The Perceived Causes of Monetary Policy Surprises

Miguel Acosta*
Federal Reserve Board

February 26, 2023

Abstract

High frequency changes in asset prices around Federal Reserve (Fed) announcements are the primary measures used for estimating the effects of monetary policy. Vis-à-vis standard theory, these measures exhibit puzzling effects on a range of macroeconomic outcomes. A debate has emerged over the source of these puzzling effects: Do they arise because the Fed simultaneously communicates about monetary policy and the economic outlook, or are they simply the result of an omitted variables bias? In this paper I decompose monetary policy announcements into a rich set of identified shocks. To sidestep the shortcomings of asset prices, I construct my shocks from the text of newspaper articles written about Fed announcements, and achieve structural identification using discrete changes in Fed communications. I find that shocks to current and expected future interest rates exhibit large, theoretically consistent effects on macroeconomic outcomes. I also find that shocks to beliefs about supply and demand factors have substantial effects even after controlling for omitted variables. This suggests that Fed announcements indeed reveal information about the economic outlook, an aspect of policy that is important to take into account.

Keywords: Monetary Policy, Communication, Natural Language Processing, High Frequency Identification

*Division of International Finance, Federal Reserve Board. miguel.acosta@frb.gov. I owe an enormous debt of gratitude to my advisors—Hassan Afrouzi, Jennifer La'O, Emi Nakamura, Jón Steinsson, and Michael Woodford—for numerous insightful discussions regarding this work. The paper has also benefited from discussions with Andrea Ajello, Paul Bouscasse, Lydia Cox, Emilien Gouin-Bonenfant, Juan Herreño, Isaiah Hull (discussant), Ellen Meade, José Luis Montiel-Olea, Christian Moser, Joseph Saia, Jesse Schreger, Yeji Sung, Pierre Yared, and seminar participants at Berkeley, Columbia, McGill, Michigan, the Federal Reserve Board, the Federal Reserve Banks of Boston and Dallas, the 14th International Conference on Computational and Financial Econometrics, the 2021 EEA Annual Congress, and the 2022 BSE Summer Forum. I thank the Federal Reserve Bank of San Francisco for support as I worked on this paper as a dissertation intern. I owe a special thanks to Yasmin Saira at the Business and Economics Library at Columbia for help accessing data, especially when physical access to Columbia's libraries was limited due to the pandemic. Hannah Bingley, Blaine Helleloid, and David Yu provided outstanding research assistance. This work was supported by financial support from an Alfred P. Sloan Foundation Pre-Doctoral Fellowship in Behavioral Macroeconomics awarded through the NBER. The views expressed here are mine only, and should not be interpreted as reflecting the views of the Board of Governors of the Federal Reserve System, or of any other person associated with the Federal Reserve System.

1 Introduction

The role of the modern-day Federal Reserve (Fed) has moved far beyond setting the level of the overnight interest rate. In the eyes of financial market participants, the Fed Chair is often seen as a fortune teller who communicates predictions regarding not only the path of future interest rates, but also forecasts of macroeconomic outcomes. An understanding of the Fed's role in the economy, then, crucially hinges on an understanding of the macroeconomic effects of both aspects of Fed communications. Two challenges arise when estimating these effects. First, communications regarding the path of interest rates and communications regarding the economic outlook are both highly endogenous with respect to economic fundamentals. Second, the Fed typically engages in both types of communications simultaneously, which complicates attempts to separately identify their effects.

In this paper I estimate the macroeconomic effects of Fed communications—both the effects of monetary policy (setting current interest rates, and explicitly communicating about future interest rates) and information provision policy (the effects of providing information about macroeconomic fundamentals). I contribute to a longstanding literature that studies the effects of monetary policy using market reactions to Fed policy announcements, and provide new estimates that overcome an important conceptual issue in the identification of these effects: Market-based measures of interest-rate expectations can respond to both types of policy, thereby identifying neither. I also provide evidence that information provision is an important component of the Fed's communication policy—rather than a statistical nuisance that challenges the identification of exogenous variation in monetary policy, shedding light on an important debate that has emerged in the literature.

To estimate the effects of both aspects of Fed communications on macroeconomic outcomes, I identify four new series of "perceived" shocks.[1] Two are monetary policy shocks: a traditional monetary policy shock (a shock to current rates), and a monetary news shock (a shock to the expected path of rates) The second two reflect communications about economic fundamentals: a perceived demand shock, and a perceived supply shock. I posit that an econometrician needs access to (at least) four measures of market reactions that respond differently to the four types of shocks in order to identify them separately. I show that short-term and long-term interest-rate forecast revisions, along with GDP and inflation forecast revisions emerge as natural candidates to accomplish the task, based on the implications of New Keynesian theory. Intuitively, because forecast revisions of these variables react differently in response to different shocks, observing their joint reactions to monetary policy announcements provides useful information for identifying the prevalence of each shock.[2]

To address the first empirical challenge—that Fed communications are endogenous with respect to economic fundamentals—I pursue a high-frequency approach to identifying changes

[1] I introduce the word "perceived" here because my shocks are based on data on expectations (as are estimates of monetary shocks in the literature that use interest rate futures data). I discuss this distinction in section 3.1.

[2] To give a familiar example, a contractionary monetary policy shock raises interest rates and lowers real GDP in standard New Keynesian models. Instead, a positive aggregate demand shock causes both variables to increase.

in expectations that are caused only by the Fed's announcements. The foundation of my shocks consists of four high-frequency measures of macroeconomic forecast revisions: short- and long-term interest rate forecast revisions, and new text-based measures of GDP and inflation forecast revisions. The first two are standard in the empirical monetary literature. The latter two I construct using newspaper articles written about each Fed policy meeting. At its core, my approach requires a mapping from newspaper text to numerical values corresponding to macroeconomic forecast revisions. I estimate this mapping by extracting features from the universe of *New York Times* articles written over the period 1987–2007, and use those features in a LASSO regression that predicts monthly changes in macroeconomic expectations from the Blue Chip survey. I then apply this mapping to articles written shortly before and shortly after each Fed policy announcement, allowing me to identify changes in expectations of GDP and inflation that are driven by the Fed's policy announcements, and not by other events that may cause the Fed's communications.

In order to recover the structural shocks from these forecast revisions, I must address the second empirical challenge: The Fed communicates simultaneously about a range of economic fundamentals. To that end, I estimate a simultaneous-equations model of how Fed watchers update their macroeconomic expectations in response to Fed announcements.[3] In the model, Fed watchers are forecasters whose model of the economy is a linear relationship between macroeconomic shocks and macroeconomic variables (here, interest rates, GDP, and inflation). Despite this simple formulation—a system of linear equations determined by exogenous shocks—this is the forecasting model implied by the dynamic stochastic general equilibrium models that permeate macroeconomic analysis.[4] This type of tight connection between identifying assumptions and an explicit theory is rare in the high-frequency identification literature.[5]

To identify the model, I use two discrete changes in the Fed's communication policy: the introduction of interest-rate forward guidance in 2003, and the 2014–2016 oil price decline that required the Fed to parse the nature of shocks in the economy. Starting with the first, as highlighted by Lunsford (2020), prior to 2003 the Fed's post-meeting policy statements primarily described the economic outlook. In August 2003 the Fed began the practice of interest-rate forward guidance when it promised to keep interest rates low "for a considerable period." To see how this can help identify the model, consider the implications of this policy change. Before 2003, learning about the Fed's economic outlook was straightforward; in contrast, inference about the path of interest rates was possible only indirectly through the Fed's discussion of the economy. This observation allows me to make my formal identification assumption: Fed announcements induced market participants to update their expectations about the future path of interest rates

[3]My forecast revision data is built from the expectations of market participants and newspaper correspondents that report about the Fed. I refer to these groups jointly as "Fed watchers."

[4]The model does not explicitly incorporate dynamics. As such, some of my identified shocks do not distinguish between current and expected future shocks. I show empirically and theoretically that this distinction matters little in my context.

[5]This is not to say that other work has ignored simultaneity, nor that theoretical evaluations of Fed communications are non-existent. Cieslak and Schrimpf (2019), Gürkaynak et al. (2021), Nakamura and Steinsson (2018), and Sastry (2021) all build models designed to interpret event-study data. Among these, only Cieslak and Schrimpf (2019) use implications of their model to design sign restrictions to identify shocks in an otherwise non-restrictive way.

more completely after 2003, relative to their expectations about economic fundamentals. My data support this assumption: From 1999–2003, interest-rate and GDP forecast revisions were weakly positively correlated: Intuitively, this regime primarily provides identification of the perceived demand shock. From 2003–2006, the series are negatively correlated: This latter regime provides identification of the monetary news shock by providing data whose variation is driven by a factor orthogonal to the information shock.[6] The third regime—from 2014–2016—encompasses a time when oil prices experienced a large decline. This episode received substantial attention in policy and academic circles, and resulted in Fed communications devoted to parsing the nature of the shock driving the decline. This regime allows me to achieve an estimate of what Fed watchers learn about supply factors. I show, however, that including this regime is not strictly necessary for identification, since Fed communications about supply factors were also present during the 2003–2006 regime. This approach is known as "identification by heteroskedasticity," first proposed by Rigobon (2003).[7] Formally, while the model is not identified within any regime separately, the additional assumption that the shocks have the same effects on observables across the regimes imposes enough parametric restrictions to jointly solve the model's implied moment conditions.[8] Note that while I estimate the model using a three-year window around the 2003 policy change and the 2014–2016 regime, I use the identified model to construct estimates of the four shocks over my entire 1999–2019 sample period.

I use a standard Bayesian VAR to estimate the macroeconomic effects of monetary policy and communications, using the four identified shocks as external instruments. I find that the monetary policy shocks I identify have effects on macroeconomic outcomes that are consistent with New Keynesian models, a conclusion I reach without imposing such consistency a priori. For both the monetary shock and the monetary news shock, shocks that raise interest rates on impact lead to economic contractions. I estimate similar effects on macroeconomic forecast revisions using simple OLS regressions. In terms of magnitudes, my results suggest that the effects of monetary news are "big," in the language of Coibion (2012), and similar to those estimated by Romer and Romer (2004): a 25 basis point increase in the policy rate leads to a roughly one percent decrease in industrial production. The effects of traditional monetary policy (i.e., changes in current-meeting interest rates) shocks are even larger, though shocks of this nature are largely a thing of the past.

The estimated effects of monetary policy that I find are large relative to previous estimates because of the reduction of measurement error provided by my new data and identifying assumptions. Because the Fed communicates simultaneously about several shocks that can have opposing macroeconomic effects, estimates that do not attempt to separate the two shocks—or do so with

[6] This regime change also identifies the (current-meeting) monetary shock, as the bulk of forecast revisions about interest rates moves farther out in the term structure after 2003.

[7] Other papers (Wright, 2012; Arai, 2017; Nakamura and Steinsson, 2018) have used the heteroskedasticity-based identification assumptions in seeking to estimate the effects of monetary policy. The approach is typically seen as a method of purging "background noise" (or latent factors, in the case of Gürkaynak et al. (2020)) from OLS regressions. I use the approach completely differently in that I am interested in estimating multiple shocks, not simply purging one shock of interest of a nuisance component. Lewis (2019) identifies similarly-named shocks using a heteroskedasticity-based approach that provides only statistical identification.

[8] This language comes from Stock and Watson (2017).

variables contaminated by other important shocks, e.g., risk premia in stock returns—will exhibit measurement error. In turn, they will be (at best) attenuated towards zero or even biased in the "wrong" direction. I discuss these alternative approaches in detail below.

Moving to the effects of communications about supply and demand factors, I find effects that are also consistent with how supply and demand shocks operate in standard models. This suggests that the Fed communicates about shocks that ultimately materialize. The effects of the perceived demand shock are substantial, and robustly so. There is no benchmark against which to compare the effects of these newly identified shocks, but the macroeconomic responses to one standard deviation shocks are comparable to those of monetary shocks. That they have any measurable effects suggests that the Fed's view is useful in forecasting economic developments. However, I find that the estimated responses to the perceived demand shock are sensitive to some of variables highlighted by Bauer and Swanson (Forthcoming). Below, I suggest that this is not evidence of the endogeneity of my measures but, instead, evidence that Bauer and Swanson have found variables that span what the Fed may believe is useful information in determining the economic outlook.

In the final part of the paper, I use a simple extension of the textbook New Keynesian model as a framework for interpreting my results. I start by showing that the impulse responses I estimate in the data are consistent with what the model predicts under a standard calibration. I then use the model to assess the assumptions that I make for identification. A standard calibration suggests that even though the assumptions do not hold exactly in the model, the bias introduced is minimal. Finally, I use the model to compare how my identification procedure stacks up against alternative approaches. There, I show that my approach allows me to identify model-simulated shocks much more closely than other leading approaches.

The notion that monetary policy announcements can convey macroeconomic information, and thus contaminate estimates of exogenous monetary shocks, was put forth by Romer and Romer (2000). Campbell et al. (2012) and Nakamura and Steinsson (2018) highlight the fact that the presence of "information effects" can contaminate traditional high-frequency estimates of monetary policy shocks. My results confirm that the contamination is substantial: A widely used high-frequency shock series is determined in large part by perceived demand shocks. Therefore, the empirical challenge became to disentangle information provision from exogenous monetary policy. My model shows that this challenge is the familiar problem of simultaneous determination encountered in supply and demand systems or structural vector autoregressions. Seen in this context, early work in this area imposed zero restrictions to identify monetary shocks, which are not warranted if Fed watchers learn about economic fundamentals from Fed announcements.[9] Nakamura and Steinsson (2018) estimate a structural model to overcome the identification challenge. In contrast to my work, their approach depends on all of the assumptions underlying the particular

[9]This is the case in Kuttner (2001) (who introduced surprises in the current-meeting interest rate) and Gürkaynak et al. (2005) (who introduced the notion of shocks to the path of rates). By not entertaining the possibility that changes in interest-rate futures could be driven by shocks other than exogenous monetary policy, the authors imposed the restriction that other shocks had zero effect on high-frequency interest-rate forecast revisions.

macroeconomic model. My work relaxes these restrictions.

Other papers cognizant of Fed information provision have sought model-free approaches to identify high-frequency monetary policy shocks. These papers typically focus on identifying a "clean" monetary shock, and do not perform detailed decompositions of other aspects of communication like I do here. Given that information effects are posited to stem from the Fed's private information, Miranda-Agrippino and Ricco (2021) and Handlan (2020) propose orthogonalizing high-frequency interest-rate surprises to the Fed's private information as captured by the Fed staff's presentation materials ("Greenbook forecasts").[10] This approach suffers from two shortcomings that I sidestep. First, it assumes that the staff's economic assessment spans that of the Fed's policy-making committee, which is ultimately tasked with policy communication. Romer and Romer (2008) show that these assessments generally do not align. Second, this approach requires choosing a set of variables that completely span the Fed's private information *when the announcement is made*. I avoid having to posit the variables over which the Fed has private information, which is difficult to know, given the vast number of indicators that inform Fed policy decisions.[11] In addition, by relying on staff-created reports, neither of the aforementioned papers can control for events that occur shortly before policy announcements, and thus both papers potentially fail to control for endogenous macroeconomic events to which the Fed might respond. My high-frequency measures avoid this concern.

Cieslak and Schrimpf (2019) and Jarociński and Karadi (2020) assume theoretically motivated sign restrictions regarding the relationship between monetary shocks, information shocks, stock returns, and interest-rate surprises. These sign restrictions allow the authors to discuss the relative importance of and identify monetary and information shocks. In high frequency, stock returns and interest-rate surprises are consistently negatively correlated, in contrast to my estimates of output expectations. This suggests a limited role for stock prices in differentiating between monetary and information shocks.[12] Examining a variable whose response to each shock differs substantially from the responses of interest rates—e.g., output expectations—instead pro-

[10]Campbell et al. (2012) also attempt to control for information effects by controlling for professional macroeconomic forecasts. As Woodford's comment to that article notes, this approach requires that the control variables span the Fed's reaction function, and suffers from the possibility that not all relevant information may be captured by the lower-frequency forecasts. Hansen and McMahon (2016) also study the effects of both types of communication in a low-frequency setting.

[11]Doh et al. (2020) and Handlan (2020) are able to control more flexibly for the Fed staff's information than Miranda-Agrippino and Ricco (2021) by using machine-learning and text-based techniques applied to the Fed's alternative policy statements. Cai et al. (2021) and Lakdawala (2019) introduce alternative methods for controlling for the Fed's *private information* (the difference between the Fed's and the public's information) which is necessary for properly "removing" information effects.

[12]This negative correlation does not serve as evidence "against" the presence of information effects—it only rejects the notion that information effects are the only shock operating when the Fed makes announcements. Put differently, without information effects, the negative correlation might be even stronger. Ultimately, the puzzle that suggested the presence of information effects was the effect of interest rate surprises on expectations of real macroeconomic variables (GDP, unemployment, etc.), not stock returns. D'Amico and King (2017) overcome this issue using a set-identification approach, However, they use lower-frequency data on expectations, which may not incorporate all relevant macroeconomic information to which the Fed responds. Andrade and Ferroni (2021) similarly overcome this issue by studying market-based measures of inflation expectations, combined with a sign-restriction approach.

vides more power for identification.[13] By studying GDP expectations, I directly address one of the main puzzles in the high-frequency literature: that both output and output expectations increase in response to positive interest-rate surprises. Additionally, since my GDP and inflation measures are not constructed using financial market data, I eliminate the potential confounding role of risk premia that pervade the "second variables" used by other authors (break-even inflation and equity returns). Finally, my identification approach does not require the a priori imposition that the identified shocks have theoretically consistent effects. That assumption produces set-identified shocks which, in contrast to my shocks, are not easily portable to other empirical settings.

Finally, Bauer and Swanson (Forthcoming) have recently argued against the presence of information effects, claiming that the wrong-sided response of macroeconomic expectations to interest rate surprises arises because both measures are driven by recent macroeconomic news (and markets systematically under-predict the Fed's response to that news). This argument is largely supported by their finding that the wrong-sided effects disappear once one controls for recent macroeconomic news.[14] My work pushes against this view. In particular, I show that their results are sensitive to reasonable and minuscule perturbations in the definition of "macroeconomic news." In appendix G, I suggest that that the Bauer and Swanson evidence is consistent with Fed announcements that emphasize different pieces of (publicly available) information than markets expect.[15] In addition, I find that my results are generally robust to including their controls, which suggests that the evidence against information effects is not as clear-cut as the authors suggest.

The rest of the paper proceeds as follows. In section 2, I present my text-based proxies of high-frequency macroeconomic forecast revisions. In section 3, I present a simple theoretical framework that explains the identification challenge, lay out my identification assumptions, and estimate the structural shocks. Section 4 contains evidence on the effects of these shocks on macroeconomic outcomes and expectations. Section 5 contains a standard model used to assess my results and identifying assumptions. Section 6 is the conclusion.

2 Measuring Expectations in High Frequency

In this section, I present my approach to measuring high-frequency, text-based measures of macroeconomic forecast revisions. In section 2.1, I present the technique for mapping from newspaper text to numerical values. Section 2.2 shows the application of this technique to articles written around Fed announcements. The remaining sections contain statistical (section 2.3) and narrative (section 2.4) exercises designed to validate the measure as a high-frequency measure of

[13]This, combined with the fact that Jarociński and Karadi's approach only provides set identification of the shocks, allows the authors to (statistically) learn very little about the effects of Fed information—their identified set of impulse responses (that are robust to the chosen prior) includes zero for most horizons. I re-examine the evidence of Jarociński and Karadi in detail and confirm the notion that stock returns are not particularly useful for identifying non-monetary shocks.

[14]In appendix G, I discuss the other evidence presented by the authors.

[15]Sastry (2021) carefully distinguishes the explanation put forth by Bauer and Swanson, the "information effects" explanation, and an "agreeing to disagree" explanation. My interpretation of the Bauer and Swanson evidence is that it is closer to the latter explanation. To the extent that this still presents unexpected communications from the Fed (and not some purely mechanical endogeneity as the authors suggest), I view my results as informative regarding the effects of the Fed's communications.

macroeconomic forecast revisions around Fed meetings. This is important to establish, given that these measures are at the heart of my identification strategy.

2.1 Mapping from Words to Expectations

At its core, my approach requires a mapping, f, from newspaper text to numerical values corresponding to macroeconomic forecast revisions. I estimate this mapping by extracting numerical values—"features" in the machine learning literature—from the universe of *New York Times* articles written over the period 1987–2007. I then use those features in a LASSO regression that predicts monthly changes in macroeconomic expectations from the Blue Chip survey. To establish notation, I estimate

$$\mathbb{E}_t^{\text{BC}}[x_\tau] - \mathbb{E}_{t-1}^{\text{BC}}[x_\tau] = f_x\left(\texttt{NYT}_t, \texttt{NYT}_{t-1}\right) + e_t \tag{1}$$

where $\mathbb{E}_t^{\text{BC}}[x_\tau]$ is the Blue Chip survey expectation of variable x (either real GDP or CPI inflation) at horizon τ, $\texttt{NYT}_t$ is a vector of numerical features extracted from the *New York Times* articles written in the month before the forecast is made, and f_x is a LASSO regression. Because f_x is estimated to predict revisions of GDP and inflation, both of which have units (percentage points), the resulting output of f_x inherits those units and, thus, has a cardinal interpretation. The approach is semi-supervised, in that I tailor the elements in $\texttt{NYT}_t$ based on my judgment about what might be useful features of the text for this context, but let f decide which features to retain. In the rest of this section, I discuss the construction in more detail.

I use the *New York Times Annotated Corpus* as my training corpus.[16] This corpus contains the universe of articles—about 1.8 million—from the New York Times between January 1987 and July 2007, each manually tagged by library scientists.[17] I use these tags to retain the 94,601 articles that were tagged as related to either economic output, prices, or labor markets.[18]

The next task is to extract the features that compose the vector $\texttt{NYT}_t$. My objective is to use the text as a proxy for macroeconomic forecast revisions, so I choose a set of word counts, and counts of word co-occurrences, that seem well suited to this task. In particular, from each document I extract counts of topic–sentiment–tense triplets, in which the tense and sentiment words must be within a three word window of the topic word (and in the same sentence).[19] A "topic-word" is a word (or bi-gram) related to the macroeconomic variable of interest. For GDP, these words are *economic growth*, *growth*, *economi*, *consumer spending*, and *output*. For inflation, the words are *inflat*, *price*, *oil prices*, *inflationari*, and *deflat*. In appendix A, I describe the process for choosing these words, alongside several other standard (but detailed) steps I take to clean the

[16] I chose this corpus in large part because of its coverage and availability. The corpus is available from the University of Pennsylvania's Linguistic Data Consortium.

[17] The corpus' documentation describes the articles as "every article published" in the newspaper, excluding wire-service articles. Each article is tagged with the broad topics discussed therein.

[18] Specifically, I retained articles labeled "economic conditions and trends", "united states economy", "prices, wages and salaries," "layoffs and job reductions," "production," or "labor."

[19] Each set of words (topic, tense, and sentiment) also includes an empty word. Put differently, I also include raw counts of each topic, tense, and sentiment word, and all pairwise doubles of these words. I explore longer word windows in robustness, though I note here for practitioners that using a shorter window allowed for a better fit of the mapping to the data.

Table 1: Mapping from Text to Forecast Revisions

(a) GDP

Coefficient	Word	Sentiment	Tense	Lag
-0.021	consumerspending	Decrease		
-0.019	output		Future	
-0.019	economi	Fall		
0.018	output	Negative	Future	✓
0.018	economicgrowth	Positive	Present	✓
-0.018	economicgrowth	Negative	Present	✓
-0.015	economicgrowth	Negative	Past	
-0.015	economicgrowth	Fall	Future	
-0.015	economicgrowth	Decrease		
-0.015	consumerspending	Fall		

(b) Inflation

Coefficient	Word	Sentiment	Tense	Lag
-0.022		Negative		
0.019	oilprices	Rise	Present	
0.019	inflationari	Decrease	Future	✓
0.018	inflationari	Weak	Past	✓
-0.018	inflationari	Decrease	Past	✓
-0.017	oilprices	Decrease	Past	✓
0.017	price	Positive		✓
0.017	oilprices	Rise	Past	
-0.015	inflationari	Rise	Future	
0.014	price	Increase	Future	✓

NOTE. This table shows the text "features" with the 10 largest (absolute) coefficients that enter the mapping from text to GDP (panel (a)) and inflation (panel (b)) forecast revisions from month $t-1$ to month t. The mapping predicts these revisions using features extracted from all relevant New York Times articles in month t and $t-1$ (the "lag" column).

text. Later in the paper, I show that varying the length of these word lists produces similar results. A "sentiment" is a word from one of the following word lists that are commonplace in the NLP literature: *positive*, *negative*, *increase*, *decrease*, *rise*, *fall*, *strong*, or *weak*.[20] If a negation is within the three-word window described above, I flip the sentiment to its converse. Finally, I tag each verb by tense: either present, past, or future.[21]

Given the text and procedure for extracting numerical data, the last step is to estimate equation (1). To do this, I construct the vector $\mathtt{NYT}_t$ by first aggregating all articles written between the 8th of month $t-1$ and the 7th of month t, since Blue Chip forecasts are typically submitted sometime in the first week of the month (Nakamura and Steinsson, 2018). I then extract the log count of each feature described above on these concatenated articles. I include features for both months t and $t-1$, to allow for both the level and (quasi-) difference to be selected. I separately estimate the LASSO regression for GDP and inflation,[22] and select the LASSO penalization parameter by 10-fold cross-validation.

Table 1 displays a summary of the estimated mapping f_x, where x is either inflation or GDP. The table contains the 10 features with the largest (absolute) coefficients. Many of the coefficients are intuitive. A *decreas*ing mention of *consumerspending* enters with a negative coefficient, as does a *fall*ing mention of *economi*. The presence of current and lagged counts of each feature can make some coefficients challenging to interpret. For example, if *output* is mentioned in lagged articles in a negative way, table 1 suggests an upward revision in GDP expectations. The bi-gram *oil prices* is quite important in shaping inflation expectations: *Ris*ing mentions enter with a positive

[20]The first two words lists are from the finance-specific dictionary of Loughran and McDonald (2011), and the remainder are from the Harvard-IV dictionaries. More details are in appendix A. I stem all words in each list using the stemmer of Porter (1980), and count all words in each list as a single "word." That is, I count the number of times that a GDP word is used with any word from the *decrease* list, without distinguishing between which word.

[21]I implement part of speech tagging using the natural language toolkit in Python. This tags each word using the parts of speech from the Penn Treebank Project, described by Taylor et al. (2003).

[22]In my baseline analysis, I average forecast revisions for the quarter of month t and the subsequent two quarters. This is akin to the summary measure of forecast revisions used by Nakamura and Steinsson (2018). I explore robustness to this choice.

coefficient, while *decreas*ing mentions enter with a negative coefficient. My results are robust to eliminating this bi-gram. While not shown, each of the five GDP and inflation topic-words are selected in some form, but with coefficients that are too small to be included in the table. Finally, the inclusion of the tense words is useful when predicting different forecast horizons of macroeconomic variables. This is visible in appendix A—there, I estimate equation (1) using a longer forecast horizon, τ, and find more future-tense words and no past-tense words.

2.2 Fed Announcements

The next step is to use the estimated f_x functions from the previous section to estimate high-frequency forecast revisions made around Fed announcements. Put simply, I achieve this by applying f_x to newspaper articles written in a relatively narrow window around Fed announcements.

Specifically, I analyze all articles written the day before, day of, and day after each Fed meeting in the *New York Times*, *Wall Street Journal*, and *Washington Post*. I collect these articles from Factiva, searching for articles with the keywords "Federal Reserve" and "FOMC."[23] The three-day window ensures that I capture the many articles written on the first day of two-day Fed meetings (these typically start "The Fed begins a two-day policy meeting today"), and the print articles written the day after the meeting. The timestamps in Factiva are unreliable, so I manually sorted all articles on the day of the Fed meeting based on whether they were written before or after the meeting took place. I also removed duplicate articles.[24]

By considering a 3-day window, I potentially include articles that are not primarily written about the policy meeting. To guard against this, I only retain articles that have *fed* in the title or have one of {*fed*, *fomc*, *federal*} within five words of one of {*meet*, *meeting*, *policy*, *decision*, *rate*} in the top quarter of the article.[25] Thus, while the window is not as high frequency as tick-level data, the ability to narrow articles to those that specifically discuss the policy meeting reduces the possibility that other events influence my analysis.[26]

Finally, I concatenate all pre- and post-meeting stemmed articles into a single document for each meeting t, given by $\texttt{PRE}_t$ and $\texttt{POST}_t$, respectively. My measure of high-frequency forecast revisions for variable x (either inflation or GDP) are given by

$$\widehat{x}_t = f_x\left(\texttt{POST}_t, \texttt{PRE}_t\right).$$

[23] *Fed* is too general, and returns articles related to food and eating. "FOMC" stands for "Federal Open Market Committee," the Fed's policy-making body. I use "Fed" and "FOMC" interchangeably in this paper.

[24] To do this, I created a vector representation of each article i, denoted by v_i, whose length was equal to the total number of unique words across all articles. The j-th element is the number of times term j appears in document i. I then calculate the pairwise distances between all articles using the cosine distance metric: For two documents v_i and v_k, this is $1 - v_i' v_k / [\sqrt{v_i' v_i} + \sqrt{v_k' v_k}]$. Plotting the distribution of cosine similarities, I found a second mode at 0.97, so I randomly chose one from each set of articles with mutual cosine distance above 0.97.

[25] Restricting this search to the top quarter is important for removing articles that only mention Fed policy meetings as an afterthought later in the article. Notably, this removes articles about the handful of other macroeconomic news releases that (rarely) occur on the same day as Fed policy announcements.

[26] Upon inspecting the articles, I found that many articles that fit the aforementioned criteria for being "relevant" instead discussed how another country's central bank might react to the Fed's announcement. Those articles almost always started with the name of a non-US city (i.e., reflecting where the article was written), so I exclude articles that have one of these non-US city names within the first 5% of the text. The list of these cities is in appendix A.1.

Figure 1: High-Frequency Forecast Revisions

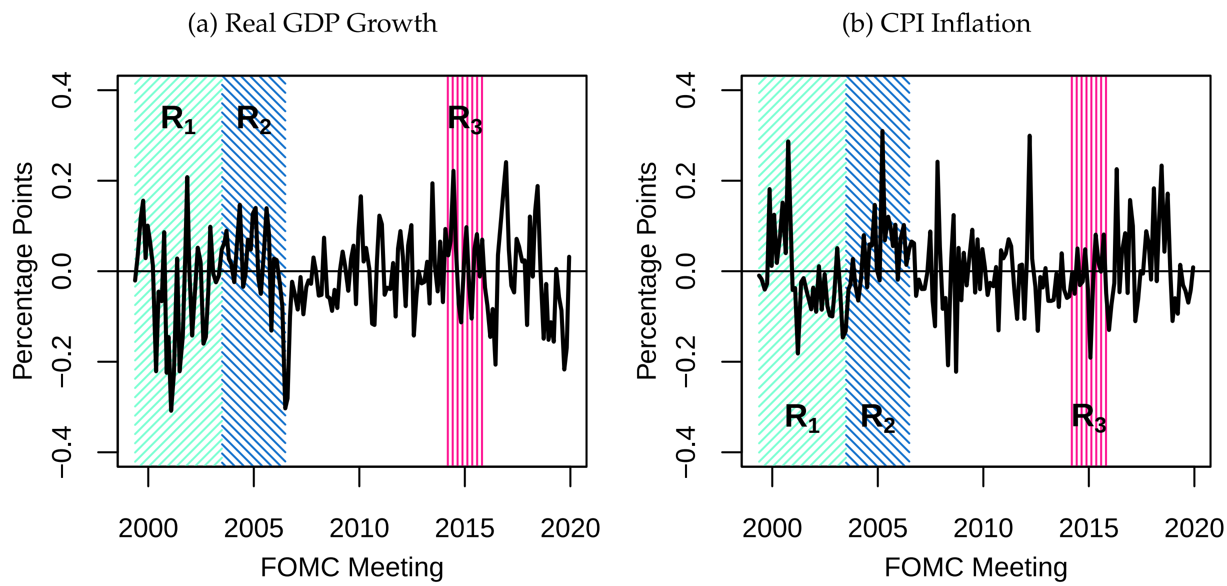

NOTE. This figure shows the high-frequency forecast revisions of real GDP growth (panel (a)) and CPI inflation (panel (b)) around Fed announcements. The sample includes all regularly scheduled meetings between May 1999 and December 2019. The three shaded regions correspond to the three regimes described in the text.

The resulting high-frequency measures are plotted in figure 1. Generally speaking, the magnitude of these surprises are quite small. Thus, it appears that most Fed announcements do not cause large forecast revisions—an unsurprising result that my manual reading of many articles, presented in section 2.4, confirms.

2.3 Statistical Validation

To build confidence in my newspaper-based measures, I discuss several validation exercises whose details are found in the appendix. My first statistical validation exercises focus on text-based changes in expectations that are not centered around Fed meetings. After all, given a credible estimate of f_x, the application to Fed announcements is a straightforward extension.

In figure 11 in appendix B, I show that, at a monthly frequency, my text-based measures of inflation and GDP forecast revisions track their Blue Chip counterparts closely. Despite being the target of estimation, there is no guarantee that the relatively simple and fixed features that I extracted from the text are rich enough to capture forecast revisions over a twenty-year period that included two recessions of very different natures, and rapid technological transformations.

A tougher test for my measures is whether their high-frequency behavior seems sensible. To test this, in appendix B, I construct a daily measure of GDP and inflation forecast revisions over the 1987–2007 period for which I have access to daily newspaper articles. I run a regression of daily changes on the surprise component of a few macroeconomic data releases, oil prices, Fed interest rate surprises, and stock returns. In levels, the surprise in the jobs report and Fed surprises are statistically-significantly correlated with the high-frequency GDP measure. GDP forecast revisions also tend to be larger on days when GDP, jobs, and Fed announcements are made. They are also larger when stock returns and changes in oil prices are larger. The inflation measure is harder to predict, responding primarily to Fed surprises and oil prices.

In figure 12 in appendix B, I show how my measures line up against the daily newspaper-based sentiment measure of Shapiro et al. (2020). Both series are positively correlated with the measure, but weakly so. This is to be expected—my measures are designed to capture "sentiment" about narrow topics, in contrast to the more-general sentiment measure. There are cases in which the measures noticeably diverge. In August 1990, my inflation measure rises noticeably, peaking in September 1990. My GDP measure and the general sentiment index both fall. This timing coincides with the start of Gulf War in August 1990, a shock that led to a spike in oil prices amidst a recession. This example highlights the fact that the sentiment measure can shed some light on the nature of economic shocks—in this example, the negative co-movement between inflation and GDP suggests a supply-driven recession. Looking at the general sentiment measure, or either of my measures in isolation, cannot reveal this type of information.

2.4 Narrative Validation

I conclude this section with a narrative description of my text-based measures. I organize the discussion around the largest realizations of my measures, which I describe in detail in appendix C. A few patterns emerge. First, the GDP and inflation forecast revisions induced by Fed announcements are broadly unremarkable. This should be no surprise—market-based measures of interest rate forecast revisions around Fed announcements are small too, typically less than five basis points.[27] Second, in order to understand the effect of the Fed's announcement, it is important to take the pre-meeting discussion into account. In some cases, post-meeting articles describe directly the surprise component of the announcement. In others, this can only be inferred by comparing post-announcement descriptions to pre-announcement descriptions of what the Fed was expected to do and say. Third, an anecdotal pattern—supported by the statistical evidence at the end of section 4.1—emerges: Stock prices around Fed announcements respond primarily (negatively) to the change in expected interest rates, regardless of the underlying "perceived cause." This is the first piece of evidence that stock returns may not be especially useful for separating monetary shocks from other components.

3 Identifying Shocks

It has long been understood that Fed announcements convey multiple dimensions of information as policymakers strive to describe their current and future actions.[28] In this context, examining how expectations are revised about any single endogenous variable in isolation presents a challenge for identifying the effects of exogenous variables. For example, forecast revisions about future interest rates can increase if the Fed commits to a series of contractionary interest rate hikes,

[27]The average absolute change in the 4-quarter Eurodollar future is 4.5 basis points. For the current-meeting Fed-Funds future, it is 1.9 basis points.

[28]The pioneering work of Gürkaynak et al. (2005) took this multidimensionality seriously, conceiving of Fed announcements as delivering independent information about both current and future interest rates.

or if the Fed's rosy economic outlook transmits to the public. I discuss the identification challenge in more detail in section 3.1. I then present a strategy for identifying the exogenous forces that Fed-watchers learn about from high-frequency forecast-revision data that depends heavily on historical knowledge of Fed communications. I present my approach to identifying these perceived structural shocks in section 3.2, and discuss their properties in section 3.3.

3.1 Simultaneity in High Frequency Data: A Motivating Example

The task of identifying structural shocks using high-frequency data is, at its core, no different from identifying structural shocks in a vector autoregression: The main challenge is to appropriately take the simultaneous determination of endogenous variables in to account. Formally, the simultaneous determination of a vector of n observable (or endogenous) variables x by n exogenous shocks ξ can be expressed as

$$x = \mathbf{M}\xi. \tag{2}$$

As a motivating example, consider a data generating process for GDP, y, and interest rates, i:

$$\begin{aligned} y &= \gamma i + \eta \\ i &= \phi y + \mu, \end{aligned}$$

where γ and ϕ are fixed coefficients, and the independent structural shocks, η (the "economic fundamental") and μ (the "monetary shock") have variances σ_η^2 and σ_μ^2 and are exogenously determined. This model provides the fairly standard result (enshrined in the IS-LM model) that output and interest rates affect one another. In turn, a structural shock to either equation affects both endogenous variables simultaneously. This presents an identification problem. The data can provide at most $n(n+1)/2$ moments from the covariance matrix of x (in the example, these three moments are the variances of y and i, and their covariance), but the model has n^2 parameters to identify (in the example, these four parameters are γ, ϕ, σ_η, and σ_μ).

Simultaneity in the form of equation (2) is a familiar problem when working with macroeconomic data, but it is also of concern when working with forecast revisions of macroeconomic variables.[29] To see this, I establish the following notation for high-frequency forecast revisions of a variable z made in response to a signal, s_t, emitted at time t:

$$\widehat{z}_t = \mathbb{E}[z \mid \mathcal{I} \cup s_t] - \mathbb{E}[z \mid \mathcal{I}]$$

where $\mathcal{I}$ is the agent's information set before the event occurs. Forecast revisions for an agent with rational expectations in this context satisfy

$$\widehat{x}_t = \mathbf{M}\widehat{\xi}_t. \tag{3}$$

Simultaneity remains an issue here, even using high-frequency data. However, one important distinction is that now the variables observable to the econometrician are determined by *perceived* structural shocks, $\widehat{\xi}_t$. Perceived structural shocks are a function of the structural shocks and any noise induced by the communication, s_t. Importantly, their variance is also a function of the variance of the shocks, and of the communication.

[29] As mentioned alongside citations above, the literature is broadly cognizant of this concern.

Intuition for the challenges present when evaluating forecast revisions around monetary policy announcements can be gleaned by considering what forecast revisions look like in the motivating example, following equation (3):

$$\widehat{i}_t = \left(\frac{1}{1-\gamma\phi}\right)\widehat{\mu}_t + \left(\frac{\phi}{1-\gamma\phi}\right)\widehat{\eta}_t \tag{4a}$$

$$\widehat{y}_t = \left(\frac{\gamma}{1-\gamma\phi}\right)\widehat{\mu}_t + \left(\frac{1}{1-\gamma\phi}\right)\widehat{\eta}_t. \tag{4b}$$

Equations (4a) and (4b) are the crux of the identification challenge. First, equation (4a) shows that in general, traditional HF measures of monetary shocks ($\widehat{i}_t$) are contaminated by what markets learn about the economic fundamental, $\widehat{\eta}_t$. There is only a special case—when the Fed has no independent knowledge about the state of the economy (so that $\widehat{\eta}_t = 0,\ \forall t$)—in which traditional measures identify the (perceived) monetary policy shock $\widehat{\varepsilon}_t$. A testable implication of the model is that if traditional estimates are not contaminated by the perceived economic fundamental, then the correlation between interest-rate and output expectations must be negative. Empirical evidence refutes this implication. Campbell et al. (2012) and Nakamura and Steinsson (2018) find a positive correlation between output and interest-rate forecast revisions using low-frequency data, which I confirm using high-frequency data.[30]

More germane to this paper is equation (4b), which shows that an estimate of output forecast revisions can neither be used to control for information effects nor to instrument for information effects. On the first point, in appendix H, I show that the residual from a regression of interest rate expectations on output expectations (i.e. a "cleaned" interest rate surprise) only identifies monetary shocks in the case that output expectations do not respond to monetary shocks. This, in turn, is only the case with full monetary neutrality (i.e. $\phi = 0$) or if there are no information effects in the first place.[31] Next, output forecast revisions are not a valid instrument for information effects, because they are not exogenous with respect to monetary policy shocks.

I take the simultaneous equation system of equation (3) seriously, but pause here to note that its simple appearance should not mask its generality. First, I have abstracted from the time horizon of the endogenous variable x. In practice, I will consider forecast revisions for variables a few quarters in the future. Empirically, I show that changing the time horizon of the endogenous variables makes little difference. This lack of a meaningful difference is also a feature of the textbook New Keynesian model that I use to evaluate my empirical results in section 5.1. Second, I have described neither the specific variables that I consider, nor the structural shocks that I seek to identify. In practice, the economic circumstances surrounding each policy announcement are never identical, and so, in principal, an announcement can convey information about an infinite-dimensional state space. To reduce the dimensionality and help to guide the thinking in well-established economic reasoning, I will lean on the textbook New Keynesian model again.

[30]Here I have not allowed the "Taylor rule coefficient," ϕ, to be stochastic. Bauer and Swanson (Forthcoming) put this forth as a potentially relevant channel. I discuss that work in detail in appendix G, but highlight here that shocks to the Taylor rule coefficient only have second-order effects in standard models.

[31]In the language of Angrist and Pischke (2008), $\widehat{y}_t$ is a "bad control" for information effects, since it is affected by monetary shocks.

In its most basic form, that model posits inflation, output, and nominal interest rates as being driven by what I call a "supply" (total factor productivity) shock, a "demand" (discount factor) shock, and an exogenous monetary policy shock. With the addition of longer-term rates and an additional shock to capture forward guidance, this is the model that I use to guide my empirics. Of course, myriad other shocks could be gleaned from the Fed's announcements, but given the relatively small sample size, I halt my exploration here, and note that this already represents an expansion of what has been considered using high-frequency data.

3.2 Overcoming Simultaneity

The system I estimate is

$$\underbrace{\begin{bmatrix} \widehat{\mathrm{FF}}_t^0 \\ \widehat{\mathrm{ED}}_t^4 \\ \widehat{y}_t \\ \widehat{\pi}_t \end{bmatrix}}_{\widehat{\mathbf{x}}_t} = \mathbf{A} \underbrace{\begin{bmatrix} \widehat{\varepsilon}_t^m \\ \widehat{\varepsilon}_t^n \\ \widehat{\varepsilon}_t^d \\ \widehat{\varepsilon}_t^s \end{bmatrix}}_{\widehat{\boldsymbol{\varepsilon}}_t}, \tag{5}$$

where t indexes Fed announcements. The four observable variables on the left-hand side are, respectively, high-frequency forecast revisions of the current-meeting Federal Funds rate, the Eurodollar rate four quarters in the future, output, and inflation.[32] I assume that these observable variables are driven by what Fed watchers learn about four structural shocks: monetary, monetary news, demand, and supply. The four-by-four matrix $\mathbf{A}$ provides the relationship between the two. My objective is to estimate $\mathbf{A}$, which will allow me to recover the vector of perceived structural shocks from the high-frequency forecast revisions at my disposal. To do this, I bring historical knowledge of how Fed communications have changed over time. This knowledge will allow me to make assumptions about the structure of the covariance matrix of the perceived structural shocks. Intuitively, this knowledge provides the additional moments needed to identify the model's parameters.

Specifically, I focus on three regimes during my 1999–2019 sample during which the historical record suggests the Fed's communications provided different amounts of information about different structural shocks. The first is the episode discussed by Lunsford (2020): the August 2003 introduction of interest-rate forward guidance. To give context, in June 2003 the Fed had lowered the Funds rate to 1%. In August 2003, the post-meeting statement declared that this "policy accommodation can be maintained for a considerable period." This was the first instance of explicit forward guidance regarding the path of interest rates.[33] This episode thus provides a natural place to split my sample for the purpose of separately identifying perceived monetary shocks from other perceived shocks ($\widehat{\varepsilon}_t^d$ and $\widehat{\varepsilon}_t^s$)). My first regime, $\mathcal{R}_1$, consists of all meetings from the start of my sample (May 1999) through June 2003. The second regime, $\mathcal{R}_2$, extends from August 2003 through

[32] I constructed the interest rate measures from tick-frequency data purchased from the Chicago Mercantile Exchange following the construction of Nakamura and Steinsson (2018).

[33] Forward guidance was used for reasons similar to what prompted its major re-emergence in 2008. In 2003, 1% was essentially seen as the effective lower bound on nominal interest rates.

Figure 2: Energy Prices

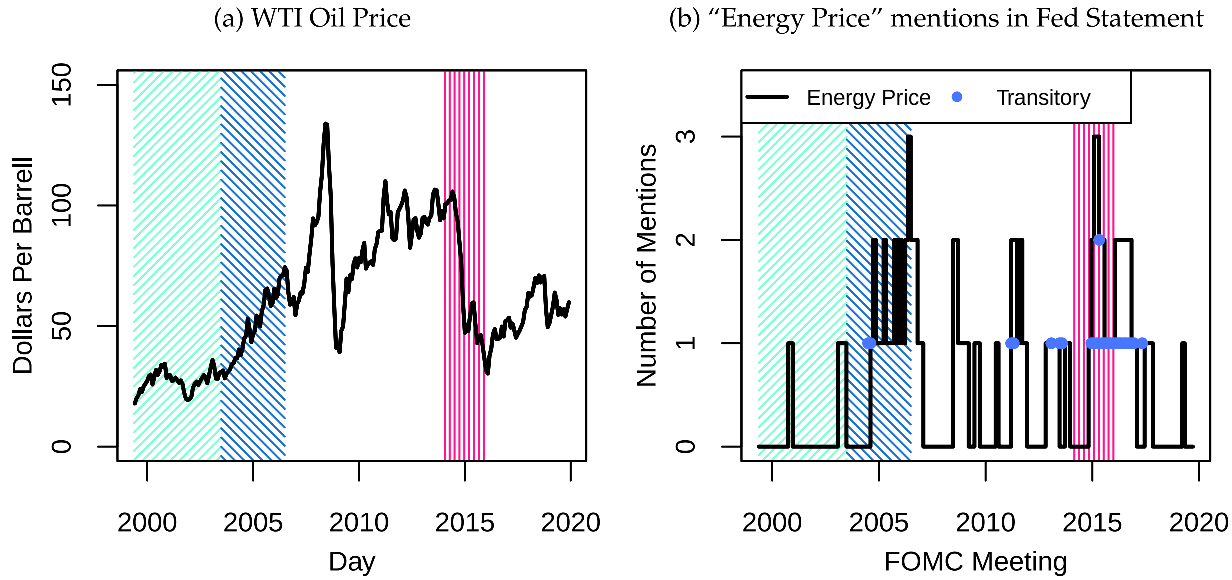

NOTE. Panel (a) shows the West Text Intermediate (WTI) oil price, taken from FRED (mnemonic WTISPLC). The black line in panel (b) shows the number of times that either "energy price," "price(s) of energy," "oil price," or "price(s) of oil" appear in each Fed statement during my 1999–2019 sample. The blue dots show the number of times the word "transitory" was used. The three shaded regions correspond to the three regimes described in the text.

May 2005.[34] These regimes are shown by the two left-most shaded regions of figure 1.

The third regime that I use for identification consists of the period 2014–2016, which I denote by $\mathcal{R}_3$. As depicted in panel (a) of figure 2, oil prices fell dramatically over this period.[35] This decline put a strain on the Fed's ability to bring inflation to its, then, two-percent target. As a consequence, a consistent part of the Fed's communications over this period was devoted to describing its outlook for energy prices. Panel (b) of figure 2 shows the number of times that energy prices have been mentioned in each Fed statement over my sample period. I also include a count of the word "transitory," since (not unlike the most-recent inflation episode) the primary way that the Fed communicated its outlook for energy prices was by describing the expected persistence of the shock. Panel (b) thus suggests that there was at least the possibility for observers to learn about the evolution of energy prices and, in turn, supply shocks from the Fed's announcements over this period.[36]

How can these regimes provide a way to identify the elements of $\mathbf{A}$? The approach—coined

[34] The first regime is a slight elongation of the sample period studied by Lunsford (2020), which I include for statistical precision—my point estimates are nearly identical using Lunsford's exact regimes, but mildly less-precisely estimated.

[35] This event has received substantial attention both in policy research circles (e.g. Grigoli et al. (2019)) and in academic circles (e.g. Baumeister and Kilian (2016)).

[36] Panel (b) also shows that statements during my second regime included discussions of energy prices—this was a result of the large uptick in prices following the start of the war in Iraq. Because of this, I show later that I can drop the 2014–2016 regime and still identify a supply shock, though less precisely so. I include the regime, however, so that I do not have to take a stand on which shock changed by more between the first and second regimes. This third regime contributes very little to the identification of the other shocks.

"identification by heteroskedasticity" in the seminal work of Rigobon (2003)—centers on the fact that variation in the variances of the underlying shocks induces different co-movements in the observable variables. Consider again the illustrative example. The covariance of output and inflation is given by $\gamma\sigma_\mu^2 + \phi\sigma_\eta^2$. Standard theory would suggest that $\gamma < 0$ and $\phi > 0$, and so an increase in the variance of the perceived monetary shock will decrease the covariance, while an increase in the variance of the information effect will do the opposite. The insight of Rigobon is that these changes in variances across regimes provides additional moments that can identify the system's parameters. Formally, I assume that $\mathbf{A}$ remains constant across regimes, while the perceived structural shocks exhibit heteroskedasticity of the form

$$\mathrm{Cov}_t\left(\widehat{\boldsymbol{\varepsilon}}_t\right) = \begin{cases} \Sigma_{(1)} & t \in \mathcal{R}_1 \\ \Sigma_{(2)} & t \in \mathcal{R}_2 \\ \Sigma_{(3)} & t \in \mathcal{R}_3 \end{cases}, \tag{6}$$

where the normalization that $\Sigma_{(1)}$ is the identity matrix, and that $\Sigma_{(2)}$ and $\Sigma_{(3)}$ are diagonal—an assumption I discus later. From an accounting perspective, these assumptions allow me to use the 10 unique elements of the covariance matrix in each of the three regimes (30 total) to identify the model's 24 parameters (16 in $\mathbf{A}$, and 4 in each of $\Sigma_{(2)}$ and $\Sigma_{(3)}$). The system could be identified with two regimes—I consider this in robustness exercises—but the third regime helps me to name the shocks, and allows me to perform tests of overidentifying restrictions. In the remainder of this section, I discuss these assumptions and the naming of the shocks.

Historical knowledge about the three regimes is necessary for naming the perceived shocks, since the heteroskedasticity that I allow in equation (6) on its own does not provide a way to name them in a way consistent with my labeling in equation (5). Recall that the variance of the perceived structural shocks should increase in the size of the structural shock, and fall with the variance of the noise accompanying the signal of that shock. The shock I call the "monetary shock" ($\widehat{\varepsilon}_t^m$), is the shock whose variance decreases the most between the first and second regime, while the variance of the "monetary news shock" ($\widehat{\varepsilon}_t^n$) increases the most. The intuition—a standard feature of Bayesian updating—is as follows. In 2003, the Fed's meeting-t communications telegraphed interest rate changes at future meetings (say $t+1$). Thus, at t, uncertainty about the Federal Funds rate decision was largely known, having been communicated at $t-1$. Thus, the introduction of forward guidance should lead to a fall in the size of *current meeting* interest rate surprises. On the other hand, signals about future rates became more precise at this time, which should allow observers to make larger forecast revisions for future interest rates. Forward guidance, then, should lead to larger monetary news shocks. The "perceived supply shock" is the shock whose variance is the largest in the third regime, since at the time the precision of the Fed's signal was arguably high given its intense focus on forecasting oil prices. Finally, the "perceived demand shock" is the remaining shock.

This historical knowledge on its own is of course not a guarantee that the procedure will "work"—there has to be some statistical evidence of heteroskedasticity in the data. Much of

Table 2: Co-Movement of HF Forecast Revisions Across Regimes

(a) No Controls

	$\widehat{\mathrm{ED}}^4_t$ $\widehat{i}_t$ and $\widehat{y}_t$	$\widehat{\mathrm{ED}}^4_t$ and $\widehat{\pi}_t$	$\widehat{\pi}_t$ and $\widehat{y}_t$
$\mathcal{R}_1$ (1999–2003)	0.054	0.334	0.089
	(0.119)	(0.108)	(0.136)
$\mathcal{R}_2$ (2003–2006)	-0.411	0.111	-0.152
	(0.161)	(0.242)	(0.189)
$\mathcal{R}_3$ (2014–2016)	-0.011	-0.056	-0.188
	(0.117)	(0.116)	(0.258)
Rest of Sample	0.087	0.165	0.195
	(0.045)	(0.060)	(0.105)
Observations	166	166	166
R^2	0.064	0.109	0.076
$p(\beta_1 = \beta_2)$	0.022	0.401	0.302
$p(\beta_1 = \beta_3)$	0.699	0.015	0.344

(b) Bauer and Swanson Controls

	$\widehat{\mathrm{ED}}^4_t$ $\widehat{i}_t$ and $\widehat{y}_t$	$\widehat{\mathrm{ED}}^4_t$ and $\widehat{\pi}_t$	$\widehat{\pi}_t$ and $\widehat{y}_t$
$\mathcal{R}_1$ (1999–2003)	0.040	0.282	0.093
	(0.103)	(0.104)	(0.123)
$\mathcal{R}_2$ (2003–2006)	-0.446	0.115	-0.186
	(0.150)	(0.240)	(0.200)
$\mathcal{R}_3$ (2014–2016)	-0.012	-0.037	-0.189
	(0.119)	(0.107)	(0.260)
Rest of Sample	0.053	0.109	0.174
	(0.043)	(0.061)	(0.106)
Observations	166	166	166
R^2	0.173	0.185	0.118
$p(\beta_1 = \beta_2)$	0.008	0.530	0.247
$p(\beta_1 = \beta_3)$	0.739	0.033	0.330

NOTE. The first column of each panel shows estimates of $\widehat{i}_t = \alpha_r + \beta_r \widehat{y}_t + e_t$, where $\widehat{i}_t$ is the HF change in Eurodollar futures (scaled to have unit effect on the daily change in 1-year treasury rates), $\widehat{y}_t$ is the measure of HF GDP forecast revisions (in percentage points), and the coefficients α_r and β_r are allowed to vary over the regimes in parentheses. The middle column shows analogous results when the HF inflation forecast revision is on the right-hand side. The right column has HF inflation revisions on the left-hand side and HF output revisions on the right-hand side. The sample includes all regularly scheduled Fed meetings between May 1999 and December 2019. The bottom two columns test for equality of the coefficients across the respective regimes. Heteroskedasticity-robust standard errors are shown in parentheses. Panel (a) contains no additional controls, while panel (b) adds the three controls of Bauer and Swanson (Forthcoming) described in appendix G.

Lunsford's work provides evidence of heteroskedasticity between $\mathcal{R}_1$ and $\mathcal{R}_2$, by showing that the correlations of interest rate forecast revisions with other variables differ across the two regimes. Table 2 provides further evidence using the variables that I consider, and also examines the third regime. Each row shows the regression coefficient between the two variables in the columns within the respective regime. From the first regime to the second, the relationship between forecast revisions of future interest rates and both GDP and inflation forecast revisions decreased, consistent with the notion that monetary news played a more-important role in the second regime.[37] From the first to the third regime, the relationship between output and inflation forecast revisions also falls, suggesting that markets learned relatively more about supply than demand factors in the third regime. That relationship is also negative in the second regime. This should not be a surprise in light of figure 2. The second regime also observed large increases in oil prices that the Fed discussed in many of its post-meeting policy statements at the time.[38] Taken together, the instability of these relationships suggests the presence of heteroskedasticity.

[37]To save space, I don't show all pairwise correlations in this table—notably missing is $\widehat{\mathrm{FF}}^0_t$. After 2003, the variance of $\widehat{\mathrm{FF}}^0_t$ falls dramatically, rendering estimates here noisy. In the first regime, $\widehat{\mathrm{FF}}^0_t$ is negatively correlated with both $\widehat{y}_t$ and $\widehat{\pi}_t$. This is not surprising: $\widehat{\mathrm{FF}}^0_t$ closely resembles the monetary shock (not the monetary news shock) that I identify. I also show, later, that other non-monetary shocks appear to feature more-prominently in longer-term interest rate forecast revisions.

[38]This is why I can drop the third regime and reach similar conclusions.

Two remaining assumptions are important to my identification—that the perceived structural shocks are uncorrelated, and that $\mathbf{A}$ is unchanged across regimes. First, while it is common to assume that structural shocks are independent, perceived structural shocks need not be if the Fed's communications are not sufficiently detailed in their discussion of each type of shock. In that case, markets will generally have to use their prior knowledge to parse the independent information revealed about each shock.[39] Empirically, this assumption can be tested with the over-identification tests that I present alongside my estimates. I also show, in the model of section 5.1, that identification by heteroskedasticity actually reverses some of the noise introduced by correlated signals, and brings the estimated perceived structural shocks closer to the true underlying shocks.

The second assumption I make is that $\mathbf{A}$ remains unchanged across the regimes. Without this assumption, there would be too many parameters to estimate relative to available moments. The over-identification (J) tests help assess the severity of this assumption. To assess the theoretical validity of this assumption, I consider the assumption within the context of linear rational (RE) expectations models. In full-information RE models there is no role for information effects, so I focus my discussion on linear RE models with imperfectly-informed agents. In appendix D, I show that the assumption is valid in a general class of dynamic linear RE models (e.g., Blanchard et al. (2013)), in which agents are symmetrically imperfectly informed about the economy's state variables.

3.3 Estimation

I estimate the parameters via GMM, and present the estimates in table 3. Panel (a) contains the mapping from shocks to observables, $\mathbf{A}$, and panels (b) and (c) show the estimates of $\Sigma_{(2)}\Sigma_{(1)}^{-1}$ and $\Sigma_{(3)}\Sigma_1^{-1}$—the variance of each of the structural shocks in $\mathcal{R}_2$ and $\mathcal{R}_3$ relative to $\mathcal{R}_1$.[40] I have named the perceived shocks consistent with the discussion above—the perceived supply shock is the shock with the largest relative variance in the third regime, while the monetary news shock is the one with the largest relative variance in regime 2. In both cases, I can safely reject the null hypothesis that their variances are smaller than the perceived demand shock. As suggested by the discussion above, markets learned more (by a statistically significant amount) about future rates than current-meeting rates following the introduction of forward guidance.

Having named the shocks based on relative variances, the estimates of $\mathbf{A}$ serve as the first check on the shocks. The perceived supply shock leads to a negative co-movement between real GDP and inflation, with minimal impact on interest rates. The perceived demand shock generates a positive co-movement in all observable variables. The monetary news shock causes a negative co-movement between future interest rates and GDP/inflation. Similarly, the monetary shock leads to a negative co-movement between interest rates and GDP/inflation. These

[39] This assumption seems less plausible if markets only observed the interest-rate decision itself. In reality, Fed announcements are composed of a post-meeting statement over my entire sample that explains the rationale behind the policy decision. Thus, the announcement is multidimensional. What's more, since the start of my sample several additional dimensions have been added to this signal: press conferences (with the opportunity to answer questions from the press) and economic forecasts.

[40] The normalization of $\Sigma_{(1)}$ to the identity matrix is standard and without loss.

Table 3: Heteroskedasticity-Based Estimates

(a) Structural Impact Matrix, $\mathbf{A}$

	Response of High-Frequency Forecast Revisions			
	Fed Funds Futures Current Meeting	Eurodollar Futures 4 Quarters Ahead	Real GDP Growth	CPI Inflation
Perceived Supply	0.03	0.09	0.86	-0.54
	(0.04)	(0.20)	(0.28)	(0.23)
Monetary	1.55	0.42	-0.52	-0.09
	(0.06)	(0.26)	(0.02)	(0.19)
Perceived Demand	0.08	0.92	0.58	0.91
	(0.19)	(0.25)	(0.23)	(0.21)
Monetary News	0.02	0.84	-0.33	-0.23
	(0.25)	(0.23)	(0.17)	(0.13)

(b) Relative Variance in Regime 2, $\Sigma_{(2)}\Sigma_{(1)}^{-1}$

Perceived Supply	Monetary Shock	Perceived Demand	Monetary News
0.51	0.02	0.47	1.71
(0.16)	(0.01)	(0.15)	(0.71)

(c) Relative Variance in Regime 3, $\Sigma_{(3)}\Sigma_{(1)}^{-1}$

Perceived Supply	Monetary Shock	Perceived Demand	Monetary News
0.93	0.02	0.26	0.59
(0.31)	(0.03)	(0.10)	(0.18)

NOTE. This table shows the estimates of the system in equations (5) and (6), estimated via two-step efficient GMM. Panel (a) shows the structural impact matrix, which shows the effect that each structural shock (in the rows) has on forecast revisions made around Fed announcements (in the columns). The shocks are normalized to have unit variance in the first regime (2000–2003). Panel (b) shows the variance of each shock in regime 2 relative to its variance in regime 1, and panel (c) shows the corresponding estimates for regime 3. Standard errors are in parentheses. The p-value for a test of the null hypothesis that the monetary news shock is smaller than the demand information shock in regime 2 is 0.038, and smaller than the monetary shock is 0.001. The p-value for a test that the perceived supply shock is smaller than the perceived demand shock in the third regime is 0.015. The Hansen (1982) J-test for overidentifying restrictions fails to reject those restrictions ($p = 0.594$)

impact responses are consistent standard New-Keynesian macroeconomic theory and the illustrative model. The time series of the estimated shocks, given by $\mathbf{A}^{-1}\widehat{\mathbf{x}}_t$, are shown in figure 3.

4 Results

Having identified the perceived shocks, I now turn to the main question of interest: What are their macroeconomic effects? In section 4.1, I present impulse responses to each shock from a Bayesian proxy SVAR. I also show OLS regressions of the effects of my shocks on Blue Chip forecast revisions and stock returns around Fed announcements. In section 4.2, I discuss several exercises designed to assess the robustness of my findings. Finally, in section 4.3, I show how my estimated effects of monetary policy compare to estimates using alternative shock measures.

Figure 3: Time Series of Estimated Perceived Structural Shocks

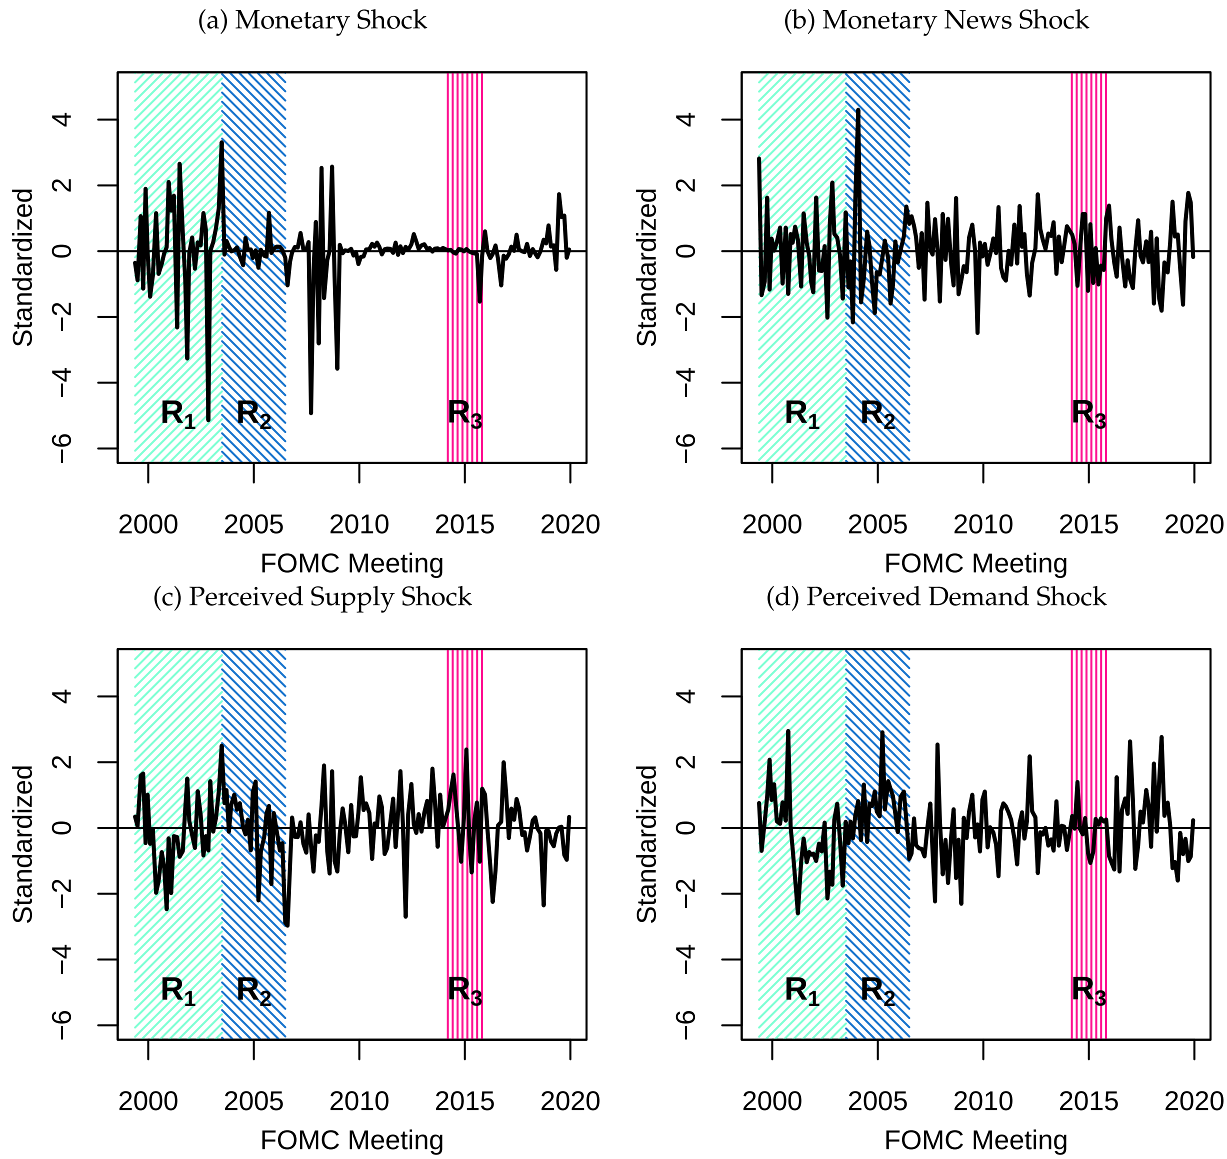

NOTE. This figure shows the time series of perceived structural shocks. Each is normalized to have a mean of zero and unit variance over the sample period shown.

4.1 Main Results

To estimate the macroeconomic effects of each of the perceived shocks, I follow the guidance of Li et al. (2022) and implement a Bayesian VAR that closely resembles that of Caldara and Herbst (2019). I instrument the reduced-form residuals of relevant variables in the system with the perceived structural shock, and plot impulse responses to a one standard deviation shock.

The VAR includes the shadow Federal Funds rate of Wu and Xia (2016), the 2-year treasury yield, the log of the industrial production index, the log of the PCE price index, and the credit spread measure used by Caldara and Herbst.[41] I estimate the VAR in first differences, and plot cumulative impulse responses. Akin to Caldara and Herbst, I impose a Minnesota prior and choose its lag length (3) and hyperparameters to maximize the VAR's marginal data density. I estimate the system's lag coefficients using a sample that starts in October 1982 (the start of Federal Funds rate targeting (Thornton, 2006)), and my proxy variables start in May 1999. I stop both samples in December 2019. Later, I discuss several perturbations to these choices, and find the results to be robust. Figures 4, 5, 6, and 7 show the impulse responses to the monetary shock, the monetary news shock, the perceived demand shock, and the perceived supply shock.

Starting with figure 4, the estimated responses to a monetary shock are roughly in line with what one would expect from a transitory monetary shock in standard models. Instrumenting the shadow rate, I find that an exogenous increase of 10 basis points leads to a fall in inflation and production, and a rise in the unemployment rate and excess bond premium. Longer term rates fall, suggesting that the contractionary effect of the shock weighs more heavily on longer-term rates than the increase in rates itself. Notably, the directions of these effects are the same as those in Caldara and Herbst, though the magnitudes for real variables (industrial production and unemployment) are substantially larger—the 10 basis point increase leads to about a one percent fall in industrial production and 15 basis increase in the unemployment rate. Comparable estimates in Caldara and Herbst are roughly 0.25 percent and 5 basis points. About half of this difference can be attributed to my extended sample period—when I estimate my system through 2008, I find results that are attenuated by about half. Those authors use $\widehat{\mathrm{FF}}_t^0$ as their instrument for the Federal Funds rate, and my monetary shock broadly resembles this instrument. To the extent that my identification procedure strips out offsetting effects by simultaneously emitted shocks, my results should incorporate less measurement error and thus be less attenuated.

Figure 5 shows the response to a monetary news shock. Here, the news shock leads to a persistent increase in the shadow rate of about 12 basis points. In sign, the statistically-significant responses of the other outcome variables largely resemble the responses to the transitory monetary shock of figure 4. Despite being much more persistent, the estimated effects are not drastically different from the response to the monetary shock. Inflation responds by about twice as much, and industrial production responds by about half as much. The larger effect on inflation may reflect the fact that much of the forward guidance about interest rates during the recovery of the Great Recession (especially later on) was targeted towards getting inflation back *up* to two percent. These responses are about twice as large as the estimates of Gertler and Karadi (2015).[42]

The responses to the perceived demand and supply shocks are shown in figures 6 and 7. Both shocks generate co-movements between nominal and real variables that is consistent with how demand and supply shocks operate in typical models. When statistically significant, the demand shock causes a positive co-movement, while the supply shock causes a negative co-

[41]Other than the shadow rate, the other variables are all from FRED and have mnemonics, respectively, of DGS2, INDPRO, PCEPI, and the difference of BAA and DGS10.

[42]The size of the responses of inflation and industrial production are about the same, but the size of the initial impulse (the interest rate response is) only half as large.

Figure 4: Response to a Monetary Shock

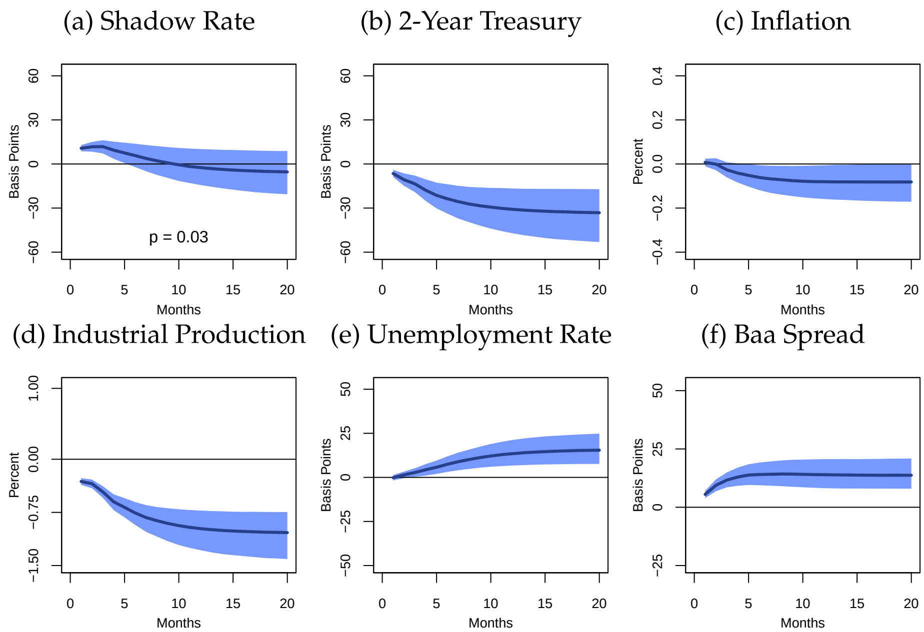

NOTE. This figure shows the impulse responses to a one standard deviation monetary shock, estimated using the baseline VAR. The p-value is shown on the variable whose reduced-form residuals are being instrumented. 90% credible regions are shaded, and the posterior median is shown with a solid line. See the text for other details.

Figure 5: Response to a Monetary News Shock

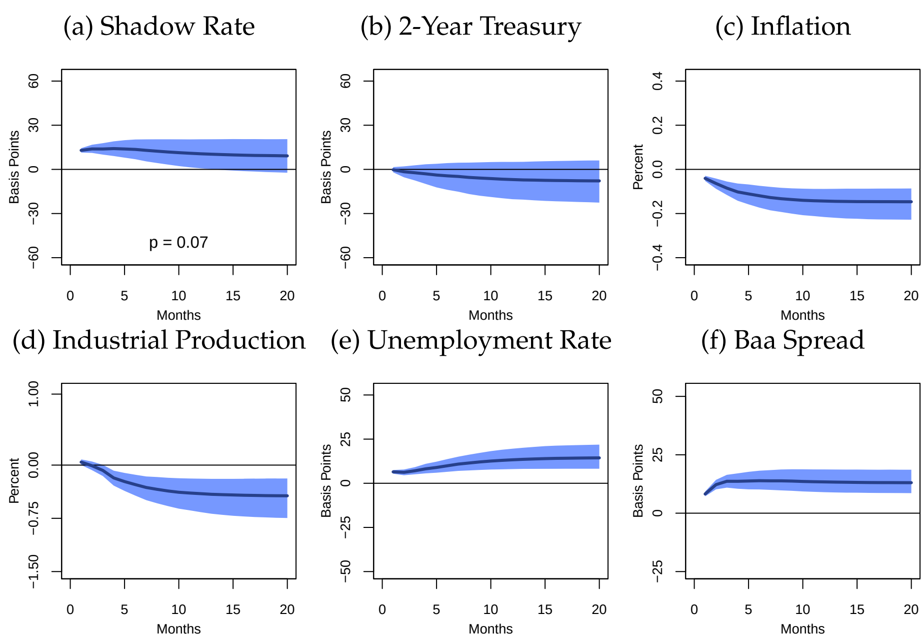

NOTE. This figure shows the impulse responses to a one standard deviation monetary news shock, estimated using the baseline VAR. The p-value is shown on the variable whose reduced-form residuals are being instrumented. 90% credible regions are shaded, and the posterior median is shown with a solid line. See the text for other details.

Figure 6: Response to a Perceived Demand Shock

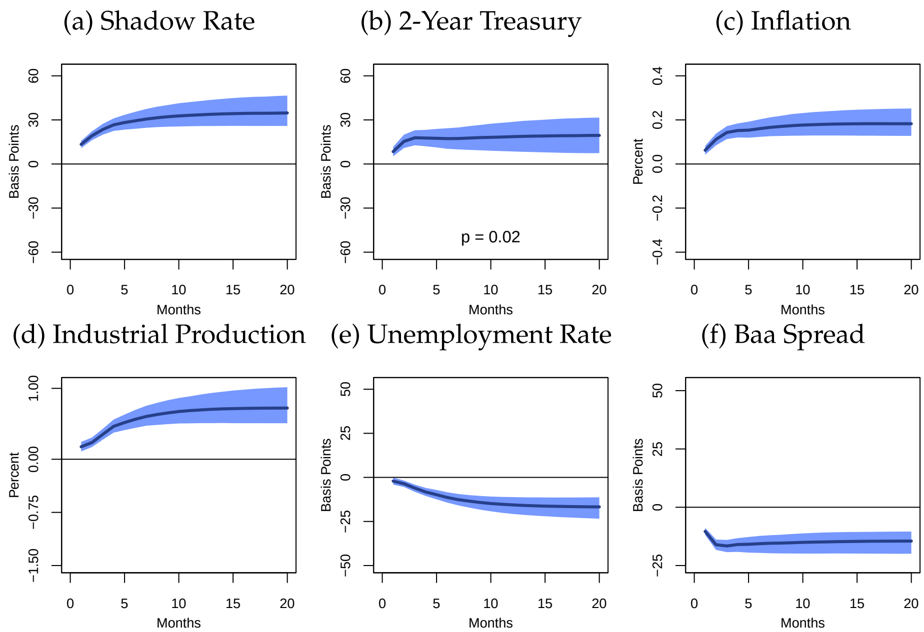

NOTE. This figure shows the impulse responses to a one standard perceived demand shock, estimated using the baseline VAR. The *p*-value is shown on the variable whose reduced-form residuals are being instrumented. 90% credible regions are shaded, and the posterior median is shown with a solid line. See the text for other details.

Figure 7: Response to a Perceived Supply Shock

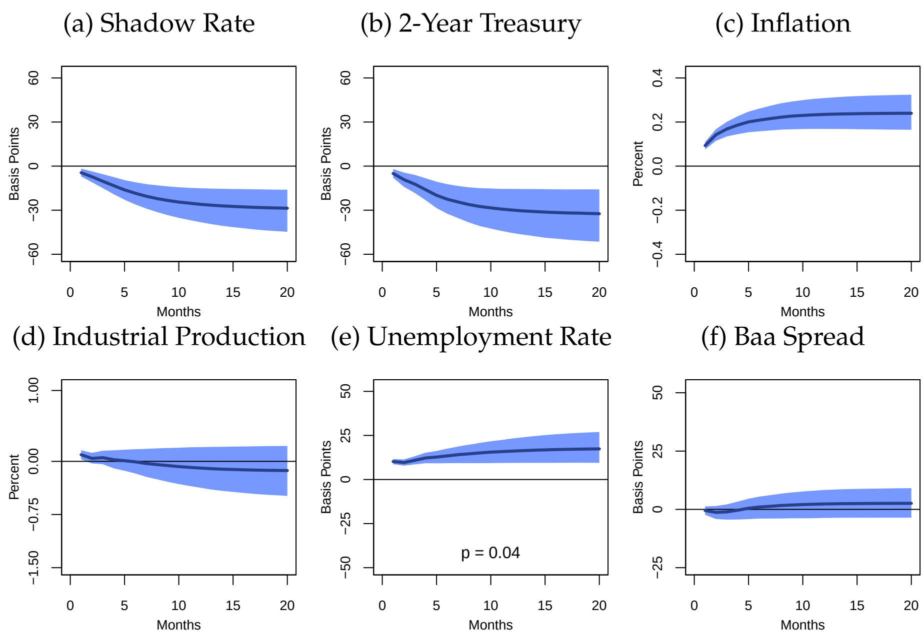

NOTE. This figure shows the impulse responses to a one standard deviation perceived supply shock, estimated using the baseline VAR. The *p*-value is shown on the variable whose reduced-form residuals are being instrumented. 90% credible regions are shaded, and the posterior median is shown with a solid line. See the text for other details.

Table 4: Effects on Other Variables (OLS)

	$\widehat{\text{FF}}_t^0$	$\widehat{\text{ED}}_t^4$	$\Delta\,\mathbb{E}^{\text{BC}}[y]$	$\Delta\,\mathbb{E}^{\text{BC}}[\pi]$	$\Delta\,\mathbb{E}^{\text{BC}}[u]$	Δ^{HF}S&P 500	$\Delta\text{nominal}_t^{10}$	Δreal_t^{10}
Monetary News	0.001	0.049	-0.018	-0.017	0.001	-0.166	0.015	0.016
	(0.000)	(0.000)	(0.016)	(0.014)	(0.008)	(0.048)	(0.005)	(0.005)
Perceived Supply	0.001	0.005	0.032	-0.016	-0.008	0.008	0.002	0.005
	(0.000)	(0.000)	(0.015)	(0.015)	(0.007)	(0.030)	(0.005)	(0.004)
Perceived Demand	0.003	0.050	0.039	0.033	-0.020	-0.104	0.026	0.025
	(0.000)	(0.000)	(0.017)	(0.016)	(0.008)	(0.040)	(0.004)	(0.005)
Monetary	0.039	0.017	-0.021	-0.046	0.005	-0.154	0.001	0.006
	(0.000)	(0.000)	(0.025)	(0.021)	(0.013)	(0.061)	(0.005)	(0.005)
Observations	165	165	139	139	139	165	165	165
R^2	1.000	1.000	0.086	0.126	0.042	0.195	0.133	0.136

NOTE. Each column shows OLS regressions in which the variable atop the column is on the left-hand side, and the four perceived shocks are on the right-hand side. Each shock is scaled to have unit standard deviation over the 165-observation sample consisting of all regularly-scheduled Fed meetings between May 1999 and December 2019, and normalized to have a positive effect on $\widehat{\text{ED}}_t^4$. The first two columns show the responses of $\widehat{\text{FF}}_t^0$ and $\widehat{\text{ED}}_t^4$ (which go into the construction of the shocks, hence the R^2 of unity). The next three columns have forecast revisions of GDP, CPI inflation, and unemployment from the Blue Chip survey on the left-hand side, all in percentage points. These regressions drop Fed meetings that occur in the first week of the month, and the forecast revisions are averaged over closest three quarters. The next column shows (100 times) the 30-minute change in the log price of E-S&P mini futures, i.e., the high-frequency return in the S&P 500. The final two columns show one-day percentage point changes in the 10 year nominal and real interest rates, respectively, from Gürkaynak et al. (2006) and Gürkaynak et al. (2008). Heteroskedasticity-robust standard errors are shown in parentheses.

movement. These results are consistent with the notion that Fed announcements reveal information about these shocks that turns out to be true.

I find broadly similar findings when I move from VAR estimates, to OLS regressions that examine the effects of my identified shocks on macroeconomic expectations. Those results are shown in table 4. The first column shows the effect of each perceived shock on the high-frequency forecast revision for $\widehat{\text{FF}}_t^0$ and $\widehat{\text{ED}}_t^4$. Consistent with the VAR evidence, monetary and monetary news shocks that raise short-term rates lead to (generally imprecisely estimated) downward revisions in GDP and inflation expectations, and increases in unemployment rate expectations. The demand and supply information shocks generate co-movements between expectations of real (unemployment and real GDP) and nominal (inflation) variables that are consistent with these revealing information about shocks that ultimately materialize. In magnitude and precision, the perceived supply and demand shocks tend to have larger effects on expectations than the monetary shocks.

One finding from table 4 that does not fit the general pattern is the response of the S&P 500. When it moves by a statistically significant amount, it co-moves negatively with longer-term interest rates. This is particularly surprising for the perceived demand shock, especially because the work of Jarociński and Karadi (2020) uses positive co-movements between the S&P 500 and Eurodollar futures to identify an "information shock" that Jarociński and Karadi posit resembles information about demand. Stock market and yields co-move tightly over the sample, so the scope for identifying an orthogonal shock seems limited, a priori. Indeed, when Jarociński and

Karadi present the identified set of the effects of monetary policy and information effects, the identified effects of their information shock seldom excludes zero. Thus, empirically the S&P 500 does not seem well suited to separately identify the effects of information and monetary shocks. In appendix I, I confirm that applying the heteroskedasticity-based assumptions to the Jarociński and Karadi data leads to poorly-identified shocks. This is also not surprising given the empirical work that documents the limited role for the stock market to predict real activity.[43]

4.2 Robustness

In appendix J, I present estimates designed to assess the robustness of my baseline VAR results. Painting broad strokes, I find robust evidence that increases in short-term interest rates caused by my monetary and monetary news shocks lead to contractionary outcomes. The perceived supply shock typically causes a negative co-movement between nominal and real variables, though in some instances it suffers from weak instrument problems. The responses to the perceived demand shock are broadly robust, with one notable exception that I discuss after briefly mentioning the other robustness exercises.

My first set of robustness checks center around the specification of the VAR. I estimate my VAR in levels, rather than first differences, and plot non-cumulative impulse response functions. I then estimate a VAR that ends in 2008, so that the shadow Fed Funds rate and effective Fed Funds rate coincide. I also consider a small VAR in the shadow rate, industrial production, and inflation. Across these specifications, the results are broadly similar to the baseline. One notable exception is the relevance of the perceived supply and demand shocks as instruments before 2008, which declines markedly, severely so for the perceived supply shock. This suggests that any effects that these shocks have in the baseline come from the post-Great Recession period. The irrelevance of the perceived supply shock as an instrument is a feature of many of the alternative specifications. This is consistent with the work of Melosi (2017), whose estimated DSGE model with Fed signalling effects suggests that relatively little is learned about supply from the Fed.

I also consider several perturbations to the shocks themselves. I estimate a system in which the text-based measures capture longer-term expectations. The results are similar to the baseline. This suggests that the exact horizon of the information (whether about future policy or economic fundamentals) is not crucial. This is not surprising given the persistence that is present in most macroeconomic data.

I then estimate shocks making only use of heteroskedasticity in $\mathcal{R}_1$ and $\mathcal{R}_2$. The main noticeable difference is that the effects of the perceived supply and demand are somewhat confounded. The perceived demand shock has a weaker effect on inflation, while the perceived supply shock no longer generates a strong negative co-movement between nominal and real variables. This should not be surprising—the third regime is included specifically to aid in the identification of the perceived supply shock. As table 3 showed, the perceived supply and demand shocks exhibit similar heteroskedasticity between the first two regimes. Adding the third regime helps to disentangle them.

[43] A recent example is the work of Chatelais et al. (2022).

The next set of robustness checks relate to the text analysis. In particular, I examine the effects of broadening the window used for determining which words modify one another. The results are little changed. Using three words for inflation and GDP, rather than five, also leads to similar results. Excluding "oil price" from the perceived supply shock weakens its effect on inflation, but otherwise makes little difference.

The final set of tests relate to the inclusion of additional controls. These exercises are motivated by the recent work of Bauer and Swanson. Briefly, those authors find that a few variables importantly modulate the response of macroeconomic expectations to high-frequency interest-rate forecast revisions, like $\widehat{\mathrm{ED}}_t^4$. I discuss the paper in more detail in appendix G but, to summarize, the authors claim that this predictability refutes the notion that the Fed's announcements reveal private information. I take a different view in this paper and in appendix G—predictability of interest rate surprises by a variable z might suggest that z reveals something useful regarding the variables which the Fed judges to be important when making an assessment of the economic outlook. I take this view in part because of the fact that the set of variables that meaningfully modulate the effects of these surprises are outside of the typical controls one might consider (the VAR, for example, contains several lags of controls for standard variables). In my robustness exercises, I orthogonalize my high-frequency forecast revisions to the controls suggested by Bauer and Swanson. I find that most of the estimated responses are unchanged, except for the responses to the perceived demand shock—this is true in the VAR (appendix figure 26) and (to a lesser extent) in my OLS estimates (appendix table 10). When I examine these controls in more detail (appendix figure 16), I find that the 13-week return on the S&P 500 is the primary source of the change. As I discuss in appendix G, this control is very important for the Bauer and Swanson findings.

4.3 Comparison with Existing Studies

I conclude my empirical exercises with a comparison of my estimated effects of monetary policy to those estimated using existing measures of monetary policy shocks. In the top row of figure 8, I show the effects of shocks that are similar to my monetary shock—the "target" factor of Gürkaynak et al. (2005) (henceforth, GSS), and the monetary shock of Romer and Romer (2004). Generally, the shocks provide similar estimates, at least in sign. That my monetary shock and the GSS target factor have similar responses again reflects the fact that they are both predominantly driven by the surprise in the Federal Funds rate.

The second row shows responses using shocks designed to capture forward guidance—my monetary news shock, the cleaned shock of Miranda-Agrippino and Ricco, and the "path" factor of GSS. My shock and that of Miranda-Agrippino and Ricco generally agree on the sign of the effect. As I show in the context of the model below, the approach of Miranda-Agrippino and Ricco is a useful way to identify monetary policy (and news) shocks. Where that approach falls somewhat short is in that it treats any residual component as a single shock. I decompose that residual here. One consequence here, even focusing only on monetary policy, is that the Miranda-Agrippino and Ricco shock conflates my monetary and monetary news shocks. My evidence suggests that the large response of industrial production, for example, likely arises more from the

Figure 8: Comparing Impulse Responses

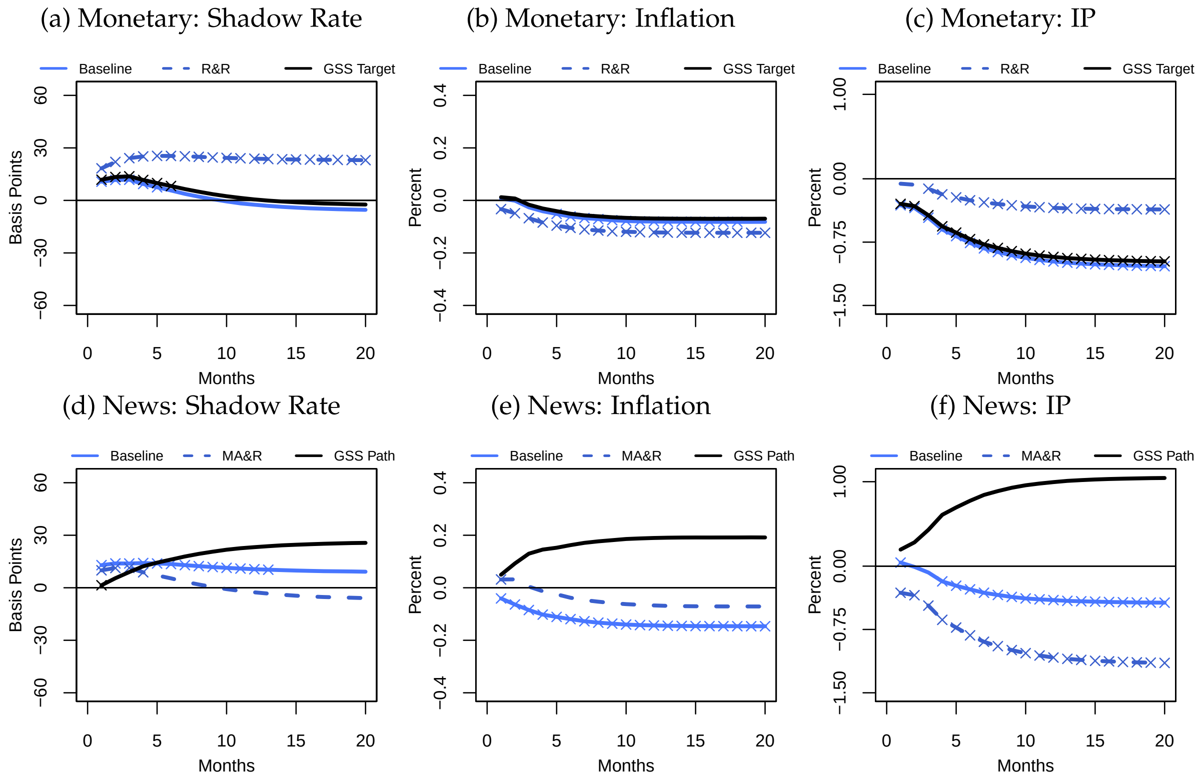

NOTE. This figure shows impulse responses of the shadow rate, inflation, and industrial production in the baseline VAR using different measures of monetary shocks as instruments. (The remaining responses are shown in figure 27 in the appendix.) The top row considers different measures of monetary shocks: my baseline estimate, the "target" factor of GSS, and the monetary shock of Romer and Romer, all updated through 2019. All three measures are used to instrument the shadow rate. The bottom row presents analogous measures using different measures of monetary news shocks: my baseline, the cleaned shock of Miranda-Agrippino and Ricco, and the "path" factor of GSS. I instrument the shadow rate with the first two, and the two-year treasury with the latter (otherwise, the latter is a weak instrument). In each figure, an "x" is drawn whenever the 90% credible set excludes zero.

former.

The most striking response in figure 8 is the one induced by the GSS path shock. In response to what would presumably be a contractionary shock, an expansion ensues. This should not be surprising—it is the type of evidence that suggested that forecast revisions of longer-term interest rates may reflect more than purely exogenous monetary policy. In appendix F, I show that the longer the interest rate maturity used, the more puzzling the responses. It does, however, illustrate the benefit of decompositions of the nature I pursue.

5 Discussion

In this section, I present a model containing two simple extensions to the textbook New Keynesian model presented in Galí (2015). The first modification is to include both transitory and persistent monetary shocks. The second is to make all agents imperfectly informed about economic shocks, and to be explicit about where they learn about those shocks—either from the Fed, or some other data release. I first describe the model and show that its impulse response functions (IRFs) are consistent with my estimates. I then use the model to assess the validity of my identification assumptions, and conclude by using the model to compare my approach to others in the literature.

5.1 Model

The economic structure of the model I consider is nearly identical to the textbook New Keynesian model of (Galí, 2015, Ch.3). The economy can be described by the following set of stochastic difference equations

$$\begin{aligned} \pi_t &= \beta\,\mathbb{E}_t[\pi_{t+1}] + \kappa \widetilde{y}_t \\ \widetilde{y}_t &= \mathbb{E}_t[\widetilde{y}_{t+1}] - \frac{1}{\sigma}(i_t - \mathbb{E}_t[\pi_{t+1}]) - \frac{1}{\sigma}\,\mathbb{E}_t[\Delta z_{t+1}] + \mathbb{E}_t[\Delta a_{t+1}] \\ i_t &= \phi_\pi \pi_t + \phi_y \widehat{y}_t + m_t + \phi_t. \end{aligned}$$

Where i_t is the nominal interest rate, $\widehat{y}_t$ is log output, $\widetilde{y}_t$ is the log output gap, and π_t is inflation, all expressed in deviations from their steady state values. The economy is driven by four mutually independent structural shocks that follow AR(1) processes: m_t is a monetary shock, ϕ_t is a persistent monetary shock, a_t is a supply (TFP) shock, and z_t is a demand (discount factor) shock. I calibrate the model's structural parameters exactly as in Galí (2015)—a full description is given in appendix E. The only exception are the monetary shocks: The autocorrelations of m_t and ϕ_t are, respectively, 0.5 and 0.9.

I enrich the information structure of the model to bring it closer to my empirics. Agents in the model have rational expectations, but not full information. I allow for the possibility that the Fed possesses an informational advantage by imposing the following timing restriction. Specifically, I assume that each discrete period, indexed by integers t, is split into the three segments. First, at time t, all shocks are realized, and the Fed privately makes its interest rate decision. Second, at time $t_F > t$, the Fed makes its policy announcement, endowing firms with independent but noisy signals of the model's four structural shocks. The interest rate decision is also announced, but its information content can also be obscured by independent noise. Third, at time $t_P > t_F$, households and firms become fully informed and make all relevant decisions (pricing, consumption, etc.). The model collapses to the textbook model as the noise components of the Fed's signals increase.

The fact that the Fed communicates about an endogenous variable complicates the solution of the model. The procedure I pursue follows that of Melosi (2017). Briefly, the solution of the model is assumed to take the form

$$x_t = M\eta_{t|t} \tag{7}$$

where x_t are the model's endogenous variables ($\widetilde{y}_t, i_t$, and π_t), and $\eta_{t|t}$ are agents' perceptions of the model's structural shocks, and high-order expectations of those shocks. Those expectations are assumed to follow a VAR(1) process

$$\eta_{t|t} = L\eta_{t-1|t-1} + N\xi_t \tag{8}$$

where ξ_t are innovations to the economy's structural shocks (including any signal noise). Agents receive signals about not only ξ_t directly,[44] but linear combinations thereof in the form of signals about endogenous variables. The matrices $M, L,$ and N—which form the model's solution—are thus jointly determined. They can be found by iterating back and forth between equations (7) and (8) until convergence.

The thick gray lines in figure 9 show impulse responses in the textbook model, i.e., the model with no Fed information advantage.[45] In this model, there is no distinction between a true and a perceived structural shock, so it serves as a useful benchmark for understanding whether the distinction is important. The responses to the monetary, supply, and demand shocks are standard—Galí (2015) provides a nice explanation. They also broadly line up with my estimates, as promised. The response to the persistent monetary shock is unusual but known—a positive (contractionary) persistent monetary shock leads to contractionary responses in output and inflation, but causes interest rates to fall to offset the contraction. However empirically implausible this may seem, it helps to explain the somewhat weak relationship between my monetary news shock and longer-term interest rates. Taken together, the model suggests that my estimates are largely consistent with the textbook theory.

5.2 Identification by Heteroskedasticity

I next use the model as a laboratory in which to assess an important heteroskedasticity-based identification (HBI) assumption that underlies my empirics: that the perceived structural shocks are uncorrelated.[46] Correlation in perceived structural shocks around the Fed's announcements will depend on the extent to which the signals of each shock are uncorrelated with each other—this is a standard consequence of Bayesian updating. This will not be the case if the *only* information conveyed at a policy announcement is the interest rate decision. Because the interest rate is an endogenous variable, changes can be driven by any of the economy's structural shocks to which the Fed is privy, and so it serves as a signal about a linear combination of fundamental shocks. This, in turn, causes forecast revisions about each shock (perceived structural shocks) to be correlated.

To assess how consequential correlated shocks may be for my results, I carry out the following experiment in the model to increase the extent of correlation. As mentioned above, I model Fed announcements as independent signals of all four structural shocks, and the release of the interest rate decision. Across ten calibrations of the model, I vary the precision of these signals in a way that increases the correlation in perceived forecast revisions. Within each calibration, I simulate 50,000 time periods. For the first 25,000 periods, all structural shocks have the same variance. During the second, I change the variances in a way consistent with my empirical exercise so that I can use HBI to identify the perceived structural shocks.[47] As in my empirics, I use HBI to estimate

[44] Agents have rational expectations, so they use the Kalman filter to update expectations. I have two separate Kalman filters set up to handle the two signaling "events:" Fed press conferences (t_F), and the later release of economic news (t_P).

[45] The information structure here is simple: all agents are symmetrically and perfectly informed and make decisions simultaneously.

[46] The other two main assumptions are that (1) there is heteroskedasticity in the data and (2) that the matrix $\mathbf{A}$ has stable coefficients across regimes. I test the first in the data, and find supportive evidence. The over-identification tests suggest that the second assumption is reasonable.

[47] Specifically, I reduce the standard deviation of innovations to the monetary shock by 95%, reduce the standard

Figure 9: HBI in the Model with Correlated Perceived Structural Shocks

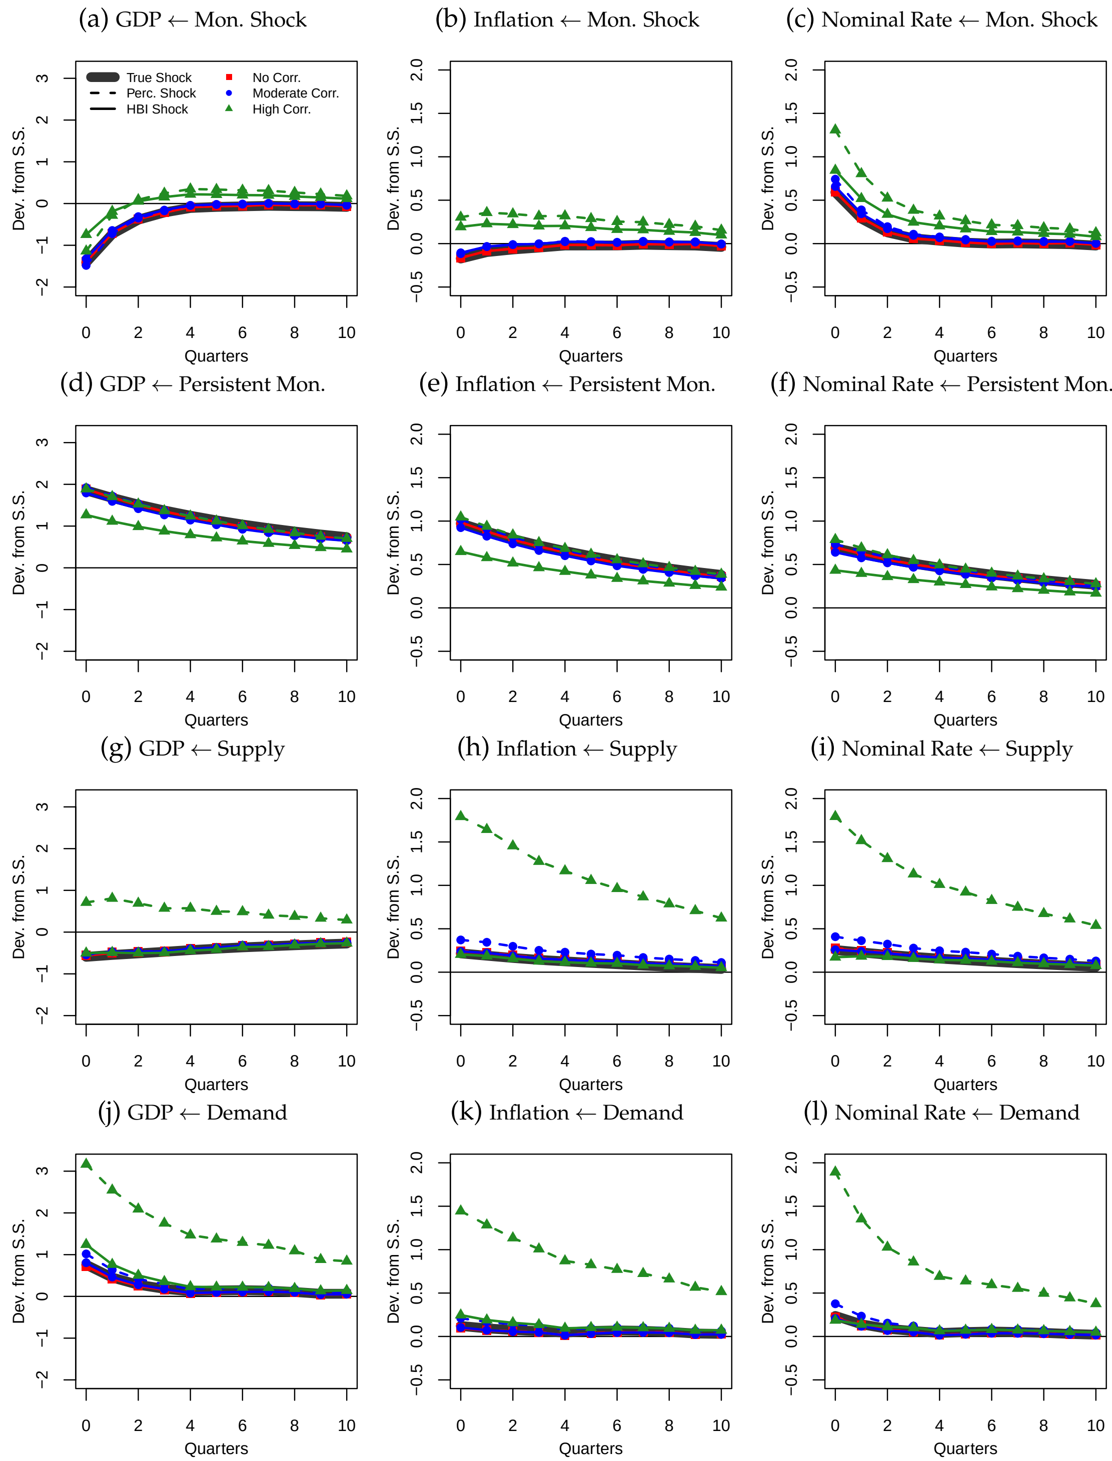

NOTE. Each panel shows the model-generated impulse responses of GDP (in percent), inflation (in percent), and the nominal interest rate (in percentage points) to each of the four structural shocks (monetary, monetary news, supply, and demand). The solid lines show the impulse response to each shock in the full-information model. Each dashed line shows the response to the perceived structural shocks under different calibrations, and the solid lines show responses to perceived shocks identified by heteroskedasticity. See the text for details on the calibration. Impulse responses are estimated using local projections on 50,000 simulated data points.

the perceived structural shocks—I call these estimates "HBI shocks" for short—using forecast revisions of output and inflation in the current period, and current- and four-quarter ahead interest rates. I then estimate local projections of the form $x_{t+k} = \alpha_k + \beta_k \varepsilon_t + e_t$, where x is either inflation, output or the nominal rate, and ε_t is one of the perceived shocks (demand, supply, or the two monetary shocks), or the perceived shocks identified by heteroskedasticity. I plot β_k for $k \in \{0, \ldots, 10\}$ in figure 9.

I present results for three representative calibrations in figure 9. In the first, labeled "No Corr.," the Fed's announcements consist of only independent signals about the structural shocks. Here, perceived, structural, and HBI shocks have identical effects to the true shocks—this is by design. In the third, "High Corr.," Fed announcements reveal the interest rate, and the noise in the other signals is such that about 70% of belief updating is done using the interest rate. Here there is some separation between the effects of the (actual and HBI) perceived shocks and the true shocks. The perceived contractionary monetary shock, for example, suggests an increase in inflation. As shown in appendix E, this arises because the perceived shock is somewhat positively correlated with an expansionary demand shock. The effects under the persistent perceived shock are somewhat attenuated, but generally consistent with the true effects. The demand and supply shocks show a dramatically different pattern. The coefficients on the perceived shocks are much larger than the true coefficients. Forecast revisions here are much smaller, so they appear to have larger effects on observables. What is remarkable, however, is that the estimated effects of the HBI shocks are closer to the true effects than the estimated effects of the perceived shocks. Table 8 in the appendix confirms that the HBI shocks tend to be more highly correlated with the true shocks (and less correlated with other shocks) than the perceived shocks. Here, HBI is serving to filter out the noise from the Fed's signal that remains embedded in the perceived shocks.[48] Thus, despite the fact that HBI is providing estimates of theoretically non-orthogonal shocks, it is effectively identifying the fundamental component that *is* uncorrelated: namely, the structural shocks.

5.3 Other Approaches to Identification

My final use of the model is to briefly comment on how my approach relates to other approaches in the literature. In figure 10, I show estimated responses to different measures of monetary shocks. First, I show an approach similar to Kuttner (2001)—simply using the forecast revisions of interest rates. Second, I use an approach similar to Nakamura and Steinsson (2018): forecast revisions of longer-term rates. Third, I orthogonalize interest rate forecast revisions to all other non-monetary structural shocks as an approximation to the approach of Miranda-Agrippino and Ricco (2021). Finally, I present the identified set of impulse responses using a sign restriction approach similar to Jarociński and Karadi (2020), in which a monetary shock causes GDP and longer-term interest

deviation of innovations to the demand shock by half, leave the supply shock variance unchanged, and double the variance of innovations to the persistent monetary shock. I could have equivalently modified the variance of the noise in the Fed's signals about each variable, but since I am varying these across different calibrations, this approach was more straightforward.

[48] This is a traditional use of HBI—see e.g. Nakamura and Steinsson (2018) or Hébert and Schreger (2017)—to remove background noise from data. The structural approach I pursue here presents a new use in the high-frequency identification literature.

Figure 10: The Effects of Monetary Policy: Comparing Identification Procedures in the Model

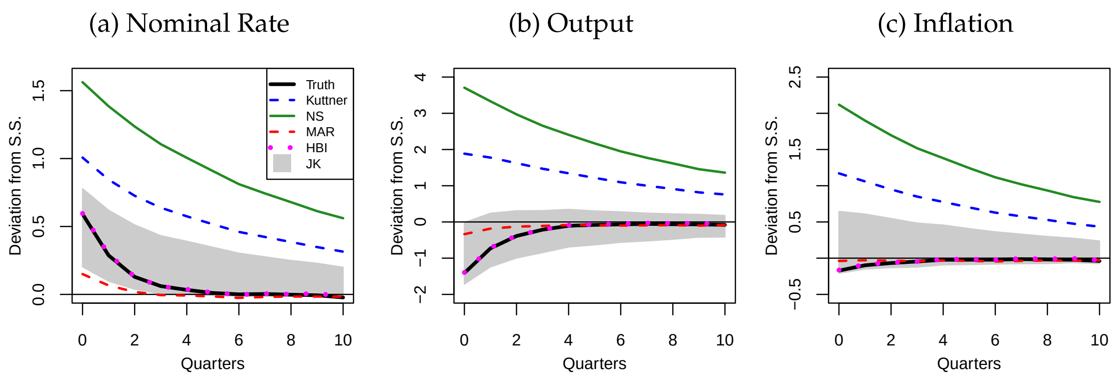

NOTE. This figure shows the impulse response of the nominal rate (in percentage points), output (in percent) and inflation (and percent) to shocks identified under different assumptions. The lines labeled "Kuttner" refer to forecast revisions of the current-period nominal rate (Kuttner, 2001), and those labeled "NS" refer to forecast revisions of interest rates four-quarters ahead (Nakamura and Steinsson, 2018). The lines labeled MAR orthogonalize the Kuttner measure with-respect to all non-monetary structural shocks, and the gray-shaded region shows the identified set of impulse responses to a shock that increases 4-quarter-ahead interest rate expectations, and decreases output expectations (Jarociński and Karadi, 2020). Impulse responses are estimated using local projections on 50,000 simulated data points.

rate forecast revisions to move in opposing directions.

I show the results in figure 10. Starting with the approaches that do not adjust for the presence of other shocks (Kuttner, 2001; Nakamura and Steinsson, 2018), the identification challenge emerges immediately. What should be a contractionary monetary shock instead resembles an expansionary demand shock. That is even more pronounced when using forecast revisions of longer-term rates, whose forecast revisions will tend to weigh more-heavily on the persistent monetary shock, whose impulse responses in turn are similar to what is seen in the graph. This pattern also holds in the data—in appendix F, I show that, in the VAR, the longer the maturity of the interest rate used as a proxy for monetary shocks, the more the estimated results resemble a demand or persistent monetary shock.

My implementation of the approach of Miranda-Agrippino and Ricco (2021) works fairly well in identifying the true effects of the shock. It is worth noting, however, that this approach works well because I am able, in the model, to exactly control for the non-monetary component of the interest rate by controlling for the other structural shocks. In practice, this can present a challenge because either one has to posit which variables span the variables over which the Fed may have private information, or else one risks over controlling, in the sense that some of the controls may themselves reflect the monetary shock to the extent that some of this is known to the staff. In any case, this approach is useful for identifying the effects of monetary policy shocks, but is not suited to identifying the effects of the other perceived shocks.

Finally, the gray-shaded region shows the identified set of impulse responses using the sign-restriction analog to the work of Jarociński and Karadi (2020). Here, because the restriction is imposed on the sign of the co-movement between output and interest rates, part of the perceived supply shock can enter since it too satisfies the sign restrictions. As such, the response of inflation is not well identified. This is the "masquerading shocks" problem of Wolf (2020). While the other identified sets contain the true impulse response, the range of responses they identify is quite wide and largely does not exclude zero, consistent with the version of the impulse responses of Jarociński and Karadi (2020) that removes the role of the prior in selecting an impulse response.

6 Conclusion

Estimating the macroeconomic effects of monetary policy is notoriously difficult, because interest rates are so highly endogenous with respect to macroeconomic variables. Despite several approaches in the literature aimed at identifying exogenous changes in interest rates (i.e., monetary policy shocks) using high-frequency data, the identification of a series that has theoretically consistent effects has proven to be quite challenging.

In this paper I provide a detailed decomposition of the information perceived by markets when the Fed makes its policy announcements. This decomposition reveals four "perceived" shocks: a monetary shock, a monetary news shock, a demand shock, and a supply shock. To separately identify these shocks, I study four data series: high-frequency interest-rate (short and long term), inflation, and GDP forecast revisions. These are the variables suggested by standard macroeconomic theory. However, the latter two did not exist, so I constructed them using newspaper articles written about Fed policy meetings.

The notion that four series do not identify four shocks, without additional identifying assumptions, is an old one in economics. I present a model of expectation revisions around Fed announcements, and show what assumptions have been made (largely implicitly), by those who have previously sought to identify monetary shocks in high frequency. Like zero restrictions in structural VARs, these are generally not supported by macroeconomic theory. My assumptions, identification by heteroskedasticity, are much less restrictive and better suited to this context. The introduction of forward guidance by the Fed, a major regime change in Fed communications, drastically changed the shocks that markets could learn about from Fed policy announcements. Paired with a regime that witnessed a large drop in oil prices to which the Fed devoted substantial time in deciphering for markets, identification by heteroskedasticity provides a tool to turn these regimes into structural identification.

Empirically, my identified shocks have macroeconomic effects that are consistent with standard theory. While other authors have identified series of monetary shocks with this feature, they have either ignored, or used inappropriate variables to identify the other perceived shocks. That the effects of the perceived demand and supply shocks have macroeconomic effects consistent with standard theory suggests that the Fed's communicated predictions ultimately materialize. Does this mean that the Fed has "better" information or a "better" reading of publicly available data? I cannot say. I present substantial evidence in this paper, however, that these non-monetary perceived shocks are not the endogenous object suggested by Bauer and Swanson. Further research should be devoted to understanding why the Fed appears to play such an important role in providing macroeconomic information, and how such communications can be further optimized.

References

Andrade, Philippe and Filippo Ferroni, "Delphic and Odyssean monetary policy shocks: Evidence from the euro area," *Journal of Monetary Economics*, 2021, *117*, 816–832.

Angrist, Joshua D and Jörn-Steffen Pischke, *Mostly harmless econometrics: An empiricist's companion*, Princeton university press, 2008.

Arai, Natsuki, "The effects of monetary policy announcements at the zero lower bound," *International Journal of Central Banking*, 2017, *13* (2), 159–196.

Bauer, Michael D. and Eric T. Swanson, "An Alternative Explanation for the "Fed Information Effect"," *American Economic Review*, Forthcoming.

Baumeister, Christiane and Lutz Kilian, "Lower Oil Prices and the U.S. Economy: Is This Time Different?," *Brookings Papers on Economic Activity*, 2016, pp. 287–336.

Blanchard, Olivier J., Jean-Paul L'Huillier, and Guido Lorenzoni, "News, Noise, and Fluctuations: An Empirical Exploration," *American Economic Review*, December 2013, *103* (7), 3045–70.

Cai, Yong, Santiago Camara, and Nicholas Capel, "It's not always about the money, sometimes it's about sending a message: Evidence of Informational Content in Monetary Policy Announcements," Working Paper 2021.

Caldara, Dario and Edward Herbst, "Monetary Policy, Real Activity, and Credit Spreads: Evidence from Bayesian Proxy SVARs," *American Economic Journal: Macroeconomics*, January 2019, *11* (1), 157–92.

Campbell, Jeffrey R, Charles L Evans, Jonas D.M. Fisher, and Alejandro Justiniano, "Macroeconomic effects of Federal Reserve forward guidance," *Brookings Papers on Economic Activity*, 2012, pp. 1–80.

Chatelais, Nicolas, Arthur Stalla-Bourdillon, and Menzie D Chinn, "Macroeconomic Forecasting using Filtered Signals from a Stock Market Cross Section," Working Paper 30305, National Bureau of Economic Research July 2022.

Cieslak, Anna and Andreas Schrimpf, "Non-monetary news in central bank communication," *Journal of International Economics*, 2019, *118*, 293 – 315.

Coibion, Olivier, "Are the effects of monetary policy shocks big or small?," *American Economic Journal: Macroeconomics*, 2012, *4* (2), 1–32.

D'Amico, Stefania and Thomas King, "What Does Anticipated Monetary Policy Do?," Working Paper 2017.

Doh, Taeyoung, Dongho Song, and Shu-Kuei Yang, "Deciphering Federal Reserve communication via text analysis of alternative FOMC statements," Working Paper 2020.

Galí, Jordi, *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*, Princeton University Press, 2015.

Gertler, Mark and Peter Karadi, "Monetary Policy Surprises, Credit Costs, and Economic Activity," *American Economic Journal: Macroeconomics*, January 2015, *7* (1), 44–76.

Grigoli, Francesco, Alexander Herman, and Andrew Swiston, "A crude shock: Explaining the short-run impact of the 2014–16 oil price decline across exporters," *Energy Economics*, 2019, *78*, 481–493.

Gürkaynak, Refet S., A. Hakan Kara, Burçin Kısacıkoğlu, and Sang Seok Lee, "Monetary policy surprises and exchange rate behavior," *Journal of International Economics*, 2021, *130*, 103443. NBER International Seminar on Macroeconomics 2020.

—, Brian P. Sack, and Eric T. Swanson, "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements," *International Journal of Central Banking*, 2005.

—, —, and Jonathan H. Wright, "The U.S. Treasury Yield Curve: 1961 to the Present," *Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series*, 2006, *28.*

—, —, and —, "The TIPS Yield Curve and Inflation Compensation," *Board of Governors of the Federal Reserve System, Finance and Economics Discussion Series*, 2008, *5.*

—, Burçin Kisacikoğlu, and Jonathan H. Wright, "Missing Events in Event Studies: Identifying the Effects of Partially Measured News Surprises," *American Economic Review*, 2020, *110* (12), 3871–3912.

Handlan, Amy, "Text Shocks and Monetary Surprises: Text Analysis of FOMC Statements with Machine Learning," Working Paper 2020.

Hansen, Lars Peter, "Large sample properties of generalized method of moments estimators," *Econometrica: Journal of the econometric society*, 1982, pp. 1029–1054.

Hansen, Stephen and Michael McMahon, "Shocking language: Understanding the Macroeconomic Effects of Central Bank Communication," *Journal of International Economics*, 2016, *99.*

Hébert, Benjamin and Jesse Schreger, "The costs of sovereign default: Evidence from argentina," *American Economic Review*, 2017, *107* (10), 3119–45.

Jarociński, Marek and Peter Karadi, "Deconstructing Monetary Policy Surprises—The Role of Information Shocks," *American Economic Journal: Macroeconomics*, April 2020, *12* (2), 1–43.

Kuttner, Kenneth N, "Monetary policy surprises and interest rates: Evidence from the Fed funds futures market," *Journal of Monetary Economics*, 2001, *47* (3), 523–544.

Lakdawala, Aeimit, "Decomposing the effects of monetary policy using an external instruments SVAR," *Journal of applied econometrics*, 2019, *34* (6), 934–950.

Lewis, Daniel J, "Announcement-Specific Decompositions of Unconventional Monetary Policy Shocks and Their Macroeconomic Effects," *FRB of New York Staff Report*, 2019, (891).

Li, Dake, Mikkel Plagborg-Møller, and Christian K Wolf, "Local projections vs. VARs: Lessons from thousands of DGPs," Working Paper, National Bureau of Economic Research 2022.

Loughran, Tim and Bill McDonald, "When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-Ks," *The Journal of Finance*, 2011, *66* (1), 35–65.

Lunsford, Kurt G, "Policy language and information effects in the early days of Federal Reserve forward guidance," *American Economic Review*, 2020, *110* (9), 2899–2934.

Melosi, Leonardo, "Signalling effects of monetary policy," *The Review of Economic Studies*, 2017, *84* (2), 853–884.

Miranda-Agrippino, Silvia and Giovanni Ricco, "The Transmission of Monetary Policy Shocks," *American Economic Journal: Macroeconomics*, July 2021, *13* (3), 74–107.

Nakamura, Emi and Jón Steinsson, "High Frequency Identification of Monetary Non-Neutrality," *Quarterly Journal of Economics*, 2018, *133* (3), 1283–1330.

Porter, Martin F, "An algorithm for suffix stripping," *Program*, 1980.

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

Romer, Christina D. and David H. Romer, "Federal Reserve Information and the Behavior of Interest Rates," *American Economic Review*, 2000, *90* (3), 429–457.

— **and** — , "A New Measure of Monetary Shocks: Derivation and Implications," *American Economic Review*, 2004, *94* (4), 1055–1084.

— **and** — , "The FOMC Versus the Staff: Where Can Monetary Policymakers Add Value?," *American Economic Review, Paper and Proceedings*, 2008, *98* (2), 230–35.

Sastry, Karthik, "Disagreement About Monetary Policy," Working Paper 2021.

Shapiro, Adam Hale, Moritz Sudhof, and Daniel J Wilson, "Measuring news sentiment," *Journal of Econometrics*, 2020.

Stock, James H and Mark W Watson, "Twenty years of time series econometrics in ten pictures," *Journal of Economic Perspectives*, 2017, *31* (2), 59–86.

Taylor, Ann, Mitchell Marcus, and Beatrice Santorini, *Treebanks: Building and Using Parsed Corpora*, Dordrecht: Springer Netherlands, 2003.

Thornton, Daniel L., "When Did the FOMC Begin Targeting the Federal Funds Rate? What the Verbatim Transcripts Tell Us," *Journal of Money, Credit and Banking*, 2006, *38* (8), 2039–2071.

Wolf, Christian K, "SVAR (Mis) identification and the Real Effects of Monetary Policy Shocks," *American Economic Journal: Macroeconomics*, 2020, *12* (4), 1–32.

Woodford, Michael, "Comment on 'Macroeconomic effects of Federal Reserve forward guidance'," *Brookings Papers on Economic Activity*, 2012, pp. 1–80.

Wright, Jonathan H., "What does monetary policy do to long-term interest rates at the zero lower bound?," *The Economic Journal*, 2012, *122* (564), F447–F466.

Wu, Jing Cynthia and Fan Dora Xia, "Measuring the macroeconomic impact of monetary policy at the zero lower bound," *Journal of Money, Credit and Banking*, 2016, *48* (2-3), 253–291.

Appendix

A Text Analysis

This appendix contains details on the text analysis.

A.1 Word Lists

I use the following word lists in my analysis. From the Harvard-IV dictionary, I use words from the increase, decrease, rise, fall, strong, and weak word lists. I also use the positive and negative words of Loughran and McDonald (2011). Those authors construct positive and negative word lists that are designed to be relevant in an economics/finance context. I note that while I count the number of uses of words in these lists, they are not required to be selected in the LASSO regression.

A.2 Cities

As described in the text, I drop any article that starts with a city/country from the following list (after having gone through and manually finding all cities/countries that are mentioned at the beginning of my articles): Basel, Beijing, Frankfurt, Hong Kong, Jakarta, London, Manila, Mexico city, Moscow, Mumbai, New delhi, Ottawa, Pretoria, Rio de Janiero, Sao Paulo, Singapore, Sydney, Tokyo, Toronto, Wellington, Xurich, Kuala Kumpur, Brasilia, Karlsruhe, Turkey.

A.3 Word List Construction

I start with a set of three "seed" words: *output*, *growth*, and *economy*. I train the word2vec algorithm of Mikolov et al. (2013) on a subset of a large corpus of newspaper articles: *The New York Times Annotated Corpus* from the University of Pennsylvania's Linguistic Data Consortium. The full corpus contains 1.8 million articles from the New York Times between 1987 and 2007, each manually tagged by library scientists. I clean the data using standard tools—I stem words using the stemmer of Porter (1980), and find bi-, tri-, and quad-grams that occur over 300 times and that are used at least 5% as often as their constituent words. The word2vec algorithm consists of constructing vector representations of words that, via a neural network, can predict a word in a set of text given the surrounding words. The algorithm is thus well suited to finding synonyms, hence its employment here. I trained the algorithm on 94,601 articles that were tagged as related to either economic output, prices, or labor markets.[49] With a vector representation of every word in the corpus, I sort words based on their distance to the average vector of my seed words.[50] The resulting list, along with the distances from the seed vector, is listed in table 5. There, I also show whether a word has its sentiment words flipped (indicated with a minus sign) when constructing features for the LASSO regression.

[49]Specifically, I retained articles labeled economic conditions and trends, united states economy, prices, wages and salaries, layoffs and job reductions, production, or labor.

[50]Mikolov et al. (2013) highlight that summing and distracting the vector representations of each word results in meaningful vectors. For example, the authors find that subtracting the vector for *man* from the vector for *king* results in a vector that is very similar to the vector for *queen*. To compute the distance of two words with vectors x and y, I compute the cosine similarity between them: $x'y/[\sqrt{(x'x)} + \sqrt{(y'y)}]$.

Table 5: Extended GDP Word List

Term	Direction	Similarity to *output* + *growth* + *economi*
economic growth	+	0.88
growth	+	0.85
economi	+	0.84
consumer spending	+	0.77
output	+	0.73
recoveri	+	0.71
consumer confidence	+	0.64
living standards	+	0.64
shrinkag	−	0.63
gross domestic product	+	0.63

A.4 The Role of Tense Words

This table shows the words used when predicting GDP for the year starting in the quarter *after* the Fed meeting.

Coefficient	Word	Sentiment	Tense	Lag
-0.020	consumerspending	Fall	Future	✓
-0.019	growth	Fall		✓
0.017	economicgrowth	Positive	Present	✓
-0.015	consumerspending	Decrease		
0.015	consumerspending	Strong		✓
-0.014	economicgrowth	Rise	Future	✓
0.012	economi	Increase	Present	
-0.012	economi	Rise	Future	
0.011	output	Rise	Present	✓
0.010	growth	Rise	Future	✓

B Text-Based Expectations: Daily and Monthly Frequency

In this appendix, I show additional exercises to validate my text-based measures of expectations using daily and monthly versions constructed from the text of the *New York Times* through 2007. Figure 11 shows the time series of the monthly Blue Chip forecast revisions used to train the algorithm, alongside the model's fit. Table 6 presents regression estimates in which the left-hand side is the daily measure of forecast revisions. In figure 12, I plot my measures of GDP and inflation forecast revisions against the daily measure of news sentiment of Shapiro et al. (2020).

Figure 11: Monthly Forecast Revisions: Blue Chip and Text-Based Proxy

(a) Real GDP Growth (b) CPI Inflation

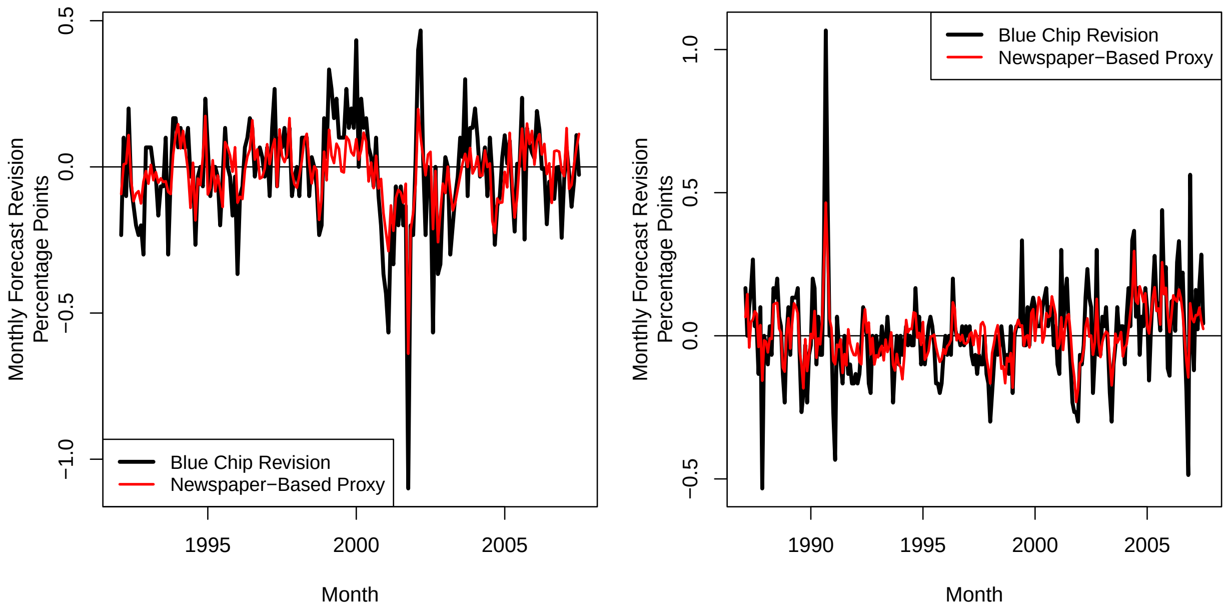

NOTE. This figure shows the actual (from the Blue Chip) and predicted (from the New York Times corpus) values of GDP (panel (a)) and inflation (panel (b)) forecast revisions. Both samples end in 2007 reflecting the end of the corpus, the GDP sample starts when the Blue Chip asks for real GDP expectations, and the inflation sample starts when the corpus begins. Respondents are asked for annualized growth rates of each variable. I average the forecast revisions for the current and following quarter of month t. All relevant articles written after the seventh day of month t and $t-1$ are included.

Table 6: Determinants of Daily Text-Based Measures of Macro Forecast Revisions

(a) Levels

	$\widehat{y}^d_{t+1}$	$\widehat{\pi}^d_{t+1}$
CPI Surprise	0.397	0.237
	(0.584)	(1.056)
GDP Surprise	0.269	0.0201
	(0.312)	(0.183)
NFPR Surprise	37.20	11.31
	(17.89)	(9.253)
$\Delta \log(\text{oil}_t)$	0.0397	1.152
	(0.850)	(0.818)
FOMC Surprise	6.997	3.243
	(2.662)	(1.330)
$\Delta \log(\text{SP500}_t)$	-0.0298	2.437
	(1.941)	(1.867)
Constant	0.00226	-0.00483
	(0.0304)	(0.0300)
Observations	3595	3595
R^2	0.0126	0.00291

(b) Size

	$\lvert\widehat{y}^d_{t+1}\rvert$	$\lvert\widehat{\pi}^d_{t+1}\rvert$
$\mathbf{1}\{\text{CPI Release}_t\}$	-0.0141	0.100
	(0.0616)	(0.0858)
$\mathbf{1}\{\text{GDP Release}_t\}$	0.320	0.0288
	(0.0983)	(0.0717)
$\mathbf{1}\{\text{NFPR Release}_t\}$	0.202	-0.0455
	(0.0947)	(0.0532)
$\lvert\Delta \log(\text{oil}_t)\rvert$	2.072	1.687
	(1.007)	(0.856)
$\mathbf{1}\{\text{FOMC Release}_t\}$	0.473	0.0914
	(0.140)	(0.0707)
$\lvert\Delta \log(\text{SP500}_t)\rvert$	4.881	1.418
	(2.008)	(1.604)
Constant	0.614	0.639
	(0.0178)	(0.0231)
Observations	3595	3595
R^2	0.0271	0.00336

NOTE. This table shows regressions of daily measures for GDP (y) and inflation (π) forecast revisions on several explanatory variables. The left-hand side in panel (a) is the level of the measures constructed using all relevant New York Times articles published on day $t+1$, standardized to have unit variance and a mean of zero. Panel (b) shows results using the absolute value of the variables in panel (a). The "release" variables are indicators for whether data on CPI, GDP, or Nonfarm Payrolls (NFPR), or an Fed policy statement were released on day t. The "surprise" variable are the surprise components of these releases—the Fed surprise is the 30-minute change in one-year ahead Eurodollar futures, and the other three surprises are taken from Bloomberg (in percentage points for CPI and GDP, and millions of jobs for NFPR). Daily changes in oil prices and the S&P 500 are taken from FRED (mnemonic DCOILWTICO) and Yahoo, respectively. The sample runs from January 1995 (when the surprise data start) through May 2007 (when the daily newspaper data end). Standard errors are calculated using heteroskedasticity and autocorrelation-consistent asymptotic standard errors with the automatic lag selection method of Newey and West (1994), as implemented by Baum et al. (2002).

Figure 12: Daily Newspaper GDP and Inflation Forecast Revisions vs. Newspaper Sentiment

(a) Real GDP Growth

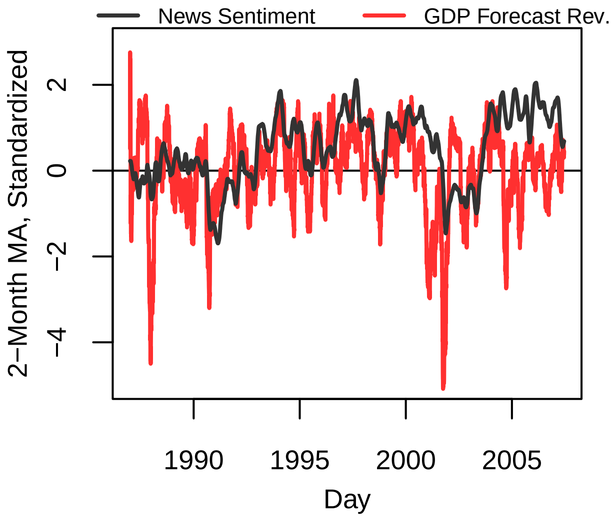

(b) CPI Inflation

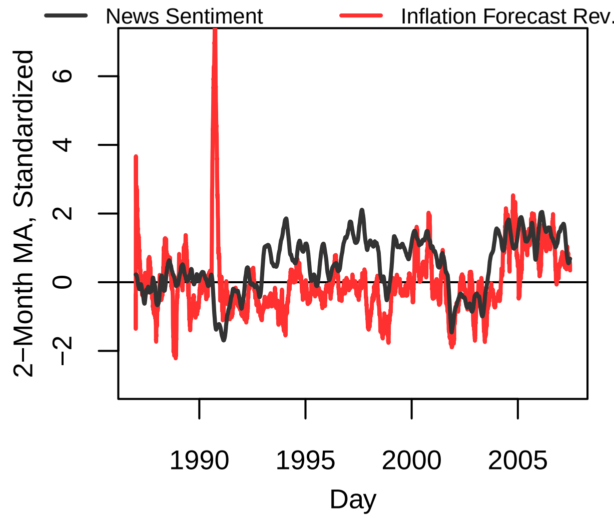

NOTE. This figure plots my newspaper-based measures of GDP and inflation forecast revisions against the newspaper-based measure of sentiment of Shapiro et al. (2020). The standardized values of 2-month moving averages are shown.

Table 7: Forecast Revisions around Selected Fed Meetings

Fed Date	$\widehat{y}_t$ b.p.	$\widehat{\pi}_t$ b.p.	$\widehat{\mathrm{FF}}_t^0$ b.p.	$\widehat{\mathrm{ED}}_t^4$ b.p.	$\Delta^{\mathrm{HF}}\mathrm{S\&P}\ 500_t$ percent
March 22, 2005	-1.6	31.0	-1.7	10.5	-0.27
January 31, 2001	-30.8	-3.7	4.0	-4.0	-0.52
June 29, 2006	-30.3	5.0	-1.5	-3.5	0.49

C Examples

This appendix contains a description of the three Fed announcements that induce the largest GDP and inflation forecast revisions. Table 7 shows how key variables change around each announcement.

March 22, 2005 This meeting induced the largest increase in inflation expectations according to my text-based measure. It is also a rare example in which the post-meeting articles paint a very clear picture of the surprising part of the Fed's announcement. One *New York Times*' headline read "Stocks Plunge As Fed Stirs Inflation Fears" while a *Wall Street Journal*'s headline read "Fed Lifts Rates, Warns on Inflation—First Concerns About Prices In 4 Years Could Presage End of 'Measured' Boosts." This surprise was apparently driven by a change in the Fed's description of inflation as "well contained" and "relatively low," in the previous statement (in February), to having "picked up in recent months."

January 31, 2001 This meeting induced the largest downward revision in GDP expectations according to my text-based measure. The newspaper articles at the time explain that the meeting's 50 basis point interest rate cut was largely expected.[51] What appears to have been less anticipated, however, was the accompanying signal about the economic outlook. A *New York Times* headline, for example, read "Fed Cuts Key Rate Half a Point, Citing Slowdown." This example highlights the usefulness of bringing in additional—non-interest-rate—data sources to parse the information gleaned from Fed announcements. Despite the fact that a *Wall Street Journal* headline read "Market's Reaction Is Anticlimactic as Rate Cut Holds No Surprises," digging deeper into articles reveal the pessimism induced by the Fed's cut. Another *Journal* article started with

> The shock and dismay at the rapid slide in the U.S. economy is palpable. Layoffs have accelerated dramatically and the Federal Reserve's full percentage point reduction in the federal-funds rate in January—with the latest installment coming yesterday—seems almost a sign of panic.

This sense of panic is not revealed in the relatively small changes in interest rates, but it is apparent in the stock market response and newspaper articles. The next example will show, however, that these two objects do not always move in tandem.

June 29, 2006 This meeting induced the second-largest downward revision in GDP according to my text-based measure. As explained in the *New York Times* before the announcement, recently-issued "stronger-than-expected economic data raised fears that the central bank would signal that more rate increases are coming." Instead, the Fed decided on a 25 basis point cut, and, as explained in the *Wall Street Journal*, "its words suggested growing prospects for a pause in rate increases." That articles goes on to explain that "the Fed expects inflation to edge lower as the economy

[51] The *Wall Street Journal*, for example, wrote the day before that the meeting was "widely expected to conclude with a 0.5 percentage-point cut in interest rates."

slows." This example highlights the importance of comparing post-announcement articles with pre-announcement articles. Pre-announcement articles indicate that markets expected the Fed to signal that the economy was running hot. Post-announcement descriptions of the Fed's outlook are unremarkable. Here, then, it is a lack of action that causes markets to revise down expectations. On the other hand, the news that the Fed would likely hold rates lower was accompanied by a decrease in stock prices. This is also an example of when my GDP measure and stock tell different stories, with stocks rising on the news that the Fed would likely not raise rates much more. This suggests that stock prices are more responsive to the interest rate itself than the underlying cause which, here, was a cooler-than-expected description of the economy.

D The Noisy Information Model

The linear dynamic RE models I consider are of the form discussed by Blanchard et al. (2013), in which agents are imperfectly informed about the economy's state variables. Models of this class take the form

$$\mathbf{A}\,\mathbf{x}_t + \mathbf{B}\,\mathbb{E}_t[\mathbf{x}_{t+1}] + \mathbf{C}\boldsymbol{\xi}_{t|t} = \mathbf{0} \tag{9}$$

where $\mathbf{x}_t$ is a vector of observable macroeconomic variables, $z_{t|\tau}$ denotes the mathematical expectation of z_t given information at time τ, and $\boldsymbol{\xi}_t$ is a vector of mutually independent structural shocks that evolves according to $\boldsymbol{\xi}_t = \mathbf{D}\boldsymbol{\xi}_{t-1} + \mathbf{F}\boldsymbol{\zeta}_t$, where $\boldsymbol{\zeta}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\zeta)$.[52] At the beginning of each period t, but before making decisions about $\mathbf{x}_t$, all agents receive the same noisy signal of the structural shock of the form $\mathbf{s}_t = \mathbf{G}\boldsymbol{\xi}_t + \mathbf{H}\boldsymbol{\nu}_t$, where $\boldsymbol{\nu}_t \sim \mathcal{N}(\mathbf{0}, \boldsymbol{\Sigma}_\nu)$. Agents use this signal to form expectations about $\boldsymbol{\xi}_t$ using the Kalman filter (i.e., agents have RE). Agents are thus imperfectly informed in a symmetric way—they share the same information set when making decisions.

Blanchard et al. (2013) show that the model in equation (9) admits a solution of the form $\mathbf{x}_t = \mathbf{Z}\boldsymbol{\xi}_{t|t}$ along with a law of motion for perceived shocks, $\boldsymbol{\xi}_{t|t}$, where $\mathbf{Z}$ depends on neither Σ_ζ not Σ_ν. That solution and the law of motion for ξ_t reveal the result: the mapping between forecast revisions of observable variables and forecast revisions of structural shocks is fixed and linear: $\mathbb{E}_t[\mathbf{x}_{t+k}] - \mathbb{E}_{t-1}[\mathbf{x}_{t+k}] = \mathbf{Z}\mathbf{D}^k[\boldsymbol{\xi}_{t|t} - \boldsymbol{\xi}_{t|t-1}]$. This relationship resembles the relationship between reduced-form and structural shocks in the parlance of structural VARs, another macroeconomic model in which the assumption of constant m_{ij} is valid.

The key ingredient for a linear RE model to feature a constant mapping between observable variables and structural shocks is that the model's structural equations can be solved independently of agents' filtering problem. The result would continue to hold, therefore, in models in which agents also observe the endogenous variables, $\mathbf{x}_t$.

[52] The assumption that the shocks follow a VAR(1) is not restrictive—any finite VARMA can be re-written as a VAR(1). Adding the vector of lagged endogenous variables to equation (9) (i.e., $\mathbf{x}_{t-1}$) slightly changes the language, but not results, of this discussion. Specifically, rather than forecast revisions of endogenous variables being related linearly to revisions of structural shocks, it's forecast revisions of *surprises* (i.e., reduced form residuals) of endogenous variables that are related linearly to revisions of structural shocks.

E Model Appendix

In this appendix, I provide additional details of the model.

Calibration I calibrate the model exactly as in chapter 3 of Galí (2015). The value of the structural parameters are

Parameter	Symbol	Value	Parameter	Symbol	Value
Discount factor	β	0.99	Risk aversion	σ	1
Inverse Frisch Elast.	φ	1	Cobb-Douglas	α	0.25
Consumption elast. of subs.	ϵ	9	Interest semielast. of mon. demand	η	4
Price stickiness	θ	0.75	Taylor rule inflation	ϕ_π	1.5
Taylor rule output	ϕ_y	0.125			

This implies the following values for the parameters used in the model

$$
\begin{aligned}
\mu &= \log\left(\frac{\epsilon}{\epsilon - 1}\right) \approx 0.117 \\
\Theta &= \frac{1-\alpha}{1-\alpha+\alpha\epsilon} = 0.25 \\
\lambda &= \frac{(1-\theta)(1-\beta\theta)}{\theta}\Theta \approx 0.021 \\
\kappa &= \lambda\left(\sigma + \frac{\varphi+\alpha}{1-\alpha}\right) \approx 0.057 \\
\Psi_{ya} &= \frac{1+\varphi}{\sigma(1-\alpha)+\varphi+\alpha} = 1.0.
\end{aligned}
$$

I set the autocorrelations of the monetary, news, supply, and demand shocks to 0.5, 0.9, 0.9, and 0.5 respectively, and set the variance of their innovations to unity.

Calibrations To increase the correlation in perceived shocks, I simulate 10 calibrations of the model. In the first 5, the interest rate is revealed only imperfectly. This is designed to remove its information content, though during these 5 calibrations I allow the structural shocks to be perfectly revealed—interest rates are thus not used when forming expectations. For the next 5, I allow the interest rate to be observed perfectly, and increase the variance of the signals of each shock. To summarize the extent to which the interest rate is used, figure 13 shows the sum of absolute elements in the Kalman gain matrix in the interest rate row over the total sum of absolute elements—a summary statistic for how much weight the interest rate receives when agents update expectations.

To have a sense of how these calibrations induce correlation in the perceived structural shocks. Table 8 shows the correlation of perceived and true shocks in the three calibrations shown in the text (1,7, and 9 in figure 13). In calibration 1, by design, the perceived shocks are uncorrelated. By calibration 9, they are quite correlated—for example, the perceived monetary and demand shocks are correlated 0.45.

Table 8: Correlations between Perceived and Structural Shocks

	Calibration 1											
	ε_t	ϕ_t	a_t	z_t	$\widehat{\varepsilon}_t$	$\widehat{\phi}_t$	$\widehat{a}_t$	$\widehat{z}_t$	$\widehat{\varepsilon}_t^H$	$\widehat{\phi}_t^H$	$\widehat{a}_t^H$	$\widehat{z}_t^H$
ε_t	1.00	-0.01	0.01	0.01	0.87	-0.00	0.00	0.01	0.87	-0.00	-0.00	0.00
ϕ_t	-0.01	1.00	-0.00	0.01	-0.00	0.44	0.01	0.01	-0.00	0.44	0.00	0.00
a_t	0.01	-0.00	1.00	-0.00	0.01	-0.00	0.44	-0.00	0.01	-0.00	0.44	-0.00
z_t	0.01	0.01	-0.00	1.00	0.01	0.01	0.01	0.87	0.01	0.01	0.00	0.87
$\widehat{\varepsilon}_t$	0.87	-0.00	0.01	0.01	1.00	0.00	0.00	0.01	1.00	0.00	0.00	0.00
$\widehat{\phi}_t$	-0.00	0.44	-0.00	0.01	0.00	1.00	0.01	0.01	0.00	1.00	0.01	0.00
$\widehat{a}_t$	0.00	0.01	0.44	0.01	0.00	0.01	1.00	0.01	0.00	0.01	1.00	0.01
$\widehat{z}_t$	0.01	0.01	-0.00	0.87	0.01	0.01	0.01	1.00	0.01	0.00	0.00	1.00
$\widehat{\varepsilon}_t^H$	0.87	-0.00	0.01	0.01	1.00	0.00	0.00	0.01	1.00	0.00	-0.00	-0.00
$\widehat{\phi}_t^H$	-0.00	0.44	-0.00	0.01	0.00	1.00	0.01	0.00	0.00	1.00	0.00	-0.00
$\widehat{a}_t^H$	-0.00	0.00	0.44	0.00	0.00	0.01	1.00	0.00	-0.00	0.00	1.00	0.00
$\widehat{z}_t^H$	0.00	0.00	-0.00	0.87	0.00	0.00	0.01	1.00	-0.00	-0.00	0.00	1.00

	Calibration 9											
	ε_t	ϕ_t	a_t	z_t	$\widehat{\varepsilon}_t$	$\widehat{\phi}_t$	$\widehat{a}_t$	$\widehat{z}_t$	$\widehat{\varepsilon}_t^H$	$\widehat{\phi}_t^H$	$\widehat{a}_t^H$	$\widehat{z}_t^H$
ε_t	1.00	-0.01	0.01	0.01	0.81	0.08	0.17	0.20	0.82	-0.12	-0.12	-0.17
ϕ_t	-0.01	1.00	-0.00	0.01	0.06	0.44	0.14	0.09	0.06	0.44	-0.01	0.03
a_t	0.01	-0.00	1.00	-0.00	0.06	0.05	0.40	0.04	0.06	0.01	0.43	-0.00
z_t	0.01	0.01	-0.00	1.00	0.11	0.07	0.08	0.77	0.11	-0.01	-0.01	0.81
$\widehat{\varepsilon}_t$	0.81	0.06	0.06	0.11	1.00	0.26	0.39	0.45	1.00	0.00	0.00	0.00
$\widehat{\phi}_t$	0.08	0.44	0.05	0.07	0.26	1.00	0.44	0.27	0.26	0.96	0.07	0.07
$\widehat{a}_t$	0.17	0.14	0.40	0.08	0.39	0.44	1.00	0.29	0.39	0.28	0.88	0.06
$\widehat{z}_t$	0.20	0.09	0.04	0.77	0.45	0.27	0.29	1.00	0.45	0.10	0.04	0.89
$\widehat{\varepsilon}_t^H$	0.82	0.06	0.06	0.11	1.00	0.26	0.39	0.45	1.00	0.00	-0.00	-0.00
$\widehat{\phi}_t^H$	-0.12	0.44	0.01	-0.01	0.00	0.96	0.28	0.10	0.00	1.00	-0.00	-0.00
$\widehat{a}_t^H$	-0.12	-0.01	0.43	-0.01	0.00	0.07	0.88	0.04	-0.00	-0.00	1.00	0.00
$\widehat{z}_t^H$	-0.17	0.03	-0.00	0.81	0.00	0.07	0.06	0.89	-0.00	-0.00	0.00	1.00

NOTE. This table shows the correlation matrix of the structural shocks (respectively, the monetary, persistent monetary, supply, and demand), their (hatted) perceived counterparts, and their (superscript-h) counterparts identified by heteroskedasticity.

Figure 13: Weight Placed on Interest Rate when Revising Expectations

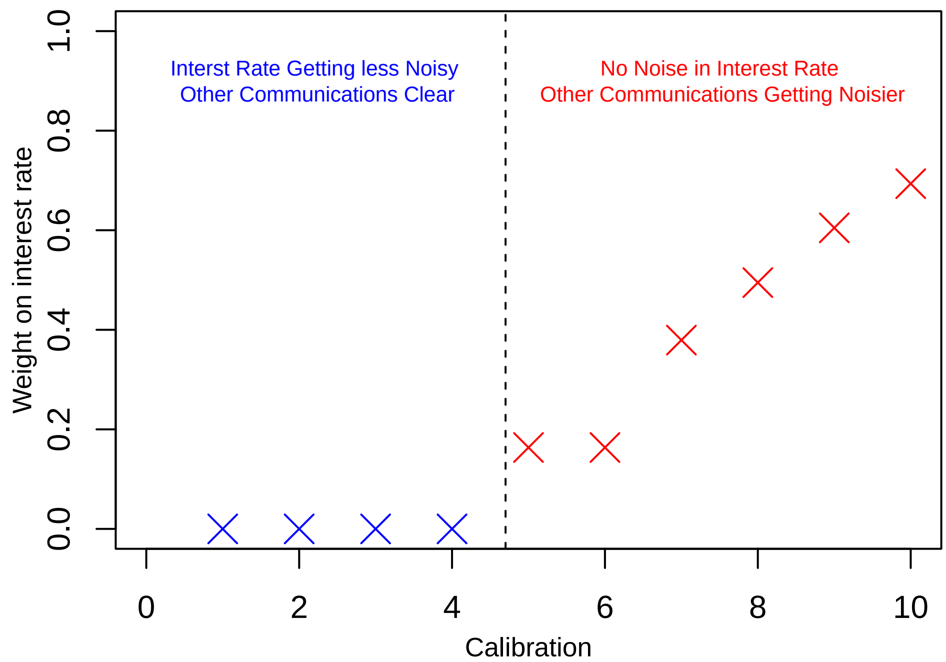

F The Role of Interest Rate Maturity

In this appendix, I show that as one increases the maturity of the interest rate used for forming high-frequency estimates of monetary shocks, the more the results look puzzling from the perspective of standard theory. To do this, I use high frequency changes in either the 1-, 2-, 3-, or 4-quarter Eurodollar future as an instrument in my baseline VAR. I show results for the shadow rate, inflation, and industrial production, but the other variables show similar features. This result can also be shown using longer- and longer-term Federal Funds futures. The results are in figure 14.

Figure 14: IRFs in the VAR Varying the Eurodollar Future Maturity

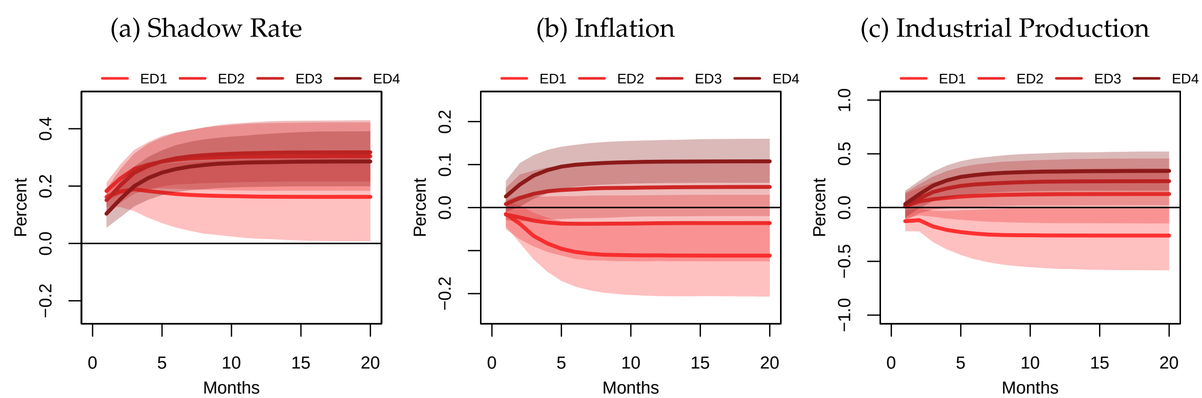

NOTE. This figure shows impulse responses of the variables in the subfigure titles in the baseline VAR. The lines labeled EDn show the estimates when the n-quarter ahead Eurodollar future is used as an instrument. 90% credible regions are shaded.

G Bauer and Swanson (Forthcoming)

The notion that high-frequency interest rate surprises contain an "information effects" component arose largely from the work of Campbell et al. (2012) and Nakamura and Steinsson (2018). Those authors estimate regressions of the form

$$\mathbb{E}_{t+1}[y_{t+k}] - \mathbb{E}_t[y_{t+k}] = \alpha + \beta\widehat{i}_t + e_t \tag{10}$$

where the left-hand side is a forecast revision of GDP, taken from the Blue Chip survey, and $\widehat{i}_t$ is a measure of high-frequency changes in interest rate expectations taken around Fed announcements (what has been traditionally called a "monetary policy shock"). Those authors find $\beta > 0$, in contrast to standard theory. This evidence is consistent with Fed announcements revealing information that is independent of purely exogenous monetary policy shocks.

The recent work of Bauer and Swanson (Forthcoming) presents several pieces of evidence that push against this conclusion. The central exercise in the paper is a demonstration that once one includes a vector of "macroeconomic news" variables as a control in equation (10), the estimate of β is attenuated or even flips sign.[53] I present evidence that this sensitivity result is itself sensitive to the variables that make up "macroeconomic news." This is noteworthy because the variables Bauer and Swanson consider are somewhat non-standard and potentially raise concerns about overfitting or spuriousness. In table 9, I revisit their regressions, but include one control variable at a time in order to understand the importance of each. I focus on the three variables that accompanied the original working-paper—the Brave et al. (2019) index of economic activity, the most-recent change in nonfarm payrolls, and the 13-week return on the S&P 500. The most-recent version of the paper includes several other controls which are even more unusual, such as "the change in the log Bloomberg total commodity price index BCOM minus 0.4 times the change in the log Bloomberg agricultural commodity price index BCOMAG."[54]

The results in table 9 suggest that two of the three "macroeconomic news" variables—the jobs number and the Brave et al. (2019) index—play essentially no role as omitted variables. What does matter, however, is the 13-week return on the S&P 500. In panel (a) of figure 15, I show that the 13-week return in particular—not 12, not 14—is the most-potent predictor of $\widehat{i}_t$. Almost every other horizon has no statistically-significant predictive power. In turn, panel (b) shows that this particular return horizon is also the most important for affecting the estimate of β. Were one to consider any other *a-priori* reasonable return horizon, the estimated coefficient would not exclude the original point estimate. Taken together, my findings suggest that the evidence against information effects is not as strong as Bauer and Swanson suggest. The skeptical view of the evidence is that "predictability" is partly a consequence of a selected set of variables on a small sample. Alternatively, the evidence shows that there are indeed a select set of variables that predict the surprises but they are somewhat idiosyncratic. The question is, then, how to interpret this evidence. Bauer and Swanson suggest that this evidence is consistent with markets persistently

[53] Bauer and Swanson include several pieces of evidence—I discuss two more here. First, they survey Blue Chip forecasters and find that they do not claim to use the Fed's communications to update their views of the state of the economy in a direction consistent with the presence of information effects. If, as suggested by my empirical analysis above, information effects are time-varying, a one-time survey may not be particularly informative if it comes at a time when information effects are not large. Second, they show—consistent with the evidence of Hoesch et al. (2020)—that the Fed's staff's forecasts are not more accurate than the private sector's. It is understood that this is not a sufficient condition for information effects to not exist, only a comment on their strength. It is worth noting, however, that if two agents are equally uncertain about some outcome but have independent signals, then each agent has much to learn from the other.

[54] The authors justify this choice by explaining that each variable separately enters their regression with a coefficient of 0.4. Even so, this choice seems oddly specific and raises questions about variable selection.

Table 9: Which Bauer and Swanson (Forthcoming) Controls Matter in Blue Chip Regressions?

NS shock	1.082	0.718	1.036	0.609	0.541
	(0.351)	(0.308)	(0.359)	(0.305)	(0.267)
BBK Index		0.105			0.112
		(0.0236)			(0.0266)
Δ Nonfarm Payrolls			0.0436		-0.223
			(0.0672)		(0.0669)
Δ S&P 500 (13-week)				1.002	0.816
				(0.157)	(0.160)
Constant	-0.0289	-0.0284	-0.0334	-0.0496	-0.0221
	(0.0116)	(0.0104)	(0.0145)	(0.0108)	(0.0128)
Observations	120	120	120	120	120
R^2	0.0771	0.239	0.0798	0.347	0.457

NOTE. This table shows estimates of equation (10) in which I add the additional control variables shown in the rows. The BBK Index is the "coincident index" of Brave et al. (2019) in the month before the Fed meeting. The change in nonfarm payrolls is the most-recent real-time estimate of the change in non-farm payrolls before the Fed meeting. The change in the log S&P price index is taken from the day before the Fed meeting to 13-weeks before. The NS shock is taken directly from the replication materials of Nakamura and Steinsson (2018). The sample matches the sample in table III of their paper and is described in the note to figure 15. Heteroskedasticity-robust standard errors are shown in parentheses.

underestimating the Fed's response to news.

An alternative consistent with Bauer and Swanson's evidence is that Fed announcements reveal information about the Fed's economic outlook that emphasizes certain non-standard variables that markets do not expect. To the extent that the public puts different weight on different variables, these variables can be observable but still predict interest rate surprises.[55] Bauer and Swanson have found a particular set of such observable non-standard variables.[56] So, do the controls provide evidence against information effects, or do they constitute yet another way to control for the Fed's alternative view of the state of the economy? The evidence is not as clear-cut as Bauer and Swanson suggest. As such, this leaves room for continued investigation into the effects of Fed communications.

[55]This sentence should not be interpreted as a statement that "markets misperceive the Fed's policy rule." It is instead consistent with a model in which markets understand the Fed's reaction to a variable x, but have different beliefs about the value of x. Sastry (2021) evaluates these alternative explanations and finds that the evidence favors this "agreeing to disagree" explanation.

[56]I say "non-standard" because these variables are outside of what one would typically include in a local projection or VAR. It is particularly noteworthy that adding the S&P 500 into my VAR does not meaningfully affect my estimates, which is additional evidence that the 13-week return is special.

Figure 15: Which horizon of S&P Returns Matters the Most?

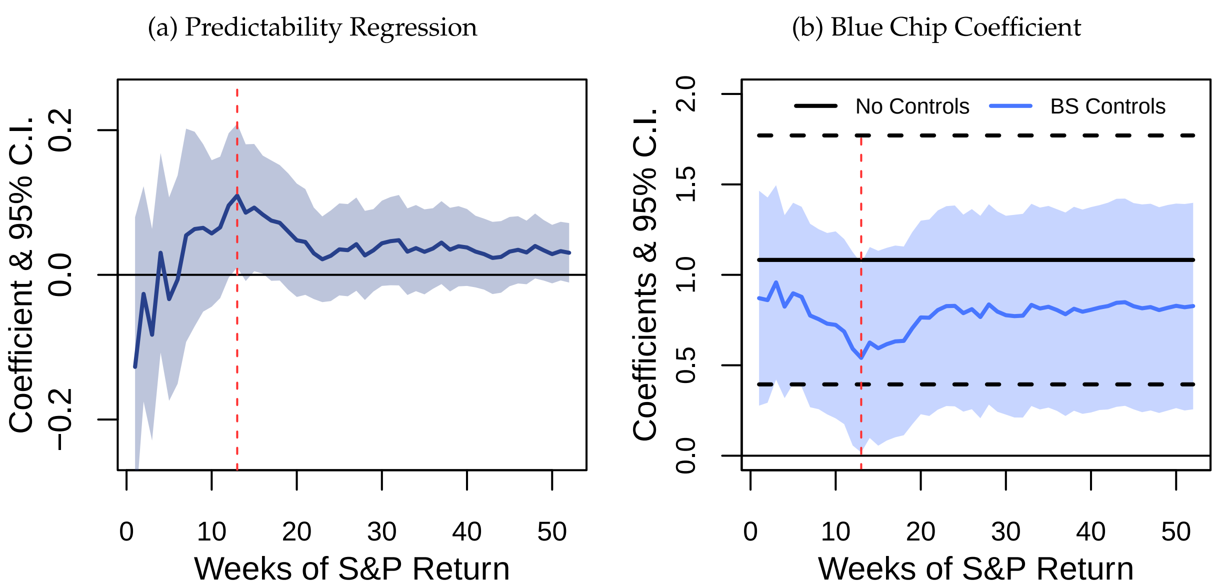

NOTE. Panel (a) presents estimates of the following regression $\widehat{i}_t^{\text{NS}} = \alpha_h + \beta_h(p_{t-1} - p_{t-7h}) + e_t$ where $\widehat{i}_t^{\text{NS}}$ is the high-frequency interest-rate surprise of Nakamura and Steinsson (2018) (from their replication materials), and p_t is the log of the S&P 500 index on day t's market close. The sample for the blue line consists of all days (t) with regularly-scheduled Fed meetings between 1995 and 2015 (excluding the July 2008–July 2009 period) and after the first week of the month (i.e., the observations used when testing for the presence of information effects with Blue Chip data above, and by Nakamura and Steinsson). The right-hand-side variable, $p_{t-1} - p_{t-7h}$, is thus the h-week return in the S&P 500 ending the day before each Fed meeting. Panel (b) presents estimates of equation (10) with black lines, and additionally (in blue) with the controls of Bauer and Swanson (described in the note to table 15) in which I vary the return window of the S&P 500 as in panel (a). I estimate both regressions for each h, which is shown on the x-axis. Bauer and Swanson present results using the 13-week return, which I highlight with a red dashed line. This return horizon has the largest effect in both regressions.

Figure 16: Bauer and Swanson Controls: What Matters?

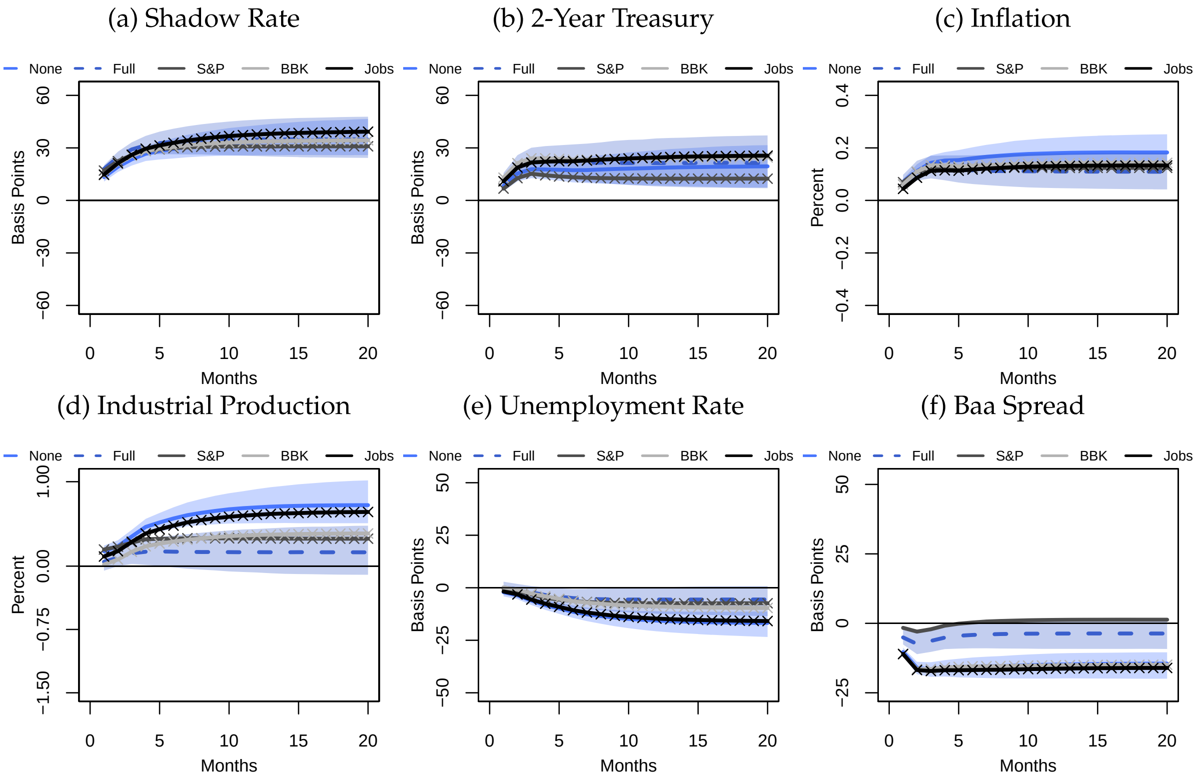

Note. This table shows responses to the perceived demand shock in my baseline VAR ("None"), and using versions of the perceived demand shock in which all HF forecast revisions have been orthogonalized to one or more of the shocks listed in the legend.

Table 10: Table 4 with Bauer and Swanson (Forthcoming) Controls

	$\widehat{\mathrm{FF}}_t^0$	$\widehat{\mathrm{ED}}_t^4$	$\Delta\,\mathbb{E}^{\mathrm{BC}}[y]$	$\Delta\,\mathbb{E}^{\mathrm{BC}}[\pi]$	$\Delta\,\mathbb{E}^{\mathrm{BC}}[u]$	Δ^{HF}S&P 500	$\Delta\mathrm{nominal}_t^{10}$	$\Delta\mathrm{real}_t^{10}$
Monetary News	0.001	0.049	-0.021	-0.017	0.007	-0.180	0.014	0.015
	(0.000)	(0.000)	(0.015)	(0.014)	(0.007)	(0.046)	(0.004)	(0.004)
Perceived Supply	0.001	0.005	0.027	-0.018	-0.001	-0.007	0.002	0.005
	(0.000)	(0.000)	(0.015)	(0.017)	(0.006)	(0.029)	(0.005)	(0.005)
Perceived Demand	0.003	0.050	0.030	0.034	0.011	-0.150	0.024	0.022
	(0.000)	(0.000)	(0.019)	(0.014)	(0.006)	(0.042)	(0.005)	(0.006)
Monetary	0.039	0.017	-0.021	-0.044	0.011	-0.161	0.000	0.005
	(0.000)	(0.000)	(0.024)	(0.021)	(0.011)	(0.068)	(0.005)	(0.005)
Observations	165	165	139	139	139	165	165	165
R^2	1.000	1.000	0.153	0.136	0.440	0.249	0.150	0.148

NOTE. This table is identical to table 10, except that I have included the controls of Bauer and Swanson (Forthcoming) in the regressions.

H Why not orthogonalize out $\widehat{y}_t$?

The illustrative model showed why $\widehat{y}_t$ is not a valid instrument for information effects—it fails the exclusion restriction by failing to be independent of the monetary shock. In addition, attempting to "purge" the monetary surprise $\widehat{i}_t$ of its information content (η_t) by orthogonalizing $\widehat{i}_t$ to $\widehat{y}_t$ will not generally identify a monetary shock. Proposition 1 states the conditions under which this procedure identifies a monetary policy shock.

Proposition 1. *Let r_t be the residual from a linear projection of $\widehat{i}_t$ on $\widehat{y}_t$. Unless $r_t = 0 \ \ \forall \ \ t$, then r_t is independent of information effects if and only if output expectations do not respond to monetary policy shocks, i.e., $\omega_\varepsilon = 0$.*

Proof. Let

$$\widehat{i}_t = \phi_\varepsilon \widehat{\varepsilon}_t + \phi_\eta \widehat{\eta}_t \tag{11}$$

$$\widehat{y}_t = \omega_\varepsilon \widehat{\varepsilon}_t + \omega_\eta \widehat{\eta}_t. \tag{12}$$

The residual r_t (i.e., the "clean" monetary shock) is

$$r_t = i_t - \widehat{\beta} y_t$$

where

$$\widehat{\beta} = \frac{\mathrm{Cov}(\widehat{y}_t, \widehat{i}_t)}{\mathrm{var}(y_t)} = \frac{\phi_\eta \omega_\eta \sigma_\eta^2 + \phi_\varepsilon \omega_\varepsilon \sigma_\varepsilon^2}{\omega_\eta^2 \sigma_y^2 + \omega_\varepsilon^2 \sigma_\varepsilon^2}$$

This residual is then a linear combination of the monetary and information shocks:

$$\begin{aligned} r_t &= \widehat{i}_t - \widehat{\beta}\widehat{y}_t \\ &= \phi_\varepsilon \, \varepsilon_t + \phi_\eta \eta_t - \widehat{\beta}(\omega_\varepsilon \, \varepsilon_t + \omega_\eta \eta_t) \\ &= \underbrace{(\phi_\eta - \widehat{\beta}\omega_\eta)}_{c_\eta} \eta_t + \underbrace{(\phi_\varepsilon - \widehat{\beta}\omega_\varepsilon)}_{c_\varepsilon} \varepsilon_t \,. \end{aligned}$$

This strategy then only provides a "clean" shock if $c_\eta \equiv \phi_\eta - \widehat{\beta}\omega_\eta = 0$.

$$0 = \phi_\eta - \widehat{\beta}\omega_\eta \iff 0 = \omega_\varepsilon \sigma_\varepsilon^2 (\phi_\eta \omega_\varepsilon - \omega_\eta \phi_\varepsilon)$$

The strategy thus provides clean shock in three cases.

1. First, the case where $\sigma_\varepsilon^2 = 0$ means that that there are no monetary shocks, so $r_t = 0 \ \ \forall \ \ t$, violating our assumptions.

2. The knife-edge case with $\phi_\eta \omega_\varepsilon = \omega_\eta \phi_\varepsilon$ also results in $r_t = 0 \ \ \forall \ \ t$. To see this, note that this assumption also implies that $c_\varepsilon = 0$:

   $$c_\varepsilon = \phi_\varepsilon - \widehat{\beta}\omega_\varepsilon = \phi_\varepsilon - \omega_\varepsilon \frac{\phi_\eta \omega_\eta \sigma_\eta^2 + \phi_\varepsilon \omega_\varepsilon \sigma_\varepsilon^2}{\omega_\eta^2 \sigma_\eta^2 + \omega_\varepsilon^2 \sigma_\varepsilon^2} = 0$$

   Thus, with $c_\eta = c_\varepsilon = 0$, we have $r_t = 0 \ \forall \ t$.

3. The final possibility is that output expectations do not respond to monetary policy shocks, i.e., $\omega_\varepsilon = 0$.

I Using Stock Returns

In this appendix, I present shocks estimated using the heteroskedasticity-based approach I pursue in the main text, but in which I replace my text-based measures with the high-frequency return in the S&P 500. Jarociński and Karadi (2020) also use stock returns as their additional variable for identifying monetary and "information" shocks. GMM estimates of the tri-variate system are shown in table 11, and impulse responses in my baseline VAR are shown in figure 17 (see appendix J for an explanation of the layout of this figure).

The monetary shock is largely unchanged, since this is largely only capturing FF_t^0, as described elsewhere in the paper. The responses to the other two variables are markedly different from my estimates, however.

The estimates in table 11 suggest that the information shock indeed captures increases in stock prices that are not associated with decreases in longer-term interest rates. The effects of that shock in the VAR, however, are small and imprecisely estimated. The exception is the effect on credit spreads, which suggests that this shock may primarily capture risk premia, though the fact that this shock is irrelevant (i.e., a weak instrument) for the variables in the system makes it hard to draw conclusions. These findings are consistent with the claim in the text that stock returns provide little information for identifying information effects. As a consequence, the system does not separate the monetary and non-monetary components from ED_t^4, leaving impulse responses that are similar to the responses to the path shock shown in figure 8.

Table 11: Heteroskedasticity-Based Estimates using the S&P 500 and no Text Measures

(a) Structural Impact Matrix, $\mathbf{A}$

	Response of High-Frequency Forecast Revisions		
	Fed Funds Futures Current Meeting	Eurodollar Futures 4 Quarters Ahead	S&P 500
Monetary News	0.03	1.24	-0.41
	(0.05)	(0.16)	(0.18)
Information	-0.21	0.08	1.07
	(0.12)	(0.41)	(0.18)
Monetary	1.57	0.46	-0.02
	(0.27)	(0.21)	(0.27)

(b) Relative Variance in Regime 2, $\Sigma_{(2)}\Sigma_{(1)}^{-1}$

Monetary News	Information	Monetary
1.01	0.24	0.04
(0.53)	(0.07)	(0.03)

NOTE. This table carries out GMM estimation of a system similar to that whose estimates are presented in table 3. The two regimes are the same as the first two regimes in the text, and the observable variables are listed in the column of panel (a).

Figure 17: Impulse Responses using Stock Returns

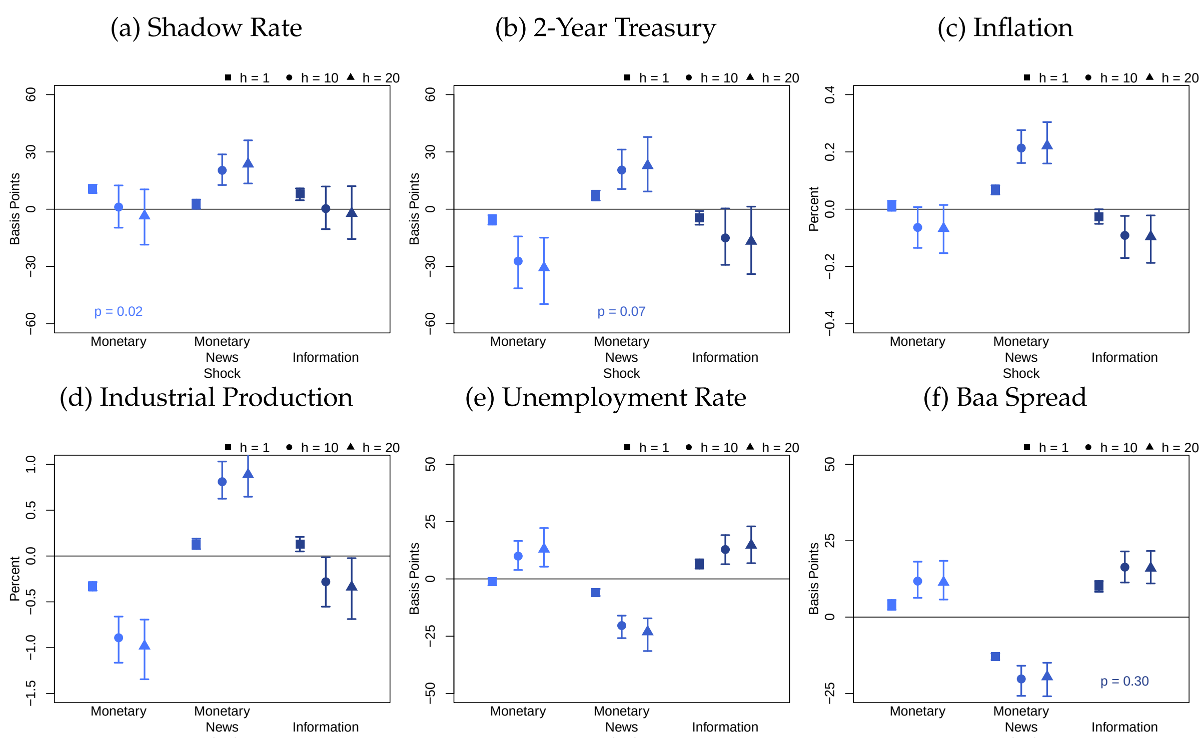

J Robustness and Additional Figure

This appendix contains several exercises designed to assess the robustness of my empirical results. Those results, in figures 18 through 26 are presented in a condensed format relative to what I present in the main text. In particular, for each shock and each outcome variable, I present the median and 90% credible region for the impact response, and the response at 10 and 20 months. As with the baseline estimates, these capture the bulk of any dynamics in my results. I also add p-values for the first-stage F-test under the variable instrumented by each shock in each figure. Rather than include notes to each figure, I describe each specification here. The first few alterations are made to the specification of the VAR itself.

- Figure 18 shows results when I estimate my VAR in levels (not in first differences) and plot the impulse responses (not the cumulative impulse response). I chose the lag lengths as described for my baseline VAR, which leaves me with 7 lags.
- Figure 19 shows results when I stop the estimation of the VAR in 2008.
- Figure 20 shows a VAR that includes the first differences of the shadow rate, 100 times the log of the industrial production index, and 100 times the log of PCE inflation.

The remaining exercises hold the VAR constant, but vary the construction of the instruments.

- Figure 21 shows results when I build my shocks using forecast revisions of GDP and inflation that are designed to reflect revisions starting and ending one quarter later than my baseline measures (from the quarter after the announcement to three quarters after).
- Figure 22 shows results when I only use the first two regimes to identify my shocks.
- Figure 23 shows results when the window I use for determining whether one word is close to another is five words, rather than three.
- Figure 24 shows results when I only use the top three GDP and inflation words.
- Figure 25 shows results when I exclude "oil prices" as an inflation word.
- Figure 26 shows results when I orthogonalize my high-frequency variables with respect to the controls of Bauer and Swanson (Forthcoming): the 13-week return on the S&P 500, the most recent real-time change in nonfarm payrolls, and the lagged Brave et al. (2019) index.

Finally, figure 27 contains the responses of the variables not shown in figure 8.

Figure 18: Robustness: VAR in Levels

(a) Shadow Rate

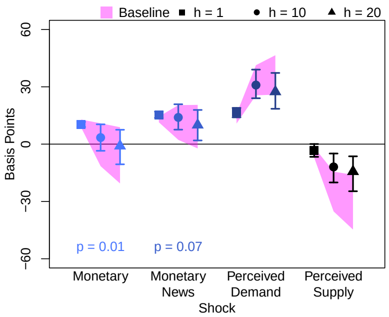

(b) 2-Year Treasury

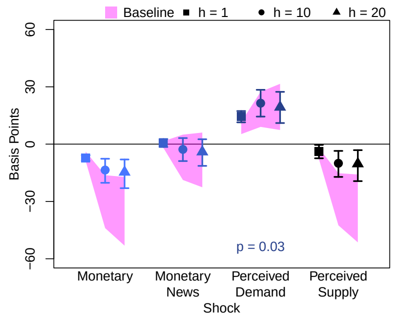

(c) Inflation

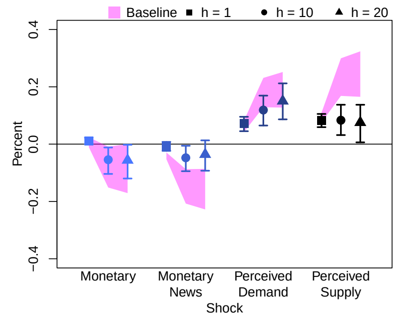

(d) Industrial Production

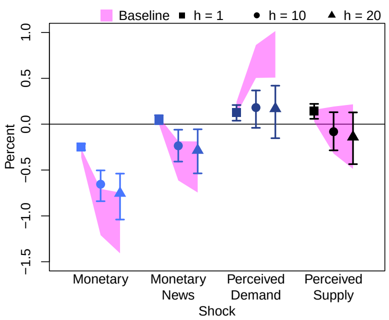

(e) Unemployment Rate

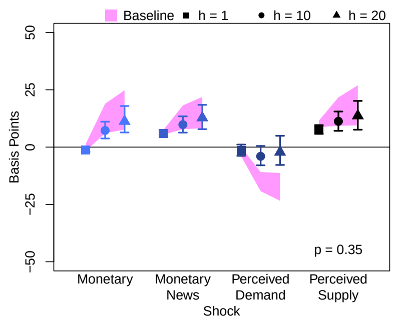

(f) Baa Spread

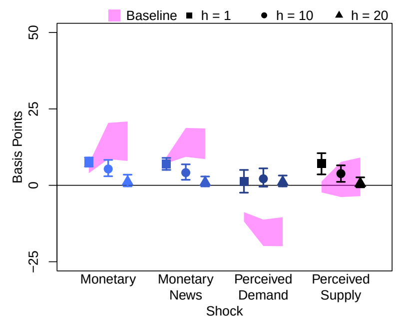

Figure 19: Robustness: End VAR in 2008

(a) Shadow Rate

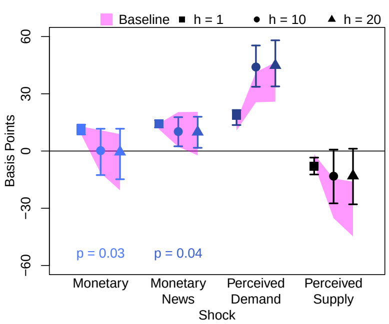

(b) 2-Year Treasury

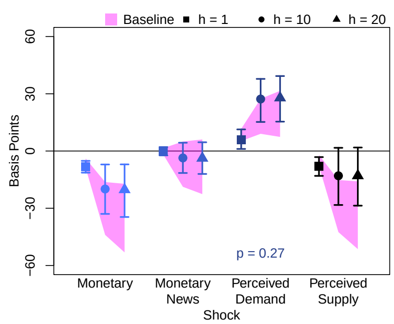

(c) Inflation

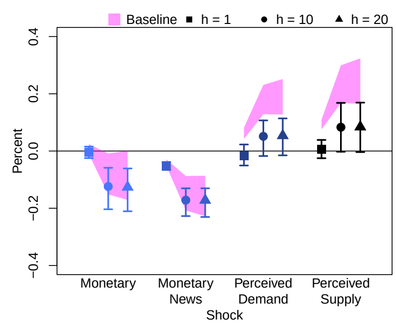

(d) Industrial Production

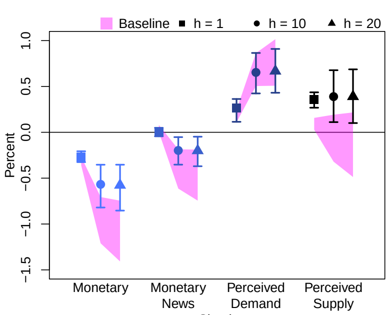

(e) Unemployment Rate

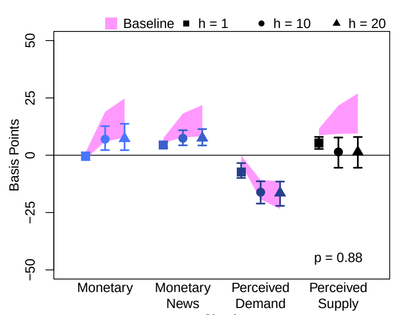

(f) Baa Spread

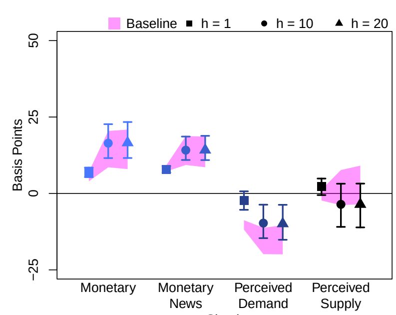

Figure 20: Robustness: Small VAR

(a) Shadow Rate (b) Inflation (c) Industrial Production

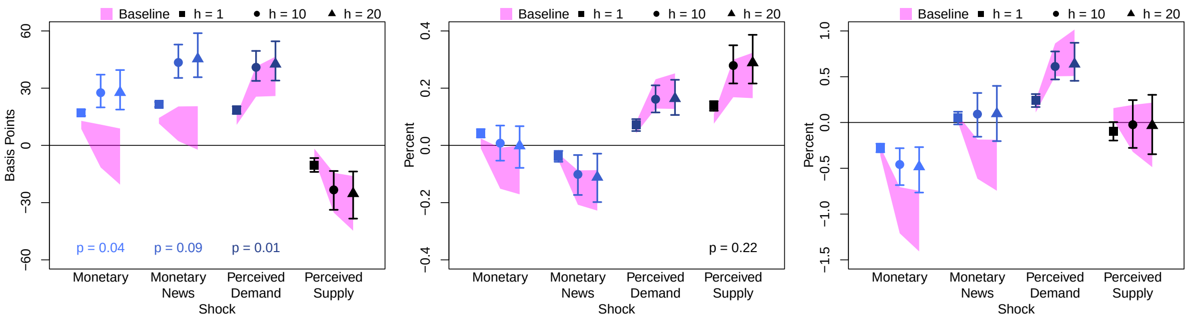

Figure 21: Robustness: Longer Forecast Horizon

(a) Shadow Rate (b) 2-Year Treasury (c) Inflation

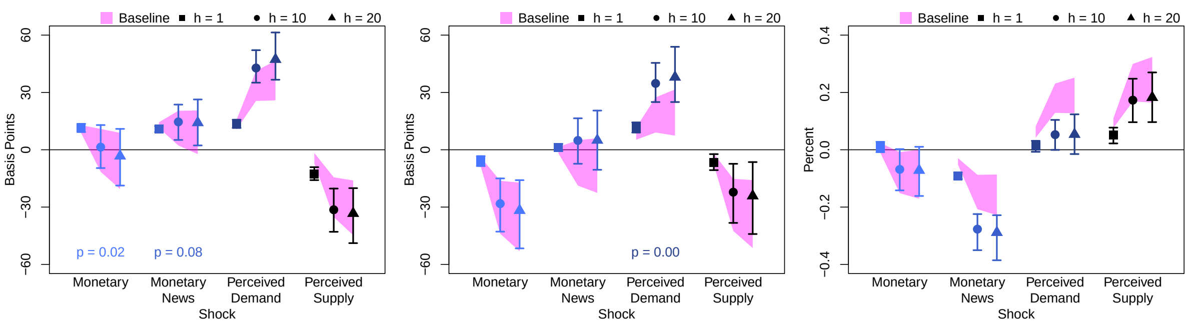

(d) Industrial Production (e) Unemployment Rate (f) Baa Spread

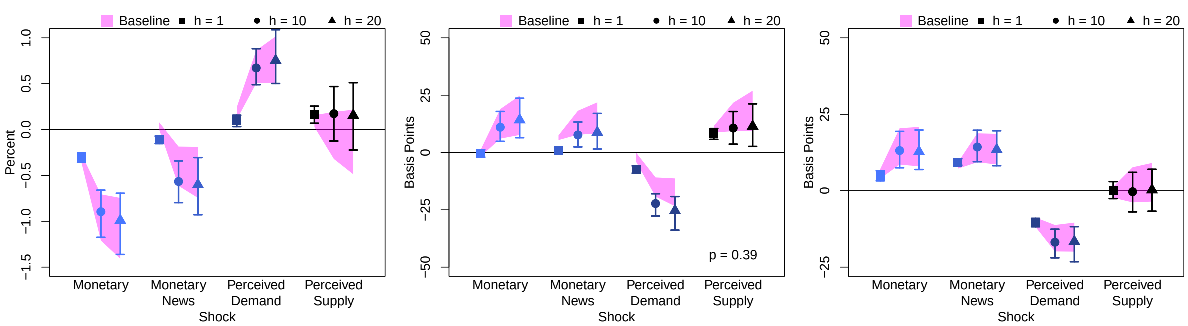

Figure 22: Robustness: 2 Regimes for Shock

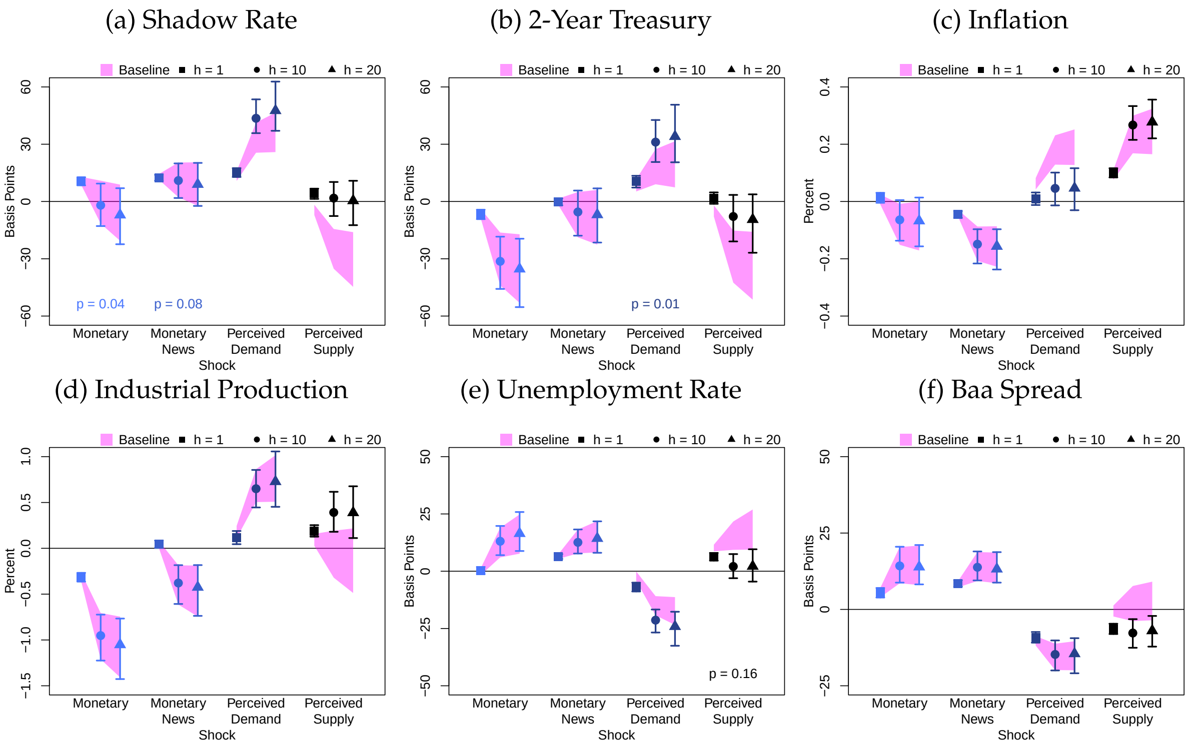

Figure 23: Robustness: 5-Word Window

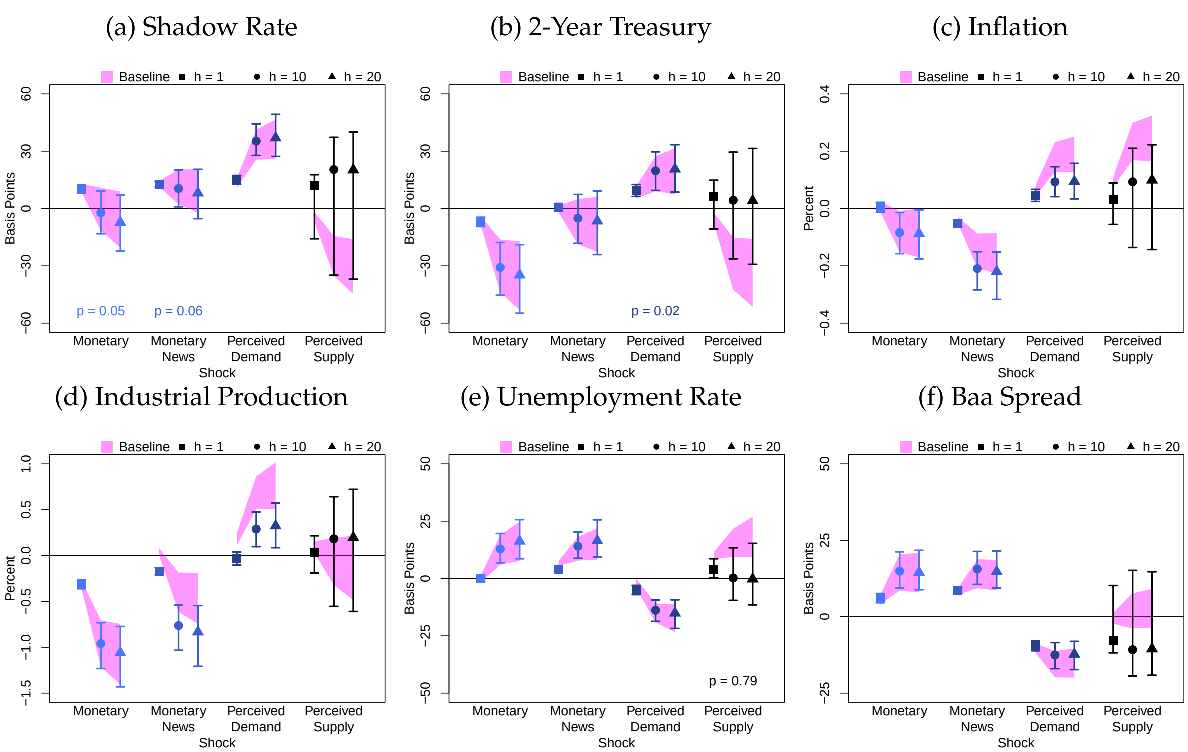

Figure 24: Robustness: 3 Words

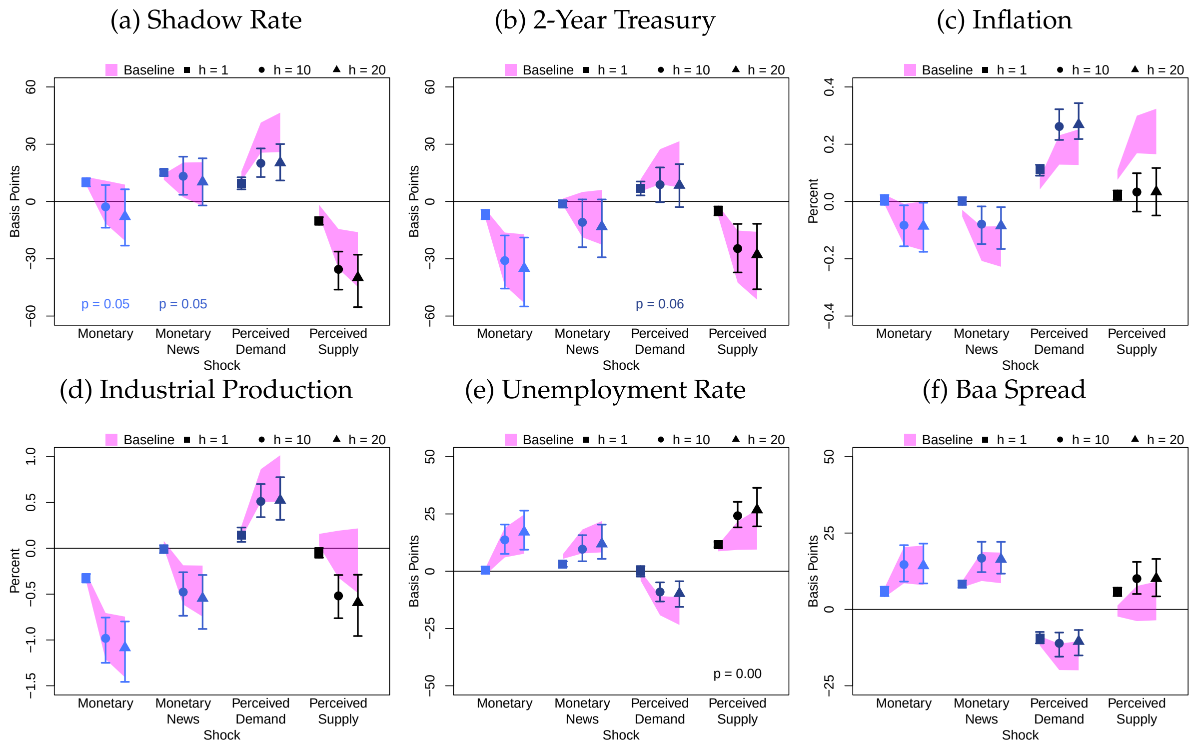

Figure 25: Robustness: Remove “Oil prices” from Inflation Words

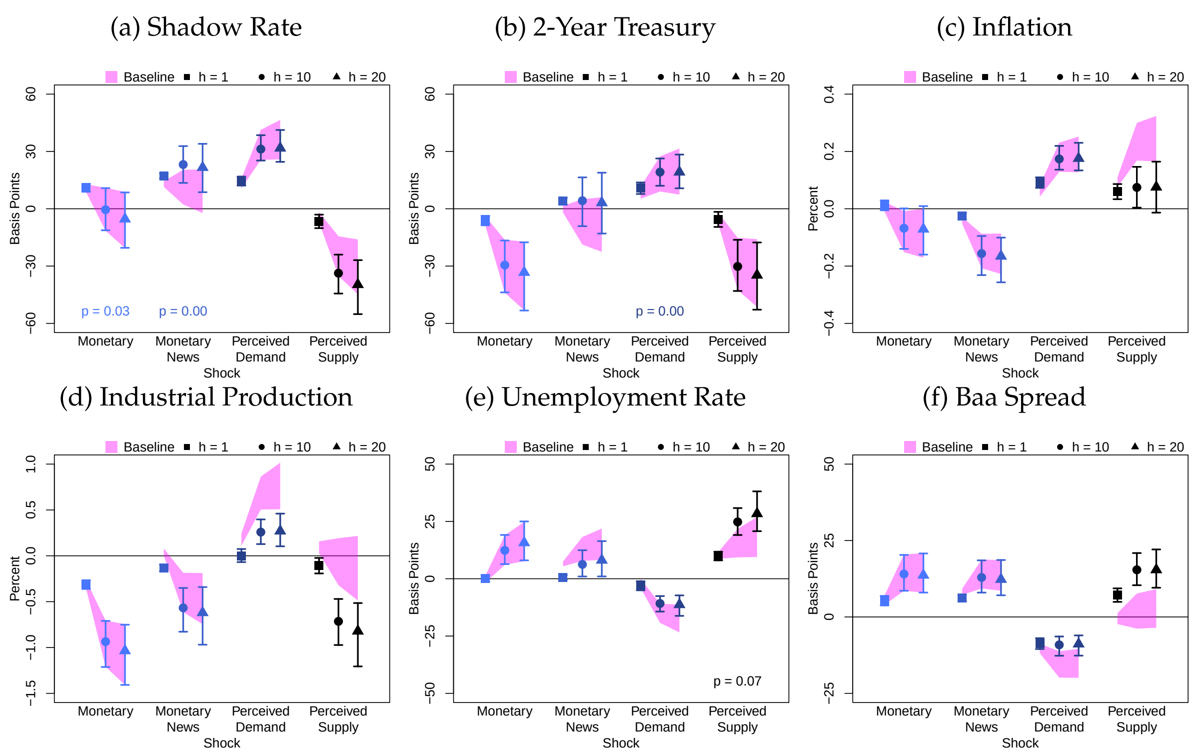

Figure 26: Robustness: All Bauer and Swanson Controls

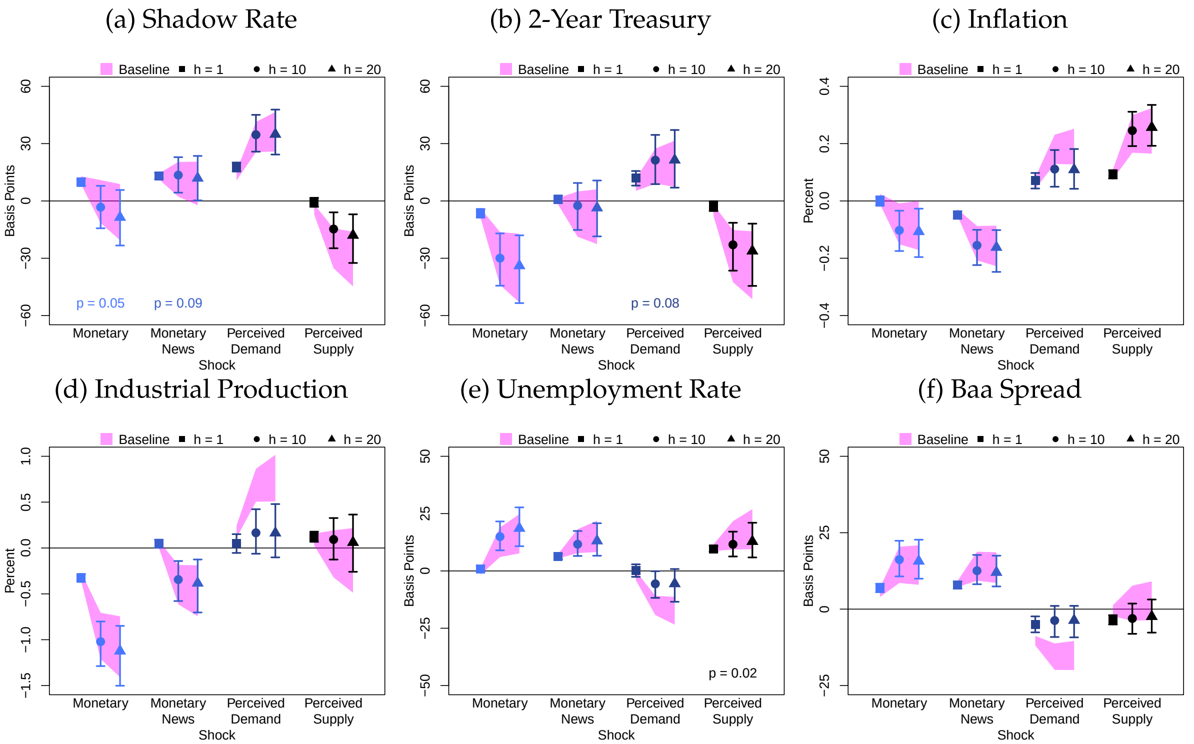

Figure 27: Remaining Responses from Figure 8

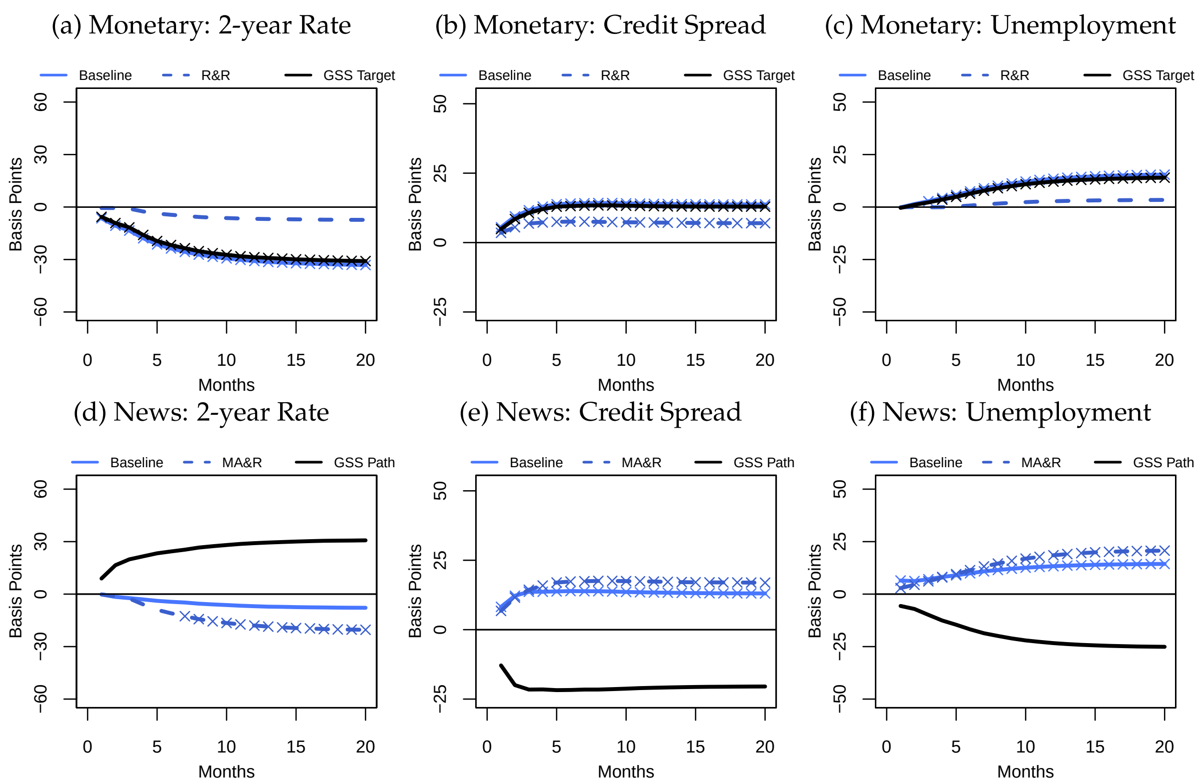

References

Bauer, Michael D. and Eric T. Swanson, "An Alternative Explanation for the "Fed Information Effect"," *American Economic Review*, Forthcoming.

Baum, Christopher F, Mark E Schaffer, and Steven Stillman, "IVREG2: Stata module for extended instrumental variables/2SLS and GMM estimation," Statistical Software Components, Boston College Department of Economics April 2002.

Blanchard, Olivier J., Jean-Paul L'Huillier, and Guido Lorenzoni, "News, Noise, and Fluctuations: An Empirical Exploration," *American Economic Review*, December 2013, *103* (7), 3045–70.

Brave, Scott A, R Andrew Butters, David Kelley et al., "A new "big data" index of US economic activity," *Economic Perspectives, FRB Chicago*, 2019, *1.*

Campbell, Jeffrey R, Charles L Evans, Jonas D.M. Fisher, and Alejandro Justiniano, "Macroeconomic effects of Federal Reserve forward guidance," *Brookings Papers on Economic Activity*, 2012, pp. 1–80.

Galí, Jordi, *Monetary policy, inflation, and the business cycle: an introduction to the new Keynesian framework and its applications*, Princeton University Press, 2015.

Hoesch, Lukas, Barbara Rossi, and Tatevik Sekhposyan, "Has the information channel of monetary policy disappeared? Revisiting the empirical evidence," Working Paper 2020.

Jarociński, Marek and Peter Karadi, "Deconstructing Monetary Policy Surprises—The Role of Information Shocks," *American Economic Journal: Macroeconomics*, April 2020, *12* (2), 1–43.

Loughran, Tim and Bill McDonald, "When is a Liability not a Liability? Textual Analysis, Dictionaries, and 10-Ks," *The Journal of Finance*, 2011, *66* (1), 35–65.

Mikolov, Tomas, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean, "Distributed representations of words and phrases and their compositionality," in "Advances in neural information processing systems" 2013, pp. 3111–3119.

Nakamura, Emi and Jón Steinsson, "High Frequency Identification of Monetary Non-Neutrality," *Quarterly Journal of Economics*, 2018, *133* (3), 1283–1330.

Newey, Whitney K and Kenneth D West, "Automatic lag selection in covariance matrix estimation," *The Review of Economic Studies*, 1994, *61* (4), 631–653.

Porter, Martin F, "An algorithm for suffix stripping," *Program*, 1980.

Sastry, Karthik, "Disagreement About Monetary Policy," Working Paper 2021.

Shapiro, Adam Hale, Moritz Sudhof, and Daniel J Wilson, "Measuring news sentiment," *Journal of Econometrics*, 2020.