Real-time Price Discovery via Verbal Communication: Method and Application to Fedspeak

Roberto Gómez-Cram Marco Grotteria*

November 2, 2020

Abstract

We advance the hypothesis and establish empirically that investors' expectations adjust slowly to Central Banks' messages. From the videos of post-FOMC-meeting press conferences, we extract the words, and timestamp them at the millisecond. We align the transcripts with high-frequency data for several financial assets to provide granular evidence on the investors' expectations formation process. When the Chairman discusses the changes between current and previous policy statement, price volatility and trading volume spike dramatically, and prices move in the same direction as they did around the statement release. Our approach allows us to quantify in monetary terms the value of information rigidity.

*We thank Jules van Binsbergen, Anna Cieslak, João Cocco, Kent Daniel, James Dow, Itamar Drechsler, Francisco Gomes, João Gomes, Christopher Hennessy, Ralph Koijen, Matteo Leombroni, Emanuel Moench, Michael McMahon, Stefan Nagel, Stephen Schaepher, Jessica Wachter, and Amir Yaron for helpful discussions and/or comments. We thank the AQR Asset Management Institute for generous financial support. Yutong Hu, Oksana Smirnova, and Tianhao Yin provided excellent research assistantship. The authors are both assistant professors of finance at the London Business School, and can be reached at rgomezcram@london.edu and mgrotteria@london.edu.

1 Introduction

When policy analysis ignores the role of expectations, and their evolution, policy conclusions are misleading (Lucas, 1972, 1973, 1976). Because of that, central bankers today give high priority to communication with financial markets in an attempt of managing the public's expectations. It is often argued that, to be effective, communication has to be clear and credible, which has historically led to a difficult trade-off between clarity and time-consistency (Kydland and Prescott, 1977; Calvo, 1978; Barro and Gordon, 1983a, 1983b; Cukierman and Meltzer, 1986; Stein, 1989). Yet, little is known on how investors actually form expectations in response to Central Bank communication. Answering such a question requires: a) tracing out investors' reaction to each specific message to avoid the confounding effect of multiple messages; b) recognizing news to the investors' information set; c) a clean identification approach.

The contribution of this paper is empirical. For the days in which the Federal Open Market Committee (FOMC) has a scheduled meeting, we scrape the videos of the Fed Chairman's post-meeting press conference. We convert the audio into interpretable text and timestamp it at the second level. We align high-frequency data for a wide range of financial assets with the exact words pronounced in each moment. This allows us to assess to which specific message the market responds. We find that the largest asset price changes occur during the minutes in which the Chairman clarifies the information added to the newly issued policy statement. In those minutes trading volume increases significantly, and asset prices move on average in the same direction as they did around the policy statement release. This generates at a coarser frequency a strong positive correlation between price changes around the statement release and the subsequent press conference.

This paper advances the hypothesis and establishes empirically that investors' expectations gradually adapt to Central Banks' messages, providing granular evidence in support of a modern macroeconomic literature on information rigidity (Lucas, 1972; Cukierman and Wachtel, 1979; Woodford, 2003; Coibion and Gorodnichenko, 2012, 2015; Bordalo, Gennaioli, Ma, and Shleifer, 2020). Our setting is natural to study how beliefs of market participants adapt to Central Banks' messages because minute-level shocks to prices are nearly ideal measures of unexpected movements in investors' expectations (Cochrane and Piazzesi, 2002; Nakamura and Steinsson, 2018).

The example of the FOMC meeting on July 31 2019 provides the intuition of our main result. Three related signals lie behind the movements in asset prices between 14:00, and 15:30 that day. First, before the conclusion of the FOMC meeting, markets expected a reduction of a quarter point in the target federal funds rate, with some possibility for a half-point cut, averaging to 35 bps. The actual rate cut was 25 bps, i.e. less than what markets expected, and market prices adjusted accordingly. Second, while investors still expected future easing, the statement included a new sentence adding uncertainty: "the Committee contemplates the future path of the target range for the federal funds rate." Third, when the press conference Q&A started, Powell was assaulted by questions on the meaning of this change in the statement. He answered "we're contemplating the future path of the target range for the federal funds rate. [...] The Committee is really thinking of this [current change] as a mid-cycle adjustment to policy." The "*mid-cycle adjustment* comment signaled there was no plan for a series of rate cuts."[1]

Figure 1 plots the intraday evolution of the interest rate implied from the 12-month Federal funds futures, the price level of SPY, i.e. the exchange-traded fund that tracks the performance of the S&P 500, and the EUR/USD exchange rate on July 31, 2019. Every reference to the new sentence in the post-meeting statement induced some investors to trade, and market prices to adjust, on average, in the same direction as they did when the statement was released. The gradual price changes and trading support the view of information rigidities in which investors update their beliefs too slowly (Hong and Stein, 1999).[2]

To measure the effect of words onto financial asset prices, we must grapple with two methodological challenges. First, we need to convert the post-meeting press conference

[1]Andrew Cinko, editor of U.S. Markets, Princeton, is the author of this sentence, which has been reported by Bloomberg on its live blog of FOMC events.

[2]Huberman and Regev (2001) provide clean evidence that financial asset prices move even on no new news. Specifically, they document the case of EntreMed's stock price which rose from $12 at the Friday close, May 1, 1998, to an open price of $85 on Monday and a close price near $52. This happened because on Sunday, May 3, 1998, the New York Times reported a recent breakthrough in cancer research which directly affected EntreMed. However, this potential breakthrough had been reported in the journal Nature, and in various popular newspapers already in November 1997. Even for public news, they show that newswatchers, á la Hong and Stein (1999), learn from different sources, in different times.

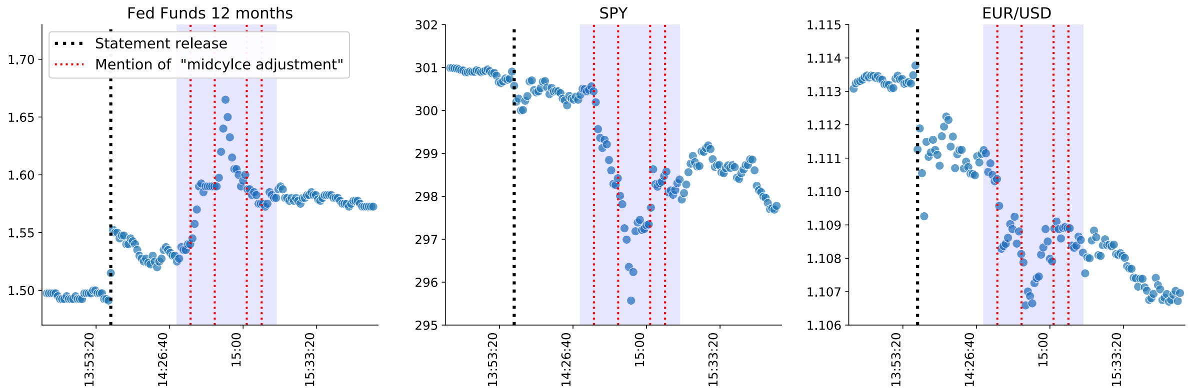

Fig. 1. *Notes*: The Figure shows the intraday evolution of the implied rate from the 12-month Federal funds futures, the SPY price level and the EUR/USD exchange rate on July 31, 2019. The black dashed vertical line highlights the time in which the FOMC statement was released (14:00). The shaded area denotes the FOMC press conference. The conference started at 14:30 and lasted for about 45 minutes. The red dotted lines highlight the time in which the Chairman mentioned "midcycle adjustment to policy."

audio into interpretable text, and time-stamp the words. We split the audio into smaller frames of around 3 seconds, which we then convert into readable text using an end-to-end deep learning algorithm for probabilistic character modeling (Hannun, Case, Casper, Catanzaro, Diamos, Elsen, Prenger, Satheesh, Sengupta, Coates, and Ng, 2014). Second, we identify statement news by tracking the words changed (added/removed) between two consecutive FOMC policy statements. We use automated text analysis to capture those sentences in the press conference text.

We find that changes in the statement are closely scrutinized for insights into what they imply for future policy rates. In the first few questions, financial reporters ask for a clarification of the statement changes and for more detail about the context of the current decision, while Fed officials in part anticipate the confusion caused by the statement.[3,4] When they talk about the statement changes (henceforth statement-related minutes) the average absolute variation in financial asset prices is larger than in the rest

[3]A key goal of post-FOMC press conferences is to clarify the decision and the related changes in the statement: "If we don't hold a press conference [...] there's a decent chance that market participants will be quite confused" (Jeffrey M. Lacker, President of the Federal Reserve Bank of Richmond, during the FOMC meeting of July 2013).

[4]The information communicated with the press conference has to be new at least for some investors to generate movements in asset prices. In fact, the press conference helps clarify the underlying motivation for the policy decision, and thereby provides news to holders of assets. The best analogy is probably with teaching: the instructor repeats the same concepts a few times, with (slightly) different words, giving more context, and trying to understand the doubts and questions from students so that eventually the concept is clear to everybody.

of the conference; trading volume goes up significantly; more importantly, prices move on average in the same direction as their initial reaction around the statement release.

Our findings are stronger for longer-maturity interest rate derivatives, stocks, and exchange rates, which highlights the link between expectations formation and forward guidance. Consistent with this hypothesis, we identify the different language patterns, and styles that characterize the minutes in which the Chairman discusses the statement. We show that those sentences tend to discuss the long-term future, and adopt a more clarifying language as defined by Pennebaker, Boyd, Jordan, and Blackburn (2015).

Price movements in statement-related minutes, once grouped together, generate the novel positive autocorrelation in price changes that we document in this paper. During FOMC days, for interest rate derivatives, stocks, and exchange rates, the response of prices to news in a 30-minute window when the FOMC announces its policy is followed by a response of the same sign and similar magnitude when the Fed Chairman explains the policy decision in a press conference. The correlation between the two price changes is large and positive, e.g. 58% for 60-month Eurodollar futures, and 44% for the S&P 500 index. This a predictive relation between changes in asset prices over two non-overlapping periods. The relation is strong and stable enough that a simple strategy that trades upon this empirical pattern gets highly profitable.

Our results on the positive autocorrelation of price changes have implications for theories of investors' learning and expectations formation. In Section 5, we evaluate several frameworks, and characterize them on the basis of whether or not they can be reconciled with our findings. Our results provide direct evidence against models in which traders are endowed with full information rational expectations (FIRE). This is important because almost every Central bank today uses FIRE-based models to guide monetary policy (Coibion, Gorodnichenko, and Kamdar, 2018). Market prices are forward looking and should already incorporate all information available to the public. So, especially at such a high frequency, they should be close to unpredictable. In addition, our findings present a puzzle to frictionless models of rational economic agents with Bayesian updating, i.e. bounded rationality or standard learning frameworks with parameter uncertainty (Lewellen and Shanken, 2002). For these models to explain our findings, it would require either implausible assumptions on investors' priors or a counterfactual positive price drift, coming from a decline in estimation risk. We also show that our results are inconsistent with theories of the *Fed put*, microstructure effects or liquidity, as well as with the idea that the positive autocorrelation of price changes is a continuation of the Lucca and Moench (2015)'s pre-announcement drift.

The models most directly consistent with our results are models of information rigidity. Throughout the manuscript, we show that a framework in which traders interpret a common signal with some trader-specific noise, i.e. observe a noisy private signal á la Lucas (1972), and Woodford (2003), naturally delivers the results in this paper. Conditional on their private signals, agents update their forecasts rationally. However, they do not know whether news to private signals reflects innovations to the monetary policy or just uninformative noise. For this reason, they decide to adjust their beliefs only slowly in response to shocks to fundamentals. Our estimates suggest that financial investors update their forecasts assigning a weight between 60% and 80% to the new signals and the rest to their previous forecasts. Given the key role of these events in the economic calendar, some could argue that these estimates represent a lower bound to the degree of information rigidity observable when studying financial markets' response to news. Moreover, our analysis does not only provide direct evidence of information rigidity, but also quantifies, via the trading strategy, in easily interpretable monetary terms the economic significance of departures from the null of full information.

Our manuscript makes a methodological contribution to the economics literature by combining video analysis with time-stamped high-frequency financial asset prices. The approach we develop contributes to an extant literature that uses textual analysis methods in different fields of economics (see for instance Tetlock, 2007; Gentzkow and Shapiro, 2010; Loughran and McDonald, 2011; Born, Ehrmann, and Fratzscher, 2014; Hansen and McMahon, 2016; Hansen, McMahon, and Prat, 2017; Gentzkow, Kelly, and Taddy, 2019a; Gentzkow, Shapiro, and Taddy, 2019b; Hassan, Hollander, van Lent, and Tahoun, 2019). Unlike these works, we are able to look at the exact moment in which each word has been pronounced. This avoids jointly gathering multiple updates together, and for our purposes it eases the understanding of which news/word the market is responding to. Our approach does not only improve on the identification of the effect of words on financial investors' beliefs, but also extends the set of questions that can be asked. The recipe that we develop can find applications in numerous settings where someone wants to bridge linguistics with economics using market prices, e.g. the field of mass communication.

Our work is also linked to a fast growing literature at the intersection of monetary policy, information transmission, and asset pricing (Gürkaynak, Sack, and Swanson, 2005; Swanson, 2017; Mueller, Tahbaz-Salehi, and Vedolin, 2017; Neuhierl and Weber, 2019; Jarociński and Karadi, 2020; Cieslak and Schrimpf, 2019; Swanson, 2020). We add to this literature by documenting large information rigidities in response to monetary policy news. The predictability in forecast revisions we document is consistent with recent works that show substantial predictability of investors' expectations about the short-term interest rates (Cieslak, 2018), and monetary policy surprises (Bauer and Swanson, 2020). They argue that this predictability is not a risk premium, but is instead due to markets having underestimated the Fed's responsiveness to the state of the economy, or more generally the "Fed response to news" channel.

A recent body of literature has focused on the European Central Bank (ECB), and, similar to our work, has analysed press conferences separately from the statement releases (see for instance Altavilla, Brugnolini, Gürkaynak, Motto, and Ragusa, 2019, 2020; Leombroni, Vedolin, Venter, and Whelan, 2020). Relative to these works, we have an exact match between the words spoken by the Central Bank Chairman in each given minute and the price of financial assets in the same minute. This improves on the identification of which specific message the market reacted too, and of the "communication surprises" in the press conference. Our machinery permits us also to identify the close connection between statement news and press conference news and to document the key role of clarification of statement news within the context of the press conference.

Finally, the literature on the signaling effects of monetary policy is among the largest in economics. Seminal contributions include Cukierman and Meltzer (1986) and Ellingsen and Soderstrom (2001). Recent contributions include Berkelmans (2011), Melosi (2016), Nakamura and Steinsson (2018). We contribute to this literature by showing the link between statement and press conference news to financial investors, and the relation between messages sent and signals received. We show how the messages communicated during the post-FOMC press conference form investors' expectations, and document the importance of those moments in which the Fed Chairman answers questions related to the interpretation of the post-meeting statement.

2 Data

Our data come from multiple sources, and their nature is twofold. On the one hand, we propose a novel way to generate and use time-stamped text as data. We apply our method to post-FOMC meeting press conferences, which are key events for financial investors worldwide. On the other hand, we have high-frequency quote-level prices for a wide range of financial assets. To the best of our knowledge, we are the first to use videos in economics by linking the time-stamped words with high-frequency financial asset prices.

2.1 FOMC meetings

Given their importance to financial investors, FOMC meetings are an ideal laboratory to study real-time price discovery. Committee members hold eight regularly scheduled meetings every year. In each of them, they set the current monetary policy actions, and discuss the likely future course of monetary policy. Starting in 1994, the decisions have been announced to the public via the release of a policy statement at 14:00 Eastern Time (EST). In April 2011, the then Chairman Ben Bernanke began the practice of holding a post-meeting press conference four times a year.[5] Since 2019 all FOMC meetings have been followed by a press conference. The overall goal of the statement and the following press conference is to increase transparency of Fed's actions, and reduce market reactions and surprises.[6]

[5]The introduction of post-FOMC-meeting press conferences in the United States was a response to the financial crisis and deep recession. In fact, clear communication is especially important when economic conditions require additional policy stimulus but the policy rate is already at its effective lower bound. The greatest public interest in the Federal Reserve's communication during this period makes our sample ideal to study the connection between the Fed Chairman's words and the movements in investors' beliefs, and the post-FOMC-meeting press conferences the perfect laboratory to analyse real-time price discovery.

[6]Janet Yellen was in charge of the subcommittee studying the rationale for moving ahead with press conferences. Quoting her words during the March 2011 FOMC meeting, "a crucial element of our mission was to consider approaches for ensuring that the public understands both the consensus of the Committee and the diversity of views among individual participants" and also "the purpose is to allow

2.1.1 *Time-stamped FOMC press conferences*

Our novel dataset is made by the audio of post-FOMC-meeting press conferences.[7] We have converted the audio into an interpretable text. We have recorded the exact time in which each word was pronounced. We have aligned the text with high-frequency asset prices.

Formally, we have converted a sequence of audio X into a sequence of words W. Let $p(W|X)$ denote the probability of a word sequence given the audio. We have obtained W^* by maximizing $p(W|X)$ over the set of all possible word sequences V, that is

$$W^* = \text{argmax}_{W \in \mathcal{V}} p(W|X). \tag{1}$$

Specifically, to obtain an estimate of W^* we have proceeded in 4 steps. First, we have split the audio into smaller frames of around 3 seconds each and preprocessed the audio clips into spectrograms. Second, we have used the end-to-end deep learning algorithm developed by Hannun et al. (2014) to optimize $p(W|X)$ directly.[8] In particular, we have used recurrent neural networks to convert the spectrograms into a sequence of characters, c, and corresponding probabilities.[9] Conditional on c, it is possible to use the Connectionist Temporal Classification algorithm of Graves, Fernández, Gomez, and Schmidhuber (2006) to draw a sequence of readable text transcriptions W. Third, once we were able to evaluate $p(W|X)$, we have followed Hannun et al. (2014) and used a beam-search algorithm to estimate W^* in (1). Fourth, we have leveraged the specific structure of our application and aligned our estimate of W^* with the text in the press-conference transcripts published by the Federal Open Market Committee. This allowed

any news to be digested into market prices."

[7] The original video files can be found at https://www.federalreserve.gov/monetarypolicy/fomccalendars.htm

[8] Alternatively, we could have applied Bayes' Theorem to obtain $p(W|X) \propto p(X|W)p(W)$ and optimize the conditional distribution $p(X|W)$ for a given language model $p(W)$. However, as explained in Hannun et al. (2014) and Amodei, Ananthanarayanan, Anubhai, Bai, Battenberg, Case, Casper, Catanzaro, Cheng, Chen, et al. (2016), estimating $p(X|W)$ separately can lead to sub-optimal results, due the lack of error propagation between the probability densities. In contrast, end-to-end methods, that optimize $p(W|X)$ directly, allow the model to learn from the data directly conditional on a sufficient large training dataset.

[9] The output of RNN will have different lengths depending, for instance, on the speed of the speaker, pronunciation, acoustic environment, spontaneous speech (e.g., "um" or "uh"), etc. Therefore, we need an additional step that maps the output of the neural network into a readable transcription. To deal with this problem we use CTC which is a state-of-the-art algorithm that addresses this issue.

Table 1. **Example of time-stamped transcription**

Start	End	Text
14:36:34.096	14:36:37.906	In terms of the rest of your question,
14:36:38.356	14:36:42.416	the Committee is really thinking of this as a way
14:36:42.416	14:36:43.526	of adjusting policy
14:36:43.526	14:36:45.576	to a somewhat more accommodative stance
14:36:46.046	14:36:48.596	to further the three objectives that I mentioned:
14:36:49.296	14:36:53.786	to insure against downside risks, to provide support
14:36:53.786	14:36:59.726	to the economy, that those factors are-where factors are
14:36:59.726	14:37:02.756	pushing down on economic growth, and then to support inflation.
14:37:02.756	14:37:05.586	So we do think it will serve all of those goals, but again,
14:37:05.586	14:37:07.696	we're thinking of it as essentially in the nature
14:37:07.696	14:37:09.526	of a midcycle adjustment to policy.

Notes: The Table reports an example of transcribed text with the starting and ending time (hours, minutes, seconds, milliseconds) within the press conference on July 31, 2019.

us to create a perfect match between the audio and the text transcription by using a combination of manual and automated procedures.

Next, we have time-stamped the text of each 3-second audio frame.[10] For each press conference, we have appended the three-second text, and aligned the beginning and the end of the press conference with the times published by Bloomberg. Table 1 shows an example of such a time-stamped transcription, and highlights the precision with which we identify the time when the words were pronounced during the press conference.

Let $\mathbf{W_j}$ be a matrix summarizing the press conference of date j. The columns correspond to the words contained in the text of the press conference, while the rows are the three-second time windows. The matrix elements are equal to one if a certain word was mentioned in a three-second window, otherwise zero. Before creating this matrix, we have performed the following pre-processing steps on the raw text: (i) lowercasing the words; (ii) removing punctuation, hyphens, and apostrophes; (iii) removing words specific to the speech-to-audio translation such as noise coming from the acoustic environment, and spontaneous speech; (iv) removing a list of very common English words (e.g., stop-words)[11]; and (v) reducing the remaining words to their root based on the Porter (1980) stemmer algorithm. Finally, for each press conference, we further recorded the start and

[10]We can increase or decrease the length of the audio frame, by modifying the length of the audio clips inputted in (1).

[11]The list of stop-words that we remove come from the Python Natural Language Toolkit.

the end of the question and answer section together with the time in which each question was asked, the name of the reporter, and her/his affiliation.

Overall, we considered all 41 press conferences, covering a sample period from April 2011 to January 2020. On average the duration of each press conference is 54 minutes and 47 seconds, where the first 10 minutes and 17 seconds correspond to the opening statement made by the Chair of the FOMC. The rest of the conference corresponds to the question and answer section that contains on average 23 different questions. After pre-processing the text, the overall vocabulary contains 7,580 unique words pronounced a total of 156,767 times, where each word is mentioned on average 20 times.

2.1.2 Extracting news from the FOMC meeting statement

The second source of information are the news contained in FOMC statements. We identify these news by tracking the sentences/words added or removed relative to the previous statement. Indeed it's common practice by Fed watchers to parse those changes to infer any new guidance on rates or variation in the economic outlook.[12]

For each press conference j, we append the changes, and build a vector $\mathbf{s_j}$. On average, each policy statement contains changes in 3.8 sentences. The average length of the changes is 7 words. Appendix A provides two examples of statement news. On purpose, we chose a statement with a large number of changes relative to the previous one, and a statement with only few variations in the text. In Section 4 and Appendix A, we show that the audience ask directly about the statement news, which allows us to link $\mathbf{s_j}$ and $\mathbf{W_j}$.

2.2 High-frequency asset prices.

After constructing a second-level time-stamped text dataset, we use high-frequency financial data to characterize the real-time price discovery. In this regard, our financial data come from three different sources.

First, we use Best-of-Book trade and quote data (BBO) data for Federal funds futures and Eurodollar futures from the Chicago Board of Trade and the Chicago Mercantile

[12]For instance, the Wall Street Journal publishes the changes between consecutive statements couple of minutes after the statement release. See `https://blogs.wsj.com/economics/tag/parsing-the-fed/`.

Exchange, respectively. As for Federal funds futures, contracts with expiration up to two years are offered, while Eurodollar futures span a longer horizon of up to seven years. At each point in time we have over 20 different maturities for Federal funds futures, and over 30 different maturities for Eurodollar futures. The advantage of using these contracts is that their prices are closely linked to investors' expectations of monetary policy actions, and target federal funds rates. For both products, our dataset contains the bid, and the ask price, the traded price, and the trading volume. Prices are reported according to the International Monetary Market Index quote convention, i.e. 100 minus the rate. For the case of Federal funds futures the rate is an arithmetic average of the daily effective rate during the contract expiration month, so a price quote of 94.25 would imply an average daily rate of 5.75 percent per annum. As to Eurodollar futures, the implied rate is the three-month London interbank offered rate for spot settlement on 3rd Wednesday of the contract expiration month. For every minute and futures maturity, we compute the implied rate estimates using mid prices.

Second, we use the Trades And Quote (TAQ) database for the intraday behavior of the S&P 500 stock market index during FOMC days as well as of its constituents. From individual stocks we form industry portfolios by combining the high-frequency prices of individual S&P 500 stock constituents with the Fama-French definition of 30 sectors. We require that in each day at least 10 stocks are present in each portfolios. At 10:00 am EST of the FOMC day we invest one dollar in each stock in the portfolio and we look at the portfolio performance during that day. The Trades And Quote database offers a complete history of trades and quotes within the U.S. National Market System. It contains the bid, and ask price and it is time-stamped at the millisecond. Within each minute, we take the median of millisecond mid prices.

Third and finally, we use spot exchange-rate quotes on seven currencies against the US Dollar: Australian Dollar, Euro, British Pound, New Zealand Dollar, Swiss Franc, Japanese Yen, Canadian Dollar. All quotes are from Dukascopy, which offers historical tick-by-tick market data for dealable interbank foreign exchange rates for each millisecond. Again, within each minute, we take the median of millisecond mid prices.

3 Statement and press conference news

We next present a model of information rigidity based on noisy signals, and derive predictions for the relationship between ex-post and ex-ante forecast revisions. We then show our empirical findings, and discuss how they relate to this model.

3.1 Noisy information model

Let x_t be the state of monetary policy at time t (from t to $t+1$). Its law of motion is described by:

$$x_t = \rho x_{t-1} + \nu_t, \tag{2}$$

where ν_t is an i.i.d. normally distributed random variable with mean zero and standard deviation σ_ν, while ρ is the autocorrelation coefficient, assumed positive without loss of generality.

There is a unit mass of atomistic agents who know the model governing monetary policy and its parameter values, but cannot observe x_t. Instead, they are endowed with their own information processing abilities to translate the Fed Chairman's words into private signals, and update their estimate in a Bayesian fashion. Specifically, they observe x_t with idiosyncratic (agent-specific) noise (á la Woodford, 2003), which creates heterogeneity in the way the public signal gets interpreted. The private signal $y_{i,t}$, received at the beginning of period t, is such that

$$y_{it} = x_t + \omega_{i,t}, \qquad \omega_{i,t} \sim N(0, \sigma_\omega^2) \tag{3}$$

where $\omega_{i,t}$ is i.i.d. both across time and agents. Let $F_{i,t}x_t$ be the forecast of agent i given her information set at time t, agent i updates the previous estimate as follows:

$$F_{i,t}x_t = F_{i,t-1}x_t + G\left(y_{i,t} - F_{i,t-1}x_t\right), \tag{4}$$

where G is the steady-state Kalman filter gain, which is a function of the parameters of the model. The Kalman gain represents the relative weight placed on the new information that agent i received at time t, i.e. $y_{i,t} - F_{i,t-1}x_t$. A value of $G = 1$ implies that the signal is fully revealing, while the presence of noise in the signal implies $G < 1$. Rewriting (4) as follows

$$F_{i,t}x_t = Gy_{i,t} + (1 - G)\, F_{i,t-1}x_t \tag{5}$$

reveals that $1 - G$ is the degree of information rigidity in this model.[13] While individuals form their forecasts rationally conditional on their information set and noisy signals, they do not know whether the new information reflects noise or innovations to the variable being forecasted, x. For this reason, they decide to adjust their beliefs only gradually in response to shocks to fundamentals.

Proposition 1. *Let x_{t+h} be the monetary policy regime effective from time $t + h$ to time $t + h + 1$, for any $h \geq 0$. Assuming expectations are noisy in the sense of* (2) *and* (3)*, then forecast revisions are autocorrelated over time:*

$$E_{t-1}[F_t x_{t+h} - F_{t-1}x_{t+h}] = \lambda[F_{t-1}x_{t+h} - F_{t-2}x_{t+h}], \tag{7}$$

where $\lambda = 1 - G$.

Proof. See Appendix F. □

We follow an extensive literature on high-frequency identification, and proxy market expectations revisions with high-frequency changes in asset prices (see for instance Nakamura and Steinsson, 2018). The standard approach in high-frequency studies is to compute price changes in narrow windows, treating those changes as measures of unexpected movements in investors' beliefs. We use those price changes in a test of the null hypothesis of full-information rational expectations. This test is commonly run

[13] Coibion and Gorodnichenko (2012, 2015) show that this model has implications consistent with the following framework by Mankiw and Reis (2002), Carroll (2003) and Reis (2006). Suppose inattentive agents update their beliefs each period with probability $G \in [0, 1]$, and when they update, they acquire full information, and have rational expectations. Forecasts are then updated according to the process

$$F_t x_{t+h} = GE_t x_{t+h} + (1 - G)F_{t-1}x_{t+h}, \tag{6}$$

where the coefficient $1 - G$ captures the degree of expectation rigidity. When G is equal to one, the model collapses to a perfectly rational expectations framework. Instead, a value of G lower than one suggests that information is slowly updated by (at least) some agents. Also this model, as the one in the text shares the feature of information rigidity, but the source is different: expectation stickiness versus noisy signals. In this paper, we do not take a stand on the source of the information rigidity, but we limit ourselves to document such a rigidity in the context of the Federal Reserve communication events.

Table 2. **Summary statistics for the changes in prices for different asset classes around FOMC post-meeting statement release and press conference.**

		Fed funds futures		Eurodollar Futures		Stocks	Forex
Δp	Event	1m-6m	9m-15m	6m-12m	24m-70m		
Average	ST	-0.19	-0.97	-0.68	-0.67	17.35	5.78
	PC	-0.08	-0.01	-0.07	-0.08	2.99	-3.68
Standard Deviation	ST	2.00	3.92	3.73	6.03	46.32	39.75
	PC	1.08	2.92	2.24	4.57	50.97	30.21
Average absolute value	ST	1.30	2.94	2.70	4.32	36.31	30.69
	PC	0.56	1.95	1.42	3.24	37.42	23.76

Notes: For a wide range of financial assets the Table reports the average value, standard deviation, and average absolute value for price changes around the times of the post-FOMC-meeting statement release (ST), and press conference (PC). In FOMC days with press conferences, the change in price around the statement is equal to the change in price between 13:50 and 14:20. The change in price around the press conference equals the change in price from the beginning to the end of the post-meeting press conference held by the Fed Chairman, which in general starts at 14:30. All values in the Table are in basis points.

using survey data. However, relative to surveys, high-frequency asset prices have several advantages: surveys might be stale or just cheap talk, while investors trade at market prices implying direct gains or losses from using them.

3.2 Variable construction

We compute changes in asset prices around two separate, non-consecutive time windows. The first is a 30-minute window around the FOMC announcement, which occurs at 14:00 EST. Following among others Gürkaynak, Sack, and Swanson (2005), Fleming and Piazzesi (2005), and Nakamura and Steinsson (2018), we compute the change in the price of financial assets between 13:50 and 14:20. The second time window is the press conference window: it starts in general at 14:30 EST, and lasts on average 55 minutes. We use the exact starting and ending time. Besides, there is no minute overlapping the two windows.

Table 2 shows the average value, standard deviation, and mean absolute value for each of the asset classes in our study both around the statement release and the press conference. We group all the assets into different buckets: short-term and medium term Federal funds futures (1–6 and 9–15 months respectively), short-term and medium term Eurodollar futures (6–12 and 24–70 months respectively), stocks, and exchange rates.

Table 3. **Press-conference shocks against statement shocks**

	Fed funds futures		Eurodollar Futures		Stocks	Forex
	1m-6m	9m-15m	6m-12m	24m-70m		
a	-0.05	0.16	-0.03	0.14	-4.30	-5.18
	[-0.41]	[0.54]	[-0.11]	[0.28]	[-0.51]	[-1.51]
λ	0.17	0.17	0.19	0.33	0.41	0.25
	[1.99]	[1.47]	[3.87]	[2.77]	[2.59]	[3.10]
R^2	8.53	5.50	8.43	20.21	14.10	11.20

Notes: For each asset bucket k, the Table reports the regression estimates for the following Equation:

$$\underbrace{\Delta p_{it,PC}}_{\Delta\text{p at press conference: e.g. 14:30–15:30}} = a_k + \lambda_k \underbrace{\Delta p_{it,ST}}_{\Delta\text{p around statement: 13:50–14:20}} + \epsilon_{it},$$

where $\Delta p_{it,PC}$ is the change in asset i's price during date-t press conference, $\Delta p_{it,ST}$ is the change in asset i's price around date-t FOMC statement release, and asset i belongs to bucket k. The two price changes, in basis points, are computed over two non-overlapping, non-consecutive time intervals. T-statistics are in square brackets. Standard errors are double clustered at date-asset level. R-squared statistics are in percentage.

For each asset within the bucket we compute the summary statistics, and then report the bucket mean value in the Table. The average value of both shocks is close to zero. Moreover, the variation of the statement and the press-conference shocks is rather similar. Both the standard deviation and the mean absolute value for press conference shocks are comparable to the ones of the statement. Hence, price changes around the press conference are of similar magnitude as price changes around the statement.[14]

3.3 Persistence in price shocks around FOMC events

Figure 2 examines the relation between price changes around the press conference and price changes around the statement release for federal funds and Eurodollar futures. Each dot corresponds to an FOMC day. The line is the regression line from a univariate linear regression model. The relation is positive across all subplots, and the slope increases with the asset maturity.

[14]In Appendix C, we compute the coefficient of variation for minute-level changes in the price of federal funds and Eurodollar futures. We report its distribution for separate maturities, and for four non-overlapping sub-periods during FOMC days. We document a large variation of these prices around the statement release, as well as during the press conference. In contrast, the variation in prices before the statement, and after the press conference is close to none. Appendix C provides further details on the exercise.

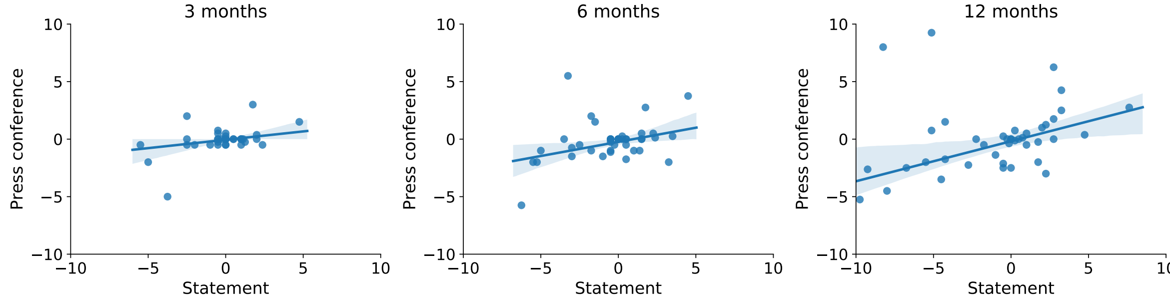

Panel A – Federal funds futures

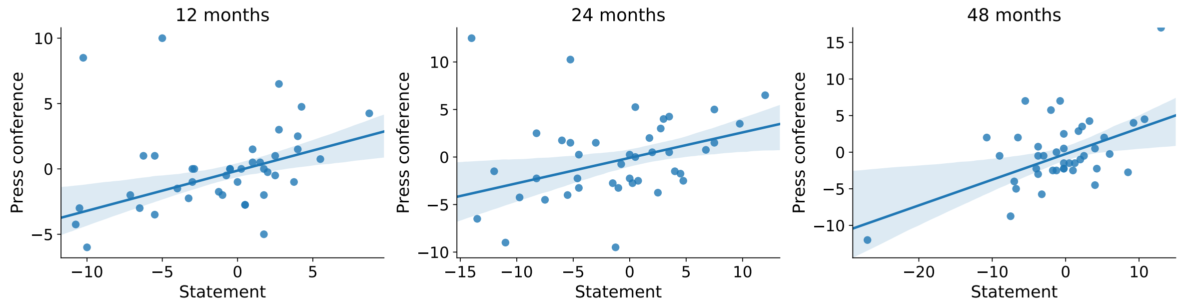

Panel B – Eurodollar futures

Fig. 2. *Notes*: The Figure shows the statement shocks on the x-axis, and the press-conference shocks on the y-axis for the 30-day Federal funds futures expiring in 3 month, 6 months and 12 months (Panel A) and Eurodollar futures expiring in 12 month, 24 months and 60 months (Panel B). The shocks are in basis points. The straight line is the regression fit line, and the dashed area around the line the 95% confidence interval bands.

To test the significance of the positive correlation between the two price changes, we run a pooled OLS where we group all the assets into different buckets as before. For each asset bucket k, we estimate the following Equation:

$$\Delta p_{it,PC} = a_k + \lambda_k \, \Delta p_{it,ST} + \epsilon_{it}, \tag{8}$$

where $\Delta p_{it,PC}$ is the change in asset i's price during date-t press conference, $\Delta p_{it,ST}$ is the change in asset i's price around date-t FOMC statement release, and asset i belongs to bucket k. We double cluster the standard errors at the date-asset level.[15]

Table 3 reports the regression results. The point estimates confirm the evidence from the scatterplots: there is a strong and statistically significant positive relation between

[15]We allow observation (i, t) to be correlated with observation (i, s) for time $s \neq t$ and with observation (j, t) for asset $j \neq i$.

the two shocks. For all asset classes, except for Federal funds futures, results are highly significant. The slope coefficient estimates are similar across all asset classes, ranging from 0.2 to 0.4. Under the noisy information model of Section 3.1, this suggests a similar degree of information rigidity for participants in different markets. Financial investors update their forecasts assigning a weight between 60% and 80% to their signals and the rest to the previous forecast.[16] Such an estimate of information rigidity is not only helpful for interpreting the economic significance of our findings, but can also be used as a moment to which calibrate monetary equilibrium models in which agents interpret noisy signals from the Central Bank. Our results on the values of information rigidity around scheduled FOMC communication events are surprising given the large attention paid by media and financial investors. Given the key role of these events in the economic calendar, our estimates can be arguably interpreted as a lower bound to the degree of information rigidity observable when studying financial markets' response to news.

Our object of interest is the degree of information rigidity, as measured by one minus the Kalman gain. However, it has become increasingly popular in the financial press and academia to evaluate the *informativeness* of the Central Bank communication through the signal-to-noise ratio, i.e. the ratio between the variance of fundamentals and the variance of the noise in private signals. Indeed, improving the signal-to-noise ratio with clearer communication has become an explicit objective for some Central Banks (Blinder, Ehrmann, Fratzscher, de Haan, and Jansen, 2008). In Appendix G, we describe and derive the relation between the Kalman gain and the signal-to-noise ratio. The key insight is that in the presence of persistent policy news, even with a high monetary policy signal-to-noise ratio, investors update their expectations slowly if they perceive information with noise. Therefore, from a policy perspective, the signal-to-noise ratio does not provide the information necessary to evaluate how fast information shapes financial investors' expectations, whereas the degree of information rigidity captures such a value directly.

[16]The R^2 in Table 3 suggests that the correlation is of the same magnitude as the slope coefficient in the regression (this is because the two shocks have similar volatilities). For instance, the correlation between price changes is 40% for medium-term Eurodollar futures, but it goes as high as 58% for the 60-month maturity. Similarly, the average correlation for all stock portfolios is 33%, while it reaches 44% for SPY.

We report additional results in Appendix D. Table D.1, and D.2 report the regression estimates of Equation 8 for each asset separately. In Table D.3, we repeat the analysis for federal funds, and Eurodollar futures, for a "placebo" event period: we use FOMC days without a press conference and compute price changes around the statement release, and around an alternative window that mimics the average press conference time (from 14:30 to 15:24 EST). We don't find any evidence of a positive autocorrelation of price changes in days without a press conference.

3.4 Economic Significance of the persistence in price shocks

To measure the economic value of information rigidity to FOMC news we implement a simple trading strategy. For each asset class, we use the half-hour returns around the FOMC statement release as a trading signal. For every asset, we take a long position at the beginning of the press conference if its price went up when the statement was released and a short position otherwise. We close the position at the end of the press conference.

We compare the results of our trading strategy with a buy-and-hold strategy, where we buy the assets regardless of the information received at the statement release. For each asset bucket k, we run the following regression:

$$\mathrm{r}_{it,MT} = \alpha_k + \beta_k \mathrm{r}_{it,B} + \epsilon_{it}, \tag{9}$$

where $\mathrm{r}_{it,MT}$ are the returns from the active strategy involving asset i in the FOMC date t, and $\mathrm{r}_{it,B}$ are the returns from a passive buy-and-hold strategy. Asset i belongs to bucket k. A positive intercept, α, implies that the active strategy has a higher Sharpe ratio relative to the passive buy-and-hold approach, i.e. higher average return scaled by the return volatility.[17]

[17] Specifically, let the Sharpe ratio for asset j be $\frac{E(r_j)-r_f}{\sigma_j}$, where r_f is the risk-free rate, and σ_j the volatility of asset j's returns. The relation between Sharpe ratio and alpha is

$$\frac{E(r_j) - r_f}{\sigma_j} = \frac{\alpha_j}{\sigma_j} + \rho_{j,B} \frac{E(r_B) - r_f}{\sigma_B},$$

where $\rho_{j,B}$ denotes the correlation parameter between r_j and r_B given by $\sqrt{R^2}$.

Table 4. **Economic Significance**

	Fed funds futures		Eurodollar Futures		Stocks	Forex
	1m-6m	9m-15m	6m-12m	24m-70m		
α	0.37	0.40	1.16	1.02	12.66	8.14
	[1.27]	[1.65]	[2.80]	[2.11]	[2.30]	[2.57]
β	0.54	0.33	0.32	0.08	-0.13	0.18
	[2.40]	[2.53]	[2.68]	[0.42]	[-0.61]	[1.12]
R^2	9.84	17.79	4.38	0.89	1.69	3.39

Notes: The Table reports the regression statistics to evaluate the economic significance of a market-timing strategy that exploits the information released around the FOMC announcement. We take a long position in the asset at the beginning of the press conference if its price went up when the statement was released and a short position otherwise. We close the position at the end of the press conference. We compare this strategy with a simple buy-and-hold strategy. For each asset bucket k, we regress the returns of the market-timing strategy onto the ones from the buy-and-hold strategy:

$$\mathrm{r}_{it,MT} = \alpha_k + \beta_k \mathrm{r}_{it,B} + \epsilon_{it},$$

where $\mathrm{r}_{it,MT}$ are the returns from a market-timing strategy involving asset i, $\mathrm{r}_{it,B}$ are the returns from a passive buy-and-hold strategy, and asset i belongs to bucket k. T-statistics are in square brackets. Standard errors are double clustered at date-asset level. The α coefficients are in basis points. R-squared statistics are in percentage.

Table 4 reports the regression results, where we double cluster the standard errors at the date-asset level. A timing strategy that exploits the information coming from the statement substantially outperforms the passive strategy. The numbers reported are not converted to reflect a lower frequency, e.g. not annualized. The intercepts are positive and statistically significant. They imply a large increase in Sharpe ratios. For instance, a timing strategy that exploits the information in the statement applied to 60-month Eurodollar futures will have a Sharpe ratio increase of 25% relative to a buy-and-hold strategy. For SPY the Sharpe ratio goes up by 34%.

Figure 3 shows the mean point-wise cumulative intraday return of the active trading strategy compared to a buy-and-hold strategy. We report results for three different assets: 60-month Eurodollar futures, SPY and EUR/USD exchange rate. The x-axis represents the minutes since the press conference started. The solid line tracks the returns from the market timing strategy, while the dotted line the returns from the passive buy-and-hold strategy. For all three assets the overperformance starts from minute 10 when the Q&A session starts. The SPY seems to react rapidly: after 20 minutes into the Q&A session the cumulative returns stabilize. For the other two assets the cumulative returns show a

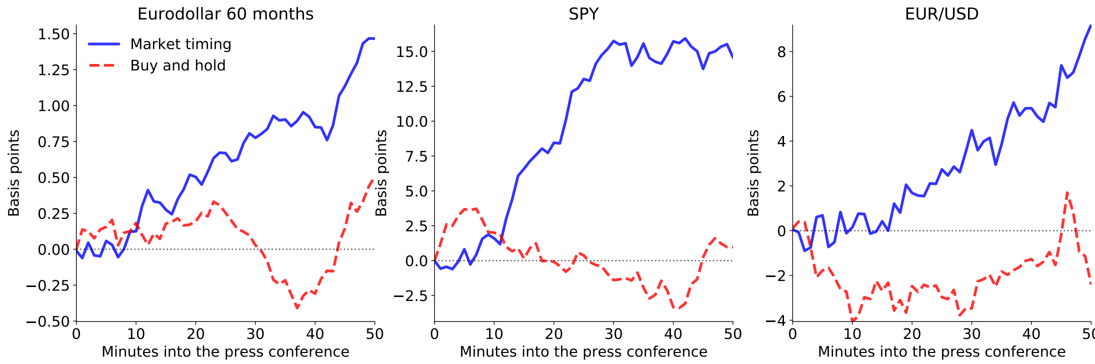

Fig. 3. *Notes*: The Figure shows the intraday evolution of the average cumulative performance, in basis points, for a timing and a passive buy-and-hold strategy. Both strategies are implemented on 60-month Eurodollar futures (left Panel), SPY (middle Panel) and EUR/USD exchange rate (right Panel). The market timing strategy exploits the information released around the FOMC announcement. We take a long position in the asset at the beginning of the press conference if its price went up when the statement was released and a short position otherwise. We close the position at the end of the press conference. We compare this strategy with a simple buy-and-hold strategy which always goes long the asset. Both strategies are implemented from the beginning to the end of the press conference.

steady growing pattern all the way from minute 10 to the end.[18]

The approach followed in this section is simple, yet provides evidence of a tight link between statement and press conference news. In the next section, we take the analysis a step further. We combine the exact words pronounced in each given minute of the press conference with higher-frequency returns, and show that the minutes in which the Fed Chairman discusses the statement news lie behind the positive autocorrelation in price changes documented so far.

4 Within press-conference analysis

We now extend the noisy information model in Section 3.1 to study the investors' expectations formation process during the press conference. We assume that: a) the Federal Reserve communicates to the public the consensus view reached during the Committee meeting, both through the post-meeting statement and the press conference;

[18]The results in this section differ from the pre-FOMC drift documented by Lucca and Moench (2015). First, we focus on a narrow 50-minute window, which is the press conference, rather than on a larger 3-day window. Second, the results of Lucca and Moench (2015) are based on the the buy-and-hold rather than the market timing strategy. Third, as they have already documented, and consistent with our results, after the FOMC announcement the returns from holding a long position on the stock market are on average zero.

b) this view does not change from the statement release to the Chairman's press conference; c) statement-related minutes are the only moments in which investors learn about the monetary policy path. This latter implies that investors update their expectations about the state variable x_t, as in (5), only during those minutes.[19]

These assumptions lead us to three new testable predictions on the high-frequency expectation formation process: 1) in every minute related to the statement some investors will update their beliefs and trade accordingly; 2) in statement-related minutes average forecasts should move more than in other minutes; 3) forecasts will move in the same direction as the initial reaction to the statement generating a positive autocorrelation in forecasts revisions.

4.1 Variable construction: linking statement and press conference news

To test the three predictions, we first need to link the statement news $\mathbf{s_j}$ computed in Section 2.1.2 with the press conference word matrix $\mathbf{W_j}$ described in Section 2.1.1. Our approach proceeds in four steps. First, we aggregate the 3-second level text in $\mathbf{W_j}$ to 1-minute frequency. Heuristically, we observe one minute is adequate to capture the asset price response to words.[20] Second, we run a part-of-speech analysis of the sentences identified in $\mathbf{s_j}$. This means splitting those sentences into nouns, adjectives, verbs, etc... Third, within the press conference text we search for all combinations that include those nouns, and verbs from $\mathbf{s_j}$. We also take their synonyms, which are reported for convenience in Table B.1 in the Appendix. To make sure that our words actually capture the link with the statement, we add the requirement that in the same sentence

[19]This assumption is without loss of generality. Calling λ the degree of information rigidity, it is rather straightforward to extend this framework to allow for different degrees of information rigidity in statement-related minutes (λ_1) and in the rest of the conference (λ_2). We subsequently test the hypothesis that $\lambda_1 < \lambda_2$ and provide empirical support for. It is just to simplify the discussion of the economic mechanism at work that we describe the case in which λ_2 is simply equal to one.

[20]The Fed video feed requires more bandwidth, and lags the audio feed by approximately three seconds. From the way we time stamp the text, we are implicitly assuming that investors listen to the audio and react to that. Aggregating the text to 1-minute frequency helps reducing any possible noise potentially coming from a 3-second difference with which different investors receive the same information. In an unreported robustness test, we have also shifted the text by three seconds, aggregated it again at the 1-minute frequency and rerun our analysis. Results were unaffected. Therefore, we can conclude that the 3-second lag between the video and audio feed is immaterial to our results.

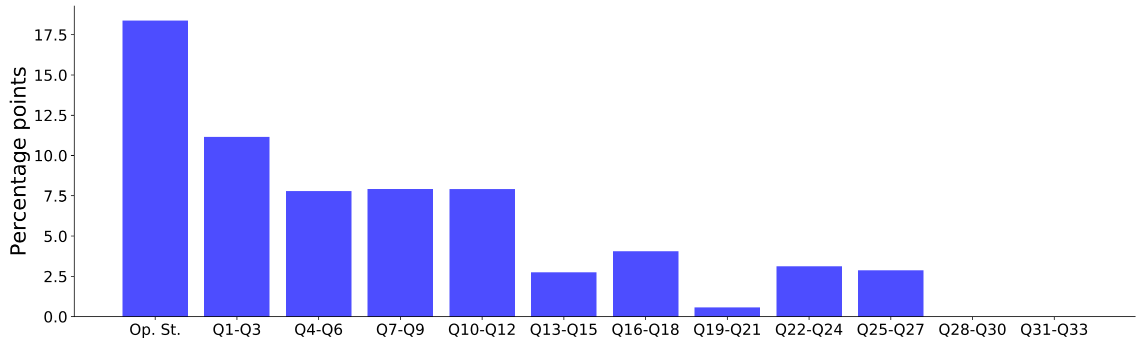

Fig. 4. *Notes*: The Figure shows the average value of the dummy, D_t, capturing statement-related minutes during different phases of the press conference. We separate the press conference into opening statement, and questions, these latter being grouped by their order.

the word "statement" should appear.[21] Fourth and finally, we create a minute-level dummy variable, D_t, equal to 1 when that combination is identified in a given minute of the press conference.

Overall, statement-related minutes account for 7.5% of press-conference minutes. However, they are not uniformly distributed. Figure 4 shows the average values of the dummy across different moments of the press conference. We separate the press conference into opening statement, and questions, these latter being grouped by their order. We document a strong connection between the average values of the dummy and the progress of the press conference. About 18% of the opening statement is devoted to statement news. During the opening statement the Chairman describes in details the policy statement: only few words/topics are actual news while the rest refers to aspects of the economic outlook or policy that have not changed relative to the previous FOMC meeting.

The figure also shows a close link between the question order and the statement news. This link fades away as time passes and more questions arrive. Fed watchers try to infer large changes in the Fed's policy from small changes in the statement's wording.[22] So they ask already in the first few questions if the Chairman could provide additional

[21]The combination of the word "statement" and the statement news is necessary to avoid false positive such as "the New York Fed's website contains a statement" or "Milton Friedman's statement."

[22]Fed watchers "try to make a living out of parsing these statements" (Peter Barnes, Fox News, April 2012)

information about the changes in the statement.[23]

We further analyse directly the linguistics of the messages to identify the different language patterns, and styles that characterize the minutes in which the Chairman talks about the statement.[24] For brevity, we report the results in Appendix E, while here we briefly describe them. Minutes in which the Chairman talks about the statement news are characterized by a larger use of future tense, relative to present or past tense; they tend to involve more insight words, such as "think" or "consider", and more relative words to qualify the statement such as "during" and "when;" finally they features more comparison words, "than" or "as," and numbers, which tend to be used almost 20% more frequently than non-statement-related minutes. Overall this suggests that messages in statement-related minutes are more specific and informative than messages in other minutes (Pennebaker, Boyd, Jordan, and Blackburn, 2015).

4.2 How the Chairman's message induce variation in investors' beliefs

We run three simple regressions that make use of the minute-level dummy variable, D_t, constructed in 4.1, to quantify the average absolute price variation, trading volume, and mean returns in statement-related minutes. Our identifying assumption is that unexpected changes in statement-related minutes for both asset prices and trading volume arise from the message communicated in that minute.

4.2.1 Absolute variation

The first regression we estimate serves to assess the average absolute variation in minute-level returns for statement-related minutes. We group all assets described in Section 2 into different buckets, and run a pooled OLS. Let $|r_{it}|$ be the absolute value of the financial returns of asset i between minutes $t-1$ and t, and D the dummy variable. For

[23] Our identifying assumption is that no other shocks influence the explanatory variable during these one-minute windows. Such an assumption is common in the literature on high-frequency identification of monetary policy. However, unlike us, they have used longer windows of one or two days, or 30 minutes: Cook and Hahn (1989); Kuttner (2001); Cochrane and Piazzesi (2002); Bernanke and Kuttner (2005) used a one- or two-day window around FOMC announcements and Nakamura and Steinsson (2018) a 30-minute window.

[24] We use a rather standard word-count strategy. The search words are categorized into language categories following the Linguistic Inquiry and Word Count (LIWC) by Pennebaker et al. (2015).

Table 5. **Absolute price variation**

	Fed funds futures		Eurodollar Futures		Stocks	Forex
	1m-6m	9m-15m	6m-12m	24m-70m		
a	0.08	0.15	0.17	0.26	3.73	2.58
	[4.54]	[19.61]	[12.08]	[35.99]	[17.36]	[13.82]
b	-0.01	0.01	0.00	0.04	0.49	0.36
	[-1.30]	[0.35]	[0.02]	[2.10]	[1.94]	[2.12]

Notes: The Table reports the regression statistics to compare the average absolute price variation in statement-related minutes to all other minutes of the press conference. For each asset bucket k, we estimate the following Equation:

$$|r_{it}| = a_k + b_k D_{t-1} + \epsilon_{it},$$

where r_{it} are the minute-level returns, in basis points, of asset i belonging to bucket k and D is the dummy variable constructed as in Section 4.1. T-statistics are in square brackets. Standard errors are double clustered at date-asset level.

each asset bucket we estimate the following Equation:

$$|r_{it}| = a_k + b_k D_{t-1} + \epsilon_{it}, \tag{10}$$

where asset i belongs to bucket k.

The assumptions in our model imply that statement-related minutes are more relevant to form investors' expectations. It follows that a positive slope coefficient, b_k, should be associated with the dummy variable. Empirically, however, it is equally likely that when the Chairman discusses the policy statement he/she simply repeats some sentences, and does not provide additional information. This would imply a negative slope coefficient.

Results are reported in Table 5. News related to the information previously released in the statement exerts a generally statistically significant influence on prices. The slope coefficient in the regression is positive and economically large. For medium-term Eurodollars, stocks and forex the slope coefficient implies that the average price variation when the Chairman discusses the policy statement news is about 14% larger than in other minutes. These results do not necessarily imply that other, non-statement minutes, do not matter, but on average they move prices less than statement-related minutes. The effect of statement-related minutes for shorter-maturity assets, federal funds or Eurodollar futures, is instead close to zero and insignificant. Consistent with the findings in Section 3, these results suggest that investors update their beliefs in statement-related

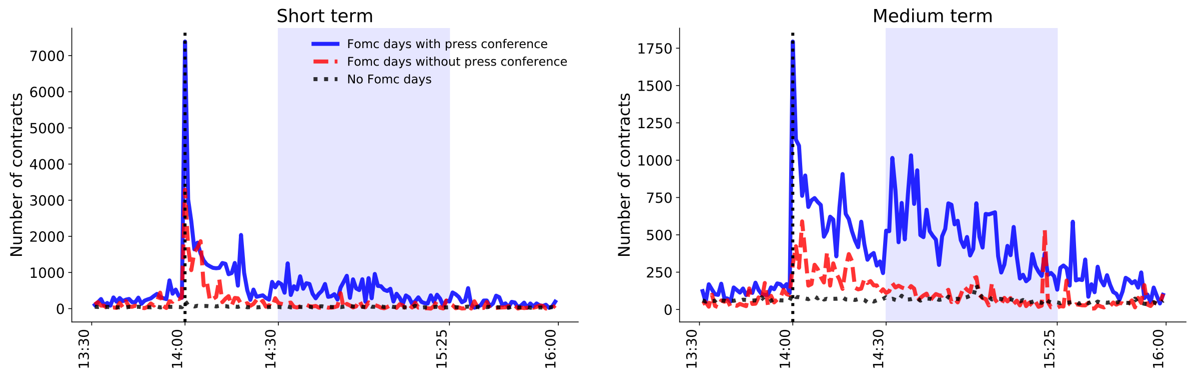

Fig. 5. *Notes*: The Figure shows the average trading volume (number of traded contracts) for the 30-day Federal funds futures for three different groups of dates. The notional amount of each contract is given by the product of the price of the futures contract times a multiplier of $4,167. The solid blue line depicts the average volume for FOMC days with a press conference, the dashed red line for FOMC days without a press conference, and the dotted black line for non-FOMC days. The left Panel shows results for Federal funds futures maturing before 9 months, while the right panel shows results for contracts with maturities above 9 months. The dashed vertical line highlights the time in which the FOMC statement is released. The shaded area denotes the FOMC press conference.

minutes mostly for the longer-term horizon.

4.2.2 Trading volume

A large literature has already documented that both trading volume and market depth increase during FOMC announcement days, and in particular in the minutes surrounding the statement release (Fleming and Piazzesi, 2005). Nevertheless, not much is known about their behavior during the Chairman's post-meeting press conference. Figure 5 shows the average trading volume for Federal funds futures in FOMC days with press conference, and compares it with the average trading volume in non-FOMC days, as well as FOMC days without a press conference (all dates starting from 2011). We plot these values from 13:30 to 16:00 for shorter-term (left Panel), and medium term (right Panel) Federal funds futures. The definition of short- and medium-term is the same as in Table 3. The basic finding is that volume jumps at announcement, and steadily decreases for shorter-term assets. In contrast, for medium-term assets the steady decrease is interrupted by the press conference, where a second jump in trading volume occurs. On average the total volume during the press conference is of the same magnitude if not higher than trading volume around the FOMC statement release. Figure D.1 in the Appendix shows this result for Federal funds and Eurodollar futures.

We then study the dynamics of the trading volume during the press conference. We scale the minute-level volume of each asset i by the total volume of the same asset during the press conference, i.e. $\frac{\text{Vol}_{ijt}}{\sum_{t\in j}\text{Vol}_{ijt}}$ where i refers to the asset, t to the minute, and j to the day. For each asset bucket k, we estimate the following Equation:

$$\frac{\text{Vol}_{ijt}}{\sum_{t\in j}\text{Vol}_{ijt}} = a_k + b_k D_t + \epsilon_{it}. \tag{11}$$

Given that the regression is estimated only on press conference days, the intercept represents the percentage of trading occurring in non-statement-related minutes during the press conference. The slope coefficient associated with the dummy is the additional average trading volume in statement-related minutes, expressed as a percentage of the total.

Ex-ante the sign of the slope coefficient could be either positive or negative. Several theoretical frameworks would predict that new information may generate trading by changing the extent of disagreement between agents (Kim and Verrecchia, 1991; Shalen, 1993). In our setting, the dummy variable capturing the discussion of statement news during the press conference would allow us to identify the moments in which investors appear to converge to similar beliefs, hence disagreement decreases. Our model would predict a positive estimate for b_k. A negative slope instead could arise when all traders interpret information in the same way, i.e. beliefs are concordant, and even in the case of heterogeneous priors, prices can move with no trading (Milgrom and Stokey, 1982; Morris, 1994; Ottaviani and Sørensen, 2015).

Table 6 reports the regression estimates. Statement-related minutes exhibit a larger trading volume which is both statistically and economically significant. The relation is stronger at longer maturities. Trading volume for Eurodollar futures between 24 and 70 months is 17% (= 0.30/1.78) larger in minutes mentioning the statement, while for Federal funds futures with maturity above 9 months that difference is almost 50% (= 0.86/1.75) larger. These values are highly economically significant given that FOMC days have been extensively shown to be among the days with the largest trading activity across several financial markets, a point reinforced by Figure 5.

Table 6. **Trading volume**

	Fed funds futures		Eurodollar futures	
	1m-6m	9m-15m	6m-12m	24m-70m
a	1.77	1.75	1.80	1.78
	[42.46]	[39.75]	[54.84]	[64.45]
b	0.48	0.86	0.04	0.30
	[1.31]	[4.96]	[0.31]	[2.57]

Notes: For each asset bucket k, the Table reports the regression statistics to compare the average trading volume in statement-related minutes to all other minutes of the press conference. We estimate the following Equation:

$$\frac{\text{Vol}_{ijt}}{\sum_{t\in j}\text{Vol}_{ijt}} = a_k + b_k D_t + \epsilon_{it}.$$

where $\frac{\text{Vol}_{ijt}}{\sum_{t\in j}\text{Vol}_{ijt}}$ is the trading volume of asset i in minute t of day j scaled by the total trading volume of asset i during the press conference in date j, and D the dummy variable constructed as in Section 4.1. Asset i belongs to bucket k. T-statistics are in square brackets. Standard errors are double clustered at date-asset level.

4.2.3 Mean returns

Finally, in this section we analyse the direction of price shocks in statement-related minutes. Our model predicts that in those minutes beliefs should be revised in the same direction as they did around the statement. We test this hypothesis and further quantify how much of the correlation documented in Section 3 is due to a discussion of the post-meeting statement news.

For each asset bucket k, we estimate the following Equation:

$$r_{it} = \begin{cases} a_k^- + b_k^- D_{t-1} + \epsilon_{it}, & \text{if } \Delta p_{ij,ST} < 0; \\ a_k^+ + b_k^+ D_{t-1} + \epsilon_{it}, & \text{if } \Delta p_{ij,ST} > 0, \end{cases} \tag{12}$$

where r_{it} are the minute-level returns, in basis points, of asset i belonging to bucket k, D is the dummy variable constructed as in Section 4.1, and $\Delta p_{ij,ST}$ is the price shock around the statement release on date j.

Table 7 reports the estimates. The Table assesses the extent to which the average changes in traders' beliefs following the Chairman's reference to the statement are of the same sign as the initial price reaction around the statement release itself.

The slope coefficient is of the same sign as the initial price reaction. The average

Table 7. **Return variation conditioning on statement news**

	Fed funds futures		Eurodollar futures		Stocks	Forex
	1m-6m	9m-15m	6m-12m	24m-70m		
Days when statement shock was negative						
a^-	-0.01	-0.01	-0.01	-0.02	-0.27	-0.17
	[-2.24]	[-0.98]	[-1.37]	[-2.52]	[-1.54]	[-1.89]
b^-	0.00	-0.03	-0.02	-0.08	-1.82	-0.55
	[0.22]	[-1.31]	[-0.52]	[-2.17]	[-2.80]	[-1.86]
Days when statement shock was positive						
a^+	0.00	0.00	0.00	0.02	0.27	0.08
	[0.43]	[0.80]	[0.15]	[1.84]	[3.31]	[0.94]
b^+	-0.00	0.01	0.05	0.04	0.58	0.22
	[-0.12]	[0.53]	[2.26]	[1.04]	[1.92]	[0.56]

Notes: For each asset bucket k, the Table reports the regression statistics to quantify how much of the correlation documented in Section 3 derives from a discussion of statement-related news. We estimate the following Equation:

$$r_{it} = \begin{cases} a_k^- + b_k^- D_{t-1} + \epsilon_{it}, & \text{if } \Delta p_{ij,ST} < 0; \\ a_k^+ + b_k^+ D_{t-1} + \epsilon_{it}, & \text{if } \Delta p_{ij,ST} > 0, \end{cases}$$

where r_{it} are the minute-level returns, in basis points, of asset i belonging to bucket k, D is the dummy variable constructed as in Section 4.1, and $\Delta p_{ij,ST}$ is the price shock around the statement release on date j. T-statistics are in square brackets. Standard errors are double clustered at date-asset level.

price movement for Eurodollars futures contracts between 24 and 70 months is -0.1 bps in statement related minutes and only -0.02 bps in the rest of the conference. The same holds for stocks, and forex where the variation is about 8 and 4 times larger in statement-related minutes, respectively. Moreover, statement-related minutes are only 7.5% of the overall minutes, yet they account for a large portion of the positive correlation in Section 3. For instance, when the initial price response around the statement was negative, statement-related minutes made 40% of the total price change over the press conference for stocks and 30% for medium-term Eurodollars futures.

This result sheds light on the learning process of financial investors after the FOMC meetings, and the Federal Reserve's communication strategy. Investors' average initial reaction when reading the statement at 14:00 is reinforced when the statement news are discussed during the press conference. Our results point out that the Chairman hints to information which both complements and better explains the statement text. A question for future research is to relate our findings to the ex-ante incentives of the Chairman in such a multi-stage signaling game.

5 Alternative explanations

There are a number of possible mechanisms linking Central bank communication with financial asset prices which differ from that examined above. In this section, we discuss six of these in the light of our results and assess their plausibility as alternative explanations.

First, as regards the Fed put, i.e. the idea that the Federal Reserve actions are excessively driven by considerations about financial markets' reactions (Cieslak, Morse, and Vissing-Jorgensen, 2019; Cieslak and Vissing-Jorgensen, 2020), we find the opposite of what the Fed put predicts. When the stock market collapses around the statement release, such a negative trend continues during the press conference. The Chairman's words reinforce the original reaction to the statement no matter what direction that reaction was.

Second, our results are inconsistent with the hypothesis that illiquidity or some microstructure effects lie behind our findings. As evidence against this hypothesis, we see no relation between price changes around the statement release and price changes in the time window between statement release and press conference or in the 30-minute window immediately after the press conference. Table 8 reports the results.

We estimate Equation (8), modifying our dependent variable to be the price change during each of the corresponding time windows. For both intervals, the slope coefficient is small and not statistically different from zero, and the R^2 close to zero. The presence of such "quiet periods" provides strong evidence against the hypothesis that our results are a mere consequence of microstructure effects.[25] It is only during the press conference that the trend we document after the statement release happens. This is why it is so important to have tick-by-tick financial data, and connect them to the words pronounced in that moment.

Third, we are able to reject that the positive autocorrelation in price changes documented in this paper is a continuation of the Lucca and Moench (2015) pre-announcement

[25] FOMC days are days in which the trading activity is at the highest (some people compare them to the final of the world cup, see for instance Figure 5) and the assets we study are among the most liquid assets in the entire financial world. So, we believe that an illiquidity story would be a very hard stretch.

Table 8. **Press-conference shocks against statement shocks**

	Fed funds futures		Eurodollar Futures		Stocks	Forex
	1m-6m	9m-15m	6m-12m	24m-70m		
Quite period 1: time window between statement release and press conference						
a	0.07	0.04	0.03	-0.32	3.14	2.87
	[1.42]	[0.30]	[0.21]	[-0.88]	[0.77]	[1.56]
λ	0.02	0.01	0.01	-0.04	0.06	0.08
	[1.01]	[0.39]	[0.41]	[-0.52]	[0.59]	[1.62]
R^2	1.24	0.21	0.08	0.74	0.84	4.12
Quite period 2: 30 minute time window after the press conference						
a	-0.05	-0.15	-0.17	-0.05	-4.11	2.72
	[-1.12]	[-1.35]	[-1.93]	[-0.25]	[-1.52]	[1.32]
λ	0.03	0.03	0.01	-0.02	0.02	0.08
	[1.74]	[1.33]	[0.65]	[-0.39]	[0.42]	[0.91]
R^2	1.98	2.34	0.17	0.71	0.17	2.48

Notes: For each asset bucket k, the Table reports the regression estimates for the following Equation:

$$\underbrace{\Delta p_{it,\text{quiet}}}_{\Delta\text{p – quiet period}} = a_k + \lambda_k \underbrace{\Delta p_{it,ST}}_{\Delta\text{p around statement}} + \epsilon_{it},$$

where $\Delta p_{it,\text{quiet}}$ is the change in asset i's price during date-t quiet period selected (either between the statement and the press conference windows or the 30 minute window right after the press conference), $\Delta p_{it,ST}$ is the change in asset i's price around date-t FOMC statement release, and asset i belongs to bucket k. The two price changes, in basis points, are computed over two non-overlapping, non-consecutive time intervals. T-statistics are in square brackets. Standard errors are double clustered at date-asset level. R-squared statistics are in percentage.

drift: we observe that when the stock market collapses around the statement release, such a negative trend continues during the press conference. This is the opposite of what the continuation of the pre-announcement drift of Lucca and Moench (2015) would imply. Moreover, again this correlation is realized only in the press-conference window and outside the press conference window we see no relation with the price change around the statement.[26]

Fourth, the results in this paper provide direct evidence against frictionless models of rational economic agents with full information. Market prices are forward looking and

[26]This is the reason why our findings cannot be explained by the post-FOMC announcement drift in U.S. bond markets documented by Brooks, Katz, and Lustig (2018). Unlike us, they study the daily response of fixed income prices for the 100 days following the FOMC announcement, and document that Treasury yields initially respond sluggishly to Fed Funds rate surprises and only after 50 days the response peaks before reverting back. We focus on the intraday variation in asset prices within FOMC days, and our results are realized only within the very short press conference window.

should already incorporate all information available to the public. So, especially at such a high frequency, they should be close to unpredictable. Imagine that investors were aware of a communication strategy of the Federal Reserve which sets that, during the press conference, the Chairman should confirm and reinforce the market interpretation of the statement. This behavior can be anticipated and exploited right away: under this framework there should be no positive autocorrelation in price changes.

Fifth, some colleagues let us notice that our findings could be generated by a model in which investors learn about the Chairman's communication strategy, political independence, policy preferences, etc... Such a model would imply that estimation risk gets reduced over time as the Chairman holds more press conferences, i.e. within a Chairman's term our results should get weaker as time passes. We propose the following test. For each Chairman in our sample, i.e. Bernanke, Yellen, and Powell, we split the sample in two halves. The first half covers the first 50% of FOMC press conferences that Chairman did, while the second half covers the 50% remaining. For each asset bucket k, we estimate the following Equation:

$$\Delta p_{it,PC} = a_k + \lambda_{1k}\, \mathbb{1}_{\text{1st half}}\, \Delta p_{it,ST} + \lambda_{2k}\, \mathbb{1}_{\text{2nd half}}\, \Delta p_{it,ST} + \epsilon_{it}, \tag{13}$$

where $\Delta p_{it,PC}$ is the change in asset i's price during date-t press conference, $\Delta p_{it,ST}$ is the change in asset i's price around date-t FOMC statement release, and asset i belongs to bucket k. We double cluster the standard errors at the date-asset level. The dummy variable $\mathbb{1}_{\text{1st half}}$ is equal to 1 if the press conference belongs to the first half of the press conferences of a Chairman, and zero otherwise, while $\mathbb{1}_{\text{2nd half}}$ is defined in the opposite way.

Table 9 reports the regression results. For all asset classes, except for stocks, the coefficient we estimate for the Chairman's second subsample is actually larger than the coefficient we estimate for the first half. This is inconsistent with a model in which traders learn about the Chairman's type. The slow reaction we have documented does not disappear as the Chairman holds more conferences, and actually it persists strongly. This is surprising given that post-FOMC meeting press conferences are events whose entire objective is a rapid and clear communication to investors.

Table 9. **Press-conference shocks against statement shocks**

	Fed funds futures		Eurodollar Futures		Stocks	Forex
	1m-6m	9m-15m	6m-12m	24m-70m		
a	-0.04	0.15	-0.04	0.16	-3.81	-5.17
	[-0.39]	[0.50]	[-0.15]	[0.30]	[-0.47]	[-1.52]
λ_{1st}	0.15	0.10	0.09	0.17	0.51	0.25
	[2.19]	[0.41]	[1.20]	[0.90]	[2.54]	[1.71]
λ_{2nd}	0.26	0.22	0.26	0.41	0.28	0.26
	[2.88]	[1.76]	[4.61]	[3.17]	[1.82]	[3.02]
R^2	14.60	6.24	10.23	22.59	15.35	11.21

Notes: For each asset bucket k, the Table reports the regression estimates for the following Equation:

$$\Delta p_{it,PC} = a_k + \lambda_{1k}\, \mathbb{1}_{\text{1st half}}\, \Delta p_{it,ST} + \lambda_{2k}\, \mathbb{1}_{\text{2nd half}}\, \Delta p_{it,ST} + \epsilon_{it},$$

where $\Delta p_{it,PC}$ is the change in asset i's price during date-t press conference, $\Delta p_{it,ST}$ is the change in asset i's price around date-t FOMC statement release, and asset i belongs to bucket k. We double cluster the standard errors at the date-asset level. The dummy variable $\mathbb{1}_{\text{1st half}}$ is equal to 1 in the first half of the sample per each Chairman, and zero otherwise, while $\mathbb{1}_{\text{2nd half}}$ is defined in the opposite way. The two price changes, in basis points, are computed over two non-overlapping, non-consecutive time intervals. T-statistics are in square brackets. Standard errors are double clustered at date-asset level. R-squared statistics are in percentage.

Sixth, we explore a standard Bayesian learning model with parameter uncertainty, á la Lewellen and Shanken (2002). Given that the hypothesis that there is learning across FOMC meetings is inconsistent with our data, here we hypothesize that every meeting is separate from the others, and investors' learning happens only within a single FOMC day. Suppose we are in the world sketched by Lewellen and Shanken (2002). For a better correspondence, we use the same notation they used, which means that we describe only how learning impacts stock prices, but the discussion can be easily generalized.

There is a riskless asset which pays a real rate r, and one risky security paying real dividend d_t, i.i.d. over time and drawn from a normal with mean δ and variance σ^2. Investors know the variance but they don't know the mean of the dividend distribution. They have some prior beliefs, centered around some δ^* with variance σ^2/h, where h is a measure of prior information (equivalent to a sample of h dividends). They update their beliefs using Bayes' rule, incorporating the information in observed dividends. With this prior, the investor's belief at time t about dividends at time $t+1$ is

$$d_{t+1} \sim \mathcal{N}\left(\frac{t}{t+h}\bar{d}_t + \frac{h}{t+h}\delta^*, \frac{t+h+1}{t+h}\sigma^2\right), \tag{14}$$

where $\bar{d}_t$ is the average dividend observed up to t. Investors are born with a constant absolute risk aversion utility with risk-aversion parameter $\gamma \geq 0$. Under the true distribution, the price of the risky asset can be shown to be

$$p_t = \frac{1}{r}\left(\frac{t}{t+h}\bar{d}_t + \frac{h}{t+h}\delta^*\right) - 2\gamma f(t+h)\sigma^2, \tag{15}$$

where $f(t)$ is a deterministic function of time. The function $f(t)$ decreases as t passes and converges to $1/r$ in the limit.[27]

Prices, and price changes, contain two terms. The first one reflects the updates in beliefs about expected dividends. The second one arises because estimation risk declines steadily over time. Consider two cases. In a framework in which investors are risk-neutral (γ is zero), the second term disappears. Because of learning, past mistakes tend to reverse and price revisions to be negatively autocorrelated. It is only when the prior is very far from the true value, and investors are confident in their wrong prior that a positive autocorrelation in price revisions may arise. This is not exactly an informational friction, but, under this hypothesis, investors must suffer from a substantial behavioral bias, starting from really wrong priors most of the times.

A model in which investors are risk-averse, and γ is positive does generate a positive covariance more easily, but at the expense of price dynamics. In fact in such a model, because estimation risk declines with time, prices tend to drift up. This is again inconsistent with our observations: we observe positive and negative shocks with almost the same frequency (e.g. for stocks), and, if anything, we observe larger shocks when shocks are negative. So, a model featuring the price drift that Equation (15) implies is also hard to square with our findings.

To sum up, we do believe there is an ongoing learning process during the press conference with journalists asking for clarifications and explanations. However, a Bayesian learning model without noisy signals for individual traders does not seem able to explain our findings.

---

[27] The function $f(t)$ takes the form $f(t) = \sum_{k=1}^{\infty} \frac{1}{(1+r)^k}\left(1 + \frac{1}{r(t+k)}\right)^2 \frac{t+k}{t+k-1}$

6 Concluding Remarks

Our paper posits a novel machinery to study financial investors' expectation formation process in events available for public observation. We apply recent advances in machine learning to scrape the videos of post-FOMC-meeting press conferences, extract the words, and timestamp these words at the millisecond. We then align the transcripts with high-frequency data for a wide range of financial assets. To the best of our knowledge, this is the first paper to add the time dimension to textual analysis, and study agent's expectation formation process at such a level of granularity.

At the moment the Chairman discusses the changes between the current and the previous policy statement, price volatility and trading volume spike dramatically, and prices move on average in the same direction as they did around the statement release, i.e. 30 minutes before the start of the press conference. This generates a strong positive correlation between price changes around the statement release and the subsequent press conference.

We have discussed a number of potential driving forces behind our results. We have examined explanations ranging from model parameters uncertainty to the *Fed put*, microstructure effects or liquidity, or learning about the Chairman's type, political independence, etc.... We have argued that it is difficult to square these explanations with all of the empirical evidence. On the other hand, the models most directly consistent with our results are models of information rigidity.

Our approach does not only improve on the identification of the effect of words on financial investors' beliefs, but also extends the set of questions that can be asked. The recipe we have developed can find applications in numerous settings where someone wants to bridge linguistics with economics using market prices. Besides, our findings suggest that noisy signals, or some other device to generate information rigidity are an important element for producing macroeconomic models with realistic information diffusion mechanisms that can capture asset price movements as well.

References

Altavilla, C., Brugnolini, L., Gürkaynak, R. S., Motto, R., Ragusa, G., 2019. Measuring euro area monetary policy. Journal of Monetary Economics 108, 162 – 179.

Altavilla, C., Brugnolini, L., Gürkaynak, R. S., Motto, R., Ragusa, G., 2020. How do financial markets react to monetary policy signals? Research Bulletin 73.

Amodei, D., Ananthanarayanan, S., Anubhai, R., Bai, J., Battenberg, E., Case, C., Casper, J., Catanzaro, B., Cheng, Q., Chen, G., et al., 2016. Deep speech 2: End-to-end speech recognition in english and mandarin. In: *International conference on machine learning*, pp. 173–182.

Barro, R. J., Gordon, D. B., 1983a. A positive theory of monetary policy in a natural rate model. Journal of Political Economy 91, 589–610.

Barro, R. J., Gordon, D. B., 1983b. Rules, discretion and reputation in a model of monetary policy. Journal of Monetary Economics 12, 101 – 121.

Bauer, M. D., Swanson, E. T., 2020. The fed's response to economic news explains the fed information effect. Working Paper 27013, National Bureau of Economic Research.

Berkelmans, L., 2011. Imperfect information, multiple shocks, and policy's signaling role. Journal of Monetary Economics 58, 373–386.

Bernanke, B. S., Kuttner, K. N., 2005. What explains the stock market's reaction to federal reserve policy? The Journal of Finance 60, 1221–1257.

Blinder, A. S., Ehrmann, M., Fratzscher, M., de Haan, J., Jansen, D.-J., 2008. Central bank communication and monetary policy: A survey of theory and evidence. Journal of Economic Literature 46, 910–945.

Bordalo, P., Gennaioli, N., Ma, Y., Shleifer, A., 2020. Overreaction in macroeconomic expectations. American Economic Review 110, 2748–82.

Born, B., Ehrmann, M., Fratzscher, M., 2014. Central bank communication on financial stability. The Economic Journal 124, 701–734.

Brooks, J., Katz, M., Lustig, H., 2018. Post-fomc announcement drift in u.s. bond markets. Working Paper 25127, National Bureau of Economic Research.

Calvo, G. A., 1978. On the time consistency of optimal policy in a monetary economy. Econometrica 46, 1411–1428.

Carroll, C. D., 2003. Macroeconomic Expectations of Households and Professional Forecasters. The Quarterly Journal of Economics 118, 269–298.

Cieslak, A., 2018. Short-Rate Expectations and Unexpected Returns in Treasury Bonds. The Review of Financial Studies 31, 3265–3306.

Cieslak, A., Morse, A., Vissing-Jorgensen, A., 2019. Stock returns over the fomc cycle. The Journal of Finance 74, 2201–2248.

Cieslak, A., Schrimpf, A., 2019. Non-monetary news in central bank communication. Journal of International Economics 118, 293–315.

Cieslak, A., Vissing-Jorgensen, A., 2020. The Economics of the Fed Put. The Review of Financial Studies (forthcoming).

Cochrane, J. H., Piazzesi, M., 2002. The fed and interest rates - a high-frequency identification. American Economic Review 92, 90–95.

Coibion, O., Gorodnichenko, Y., 2012. What can survey forecasts tell us about information rigidities? Journal of Political Economy 120, 116 – 159.

Coibion, O., Gorodnichenko, Y., 2015. Information rigidity and the expectations formation process: A simple framework and new facts. The American Economic Review 105, 2644–2678.

Coibion, O., Gorodnichenko, Y., Kamdar, R., 2018. The formation of expectations, inflation, and the phillips curve. Journal of Economic Literature 56, 1447–91.

Cook, T., Hahn, T., 1989. The effect of changes in the federal funds rate target on market interest rates in the 1970s. Journal of Monetary Economics 24, 331 – 351.

Cukierman, A., Meltzer, A. H., 1986. A theory of ambiguity, credibility, and inflation under discretion and asymmetric information. Econometrica 54, 1099–1128.

Cukierman, A., Wachtel, P., 1979. Differential inflationary expectations and the variability of the rate of inflation: Theory and evidence. American Economic Review 69, 595–609.

Ellingsen, T., Soderstrom, U., 2001. Monetary policy and market interest rates. American Economic Review 91, 1594–1607.

Fleming, M. J., Piazzesi, M., 2005. Monetary policy tick-by-tick. Working paper, Stanford University.

Gentzkow, M., Kelly, B., Taddy, M., 2019a. Text as data. Journal of Economic Literature 57, 535–74.

Gentzkow, M., Shapiro, J. M., 2010. What drives media slant? evidence from u.s. daily newspapers. Econometrica 78, 35–71.

Gentzkow, M., Shapiro, J. M., Taddy, M., 2019b. Measuring group differences in high-dimensional choices: Method and application to congressional speech. Econometrica 87, 1307–1340.

Graves, A., Fernández, S., Gomez, F., Schmidhuber, J., 2006. Connectionist temporal classification: labelling unsegmented sequence data with recurrent neural networks. In: *Proceedings of the 23rd international conference on Machine learning*, pp. 369–376.

Gürkaynak, R. S., Sack, B., Swanson, E., 2005. Do Actions Speak Louder Than Words?

The Response of Asset Prices to Monetary Policy Actions and Statements. International Journal of Central Banking 1, 55 – 93.

Hannun, A. Y., Case, C., Casper, J., Catanzaro, B., Diamos, G., Elsen, E., Prenger, R., Satheesh, S., Sengupta, S., Coates, A., Ng, A. Y., 2014. Deep speech: Scaling up end-to-end speech recognition. CoRR abs/1412.5567.

Hansen, S., McMahon, M., 2016. Shocking language: Understanding the macroeconomic effects of central bank communication. Journal of International Economics 99, S114–S133.

Hansen, S., McMahon, M., Prat, A., 2017. Transparency and Deliberation Within the FOMC: A Computational Linguistics Approach. The Quarterly Journal of Economics 133, 801–870.

Hassan, T. A., Hollander, S., van Lent, L., Tahoun, A., 2019. Firm-Level Political Risk: Measurement and Effects. The Quarterly Journal of Economics 134, 2135–2202.

Hong, H., Stein, J. C., 1999. A Unified Theory of Underreaction, Momentum Trading, and Overreaction in Asset Markets. Journal of Finance 54, 2143–2184.

Huberman, G., Regev, T., 2001. Contagious speculation and a cure for cancer: A nonevent that made stock prices soar. The Journal of Finance 56, 387–396.

Jarociński, M., Karadi, P., 2020. Deconstructing monetary policy surprises—the role of information shocks. American Economic Journal: Macroeconomics 12, 1–43.

Kim, O., Verrecchia, R., 1991. Trading volume and price reactions to public announcements. Journal of Accounting Research 29, 302–321.

Kuttner, K. N., 2001. Monetary policy surprises and interest rates: Evidence from the fed funds futures market. Journal of Monetary Economics 47, 523 – 544.

Kydland, F. E., Prescott, E. C., 1977. Rules rather than discretion: The inconsistency of optimal plans. Journal of Political Economy 85, 473–491.

Leombroni, M., Vedolin, A., Venter, G., Whelan, P., 2020. Central Bank Communication and the Yield Curve. Journal of Financial Economics (forthcoming).

Lewellen, J., Shanken, J., 2002. Learning, asset-pricing tests, and market efficiency. The Journal of Finance 57, 1113–1145.

Loughran, T., McDonald, B., 2011. When is a liability not a liability? textual analysis, dictionaries, and 10-ks. The Journal of Finance 66, 35–65.

Lucas, R. E., 1972. Expectations and the neutrality of money. Journal of Economic Theory 4, 103 – 124.

Lucas, R. E., 1973. Some international evidence on output-inflation tradeoffs. The American Economic Review 63, 326–334.

Lucas, R. E., 1976. Econometric policy evaluation: A critique. Carnegie-Rochester Conference Series on Public Policy 1, 19 – 46.

Lucca, D. O., Moench, E., 2015. The pre-fomc announcement drift. The Journal of Finance 70, 329–371.

Mankiw, N. G., Reis, R., 2002. Sticky Information versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve. The Quarterly Journal of Economics 117, 1295–1328.

Melosi, L., 2016. Signalling Effects of Monetary Policy. The Review of Economic Studies 84, 853–884.

Milgrom, P., Stokey, N., 1982. Information, trade and common knowledge. Journal of Economic Theory 26, 17 – 27.

Morris, S., 1994. Trade with heterogeneous prior beliefs and asymmetric information. Econometrica 62, 1327–47.

Mueller, P., Tahbaz-Salehi, A., Vedolin, A., 2017. Exchange rates and monetary policy uncertainty. The Journal of Finance 72, 1213–1252.

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

Neuhierl, A., Weber, M., 2019. Monetary policy communication, policy slope, and the stock market. Journal of Monetary Economics 108, 140–155.

Ottaviani, M., Sørensen, P. N., 2015. Price reaction to information with heterogeneous beliefs and wealth effects: Underreaction, momentum, and reversal. American Economic Review 105, 1–34.

Pennebaker, J., Boyd, R., Jordan, K., Blackburn, K., 2015. The development and psychometric properties of LIWC2015. University of Texas at Austin.

Porter, M. F., 1980. An algorithm for suffix stripping. Program 14, 130–137.

Reis, R., 2006. Inattentive consumers. Journal of Monetary Economics 53, 1761–1800.

Shalen, C. T., 1993. Volume, Volatility, and the Dispersion of Beliefs. The Review of Financial Studies 6, 405–434.

Stein, J. C., 1989. Cheap talk and the fed: A theory of imprecise policy announcments. Amercian Economic Review 79, 32–42.

Swanson, E., 2020. Measuring the effects of federal reserve forward guidance and asset purchases on financial markets. Working paper, University of California, Irvine.

Swanson, E. T., 2017. Measuring the effects of federal reserve forward guidance and asset purchases on financial markets. Working Paper 23311, National Bureau of Economic Research.

Tetlock, P. C., 2007. Giving content to investor sentiment: The role of media in the stock market. Journal of Finance 62, 1139–1168.

Woodford, M., 2003. Imperfect common knowledge and the effects of monetary policy. In: Aghion, P., Frydman, R., Stiglitz, J., Woodford, M. (eds.), *Knowledge, Information, and Expectations in Modern Macroeconomics: In Honor of Edmund S. Phelps*, Princeton University Press, Princeton, NJ.

Appendix A Example Statement News

We report in red the words that were present in the previous statement and that got removed in the new statement. We report in green the words added relative to previous statement. We highlight the words to which the question reported below the statement refers.

A.1 January 2019

A.1.1 Statement

Information received since the Federal Open Market Committee met in December November indicates that the labor market has continued to strengthen and that economic activity has been rising at a solid strong rate. Job gains have been strong, on average, in recent months, and the unemployment rate has remained low. Household spending has continued to grow strongly, while growth of business fixed investment has moderated from its rapid pace earlier last in the year. On a 12-month basis, both overall inflation and inflation for items other than food and energy remain near 2 percent. Although market-based measures of inflation compensation have moved lower in recent months, survey-based measures Indicators of longer-term inflation expectations are little changed, on balance.Consistent with its statutory mandate, the Committee seeks to foster maximum employment and price stability. In support of these goals, The Committee decided to maintain judges that some further gradual increases in the target range for the federal funds rate at 2-1/4 to 2-1/2 percent. The Committee continues to view will be consistent with sustained expansion of economic activity, strong labor market conditions, and inflation near the Committee's symmetric 2 percent objective as the most likely outcomes over medium term. The Committee judges that risks to the economic outlook are roughly balanced, but will continue to monitor In light of global economic and financial developments and muted inflation pressures, assess their implications for the economic outlook. In view of realized and expected labor market conditions and inflation, the Committee will be patient as it determines what future adjustments decided to raise the target range for the federal funds rate may be appropriate to support these outcomes 2-1/4 to 2-1/2 percent. In determining the timing and size of future adjustments to the target range for the federal funds rate, the Committee will assess realized and expected economic conditions relative to its maximum employment objective and its symmetric 2 percent inflation objective. This assessment will take into account a wide range of information, including measures of labor market conditions, indicators of inflation pressures and inflation expectations, and readings on financial and international developments.

A.1.2 Question

Second question in the press conference:

HEATHER LONG. Heather Long from the Washington Post. Last week, the IMF said risks are clearly skewed to the downside for the U.S. and global economy. Can you clarify-does the FOMC see risks as skewed to the downside, particularly after you removed the statement about risks being balanced?

CHAIRMAN POWELL. We had an extensive discussion of the baseline and also of the risks to the baseline, and the risks are, of course, the fact that financial conditions have tightened, that global growth has slowed, as well as some, let's say, government-related risks like Brexit and trade discussions, and also the effects and ultimate disposition of the shutdown. So we looked at-we look at those, and the way we think of it is that policy-we will use our policy, and we have, to offset risks to the baseline. So we view the baseline as still solid, and part of that is the way we adjusted our baseline to address those risks. So that's the way we're thinking about that now.

A.2 January 2020

A.2.1 Statement

Information received since the Federal Open Market Committee met in December October indicates that the labor market remains strong and that economic activity has been rising at a moderate rate. Job gains have been solid, on average, in recent months, and the unemployment rate has remained low. Although household spending has been rising at a moderate strong pace, business fixed investment and exports remain weak. On a 12-month basis, overall inflation and inflation for items other than food and energy are running below 2 percent. Market-based measures of inflation compensation remain low; survey-based measures of longer-term inflation expectations are little changed.Consistent with its statutory mandate, the Committee seeks to foster maximum employment and price stability. The Committee decided to maintain the target range for the federal funds rate at 1-1/2 to 1-3/4 percent. The Committee judges that the current stance of monetary policy is appropriate to support sustained expansion of economic activity, strong labor market conditions, and inflation returning to near the Committee's symmetric 2 percent objective. The Committee will continue to monitor the implications of incoming information for the economic outlook, including global developments and muted inflation pressures, as it assesses the appropriate path of the target range for the federal funds rate. In determining the timing and size of future adjustments to the target range for the federal funds rate, the Committee will assess realized and expected economic conditions relative to its maximum employment objective and its symmetric 2 percent inflation objective. This assessment will take into account a wide range of information, including measures of labor market conditions, indicators of inflation pressures and inflation expectations, and readings on financial and international developments.

A.2.2 Question

First question in the press conference:

CHRISTOPHER CONDON. Thank you. Chris Condon, Bloomberg News. Mr. Chairman, I would like you to comment in a little bit more depth about one small change I've noted in the statement. It notes that policy will be appropriate to bring-the Committee believes-inflation back to the Committee's 2 percent symmetric inflation objective. That's a slight change from the last time, when you were expecting it to bring inflation outcomes back near the objective. And I would put this also in the context of a comment you made at the last press conference where you drew attention to the fact that a number of *policymakers had projected inflation overshoots two or three years out under appropriate monetary policy. Should we take all of this together to mean simply that the Committee is more confident that a 2 percent outcome for inflation is already baked in the cake, or that this is a signal that the Committee has a stronger resolution to bring inflation at least to the 2 percent objective and put-bring into play an informal makeup strategy for inflation?*

CHAIRMAN POWELL. Yes. So, in making that change, our goal was, really-that was, changing "near" to "returning to"-was to avoid possible misinterpretation. So you may remember, in the December minutes we noted that a few Committee members suggested that the language that stated that monetary policy would support inflation near 2 percent could be misinterpreted as suggesting that policymakers were comfortable with inflation running below that level. So we thought about that in the intermeeting period and concluded that it would be appropriate to adjust that language to send a clearer signal that we're not comfortable with inflation rising persistently-running persistently below our 2 percent symmetric objective. So, yes, there is something in that. It's just that we wanted to underscore our commitment to 2 percent not being a ceiling to inflation running around-symmetrically around 2 percent, and that we're not satisfied with inflation running below 2 percent, particularly at a time such as now where we're a long way into an expansion and a long way into a period of very low unemployment when, in theory, inflation should be moving up.

Appendix B Complementary Dictionary

Table B.1. **Complementary Dictionary**

Press conference	Statement
mandate consistent	consistent committee's dual mandate
moderate-growth	moderate pace economic growth
statutory mandate	dual mandate
inflation goal	inflation run level consistent committee's dual mandate
economic outlook	economic growth
short-term securities	treasury security remaining maturity approximately 3 year less
maintain accommodation	maintain highly accommodative stance
highly accommodative policy stance	highly accommodative stance
the exit strategy put consistent statement today	begin remove policy accommodation
inflation readings	inflation running
fiscal issues	fiscal retrenchment
unemployment come down	improvement labor conditions
weakness economy	strength broader economy
hold funds rate	keeping target fund rate level
end bond purchases	asset purchase program end
decided make another reduction pace asset purchases	committee end current program asset purchase
inflation fomc's objectives	inflation running committee's longer-run objective
normalizing policy	committee end current program asset purchase
raise funds rate target range	increase fund rate
2 percent target	2 percent objective
despite risks abroad	global economic development pose risks
labor expected tighten	strengthening labor market
risks outlook	global economic development pose risks
global economic developments	net export
economy growing roughly trend	economic activity expand moderate pace
labor conditions continued improve	labor indicator strengthen
wage growth	labor indicator strengthen
case rate increase strengthened	case increase fund rate strengthened
bring inflation back 2 percent	inflation stabilize around 2 percent
undertake beginning plan	begin implementing balance sheet normalization program
broader measures labor utilization continued strengthen	solid labor condition remain strong
raising target range	increase target range
gradual increases rate	expects gradual increase target range
conditions likely call three rate increases	committee judge gradual increase target range
future adjustment policy	adjustment target range fund rate appropriate
policy stance appropriate	adjustment target range fund rate appropriate
weakness global growth trade	export weakened
particularly weak global growth trade	export weakened
understanding word "appropriate"	appropriate path target range fund rate
"appropriate" statement	appropriate path target range fund rate
believes-inflation back committee's 2 percent symmetric inflation objective	inflation returning committee's symmetric 2 percent objective
stance monetary policy remains accommodative	monetary policy accommodative
core inflation running 2	inflation declined running 2 percent

Notes: The Table reports the dictionary that we use to complement the algorithm described in Section 4.1 to identify policy-statement news within the press conference.

Appendix C Variation in FOMC days

We separate time in four non-overlapping periods. The pre-statement period starts 45 minutes before the FOMC statement release and ends one second before the FOMC statement release. The statement period starts with the FOMC statement release and ends one second before the FOMC press conference starts. The press conference period considers the entire duration of the press conference. The post press conference period starts one second after the FOMC press conference and ends 45 minutes after.

For each FOMC press conference day and each of the sub-periods we compute the ratio between standard deviation and average (coefficient of variation) of minutely-level changes in the rates implied from federal funds and Eurodollar futures contracts. Figure C.1 displays the distribution across all FOMC press conferences of the coefficient of variation so computed for each of the four sub-periods. Panel A shows the results for federal funds rate futures, while Panel B for Eurodollar futures. They both include a large range of maturities.

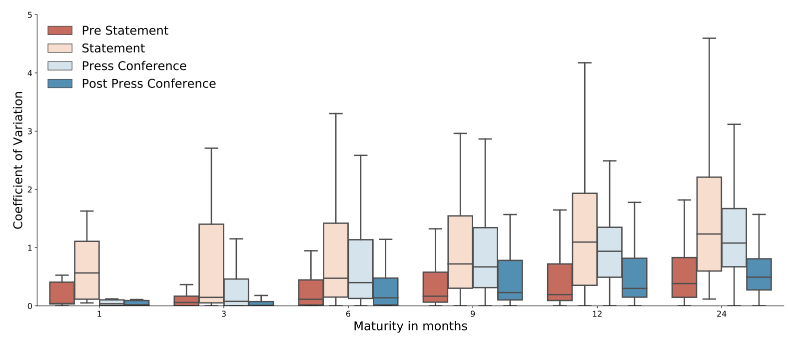

Panel A –Federal funds futures

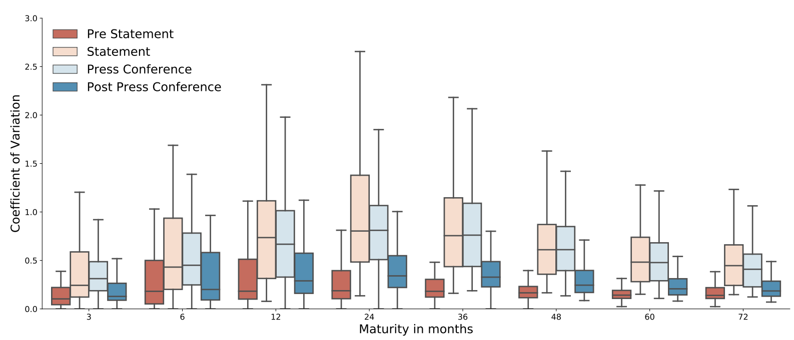

Panel B – Eurodollar futures

Fig. C.1. *Notes*: The Figure shows the distribution of the coefficient of variation (CV), i.e. the ratio of standard deviation to average value, for the Federal funds futures and Eurodollar futures for different expiration dates. We present the CV distribution over four non-overlapping sub-periods. The sample consists of all 41 FOMC meetings containing a press conference from January 2011 to January 2020.

Appendix D Additional Figures and Tables

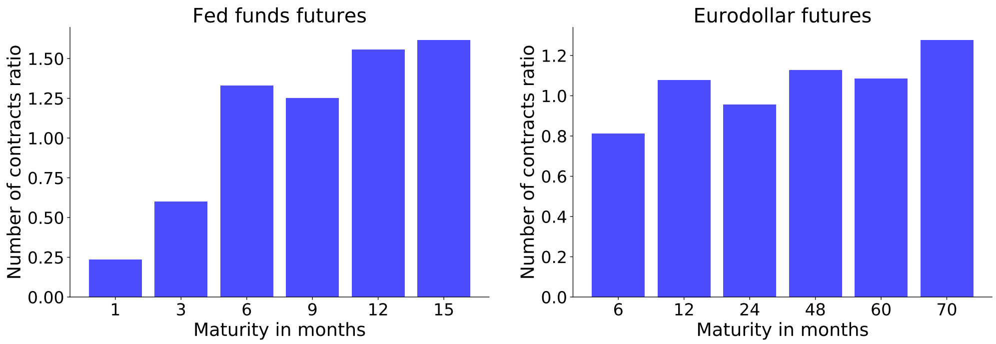

Fig. D.1. *Notes*: For each FOMC day with a press conference we compute the ratio between the trading volume (number of traded contracts) during the press conference and the trading volume around the statement release. The windows to compute price changes are the same as in Section 3. The Figure shows the average ratio across all FOMC days with a press conference for several contract maturities. The left Panel reports results for the 30-day Federal funds futures, while the right Panel reports results for Eurodollar futures.

Table D.1. **Press-conference shocks against statement shocks – Federal funds futures, Eurodollar futures**

Federal funds futures				Eurodollar futures			
Maturity	a	λ	R^2	Maturity	a	λ	R^2
1 months	0.088	0.044	2.916	3 months	-0.117	0.338	37.579
	[0.700]	[0.775]	–		[-0.493]	[3.801]	–
3 months	-0.003	0.293	24.138	6 months	-0.459	0.335	33.177
	[-0.020]	[3.385]	–		[-1.475]	[3.071]	–
6 months	-0.081	0.262	13.915	12 months	0.221	0.174	6.162
	[-0.300]	[2.446]	–		[0.428]	[1.600]	–
9 months	0.082	0.201	6.995	24 months	0.114	0.165	5.439
	[0.208]	[1.691]	–		[0.164]	[1.498]	–
12 months	0.218	0.180	5.808	36 months	0.221	0.277	18.319
	[0.456]	[1.531]	–		[0.316]	[2.957]	–
15 months	0.199	0.160	4.700	48 months	0.190	0.386	30.064
	[0.346]	[1.333]	–		[0.302]	[4.095]	–
18 months	0.344	0.138	3.731	60 months	0.078	0.497	33.450
	[0.514]	[1.131]	–		[0.133]	[4.427]	–
24 months	0.190	0.061	0.697	70 months	0.010	0.510	27.055
	[0.150]	[0.325]	–		[0.016]	[3.704]	–

Notes: For each asset bucket k, the Table reports the regression estimates for the following Equation:

$$\underbrace{\Delta p_{it,PC}}_{\Delta\text{p at press conference: e.g. 14:30–15:30}} = a_k + \lambda_k \underbrace{\Delta p_{it,ST}}_{\Delta\text{p around statement: 13:50–14:20}} + \epsilon_{it},$$

where $\Delta p_{it,PC}$ is the change in asset i's price during date-t press conference, $\Delta p_{it,ST}$ is the change in asset i's price around date-t FOMC statement release, and asset i belongs to bucket k. The two price changes, in basis points, are computed over two non-overlapping, non-consecutive time intervals. T-statistics are in square brackets. Standard errors are double clustered at date-asset level. R-squared statistics are in percentage.

Table D.2. **Press-conference shocks against statement shocks – Stock portfolios, Foreign Exchange rates**

Stocks				Foreign Exchange			
Portfolio	a	λ	R^2	- vs. usd	a	λ	R^2
SPY	-7.115	0.515	19.635	aud	-8.005	0.231	9.086
	[-0.988]	[3.087]	–		[-1.390]	[1.949]	–
Mining	-6.236	0.333	8.555	eur	-6.303	0.225	9.954
	[-0.517]	[1.885]	–		[-1.415]	[2.050]	–
Utilities	-4.133	0.387	27.675	gbp	-2.269	0.252	13.157
	[-0.641]	[3.712]	–		[-0.637]	[2.399]	–
Manufacturing	-3.299	0.524	19.962	nzd	-7.771	0.235	9.851
	[-0.403]	[3.079]	–		[-1.316]	[2.038]	–
Fabricated Metal	-8.296	0.578	26.110	chf	-5.813	0.271	13.231
	[-1.054]	[3.517]	–		[-1.313]	[2.407]	–
Retail	-3.310	0.403	12.041	jpy	-2.099	0.323	16.099
	[-0.471]	[2.281]	–		[-0.497]	[2.700]	–
Information	-3.402	0.445	14.275	cad	-5.193	0.361	17.491
	[-0.406]	[2.516]	–		[-1.230]	[2.838]	–
Finance and Insurance	-1.758	0.243	3.650				
	[-0.200]	[1.200]	–				

Notes: For each asset bucket k, the Table reports the regression estimates for the following Equation:

$$\underbrace{\Delta p_{it,PC}}_{\Delta\text{p at press conference: e.g. 14:30–15:30}} = a_k + \lambda_k \underbrace{\Delta p_{it,ST}}_{\Delta\text{p around statement: 13:50–14:20}} + \epsilon_{it},$$

where $\Delta p_{it,PC}$ is the change in asset i's price during date-t press conference, $\Delta p_{it,ST}$ is the change in asset i's price around date-t FOMC statement release, and asset i belongs to bucket k. The two price changes, in basis points, are computed over two non-overlapping, non-consecutive time intervals. T-statistics are in square brackets. Standard errors are double clustered at date-asset level. R-squared statistics are in percentage.

Table D.3. **Regression estimates: Placebo dates – FOMC days without a press conference**

	Federal funds futures		Eurodollar futures	
	1m-6m	9m-15m	6m-12m	24m-70m
a	0.16	0.15	-0.04	-0.18
	[2.17]	[2.15]	[-0.16]	[-0.70]
λ	0.08	0.05	0.01	0.04
	[1.49]	[1.19]	[0.16]	[0.33]
R^2	0.35	1.30	0.01	0.55

Notes: In this Table, we repeat the analysis of Table 3, but for a "placebo" event period: we use FOMC days without a press conference and compute price changes around the statement release, and around an alternative window that mimicks the average press conference time (from 14:30 to 15:24 EST). The columns in the Table report the estimates of the following Equation for each asset bucket k:

$$\underbrace{\Delta p_{it,AW}}_{\Delta\text{p in alternative window}} = a_k + \lambda_k \underbrace{\Delta p_{it,ST}}_{\Delta\text{p around statement: 13:50–14:20}} + \epsilon_{it},$$

where $\Delta p_{it,AW}$ is the change in asset i's price during date-t alternative window mimicking the press conference time, and $\Delta p_{it,ST}$ is the change in asset i's price around date-t FOMC statement release. The two price changes are computed over two non-overlapping, non-consecutive time intervals. T-statistics are in square brackets. Standard errors are double clustered at date-asset level. R-squared statistics are in percentage.

Appendix E Linguistic analysis of the statement message

In the main text, we have shown how financial investors perceive statement-related messages. In this section, we analyse the linguistics of the messages directly. The goal is to identify the different language patterns, and styles that characterize the minutes in which the Chairman talks about the statement. We use a rather standard word-count strategy. The search words are categorized into language categories following the Linguistic Inquiry and Word Count (LIWC) by Pennebaker et al. (2015).

For each minute, we count the words in a given category and divide it by the total number of words in that minute. We then regress the frequency variable onto the dummy variable contructed as in Section 4.1. Let D_t be such a dummy, and Freq_{it} the frequency value for the semantic category i in minute t. We estimate the following Equation:

$$\text{Freq}_{it} = a + bD_t + \epsilon_{it}. \tag{E.1}$$

The intercept represents the average frequency in all non-statement-related minutes. The slope coefficient represents by how much the frequency value changes on average in statement-related minutes. Table E.1 reproduces the estimates.

First, statement-related minutes exhibit fewer negations (16.5% less). From a linguistic perspective this is important. Psychology literature suggests that there is a fundamental asymmetry between negative and affirmative propositions in natural language: negative sentences are less valuable than affirmative ones, less specific and less informative. The fewer negations may suggest a larger informational value of statement-related minutes.

Second, we find more comparison words such as “than” or “as” (about 10% more) and numbers (almost 20% more number usage). Beyond providing more tangible information, comparisons and numbers serve to give the sense of the ideas expressed. Besides, numbers are often considered a neutral, and transparent sign of the reality.[28]

Third, both relativity words, such as “during” or “when,” and insight words, such as “consider,” “know”, or “think,” increase (18% and 14% respectively). For an intuition as to why this adds nuances to the statement, imagine the following sentence: “I think you are wrong in this instance.” Removing signals of insight, we have: “you are wrong in this instance,” and then removing relativity we have: “you are wrong.” Relativity and insight words tend to moderate the sentence meaning. They tend to reflect opinions rather than accepted truths, and qualify that opinion to a specific environment, case, time relative to which the statement holds.

Fourth and finally, in terms of the usage of words, it is interesting to notice that words describing motion (e.g. where the economy is *going*), or time (past, present and future) go up by 20% and 33%, respectively.

Given that time is an more important dimension of statement-related minutes, we analyse the tense, and the time focus of those minutes. The last three columns of Table E.1 show that in the minutes in which the statement is discussed, messages feature fewer words related to the past or present. On the other hand, the attention to the future

[28] The use of numbers has also been found to have a positive impact on the audience perceptions because it suggests competence and skills.

Table E.1. **Linguistic Analysis: part of the speech**

	Negations	Comparisons	Numbers	Insight	Relativity	Motion	Time	Past	Present	Future
a	0.97	2.48	1.66	2.30	11.84	1.92	3.50	2.43	16.9	1.69
	[46.68]	[71.65]	[30.68]	[67.73]	[141.97]	[63.59]	[70.57]	[59.88]	[146.14]	[53.92]
b	-0.16	0.24	0.31	0.42	1.65	0.38	1.15	-0.15	-1.93	0.37
	[-2.26]	[2.02]	[1.62]	[3.52]	[5.7]	[3.65]	[6.65]	[-1.05]	[-4.8]	[3.42]

Notes: The Table reports the regression estimates for the following Equation:

$$\text{Freq}_{it} = a + bD_t + \epsilon_{it},$$

where Freq_{it} is the frequency value for words in category i in minute t, and D_t is the dummy variable constructed as in Section 4.1. T-statistics are in square brackets. The search words are categorized into language categories following the Linguistic Inquiry and Word Count (LIWC) by Pennebaker et al. (2015). Each column corresponds to one category.

(captured again by the word frequency variable) goes up by 22%.

To sum up, during statement-related minutes messages talk more about the future, and less about the present or past, have a larger informational value captured by fewer negations, more comparisons, and more numbers, and also add qualifying phrases that may reflect the fact that the Chairman's tends to mention his own view or that may add nuances talking about a specific situation to which that statement applies.

Appendix F Proofs

Proof of Proposition 1. Conditional on $t-1$ information, the forecast of agent i about the state x_t is given by

$$F_{i,t-1}x_t = \rho F_{i,t-1}x_{t-1}. \tag{F.1}$$

After receiving signal $y_{i,t}$, agent i updates the previous estimate by incorporating this new information:

$$F_{i,t}x_t = F_{i,t-1}x_t + G\left(y_{i,t} - F_{i,t-1}x_t\right) \tag{F.2}$$

where G is the steady-state Kalman filter gain.

Given the prediction and update step described by equations (F.1) and (F.2) respectively, we can write revisions in expectations as

$$\begin{aligned}
F_{i,t}x_{t+h} - F_{i,t-1}x_{t+h} &= \rho^h\left[F_{i,t}x_t - F_{i,t-1}x_t\right] \\
&= \rho^h\left[Gy_{i,t} + (1-G)F_{i,t-1}x_t - F_{i,t-1}x_t\right] \\
&= \rho^h\left[Gy_{i,t} - GF_{i,t-1}x_t\right] \\
&= \rho^h\left[Gy_{i,t} - \rho GF_{i,t-1}x_{t-1}\right] \\
&= \rho^h\left[Gy_{i,t} - \rho G\left(Gy_{i,t-1} + (1-G)F_{i,t-2}x_{t-1}\right)\right] \\
&= \rho^h\left[Gy_{i,t} - \rho G\left(Gy_{i,t-1} + \rho(1-G)F_{i,t-2}x_{t-2}\right)\right] \\
&= \rho^h\left[Gy_{i,t} - \rho G^2 y_{i,t-1} - \rho^2 G(1-G)F_{i,t-2}x_{t-2}\right]
\end{aligned}$$

Defining $F_t x_{t+h} = E_i\left[F_{i,t}x_{t+h}\right]$ the average forecast across agents and using $E_i\left[\omega_{i,t}\right] = 0$, i.e. on average signals are able to fully recover the state x_t, it follows

$$F_t x_{t+h} - F_{t-1}x_{t+h} = \rho^h\left[Gx_t - \rho G^2 x_{t-1} - \rho^2 G(1-G)F_{t-2}x_{t-2}\right] \tag{F.3}$$

Using the law of motion for x_t in (2) and substituting in (F.3) we get:

$$\begin{aligned}
F_t x_{t+h} - F_{t-1}x_{t+h} &= \rho^h\left[G\left(\rho x_{t-1} + \nu_t\right) - \rho G^2 x_{t-1} - \rho^2 G(1-G)F_{t-2}x_{t-2}\right] \\
&= \rho^h\left[\rho G(1-G)x_{t-1} + G\nu_t - \rho^2 G(1-G)F_{t-2}x_{t-2}\right] \\
&= \rho^h(1-G)\left[\rho Gx_{t-1} - \rho^2 GF_{t-2}x_{t-2}\right] + \rho^h G\nu_t \\
&= (1-G)\left(\rho^h\left[\rho Gx_{t-1} - \rho^2 GF_{t-2}x_{t-2}\right]\right) + \rho^h G\nu_t \\
&= (1-G)\left(F_{t-1}x_{t+h} - F_{t-2}x_{t+h}\right) + \rho^h G\nu_t \\
&= \lambda\left(F_{t-1}x_{t+h} - F_{t-2}x_{t+h}\right) + \rho^h G\nu_t
\end{aligned} \tag{F.4}$$

where the penultimate line follows from

$$\begin{aligned}
F_{i,t-1}x_{t+h} - F_{i,t-2}x_{t+h} &= \rho^h\left(\rho F_{i,t-1}x_{t-1} - \rho^2 F_{i,t-2}x_{t-2}\right) \\
&= \rho^h\left(\rho Gy_{i,t-1} + \rho^2(1-G)F_{i,t-2}x_{t-2} - \rho^2 F_{i,t-2}x_{t-2}\right) \\
&= \rho^h\left(\rho Gy_{i,t-1} - \rho^2 GF_{i,t-2}x_{t-2}\right)
\end{aligned} \tag{F.5}$$

after averaging across agents we then have $F_{t-1}x_{t+h} - F_{t-2}x_{t+h} = \rho^h\left(\rho Gx_{t-1} - \rho^2 GF_{t-2}x_{t-2}\right)$.

□

Appendix G Relation between Kalman gain and signal-to-noise ratio

Table G.1. **Implied estimates of agent-specific noise**

	Kalman gain			Kalman gain		
	0.6	0.7	0.8	0.6	0.7	0.8
ρ	Signal-to-noise ratio			Coefficient of determination		
0.00	1.22	1.53	2.00	0.60	0.70	0.80
0.10	1.23	1.53	2.01	0.60	0.70	0.80
0.20	1.24	1.55	2.03	0.61	0.71	0.81
0.40	1.29	1.63	2.15	0.63	0.73	0.82
0.60	1.42	1.80	2.41	0.67	0.76	0.85
0.80	1.76	2.29	3.11	0.76	0.84	0.91
0.90	2.31	3.05	4.20	0.84	0.90	0.95
0.99	6.77	9.10	12.71	0.98	0.99	0.99

Notes: The table reports the signal-to-noise ratio and the coefficient of determination implied by the estimates of the Kalman gain G. The signal-to-noise ratio is defined as the variance of x_t relative to the variance of the noise in the private signal in equation (3), σ_ω^2, while the coefficient of determination is defined as the proportion of variance in y_{it} explained by the variance in x_t. Section G.1 in Appendix G provides the details on how to compute these values.

In this section we describe the relation between the Kalman gain and the signal-to-noise ratio, which we derive in Section G.1. Let φ_x denote the square root of the "signal-to-noise ratio", defined as the variance of x_t relative to the variance of the noise in the private signal in (3), σ_ω^2. The signal is more informative, the larger is the variance of the hidden state x_t relative to the private noise. For any combination of the Kalman gain and the autocorrelation parameter of the hidden state in (2), ρ, we solve for the value of the implied signal-to-noise ratio. We consider eight possible values for ρ, including values that imply i.i.d. hidden state as well as values for which x_t is highly persistent, while the Kalman gain estimates are from Table 3.

The limitation when using the signal-to-noise ratio is that it is not a priori clear how large the signal-to-noise ratio needs to be for a signal to be considered *informative*. So in Table G.1 we add the easier-to-interpret coefficient of determination, namely the proportion of variance in y_{it} explained by the variance in x_t. Equation (G.1) shows a one-to-one mapping between the signal-to-noise and the coefficient of determination:

$$\frac{Var(x_t)}{Var(y_{it})} = \frac{Var(x_t)}{Var(x_t) + \sigma_\omega^2} = \frac{\varphi_x^2}{\varphi_x^2 + 1}. \tag{G.1}$$

Table G.1 reports the values of the signal-to-noise ratios and the coefficient of determination implied by various combinations of ρ and the Kalman gain estimates from Table 3. Note that when ρ is equal to zero, the hidden state is a white noise process, and the coefficient of determination equals the Kalman gain. The coefficient of determination converges monotonically to 1 as ρ moves towards 1: the more the signals reflect transitory rather than permanent policy changes, the faster agents react to these signals by updating their expectations. So even if the monetary policy signal-to-noise ratio is high, and the variation in the signals can be attributed almost entirely to variation in fundamentals, agents will react to signals slowly if policy changes are long lasting. This suggests that, from a policy perspective, increasing the signal-to-noise ratio is not enough if investors perceive policy decisions to last longer. The degree of information rigidity summarizes both properties, providing a direct description of how fast information gets used in the financial markets.

G.1 Deriving the relation between G, φ_x, and ρ

Let G be the Kalman gain, φ_x the square-root of the signal-to-noise ratio. Assume expectations are noisy in the sense of (2) and (3). Moreover, call $y_{i,t}$ the private signal, x_t the hidden state, ρ the autocorrelation parameter of the hidden state, and σ_ω the standard deviation of the noise in the signal Equation (3).

The Kalman filtering equations for each step $t = 1, 2, 3, \dots$ are:

$$\begin{aligned} P_{t-1|t} &= \rho^2 P_{t-1|t-1} + \varphi_x^2(1-\rho^2)\sigma_\omega^2 \\ G_t &= P_{t-1|t}\left(P_{t-1|t} + \sigma_\omega^2\right)^{-1} \\ P_{t|t} &= (1 - G_t)P_{t-1|t} \\ x_{t-1|t} &= \rho x_{t-1|t-1} \\ x_{t|t} &= x_{t-1|t} + G_t\left(y_{i,t} - x_{t-1|t}\right) \end{aligned} \tag{G.2}$$

Note that $P_{t-1|t}$, G_t, and $P_{t|t}$ do not depend on the values of the private signal $y_{i,t}$, but only on the model parameters ρ, σ_ω, and φ_x (given $P_{0|0}$).

To provide the intuition on how φ_x affects the step-t Kalman gain G_t, it is useful to show the explicit solution for the first step:

$$\begin{aligned} P_{0|1} &= \rho^2 P_{0|0} + \varphi_x^2(1-\rho^2)\sigma_\omega^2 \\ &= \rho^2\varphi_x^2\sigma_\omega^2 + \varphi_x^2(1-\rho^2)\sigma_\omega^2 \\ &= \varphi_x^2\sigma_\omega^2 \end{aligned} \tag{G.3}$$

where we use the unconditional variance for x_t as a starting point for $P_{0|0} = E\left[(x_0 - x_{0|0})(x_0 - x_{0|0})\right]$. Given, $P_{0|1}$ the Kalman gain at time 1 is

$$G_1 = P_{0|1}\left(P_{0|1} + \sigma_\omega^2\right)^{-1} = \varphi_x^2\sigma_\omega^2\left(\varphi_x^2\sigma_\omega^2 + \sigma_\omega^2\right)^{-1} = \frac{\varphi_x^2}{\varphi_x^2 + 1}, \tag{G.4}$$

which is the coefficient of determination computed in Equation (G.1).

Manipulating the equations for $P_{t-1|t}$, G_t, and $P_{t|t}$ we obtain the following Ricatti equation:

$$P = \rho^2\left(P - P(P + \sigma_\omega^2)^{-1}P\right) + \varphi_x^2(1-\rho^2)\sigma_\omega^2. \tag{G.5}$$

Once, we solve for P we can substitute it for $P_{t-1|t}$ in the Kalman gain formula for G_t in

(G.2) to obtain the steady-state Kalman gain:

$$G = P\left(P + \sigma_\omega^2\right)^{-1}. \tag{G.6}$$

Unfortunately, G has no closed-form solution. However, by fixing ρ to some value $\bar{\rho}$, and normalizing σ_ω^2 to one, G becomes a function of φ_x only, for which can solve numerically.