E}$

To confirm our panel estimates for elasticity, we explore subsets of the data. Particularly, we assess the time series dimension by using a rolling window of 12 months. We use 12 months because, although there are approximately 380 assets in the cross-section, the estimates are still noisy when using a smaller window. Hence, let $m \in \mathcal{M}$ be a month in the sample. For each month m, we consider τ_m as the set of events in that month. Denote rw as the month rolling window to estimate $\mathcal{E}_t$

$$\Delta log(\widetilde{m_{\tau_{m,m+rw},j}}) = aD_{\tau_{m,m+rw}} + (b_0 D_{\tau_{m,m+rw}} + b_1 X_{\tau_{m,m+rw},j}) \Delta log(\widetilde{\sigma_{\tau_{m,m+rw},j}}) \quad (3.21)$$

Figure 3.11 reports the estimates from using $rw = 11$ and the same explanatory variables used in Table 3.6 for the panel estimation. The estimates for each subsample confirm the signs observed in the panel. The explanatory variables are never statistically greater than zero and most of the time we have statistically negative estimates.

In addition, we can study how the elasticity changes over time. Particularly, we can evaluate the effect of changes in overall disagreement from the changes in $\hat{b}_0$. This coefficient is the elasticity remainder after controlling for the cross-section changes. Figure 3.12 plots the parameter estimates. First, it is interesting that, even tough we did not impose any constraint on the elasticity estimate, its values are all bounded between 0 and 1. Additionally, we observe sharp declines around the Enron scandal in 2002, the financial crisis in 2008, the European debt crisis in 2011, and the US government shutdown in 2013. Those periods are associated with a low investor confidence which, in turn, should imply a positive effect on elasticity because investors have less trust in their own interpretation of announcements resulting in less disagreement. To confirm this hypothesis, the correlation between the baseline elasticity and the composite investor sentiment measure from Baker and Wurgler (2006) is -0.4. We do not observe a correlation closer to -1 because at the same time that investor confidence is low, the aggregate uncertainty might be high and, consequently, offsets some of the impact coming investor sentiment in disagreement.

To further address the use of volume-volatility elasticity to understand how investors disagree when interpreting new information, we evaluate a different design. This design will highlight the asymmetric role of institutional ownership for a mispriced stock.

FIGURE 3.11: Elasticity Parameters Over Time

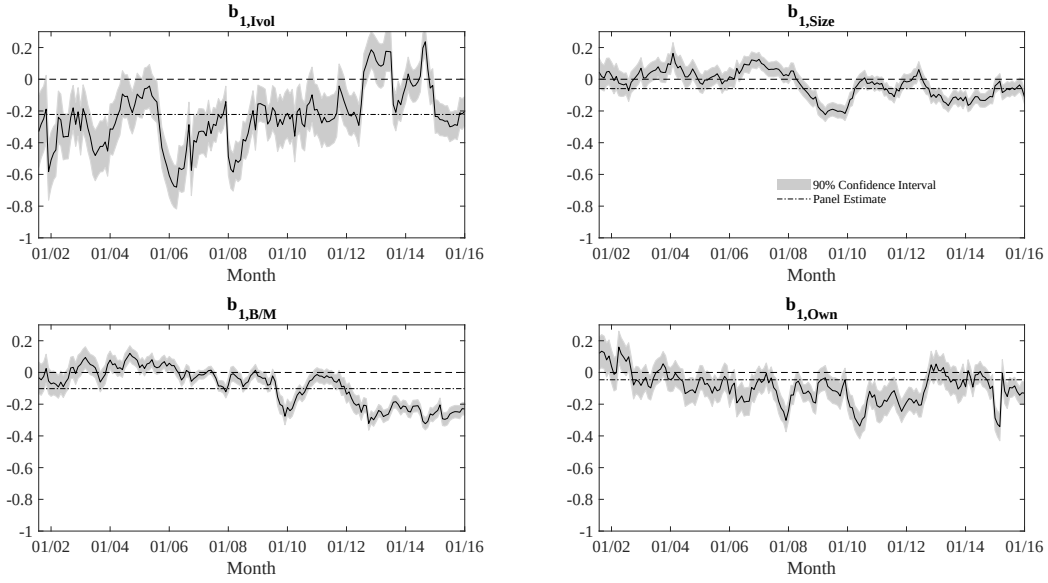

Note: In this figure, we report the estimated elasticity coefficients for the explanatory variables from regression $\Delta log(\widetilde{m_{\tau_{m,m+rw},j}}) = aD_{\tau_{m,m+rw}} + (b_0 D_{\tau_{m,m+rw}} + b_1 X_{\tau_{m,m+rw},j})\Delta log(\widetilde{\sigma_{\tau_{m,m+rw},j}})$ *with its* 90% *confidence interval. We use a monthly rolling window consisting of 11 previous and the current month to get estimates.*

3.5.3 Misprice and Disagreement

In this section, we focus on the make-up of the pool of investors for a particular stock. For example, consider two types of investor: noise and informed. They disagree on how to evaluate the stock and the informed agent faces the risk that, in the short term, he could have losses with the price moving on the opposite direction. Consequently, we should observe a different composition of investors depending on the transaction cost for each stock.

Moreover, based on the arbitrage literature, a stock will be mispriced when there are elevated costs for informed investors to take the appropriate position on the market. Usually, there are more costs to take a short position than long because the

FIGURE 3.12: Estimated Baseline Elasticity (FOMC)

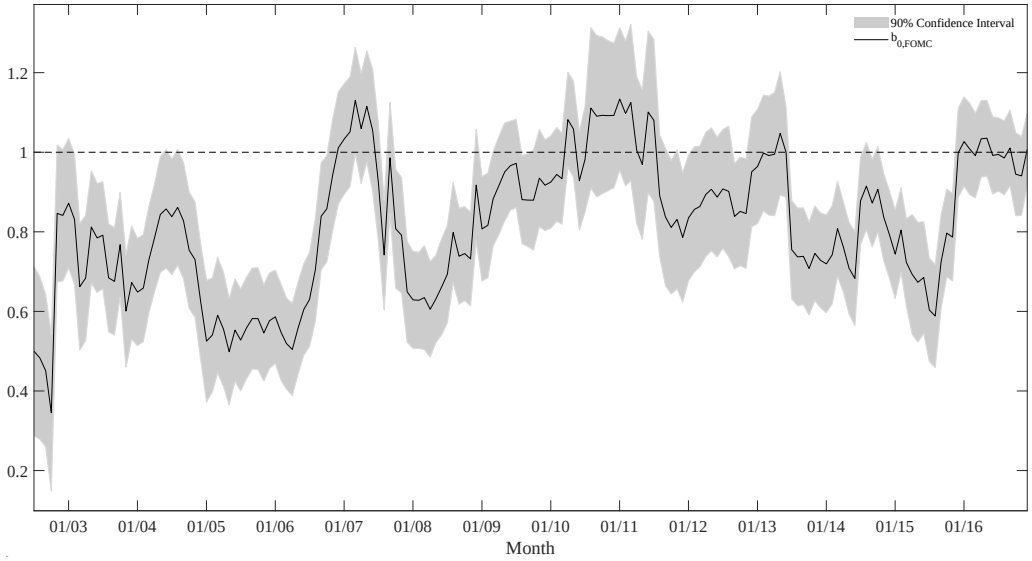

Note: In this figure, we report the estimated baseline elasticity (b_0) *associated with the FOMC dummy from the regression* $\Delta log(\widetilde{m_{\tau_{m,m+rw},j}}) = aD_{\tau_{m,m+rw}} + (b_0 D_{\tau_{m,m+rw}} + b_1 X_{\tau_{m,m+rw},j})\Delta log(\widetilde{\sigma_{\tau_{m,m+rw},j}})$ *and its* 90% *confidence interval. We use a monthly rolling window consisting of the 11 previous months and the current month to get each estimate.*

investor will have to pay a fee to borrow the stock. Hence, it seems plausible that mostly noise traders might own overpriced stocks, consequently, creating a more homogeneous pool of investors which leads to less observed disagreement. Conversely, underpriced stocks should have a more diverse composition between noise and informed with greater disagreement reflected in the elasticity.

Therefore, to properly verify this claim on *observed* disagreement, we need a measure for over/ underpriced stocks. Stambaugh et al. (2015) constructed an interesting measure to quantify a stock misprice. They examined the literature for possible anomalies that survive after controlling for the Fama-French factor model and found 11 different misprice indicators. Then, for each of these anomalies, they built a rank

with the stocks in which the higher the ranking, the lower the possible abnormal future returns because the stock is overpriced today. After ranking for each anomaly, they associated for a stock its respective quantile in each category and the misprice score was the average across the quantiles.

This misprice score was obtained at Robert Stambaugh's website[7]. It is available monthly from July 1965 to December 2016, and the number of stocks in each month is not fixed.

For our study, we are interested only in the direction of misprice. So, we simply define:

$$\begin{cases} \text{Underpriced: } Misp_i < 30 \\ \text{Overpriced: } Misp_i > 70 \end{cases}$$

The underprice cutoff is based on evidence from Stambaugh et al. (2015) and the overprice was chosen to preserve symmetry. Figure 3.13 plots the number of stocks in each misprice category in the sample over time.

For our analysis, we use the misprice score without normalization because it is already a bounded measure based on a cross section ranking.

With the misprice direction defined, we expect to capture the asymmetric dynamic of disagreement from the estimation of $\mathcal{E}$. Particularly, the effect of each misprice direction on the volume-volatility elasticity can be captured using a new combination of explanatory variables. We still have the institutional ownership measured as small to high percent and the misprice score. What will be new is the interaction between the misprice direction and the indicator functions using the median value for each event to signal high/ low ownership.

[7] http://finance.wharton.upenn.edu/ stambaug/

FIGURE 3.13: Stock Misprice in Sample

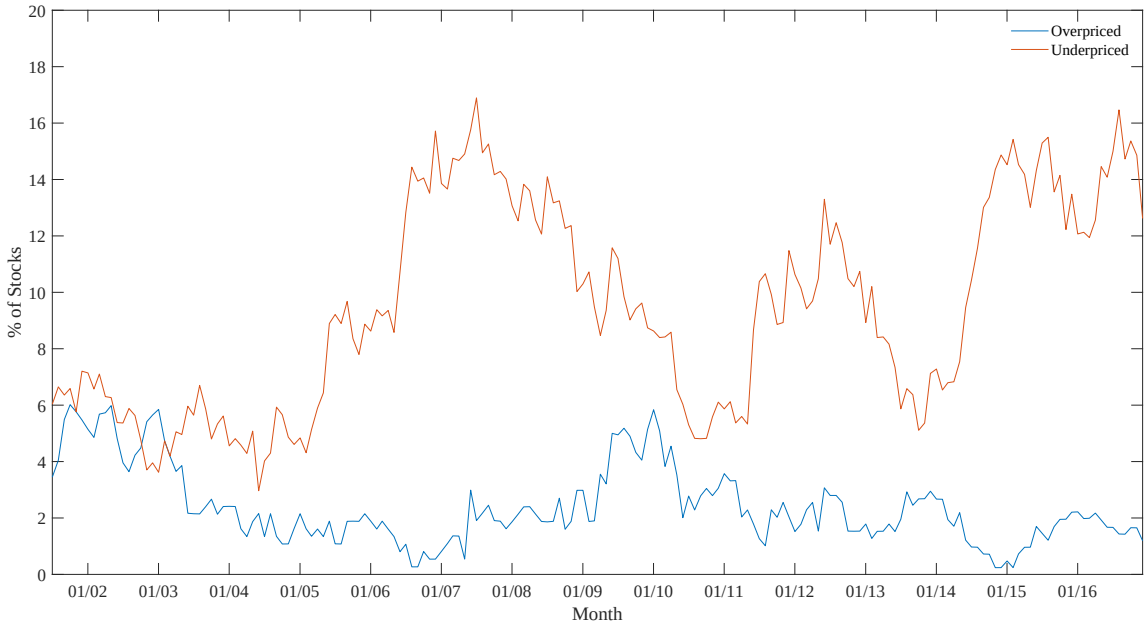

Note: In this figure, we report the changes over time in the percent of stocks within our sample that were classified as overpriced (blue line) or underpriced (red line).

To illustrate, we present the specification used to study whether institutional ownership has asymmetric impact on overpriced stocks:

$$\widetilde{\Delta log(m_{\tau,j})} = aD_\tau + \Big(b_0 D_\tau + b_{1,1} O_{\tau,j} + b_{1,2} 1_{[M_{\tau,j}>70 \,\cap\, O_{\tau,j}<q_{0.5}(O_{\tau,j})]} M_{\tau,j} O_{\tau,j} + \cdots$$
$$b_{1,3} 1_{[M_{\tau,j}>70 \,\cap\, O_{\tau,j}>q_{0.5}(O_{\tau,j})]} M_{\tau,j} O_{\tau,j} \Big) \widetilde{\Delta log(\sigma_{\tau,j})} \tag{3.22}$$

where the estimates of $b_{1,2}$ and $b_{1,3}$ should provide information about the impact of institutional ownership on the elasticity of overpriced stocks. It consists of a simple

comparison,

$$\text{If } b_{1,2} = b_{1,3} \implies \text{no asymmetry}$$

$$\text{If } b_{1,2} = b_{1,3} \implies \text{asymmetry}$$

In Table 3.7, we present the estimates for equation (3.22) using the misprice score interacted with both over and underprice dummies. We observe a different effect on elasticity depending on the misprice direction. If the stock is overpriced, $b_{1,2} \neq b_{1,3}$, and they have different signs, meaning that the cost of a short position keeps informed investors outside the market. If the stock is underpriced, $b_{1,2} = b_{1,3}$, which means that institutional ownership might not be relevant for investors to decide to enter the market. This interpretation of the coefficients is in line with the literature because, for underpriced stocks, an informed investor should take long position, and we are using a measure for the cost of a short position.

As a result, an overpriced stock should have less informed investors when it has low ownership than when it has high ownership. For underpriced stocks, there should not have be distinction on the composition between informed and noise traders based on institutional ownership.

3.6 Conclusion

In this chapter we explored how the volume-volatility elasticity can help us understand the disagreement among investors in the cross-section of stocks. To estimate the volume-volatility elasticity, we used high-frequency econometric tools to deliver a reliable estimate both in finite samples and asymptotically. The elasticity estimates showed that the most important components that drive disagreement are the idiosyncratic volatility, market size, book-to-market ratio, and the institutional ownership.

Table 3.7: Volume-Volatility Elasticity and Misprice

Baseline Elasticity by Event (b_0):					
CC	**0.347**	**0.298**	**0.347**	**0.344**	**0.344**
	(0.012)	(0.006)	(0.012)	(0.012)	(0.012)
ISMM	**0.316**	**0.268**	**0.316**	**0.313**	**0.313**
	(0.012)	(0.006)	(0.012)	(0.012)	(0.012)
FOMC	**0.544**	**0.497**	**0.544**	**0.541**	**0.541**
	(0.016)	(0.013)	(0.016)	(0.016)	(0.017)
ISMNM	**0.331**	**0.283**	**0.332**	**0.329**	**0.329**
	(0.012)	(0.005)	(0.012)	(0.012)	(0.012)
Explanatory Variables (b_1):					
Institutional Ownership ($b_{1,1}$)	**-0.082**		**-0.080**	**-0.078**	**-0.075**
	(0.017)		(0.017)	(0.017)	(0.018)
Overpriced ($b_{1,1}$)		0.023			
		(0.013)			
Underpriced ($b_{1,1}$)		**-0.054**			
		(0.018)			
Overpriced x Low Ownership ($b_{1,2}$)				**0.163**	**0.160**
				(0.010)	(0.034)
Overpriced x High Ownership ($b_{1,3}$)				**-0.058**	-0.061
				(0.014)	(0.029)
Underpriced x Low Ownership ($b_{1,2}$)			**-0.092**		-0.086
			(0.032)		(0.045)
Underpriced x High Ownership ($b_{1,3}$)			**-0.099**		**-0.099**
			(0.029)		(0.039)

Note: This table reports estimates and bootstrap standard deviation using the specification from Equation (3.22). The estimates are based on approximately 385 stocks per month and announcements from Jul.2001 until Dec.2016. The coefficients in bold are significant at 1% *level.*

Although the first three characteristics capture how complex it is to process information about a specific stock, the last tells us more about the composition of investors on the market for a particular stock.

This result on the disagreement components might lead to further investigation. The characteristics we identified are the most standard factors in the asset pricing

literature, and disagreement is sometimes connected with stock misprice (overvalued) and asset bubbles. Our results, we hope, improve our understanding of how information is processed in financial markets. Particularly, by using the cross section of stocks, we are able to capture systematic components that almost entirely explain the existing disagreement among investors. As hypothesized in the literature, disagreement among agents can drive a stock overprice (Miller (1977) and Hong and Stein (2007)), and our finding on disagreement should lead to further asset pricing investigations.

Another interesting result combines the institutional ownership with misprice direction. Although the institutional ownership proxies for the cost of a short position on a stock, the misprice indicates what position an informed trader should take on the stock. Sometimes, this arbitrage cost can deter informed traders from taking a position on the market. What we find confirms this view because there is an asymmetric pattern for the volume-volatility elasticity when the stock is overpriced and we compare high and low ownership.

4

Conclusion

In this dissertation we expanded the knowledge of how the information from macroeconomic announcements is processed by market participants. After the information is released, investors almost immediately rebalance their portfolios to reflect their updated priors. regarding the actual value of the asset. These reactions can be observed through the jumps in volume and volatility and that is what we explore to investigate whether investors disagree when interpreting the news from macro announcements.

In Chapter 2, we have laid out the necessary econometric theory to estimate the relation between jumps in different stochastic processes using a general loss function. In addition, we have provided a new method to overcome a few econometric challenges that have arisen in this setting. First, we do not observe the continuous path of the variables and we need to nonparametrically estimate the jumps which adds more uncertainty causing a finite sample bias in the relation of interest. To overcome it

we propose a bootstrap method where the objective function is slightly different to generate an estimate of the bias which we use to correct the sample version. The second challenge arises from the nonstandard asymptotic distribution which makes the usual inference procedures unreliable. The solution we propose is to consider a local bootstrap procedure to make inference.

Then, in Chapter 3, we apply the econometric tools from the previous chapter to provide an empirical analysis of investor disagreement in the stocks comprising the $S\&P500$. There we provide three different contributions. First, we propose a model that captures the systemic disagreement in the cross section of stocks. It comprises 4 explanatory variables which are: the idiosyncratic volatility, the market capitalization, the book-to-market ratio, and the institutional ownership. These characteristics are associated with how complex it is to understand the implications of the event for the future cash-flows of the asset and how diverse is the pool of investors in the market of the particular stock. With this model for the cross section of stocks we provide a market-based time series measure of how investor disagreement and investor confidence varies over time. It does have a reasonable interpretation with abrupt falls during times of high uncertainty. For last, we explore the stock misprice direction interacted with its degree of arbitrage constraint to document that not all disagreement is observed through market variables because a few investors may opt to stay outside the market whenever there is a high arbitrage constraint.

As extensions, it would be interesting to make connections with the asset pricing literature by actually matching the estimated investor disagreement with expected returns in the cross section of stocks. Also, the time series measure for investor disagreement might be able to explain some of the conditional returns on the market.

Bibliography

Aït-Sahalia, Y., J. Fan, R. J. A. Laeven, C. D. Wang, and X. Yang (2017). Estimation of continuous and discontinuous leverage effects. *Journal of the American Statistical Association 112*(520), 1744–1758.

Aït-Sahalia, Y. and J. Jacod (2014). *High-Frequency Financial Econometrics*. Princeton University Press.

Aït-Sahalia, Y., I. Kalnina, and D. Xiu (2020). High-frequency factor models and regressions. *Journal of Econometrics*.

Alexeev, V., M. Dungey, and W. Yao (2017). Time-varying continuous and jump betas: The role of firm characteristics and periods of stress. *Journal of Empirical Finance 40*(C), 1–19.

Andersen, T. G. (1996). Return volatility and trading volume: An information flow interpretation of stochastic volatility. *Journal of Finance 51*(1), 169–204.

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

Andersen, T. G. and T. Bollerslev (1998). Deutsche mark–dollar volatility: intraday activity patterns, macroeconomic announcements, and longer run dependencies. *the Journal of Finance 53*(1), 219–265.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and C. Vega (2003). Micro effects of macro announcements: Real-time price discovery in foreign exchange. *American Economic Review 93*(1), 38–62.

Andersen, T. G., T. Bollerslev, F. X. Diebold, and C. Vega (2007). Real-time price discovery in stock, bond and foreign exchange markets. *Journal of International Economics 73*(1), 251–277.

Angelis, D. D., P. Hall, and G. A. Young (1993). Analytical and bootstrap approximations to estimator distributions in L^1 regression. *Journal of the American Statistical Association 88*(424), 1310–1316.

Baker, M. and J. Wurgler (2006). Investor sentiment and the cross-section of stock returns. *The journal of Finance 61*(4), 1645–1680.

Baker, S. R., N. Bloom, and S. J. Davis (2016). Measuring economic policy uncertainty. *The Quarterly Journal of Economics 131*(4), 1593–1636.

Barndorff-Nielsen, O. E., P. R. Hansen, A. Lunde, and N. Shephard (2009). Realized kernels in practice: Trades and quotes. *The Econometrics Journal 12*(3), C1–C32.

Ben-David, I., F. Franzoni, R. Moussawi, and J. Sedunov (2016). The granular nature

of large institutional investors. Technical report, National Bureau of Economic Research.

Bibinger, M., C. J. Neely, and L. Winkelmann (2017, April). Estimation of the discontinuous leverage effect: Evidence from the NASDAQ order book. Working Papers 2017-12, Federal Reserve Bank of St. Louis.

Bibinger, M. and L. Winkelmann (2015). Econometrics of co-jumps in high-frequency data with noise. *Journal of Econometrics 184*(2), 361 – 378.

Bollerslev, T., J. Li, and L. S. S. Chaves (2020). Generalized jump regressions for local moments. *Journal of Business & Economic Statistics*, 1–11.

Bollerslev, T., J. Li, and Y. Xue (2018). Volume, volatility and public news announcements. *Review of Economic Studies 85*(4), 2005–2041.

Brownlees, C. T. and G. M. Gallo (2006). Financial econometric analysis at ultra-high frequency: Data handling concerns. *Computational Statistics & Data Analysis 51*(4), 2232–2245.

Carhart, M. M. (1997). On persistence in mutual fund performance. *The Journal of finance 52*(1), 57–82.

Carlstein, E. (1986). The use of subseries values for estimating the variance of a general statistic from a stationary sequence. *Annals of Statistics 4*(3), 1171–1179.

Cieslak, A. and A. Vissing-Jorgensen (2020). The economics of the fed put. Technical report, National Bureau of Economic Research.

Clark, P. K. (1973). A subordinated stochastic process model with finite variance for speculative prices. *Econometrica 41*(1), 135–155.

Davison, A. C. and D. V. Hinkley (1997). *Bootstrap Methods and their Application*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press.

D'avolio, G. (2002). The market for borrowing stock. *Journal of financial economics 66*(2-3), 271–306.

De Long, J. B., A. Shleifer, L. H. Summers, and R. J. Waldmann (1990). Noise trader risk in financial markets. *Journal of political Economy 98*(4), 703–738.

Dávila, E. and C. Parlatore (2019, March). Trading costs and informational efficiency. Working Paper 25662, National Bureau of Economic Research.

Efron, B. and R. Tibshirani (1994). *An Introduction to the Bootstrap*. Chapman and Hall/CRC Monographs on Statistics and Applied Probability. Taylor & Francis.

Fama, E. F. and K. R. French (1992, 6). The cross-section of expected stock returns. *The Journal of Finance 47*(2), 427–465.

Fama, E. F. and K. R. French (2015). A five-factor asset pricing model. *Journal of*

Financial Economics 116(1), 1 – 22.

Fedyk, A. (2021, 03). Disagreement after News: Gradual Information Diffusion or Differences of Opinion? *The Review of Asset Pricing Studies*. raab008.

Gonçalves, S. and N. Meddahi (2009). Bootstrapping realized volatility. *Econometrica 77*, 283–306.

Hahn, J. (1995). Bootstrapping quantile regression estimators. *Econometric Theory 11*(1), 105–121.

Hall, P. (1997). *The Bootstrap and Edgeworth Expansion*. Springer Series in Statistics. Springer New York.

Hong, H. and J. C. Stein (2007). Disagreement and the stock market. *Journal of Economic perspectives 21*(2), 109–128.

Horowitz, J. L. (2001). The bootstrap. In *Handbook of Econometrics*, Volume 5. Elsevier.

Huber, P. J. (1967). The behavior of maximum likelihood estimates under nonstandard conditions. In *Proceedings of the Fifth Berkeley Symposium on Mathematical Statistics and Probaility, 1967*, Volume 1. University of California Press.

Huber, P. J. and E. M. Ronchetti (2009). *Robust Statistics*. John Wiley & Sons Inc.

Jacod, J. and P. Protter (2012). *Discretization of Processes*. Springer Verlag.

Jacod, J. and V. Todorov (2010, 08). Do price and volatility jump together? *The Annals of Applied Probability 20*(4), 1425–1469.

Kalina, I. and D. Xiu (2017). Nonparametric estimation of the leverage effect: A trade-off between robustness and efficiency. *Journal of the American Statistical Association 112*(517), 384–396.

Kandel, E. and N. D. Pearson (1995). Differential interpretation of public signals and trade in speculative markets. *Journal of Political Economy 103*(4), 831–872.

Kim, O. and R. E. Verrecchia (1991). Trading volume and price reactions to public announcements. *Journal of accounting research*, 302–321.

Knight, K. (1989). Limit theory for autoregressive-parameter estimates in an infinite-variance random walk. *Canadian Journal of Statistics 17*(3), 261–278.

Knight, K. (1998). Limiting distributions for L_1 regression estimators under general conditions. *Annals of Statistics 26*(2), 755–770.

Koenker, R. (2005). *Quantile Regression*. Econometric Society Monographs. Cambridge University Press.

Koenker, R. and G. Bassett (1978). Regression quantiles. *Econometrica 46*(1), 33–50.

Koenker, R. and G. Bassett (1982). Robust tests for heteroscedasticity based on regression quantiles. *Econometrica 50*(1), 43–61.

Kumar, A. (2009). Hard-to-value stocks, behavioral biases, and informed trading. *Journal of Financial and Quantitative Analysis 44*(6), 1375–1401.

Kunsch, H. R. (1989). The jackknife and the bootstrap for general stationary observations. *Annals of Statistics 17*(3), 1217–1241.

Kyle, A. S. (1985). Continuous auctions and insider trading. *Econometrica 53*(6), 1315–1335.

Law, T.-H., D. Song, and A. Yaron (2018). Fearing the fed: How wall street reads main street. Technical report, The Wharton School, University of Pennsylvania.

Li, J., V. Todorov, and G. Tauchen (2017a). Jump regressions. *Econometrica 85*(1), 173–195.

Li, J., V. Todorov, and G. Tauchen (2017b). Robust jump regressions. *Journal of the American Statistical Association 112*(517), 332–341.

Li, J. and D. Xiu (2016). Generalized method of integrated moments for high-frequency data. *Econometrica 84*(4), 1613–1633.

Li, S. Z. and M. Schwartz-Ziv (2019). How are shareholder votes and trades related? Working paper, ECGI.

Lucca, D. O. and E. Moench (2015). The pre-FOMC announcement drift. *Journal of Finance 70*, pp. 329–371.

Mancini, C. (2001). Disentangling the jumps of the diffusion in a geometric jumping Brownian motion. *Giornale dell'Istituto Italiano degli Attuari LXIV*, 19–47.

Miller, E. M. (1977). Risk, uncertainty, and divergence of opinion. *The Journal of finance 32*(4), 1151–1168.

Nakamura, E. and J. Steinsson (2018). High-frequency identification of monetary non-neutrality: the information effect. *The Quarterly Journal of Economics 133*(3), 1283–1330.

Patton, A. J. and A. Timmermann (2010). Why do forecasters disagree? lessons from the term structure of cross-sectional dispersion. *Journal of Monetary Economics 57*(7), 803–820.

Patton, A. J. and M. Verardo (2012). Does beta move with news? firm-specific information flows and learning about profitability. *The Review of Financial Studies 25*(9), 2789–2839.

Peng, L. and W. Xiong (2006). Investor attention, overconfidence and category learning. *Journal of Financial Economics 80*(3), 563–602.

Pollard, D. (1985). New ways to prove central limit theorems. *Econometric Theory 1*(3), 295–313.

Savor, P. and M. Wilson (2014). Asset pricing: A tale of two days. *Journal of Financial Economics 113*, pp. 171–201.

Shleifer, A. and R. W. Vishny (1997). The limits of arbitrage. *The Journal of Finance 52*(1), 35–55.

Stambaugh, R. F., J. Yu, and Y. Yuan (2015). Arbitrage asymmetry and the idiosyncratic volatility puzzle. *The Journal of Finance 70*(5), 1903–1948.

Tauchen, G. and M. Pitts (1983). The price variability-volume relationship on speculative markets. *Econometrica 51*(2), 485–505.

Wang, C. D. and P. Mykland (2015). The estimation of leverage effect with high-frequency data. *Journal of the American Statistical Association 109*(505), 197–215.

Wood, R. A., T. H. McInish, and J. K. Ord (1985). An investigation of transactions data for NYSE stocks. *Journal of Finance 40*(3), 723–739.

Wu, C. F. J. (1986, 12). Jackknife, bootstrap and other resampling methods in regression analysis. *The Annals of Statistics 14*(4), 1261–1295.

Zhang, L., P. A. Mykland, and Y. Ait-Sahalia (2005). A tale of two time scales: Determining integrated volatility with noisy high-frequency data. *Journal of the American Statistical Association 100*(472), 1394–1411.

Biography

Leonardo Salim Saker Chaves is a Brazilian economist. In 2012, he concluded his Bachelor degree in Economics from the Federal University of Rio de Janeiro (UFRJ, Brazil) with a project studying the foundation of the Brazilian Central Bank (BCB). In 2013, he started his Masters in Economics at Getulio Vargas Foundation (EPGE - FGV/RJ) where he pursued research in econometric theory focused on hypothesis testing under weak identification. In 2015, he was accepted into the PhD program of Economics at Duke and, in 2016, he started working as research assistant in a financial econometrics project to develop new tools to understand how investors process macroeconomic events. This work later developed into his first publication in 2020 titled "Generalized Jump Regressions for Local Moments" in the Journal of Business and Economic Statistics (JBES). Also, this relation developed with the Professors was so fruitful that they later advised him in this dissertation aimed to improve the understanding of how macroeconomic information is processed by investors. His main interests are econometric theory, finance, and macroeconomics.