

The Macroeconomic Effects of the Federal Reserve's Conventional and Unconventional Monetary Policies

Eric T. Swanson
University of California, Irvine
eric.swanson@uci.edu

Abstract

I separately identify and estimate the effects of the Federal Reserve's federal funds rate, forward guidance, and large-scale asset purchase (LSAP) policies on the U.S. economy. I extend the high-frequency identification strategy of Bauer and Swanson (2023b) for monetary policy VARs by allowing each of the above policies to have possibly different economic effects. I follow Swanson (2021) and Swanson and Jayawickrema (2023) to separately identify federal funds rate, forward guidance, and LSAP components of monetary policy announcements using high-frequency interest rate changes around FOMC announcements, post-FOMC press conferences, FOMC meeting minutes releases, and speeches and testimony by the Fed Chair and Vice Chair. I find that changes in the federal funds rate have had the most powerful effects on the U.S. economy, suggesting that the Fed and other central banks should continue to focus on short-term interest rates as their primary monetary policy tool going forward.

JEL Classification: E52, E58

Version 1.0

August 16, 2023

I thank Juan Rubio-Ramirez, Thomas Drechsel, and seminar participants at UC San Diego, the Renmin University of China Macroeconomics and Finance Workshop, and the IMF/Norges Bank Conference on the Future of Macroeconomic Policy for helpful discussions, comments, and suggestions. The views expressed in this paper, and all errors and omissions, are my own and are not necessarily those of any other individuals or groups listed above.

1. Introduction

One of the most important topics in Monetary Economics is the estimation of monetary policy’s effects on the economy. Prior to 2008, empirical estimates of monetary policy’s effects typically focused on changes in the federal funds rate, the Federal Reserve’s conventional monetary policy tool (e.g., Christiano, Eichenbaum, and Evans, 1999; Romer and Romer, 2004; Faust, Swanson, and Wright, 2004). However, in December 2008, the Federal Reserve’s Federal Open Market Committee (FOMC) lowered the federal funds rate to essentially zero and instead began conducting monetary policy using the “unconventional” tools of forward guidance—communication by the FOMC about the likely future path of the federal funds rate—and large-scale asset purchases, or LSAPs—purchases by the Federal Reserve of hundreds of billions of dollars of longer-term U.S. Treasuries and mortgage-backed securities. The goal of both of these unconventional policies was to stimulate the economy by lowering medium- and longer-term interest rates, even if the federal funds rate was held at essentially zero due to the Fed’s concerns about lowering the federal funds rate below zero—the zero lower bound.

The Fed’s pivot to forward guidance and LSAPs during the 2009–15 U.S. zero lower bound period makes it very difficult for empirical work to ignore these policies after 2008. Thus, more recent work on monetary policy’s effects has focused on interest rates with a longer maturity than the overnight federal funds rate, to better capture some of the effects of forward guidance. For example, Gertler and Karadi (2015) use the one-year Treasury yield and the three-month-ahead federal funds futures rate as their measures of monetary policy, while Gürkaynak, Sack, and Swanson (2005a), Nakamura and Steinsson (2018), and Bauer and Swanson (2023b) extract principal components from federal funds futures and Eurodollar futures with maturities of up to one year. A drawback of this approach, however, is that it does not distinguish between changes in the federal funds rate and changes in forward guidance—and does not consider LSAPs at all—so there is no way to compare the effects of these different policies.

In this paper, I separately identify and estimate the effects of changes in the federal funds rate, forward guidance, and large-scale asset purchases on the U.S. economy. I first follow Swanson (2021) and Swanson and Jayawickrema (2023) to separately identify federal funds rate, forward guidance, and LSAP components of U.S. monetary policy announcements using high-frequency changes in short-, medium-, and long-term interest rates around FOMC announcements, post-FOMC press conferences, FOMC meeting minutes releases, and speeches and testimony by the Federal Reserve Chair and Vice Chair. This gives me a much more comprehensive measure of

changes in forward guidance and LSAPs than has been available to previous authors, because information about the Fed’s forward guidance and LSAPs has often been released to the public through these other channels, especially speeches by the Fed Chair (Swanson and Jayawickrema, 2023; Swanson, 2023). In fact, Swanson and Jayawickrema show that speeches by the Fed Chair have been even more important than FOMC announcements for stocks, bonds, and all but the very shortest-maturity interest rate futures.

I then estimate the effects of these conventional and unconventional monetary policies on the U.S. economy using high-frequency identification of a structural vector autoregression (VAR), following Stock and Watson (2012, 2018), Gertler and Karadi (2015), Ramey (2016), and Bauer and Swanson (2023b). The idea is that high-frequency changes in interest rates around FOMC and other major monetary policy announcements can be used as an instrumental variable in the VAR to estimate the effects of exogenous interest rate changes on the economy. While previous authors considered a single monetary policy instrument—such as the three-month-ahead federal funds futures rate as in Gertler and Karadi (2015)—here I consider the three separate monetary policy instruments described above: changes in the federal funds rate, forward guidance, and LSAPs. I am able to conduct this analysis at a more detailed level than previous authors because I have a much better set of instrumental variables for these three types of monetary policies.

When conducting my high-frequency identification, I take particular care to account for the “Fed Response to News” bias documented by Bauer and Swanson (2023a,b). Under the standard assumptions of full information and rational expectations (FIRE) in financial markets, high-frequency changes in interest rates around monetary policy announcements should be uncorrelated with any economic or financial news that is released prior to those announcements. Intuitively, if there were any correlation, financial market participants would be able to trade profitably on that predictability and drive it away in the process. However, Bauer and Swanson (2023a,b) and others have shown that this orthogonality does not seem to hold in practice; for example, the most recent surprise in the nonfarm payrolls data release is positively correlated with the subsequent high-frequency monetary policy surprise. Bauer and Swanson (2023a,b) present a model and evidence that this correlation is due to a “Fed Response to News” channel, according to which financial markets did not have full information and the Fed responded to incoming economic news more aggressively than the markets had expected.

Bauer and Swanson (2023b) show that the correlation between the high-frequency instrument and past economic and financial news leads to substantially biased estimates in a structural

VAR if the high-frequency instrument is not corrected for that correlation. Intuitively, if positive news about output leads the Fed to raise interest rates by more than the market expected, resulting in the appearance of a surprise monetary policy tightening in the high-frequency data, then the structural VAR will attribute part of the resulting increase in output to tighter monetary policy—exactly opposite to what theoretical models would predict for an exogenous monetary policy tightening. Bauer and Swanson show that this bias substantially attenuates or even reverses the sign of the estimated impulse response functions of output, inflation, and credit spreads to a monetary policy shock in a VAR. I thus follow the Bauer and Swanson (2023a,b) prescription to avoid this bias and orthogonalize my high-frequency instruments with respect to economic and financial market news released in the days and weeks before each monetary policy announcement.

Overall, I find that changes in the federal funds rate had the largest effects on the U.S. economy, followed by changes in forward guidance and, lastly, changes in LSAPs. This is surprising, because Swanson (2021) finds that changes in forward guidance and LSAPs were about equally as powerful as changes in the federal funds rate at moving financial market variables such as Treasury yields, corporate bond yields, and exchange rates. However, I show here that commodity prices respond more strongly to the federal funds rate, and Swanson (2021) shows that the stock market responds more strongly to the federal funds rate, so these outsized effects may be contributing to the greater effects of the federal funds rate on economic variables like inflation and output.

My results thus suggest that central banks should continue to focus on short-term interest rates—conventional monetary policy—as their primary monetary policy tool going forward. Forward guidance and LSAPs are certainly effective at moving financial markets (Swanson, 2021), and appear to be somewhat effective at moving output and inflation, but the effects of conventional monetary policy on the economy appear to be larger and more robust.

After briefly reviewing the literature, the remainder of the paper proceeds as follows. In Section 2, I describe the Swanson and Jayawickrema (2023) high-frequency monetary policy announcement data, including the types of monetary policy announcements considered and how those high-frequency interest rate changes are used to identify changes in the federal funds rate, forward guidance, and large-scale asset purchases. In Section 3, I present my structural VAR analysis with high-frequency identification and describe the Bauer-Swanson (2023a,b) Fed Response to News bias and bias correction. In Section 4, I perform the analysis and present the results. Section 5 discusses the results and concludes.

Related Literature

The present paper is closely related to several strands of the literature. First, there are many studies that use high-frequency changes in interest rates around FOMC announcements to measure changes in monetary policy.¹ Kuttner’s (2001) seminal paper and Bernanke and Kuttner (2005) use this approach, but focus only on changes in the federal funds rate using very short-term federal funds futures. Gürkaynak, Sack, and Swanson (2005a) extend Kuttner’s analysis to decompose the Fed’s FOMC announcements into changes in the federal funds rate and forward guidance, using a range of federal funds futures and Eurodollar futures rates out to a horizon of about one year. Swanson (2021) extends these methods further to separately identify changes in LSAPs as well, but restricts attention to FOMC announcements. Swanson and Jayawickrema (2023) use the methods of Swanson (2021), Rogers, Scotti, and Wright (2018) and Gilchrist, Yue, and Zakrajsek (2019) to estimate federal funds rate, forward guidance, and LSAP components of FOMC and other major monetary policy announcements.

A second strand of the literature looks beyond FOMC announcements to measure high-frequency interest rate changes around other types of monetary policy announcements as well. Gagnon et al. (2011) analyze FOMC announcements in 2009–10 and include one speech by Fed Chair Bernanke; Wright (2012) considers FOMC announcements from 2008–11 and four speeches by Fed Chair Bernanke; Cieslak and Schrimpf (2019) include FOMC announcements and post-FOMC press conferences from 1997–2017, but no other speeches by the Fed Chair; and Kim, Laubach, and Wei (2020) include FOMC announcements from 1991–2015, but just “a few” Fed speeches. In contrast to these studies, Swanson and Jayawickrema (2023), whose data I use here, include all FOMC announcements, post-FOMC press conferences, FOMC meeting minutes releases, and speeches and Congressional testimony by the Fed Chair and Vice Chair from 1988 to 2019, a much more comprehensive set of monetary policy announcements. Bauer and Swanson (2023b) used an earlier version of this dataset that included just FOMC announcements, post-FOMC press conferences, and a subset of the most important speeches by the Fed Chair from 1988–2019.

A third strand of the literature uses high-frequency monetary policy surprises to help identify

¹ An alternative approach to measuring monetary policy shocks, pioneered by Romer and Romer (2004), takes the residuals from a regression of the FOMC’s federal funds rate decisions on information in the Fed’s internal Greenbook forecast of the economy. Aruoba and Drechsel (2023) extend this alternative approach to include textual analysis of the Greenbook and related Fed documents and show that their shock measure is highly correlated with the high-frequency federal funds rate surprises in Swanson (2021).

and estimate the effects of monetary policy on macroeconomic variables in a structural VAR. Early examples are Cochrane and Piazzesi (2002), Faust et al. (2003), and Faust, Swanson, and Wright (2004). Stock and Watson (2012, 2018) discuss how to use high-frequency monetary policy surprises as an external instrument to identify the effects of monetary policy in a VAR, and Gertler and Karadi (2015) and Ramey (2016) follow this approach to obtain estimates that are regarded as benchmarks. Bauer and Swanson (2023b) extend these analyses to assess their robustness with respect to the “Fed Response to News” channel discussed above, and also extend the set of high-frequency monetary policy surprises to include post-FOMC press conferences and speeches by the Fed Chair to increase the power of their high-frequency instrument.

The three papers most closely related to the present paper use the high-frequency identification approach discussed above to estimate the effects of forward guidance and LSAPs on the U.S. economy. Miranda-Agrippino and Ricco (2023) use the Swanson (2021) measures of the federal funds rate, forward guidance, and LSAP components of FOMC announcements to estimate their effects in a structural VAR. They focus primarily on changes in the federal funds rate because their results for forward guidance and LSAPs are puzzling and fragile, as they note. Kim, Laubach, and Wei (2020) obtain very imprecise estimates of the effects of forward guidance, probably due to their short sample and inclusion of only a very small number of non-FOMC announcements in their instrument. Their estimates for the effects of LSAPs are better, but still generally not statistically significant. Eberly, Stock, and Wright (2020) identify the effects of the federal funds rate and a “slope shock”—basically the union of forward guidance and LSAP components of FOMC announcements—in a monetary policy VAR. In contrast to these studies, I separately identify the effects of the federal funds rate, forward guidance, and LSAPs and I have a much more powerful high-frequency instrument for each of these policies, due to my inclusion of speeches by the Fed Chair and Vice Chair, press conference, and minutes releases as additional monetary policy announcements. Finally, I correct my high-frequency instruments for the Fed Response to New bias documented by Bauer and Swanson (2023a,b). As a result, my estimates should be substantially less biased, more precise, and more robust.

Finally, there is a literature on the “Fed Information Effect” and “Fed Response to News” channels that finds that high-frequency changes in interest rates around FOMC announcements are correlated with economic and financial data that pre-date the FOMC announcements. For example, Cieslak (2018) documents correlation with the lagged federal funds rate and employment growth; Miranda-Agrippino (2017) and Miranda-Agrippino and Ricco (2021) with macroeconomic

factors from a dynamic factor model; Bauer and Swanson (2023a,b) with major macroeconomic data release surprises (e.g., nonfarm payrolls) and changes in financial markets (e.g., S&P 500 and commodity prices); Karnaukh and Vokata (2022) with the most recent Blue Chip GDP forecast revisions; Bauer and Chernov (2023) with option-implied skewness of Treasury yields; and Sastry (2021) with the consumer sentiment release, S&P 500 stock returns, and the most recent Blue Chip GDP forecast. Some authors, such as Romer and Romer (2000), Nakamura and Steinsson (2018), and Miranda-Agrippino and Ricco (2021), have argued for the existence of a “Fed Information Effect”, according to which the FOMC’s announcements release information to the private sector about the state of the economy (such as output or inflation) that the private sector didn’t previously have. However, Bauer and Swanson (2023a,b) present evidence that these empirical results are better explained by a “Fed Response to News” channel, where the Fed has simply responded to incoming publicly available economic news by more than the private sector expected. In this paper, I show that changes in the federal funds rate, forward guidance, and LSAPs have the same predictability issues found by the studies cited above. Bauer and Swanson (2023b) show that this correlation significantly biases estimated impulse response functions from a VAR, so I follow their prescriptions to orthogonalize the high-frequency instruments with respect to that news to reduce or eliminate the bias.

2. High-Frequency Monetary Policy Announcement Data

I first construct high-frequency measures of the Fed’s changes in the federal funds rate, forward guidance, and large-scale asset purchases (LSAPs) from 1988–2019, following Swanson and Jayawickrema (2023). I summarize that data construction here, and refer the reader to Swanson and Jayawickrema (2023) for additional details.

2.1 Types of Monetary Policy Announcements Considered

I consider five types of monetary policy announcements by the Federal Reserve: FOMC announcements, post-FOMC press conferences, FOMC meeting minutes releases, speeches and testimony by the Fed Chair, and speeches and testimony by the Federal Reserve Board Vice Chair.

The FOMC holds eight scheduled meetings per year at which it decides what the federal funds rate target will be, and the outcome of those decisions is announced following the end of

each meeting.² In addition, the FOMC sometimes changes its target for the federal funds rate in between scheduled meetings—typically when economic conditions deteriorate rapidly and the FOMC does not want to wait several weeks for the next scheduled meeting—and announces its decision shortly afterward. These are referred to as “unscheduled” or “intermeeting” FOMC announcements. Unless otherwise specified, the term “FOMC announcement” includes both types: scheduled and unscheduled. Since 1994, these announcements have typically been accompanied by an FOMC statement that explains the rationale for the decision; these statements have gradually grown in length over time and currently span about six paragraphs.

Beginning in April 2011, the Federal Reserve Chair holds a press conference in the afternoon after approximately every other FOMC meeting (and after every FOMC meeting beginning in 2019) to answer questions from the press about the FOMC’s decision, the FOMC statement, the rationale for its decision, and monetary policy and the economy more generally.

A few weeks after each FOMC meeting, the FOMC approves the minutes of the meeting and those minutes are released to the public. The minutes summarize all of the discussion that took place at the meeting, including issues related to the U.S. and global economies, U.S. and global financial markets, and the rationale for the FOMC’s monetary policy decision, including any debates or disagreement about that decision.

In addition to the official FOMC communication above, individual FOMC members often give speeches to the public or testimony to Congress in which they discuss their views of the U.S. economy and monetary policy and answer questions from the audience. Financial market participants read and watch these speeches very carefully to look for hints about future U.S. monetary policy. Ideally, we would include every speech by every FOMC member in our analysis, but there are 19 members, each of whom gives about 20–40 speeches per year. To keep the set of speeches down to a more manageable number, Swanson and Jayawickrema (2023) focus on two of the most influential members of the FOMC: the Federal Reserve Board Chair and Vice Chair.

The Federal Reserve Board Chair is also the Chair of the FOMC and is by far the most influential member of the Committee. The Chair sets the agenda for each FOMC meeting, determines the order in which the Committee members present their views, presents his or her own views at the end, and has never been on the losing side of an FOMC vote. While financial

²The FOMC has in effect explicitly announced its decisions for the federal funds rate target after each FOMC meeting since the beginning of 1994. Prior to 1994, the FOMC effectively announced its decisions for the federal funds rate target through the size and type of open market operation conducted in the federal funds market the morning following the FOMC meeting. See Swanson and Jayawickrema (2023) for additional details.

market participants closely watch speeches and testimony by all FOMC members, those by the Fed Chair are given by far the most attention due to the Chair's outsized influence on the Committee.

The Federal Reserve Board Vice Chair is less influential than the Chair, but is more influential than the other Federal Reserve Board Governors and Bank Presidents, with the possible exception of the Federal Reserve Bank of New York President. For example, the Board Vice Chair, like the Chair, frequently testifies before Congress, which other Governors and Bank Presidents rarely do. The Board Vice Chair is also located in the same building as the Chair, is typically in frequent communication with the Chair, and has never voted against the Chair's position at an FOMC meeting. Thus, speeches and testimony by the Vice Chair are also given very high weight by market participants.

2.2 Intradaily Yield Changes around Monetary Policy Announcements

Swanson and Jayawickrema (2023) collect the dates and times of every monetary policy announcement from 1988 to 2019 that falls into one of the five categories above. They then compute the high-frequency, intradaily change in interest rates in a narrow window of time around each of those announcements.

From 1988 to 2019, there are 256 scheduled FOMC announcements, plus an additional 68 unscheduled announcements, for a total of 324 FOMC announcements. However, one of those announcements, September 17, 2001, occurred before financial markets opened that day and after they had been closed for several days following the September 11 terrorist attacks, which makes it impossible to get high-frequency measures of the financial market responses to that announcement that exclude the effects of the terrorist attack itself. I thus exclude that announcement from my analysis, as is standard in the literature, leaving 323 FOMC announcements total. Jayawickrema and Swanson (2023) follow Gürkaynak et al. (2005a) and measure the change in financial markets using an intradaily window beginning 10 minutes before each FOMC announcement and ending 20 minutes after, for a total intradaily window length of 30 minutes.

From 2011 to 2019, there are 40 post-FOMC-meeting press conferences—four each year from 2011 to 2018, and eight in 2019. Post-FOMC press conferences typically last for about one hour, so Swanson and Jayawickrema (2023) begin the intra-daily window 10 minutes before the start of the press conference and end it 80 minutes after, for a total window length of 90 minutes.

FOMC minutes are released eight times per year, a few weeks after each regularly scheduled FOMC meeting. However, Swanson and Jayawickrema (2023) found that FOMC minutes releases

before 1997 essentially never had a significant effect on financial markets, partly because they were released after the market close on Friday afternoons. Thus, I follow Swanson and Jayawickrema and drop the pre-1997 minutes releases from the rest of my analysis, leaving 184 minutes releases—eight per year from 1997 to 2019. The FOMC meeting minutes are much longer than an FOMC statement, comprising about 10–20 pages of text, so Swanson and Jayawickrema (2023) use a longer intradaily window for those announcements, beginning 10 minutes before the minutes release and ending 50 minutes after, for a total window length of 60 minutes.

From 1988 to 2019, the Fed Chair gave 847 speeches and Congressional testimony, not counting the 40 post-FOMC press conferences described above. The Vice Chair gave 310 speeches and Congressional testimony over the same period. (For brevity, I will henceforth use the term “speeches” to refer to both speeches and Congressional testimony.) However, the Fed Chair and Vice Chair often give speeches that are either ceremonial (e.g., commencement or dedication speeches) or are on topics other than monetary policy, such as bank regulation, securities market regulation, fiscal policy, Social Security, the stock market, the exchange rate, check clearing, and other economic and financial issues of national importance. To reduce potential noise and identify those speeches that did contain information about monetary policy, Swanson and Jayawickrema (2023) read the market commentary in *The Wall Street Journal* or *The New York Times* following each speech. This resulted in 364 Fed Chair speeches and 102 Vice Chair speeches that contained enough information about monetary policy to be mentioned as having possible implications for interest rates in the market commentary.

Speeches by the Fed Chair and Vice Chair to the public are typically 30–60 minutes long and can be followed by as much as 30 minutes of answering questions from the audience. For these speeches, Swanson and Jayawickrema (2023) use an intradaily window of 2 hours, beginning 15 minutes before the start of the speech and ending 1 hour and 45 minutes after. Congressional testimony is typically even longer, often consisting of a 30–60 minute opening statement followed by two hours of answering questions from members of Congress. Thus, Swanson and Jayawickrema use a 3.5-hour intradaily window for Congressional testimony, beginning 15 minutes before the start of the testimony and ending 3 hours and 15 minutes after.

Finally, Swanson and Jayawickrema (2023) check whether the intradaily windows around any of the announcements overlap with a macroeconomic data release or other market-moving event such as a Treasury auction. When such an overlap occurs, they read the market commentary in *The Wall Street Journal* or *The New York Times* to determine whether the data release was a

significant mover of financial markets that day and adjust the window start or end time to avoid overlapping with the data release if necessary.

2.3 Federal Funds Rate, Forward Guidance, and LSAPs

I follow Swanson and Jayawickrema (2023) and decompose each of the monetary policy announcements above into federal funds rate, forward guidance, and large-scale asset purchase components, as follows.

Changes in the FOMC’s federal funds rate target are always accompanied by an FOMC announcement, where the term “FOMC announcement” includes notable pre-1994 open market operations, as discussed above. For all other monetary policy announcements (i.e., non-FOMC announcements), I define the surprise change in the federal funds rate target to be zero.

For FOMC announcements, I use the methods of Gürkaynak, Sack, and Swanson (2005a) to separate forward guidance from changes in the federal funds rate. Briefly, define X^{FOMC} to be a matrix of short- and medium-term interest rate futures responses to FOMC announcements and extract the first two principal components of X^{FOMC} . Rotate those two principal components so that the second one has no effect on the current federal funds rate. The first of these two rotated factors thus corresponds to the surprise change in the federal fund rate target, and the second to the surprise change in forward guidance (because it causes interest rate futures to change for reasons other than changes in the current federal funds rate target). See Gürkaynak et al. (2005a) and Swanson and Jayawickrema (2023) for additional details and discussion.

For non-FOMC announcements (press conferences, minutes releases, Fed Chair speeches, and Vice Chair speeches), I use essentially the same method, except that there is no federal funds rate change. I define X^{type} to be a matrix of short- and medium-term interest rate futures responses for monetary policy announcements of a given type (press conferences, minutes releases, etc.). I then extract the first principal component of X^{type} and define it to be the change in forward guidance for that announcement type—this is analogous to the definition of forward guidance for FOMC announcements, above, because there are no changes in the federal funds rate. I rescale each of these forward guidance estimates (for FOMC and non-FOMC announcements) so that a 1-unit change in forward guidance has a 1-percent effect on the 2-year Treasury yield on average. Finally, I treat the union of these forward guidance announcements as a single forward guidance series. Swanson and Jayawickrema (2023) test whether forward guidance from these different announcement types has the same effects on financial markets and find that they do, consistent

with this unified treatment.

To identify LSAPs, I follow Swanson and Jayawickrema (2023) and define the change in LSAPs prior to 2009 to be zero, and from 2009 onward to be the change in long-term Treasury bond yields around each announcement, orthogonalized with respect to the federal funds rate and forward guidance defined above.³ This last identifying assumption is intuitive and is essentially the same one used by Rogers, Scotti, and Wright (2018) and Gilchrist, Yue, and Zakrajsek (2019). It is also simpler than the one in Swanson (2021) and can be used for all five monetary policy announcement types, even those that begin after 2008 (like press conferences) or have very little variation before 2008 (such as speeches by the Fed Vice Chair).⁴ I normalize the scale of the LSAP factor so that a 1-unit change in LSAPs lowers the 10-year Treasury yield by 1 percent.

3. Structural VAR Framework and High-Frequency Identification

I estimate the effects of the federal fund rate, forward guidance, and LSAPs on the U.S. economy using high-frequency identification of a structural VAR. Many recent papers have used a one-dimensional high-frequency measure of interest rate changes around FOMC announcements to estimate the effects of monetary policy in a VAR—see, e.g., Cochrane and Piazzesi (2002), Faust et al. (2003), Faust, Swanson, and Wright (2004), Stock and Watson (2012, 2018), Gertler and Karadi (2015), Ramey (2016), Miranda-Agrippino and Ricco (2021), and Bauer and Swanson (2023b). High-frequency interest rate changes around monetary policy announcements are appealing in these applications because they plausibly rule out reverse causality and other endogeneity problems, as I discuss below. I extend these methods to separately analyze the effects of the federal funds rate, forward guidance, and LSAPs on the U.S. economy.

3.1 Structural VAR

The VAR in this analysis includes seven monthly macroeconomic variables: the log of industrial production, the log of the consumer price index, the log of the Commodity Research Bureau index

³I define the change in the long-term Treasury bond yield to be the average of the change in the 10-year Treasury yield and the 30-year Treasury yield.

⁴Swanson (2021) includes a range of Treasury yield responses in the matrix X^{FOMC} and extracts the first three principal components from that matrix. He imposes identifying restrictions that the second and third factors do not systematically affect the current federal funds rate target, and that the third factor (LSAPs) has minimum variance over the pre-2008 sample. These identifying assumptions work well for FOMC announcements, but cannot be applied to press conferences, which begin in 2011, and work poorly for FOMC minutes releases and Vice Chair speeches, which have relatively little variation before 2008.

of commodity prices, the Gilchrist-Zakrajsek (2012) credit spread, the Wu-Xia (2016) shadow federal funds rate, the two-year Treasury yield, and the 10-year Treasury yield.⁵ Industrial production is a standard measure of output and the CPI is a standard measure of prices. The CRB commodity price index is not essential but helps to describe the response of prices in the VAR, and Bauer and Swanson (2023b) include it in their “best practice” estimates. I include the GZ credit spread because Caldara and Herbst (2019) found credit spreads to be important for the estimation of monetary policy VARs. The Wu-Xia (2016) shadow federal funds rate is not essential, but helps to distinguish between the effects of changes in the federal funds rate vs. forward guidance. The regular federal funds rate is severely constrained by the zero lower bound from 2009–15, which makes it invalid in a linear VAR specification; the Wu-Xia (2016) shadow federal funds rate avoids this problem by estimating an equivalent value for the federal funds rate—which can be negative—based on the rest of the yield curve in those years. The two-year Treasury yield is a good measure of the overall stance of monetary policy, is very sensitive to changes in forward guidance (Gürkaynak et al., 2005a), and was largely unconstrained by the zero lower bound from 2009–15 (Swanson and Williams, 2014; Swanson, 2018). Finally, the 10-year Treasury yield is very sensitive to changes in LSAPs (Swanson, 2021), making it useful for estimating the effects of LSAPs on the economy.

This VAR specification is very similar to Bauer and Swanson (2023b), except that it includes three interest rate measures—the shadow federal funds rate, two-year Treasury yield, and 10-year Treasury yield—rather than just one. Having three interest rates in the VAR is useful because I am interested in estimating the effects of three different types of monetary policies, which are most easily distinguished by their effects on these three different interest rates.⁶

I stack these seven variables into a vector Y_t and estimate the reduced-form VAR,

$$Y_t = \alpha + B(L)Y_{t-1} + u_t, \tag{1}$$

where α is a constant, $B(L)$ a matrix polynomial in the lag operator with 6 monthly lags, and

⁵Industrial production and the CPI are from the Federal Reserve Bank of St. Louis FRED database. The CRB commodity price index is from Bloomberg. An updated version of the Gilchrist-Zakrajsek (2012) credit spread is from the Federal Reserve Board. An updated version of the Wu-Xia (2016) shadow federal funds rate is from Cynthia Wu’s website. The 2-year and 10-year Treasury yields are zero-coupon yields from the updated Gürkaynak, Sack, and Wright (2007) database at the Federal Reserve Board.

⁶Note that it is not necessary to have three different interest rates in the VAR to estimate the effects of the three different monetary policy shocks. The external instruments approach below works just as well in theory with any collection of variables in the VAR, as long as the reduced-form VAR is well specified. The advantage of including three different interest rates is that it helps to clearly illustrate the different effects of the three different monetary policy shocks I consider.

u_t is a 7×1 vector of serially uncorrelated regression residuals, with $\text{Var}(u_t) = \Omega$. I estimate regression (1) from January 1973 to February 2020 via ordinary least squares, as in Bauer and Swanson (2023b). The Gilchrist-Zakrajsek (2012) credit spread data begin in 1973, which prevents me from beginning the sample earlier, and I end the sample in February 2020 to avoid the large swings in the macroeconomic data due to the Covid pandemic.⁷

I follow standard practice and assume that the economy is driven by a set of serially uncorrelated structural shocks, ε_t , with $\text{Var}(\varepsilon_t) = I$ (see, e.g., Ramey, 2016). Since the dynamics of the economy are determined by $B(L)$, the effects of different structural shocks ε_t on Y_t are completely determined by differences in their impact effects on Y_t in period t —that is, by their effects on u_t ,

$$u_t = S\varepsilon_t, \quad (2)$$

which I assume are linear, with S a matrix of appropriate dimensions.⁸ I assume that one of the structural shocks is a “federal funds rate shock”, a second is a “forward guidance shock”, and a third is an “LSAP shock”, and I denote those shocks by ε_t^{ff} , ε_t^{fg} , and ε_t^{lsap} , respectively, and order them first, second, and third in the vector ε_t . The first three columns of S , denoted s_1 , s_2 , and s_3 , then describe the impact effects of each of the three structural monetary policy shocks ε_t^{ff} , ε_t^{fg} , and ε_t^{lsap} on u_t and hence Y_t .

3.2 High-Frequency Identification

To identify the impact effects s_1 , s_2 , and s_3 , I use high-frequency identification. Let $i \in \{\text{federal funds rate, forward guidance, LSAPs}\}$ denote one of the three monetary policies being studied, and let \tilde{z}_t^i denote the set of high-frequency changes in monetary policy instrument i around all of the monetary policy announcements above (FOMC announcements, press conferences, minutes releases, Fed Chair speeches, and Vice Chair speeches) from 1988–2019. Let z_t^i denote the monthly version of \tilde{z}_t^i , obtained by summing over all of the high-frequency changes in policy instrument i within each month.

The idea is that z_t^i is a good instrument for ε_t^i for each of the three monetary policy tools i . As discussed by Stock and Watson (2012, 2018), in order for z_t^i to be a valid instrument for ε_t^i , it

⁷The 10-year Treasury yield data begin in 1971, so it would be difficult to extend the VAR estimation back much further even without the GZ credit spread.

⁸It is common in the literature to assume that S is a nonsingular square matrix, which is equivalent to assuming that the VAR is invertible (see Stock and Watson, 2018, and Plagborg-Møller and Wolf, 2021, for discussions of invertibility). I do not require that assumption here, since I am only interested in estimating the effects of the first three shocks, ε_t^{ff} , ε_t^{fg} , and ε_t^{lsap} . In principle, the number of different structural shocks in ε_t could be large.

must satisfy an instrument *relevance* condition,

$$E[z_t^i \varepsilon_t^i] \neq 0, \quad (3)$$

and an instrument *exogeneity* condition,

$$E[z_t^i \varepsilon_t^{-i}] = 0, \quad (4)$$

where ε_t^{-i} denotes any element of ε_t other than ε_t^i .⁹

The appeal of high-frequency interest rate changes around monetary policy announcements is that they very plausibly satisfy conditions (3)–(4). First, FOMC announcements, Fed Chair speeches, and the other announcements in my data set are a very important part of the news about monetary policy each month, so the correlation between z_t^i and ε_t^i in (3) should be positive and large.¹⁰ Importantly, including non-FOMC announcements such as Fed Chair speeches provides me with a *much* more relevant instrument than using FOMC announcements alone, as shown by Bauer and Swanson (2023b) and as I also verify below. Second, because z_t^i consists of high-frequency interest rate changes in narrow windows of time around monetary policy announcements, it is very unlikely that other structural shocks in ε_t^{-i} are significantly affecting financial markets at exactly the same time, so these other shocks should be uncorrelated with z_t^i , implying (4). (And the three monetary policy instruments i are already orthogonal with respect to each other by construction, so condition (4) is also satisfied for them.) Stock and Watson (2012, 2018) refer to high-frequency instruments like z_t^i as *external instruments* because they are not contained within the VAR data Y_t .

Given the external instruments z_t^i , I estimate the impact effects s_1 , s_2 , and s_3 in the VAR as described in Stock and Watson (2012, 2018), Gertler and Karadi (2015), and Bauer and Swanson (2023b), considering each of the three monetary policy shocks and instruments separately, one at a time. (Because the monetary policy shocks are mutually orthogonal and the instruments z_t^i are mutually orthogonal, there is no difference between estimating the effects of the shocks individually vs. jointly.)¹¹ For example, consider first z_t^{ff} and s_1 . Denote the shadow federal

⁹ Strictly speaking, I have not assumed invertibility of the VAR, so the instrument z_t^i must also satisfy a *lead-lag exogeneity* condition, $E[z_t^i \varepsilon_\tau^i] = 0$ for $\tau \neq t$, discussed below. See Stock and Watson (2018).

¹⁰ Note that $z_t^i \neq \varepsilon_t^i$ in general, because not all the news about monetary policy each month is released in FOMC announcements, press conferences, FOMC minutes, Fed Chair speeches, or Vice Chair speeches. For example, speeches by other FOMC members also contain information about monetary policy.

¹¹ Arias, Rubio-Ramírez, and Waggoner (2021) discuss individual vs. joint identification with external instruments and show how the results can differ across the two methods when the instruments are not orthogonal.

funds rate in Y_t by Y_t^{ff} , and the corresponding reduced-form residuals by u_t^{ff} . We then estimate the vector s_1 by running the regression

$$Y_t = \tilde{\alpha} + \tilde{B}(L)Y_{t-1} + s_1 Y_t^{ff} + \tilde{u}_t \quad (5)$$

via equation-by-equation two-stage least squares, where $\tilde{B}(L)$ has the same number of lags as $B(L)$ and z_t^{ff} is used as the instrument for Y_t^{ff} . Note that one can obtain the same point estimates for s_1 by regressing the reduced-form residuals u_t from (1) on u_t^{ff} using z_t^{ff} as the instrument; Stock and Watson (2012) recommend using specification (5) to avoid generated regressors and correctly estimate the standard errors for s_1 . Note also that the sample for the two-stage least squares regression (5) used to estimate s_1 does not have to be the same as for the reduced-form VAR regression (1) used to estimate α and $B(L)$. In fact, my high-frequency interest rate change data is available only from 1988–2019, while I can estimate the reduced-form VAR coefficients α and $B(L)$ over the longer sample from 1973:1–2020:2.

It is straightforward to show that (3)–(4) imply regression (5) produces an unbiased and consistent estimate of s_1 , with the impact effect on Y_t^{ff} normalized to unity. In my empirical results below, I rescale s_1 so that the impact effect on Y_t^{ff} is 25 basis points (bp), rather than 1 percentage point.

An important statistic for assessing the quality of the instrument z_t^i is its first-stage F -statistic. In the first-stage regression,

$$Y_t^{ff} = \gamma + C(L)Y_{t-1} + \theta z_t^{ff} + \eta_t, \quad (6)$$

the lag polynomial $C(L)$ has the same number of lags as $B(L)$ and the first-stage F -statistic is the squared value of the t -statistic for the estimated coefficient $\hat{\theta}$.¹²

The procedure for estimating the impact effects s_2 and s_3 for forward guidance and LSAPs is essentially the same. For forward guidance, it is natural to use the two-year Treasury yield in the VAR as the reference variable for the two-stage least squares regression, so (5) instead takes the form

$$Y_t = \check{\alpha} + \check{B}(L)Y_{t-1} + s_2 Y_t^{2y} + \check{u}_t, \quad (7)$$

¹²In theory, z_t^{ff} should be orthogonal to Y_{t-1} and the constant term, so (6) can be run without the lagged controls, and then the first-stage F -statistic is the exclusion F -statistic for the whole regression. Note, however, that the first-stage F -statistic is not the exclusion F -statistic for the regression (6) with the lagged controls included, because that will typically be a large number even if the instrument z_t^{ff} had zero relevance.

where Y_t^{2y} denotes the two-year Treasury yield in the VAR, and z_t^{fg} is the instrument for Y_t^{2y} . For LSAPs, the natural reference variable is the 10-year Treasury yield, so the two-stage least squares regression is

$$Y_t = \tilde{\alpha} + \tilde{B}(L)Y_{t-1} + s_3 Y_t^{10y} + \tilde{u}_t, \quad (8)$$

where Y_t^{10y} denotes the 10-year Treasury yield in the VAR, and z_t^{lsap} is the instrument for Y_t^{10y} .

Given the estimated impact effects s_1 , s_2 , and s_3 , it is then straightforward to use the estimated matrix lag polynomial $B(L)$ from regression (1) to compute the impulse response functions for Y_t to each of the structural shocks ε_t^{ff} , ε_t^{fg} , and ε_t^{lsap} .

3.3 The “Fed Response to News” Channel and Bias

One of the appeals of high-frequency identification is that high-frequency interest rate changes around monetary policy announcements are typically viewed as being orthogonal to macroeconomic and financial data that predate the announcement. This view is supported by the standard argument that, otherwise, financial market participants would be able to trade profitably on that predictability and drive it away in the process. The orthogonality is important for the identification, because it implies that z_t^i is orthogonal to shocks to other variables earlier in the same month (exogeneity condition (4)) and to lagged variables in the VAR (the lead-lag exogeneity condition discussed in Stock and Watson, 2018, and footnote 8, above).

However, a number of recent studies have found that high-frequency interest rate changes such as z_t^i are not orthogonal to past economic and financial news in practice. For example, Bauer and Swanson (2023a,b) show that the most recent surprise in the nonfarm payrolls data release is positively correlated with the high-frequency monetary policy surprise for the next FOMC announcement, and Cieslak (2018), Miranda-Agrippino and Ricco (2021), Bauer and Swanson (2023a,b), and Sastry (2021) document substantial additional correlation with a number of other economic and financial news series that predate those monetary policy announcements.

Bauer and Swanson (2023a,b) present a model and empirical evidence that these correlations are due to a “Fed Response to News” channel, where the private sector has imperfect information about the Fed’s monetary policy reaction function and the Fed responded to economic news more aggressively than the private sector had expected. Bauer and Swanson (2023b) also present evidence that the Fed’s responsiveness to the economy has increased over time, which provides a plausible explanation for why the private sector’s estimate of the Fed’s responsiveness has lagged behind, leading to the private sector’s underestimation on average over the 1988–2019 sample.

Bauer and Swanson (2023b) show that the correlation between the high-frequency instrument z_t^i and past economic and financial news will lead to substantially biased estimates in a structural VAR if the high-frequency instrument is not purged of that correlation. Intuitively, if positive news about output leads the Fed to raise interest rates by more than the market expected, resulting in the appearance of a surprise monetary policy tightening in the high-frequency data, then the structural VAR will attribute part of the resulting increase in output to tighter monetary policy—exactly opposite to what standard theory would predict and leading to biased estimates. Bauer and Swanson (2023b) show that this bias greatly attenuates or even reverses the sign of the estimated impulse response functions of output, inflation, and credit spreads to a monetary policy shock in a standard monetary policy VAR.

To eliminate the bias, Bauer and Swanson (2023a,b) recommend orthogonalizing the high-frequency instruments z_t^i with respect to economic and financial data that pre-date the monetary policy announcements. In other words, run regressions of the form

$$z_t^i = \delta + \psi' X_{t-} + z_t^{i\perp}, \quad (9)$$

where the regressors X_{t-} are measures of economic and financial news that pre-date the monetary policy announcements in z_t^i each month, δ and ψ are parameters, and the residuals from the regression, $z_t^{i\perp}$, are orthogonal to the data in X_{t-} . If X_{t-} is sufficiently comprehensive, then the orthogonalized instrument $z_t^{i\perp}$ should produce impulse response function estimates that are unbiased.

In my analysis below, I show that the high-frequency instruments for all three monetary policy tools—the federal funds rate, forward guidance, and LSAPs—are correlated with past macroeconomic and financial news that predate the monetary policy announcements, extending the findings in Bauer and Swanson (2023a,b). It is thus very important to orthogonalize these instruments to avoid biased impulse response function estimates.

4. Results

I now perform the structural VAR analysis described in Section 3 using the high-frequency monetary policy surprise data for the federal funds rate, forward guidance, and large-scale asset purchases described in Section 2. I first check the importance of the Fed Response to News channel documented by Bauer and Swanson (2023a,b) for the high-frequency monetary policy instruments, and then report the results for each of these three monetary policies in turn.

4.1 The Fed Response to News Channel

I first check whether the high-frequency instruments z_t^{ff} , z_t^{fg} , and z_t^{lsap} are correlated with economic and financial news that predate the monetary policy announcements, as documented by Bauer and Swanson (2023a,b) for their high-frequency monetary policy surprise measure. Table 1 reports results from regressing these instruments on a number of variables suggested by Bauer and Swanson. For macroeconomic variables, I consider the most recent surprise in the nonfarm payrolls release, the unemployment rate release, the GDP release, and the core CPI release.¹³ For financial variables, I consider the percent change in the S&P 500 stock index from 3 months before the monetary policy announcement to the day before the monetary policy announcement, the change in the shadow federal funds rate, 2-year Treasury yield, and 10-year Treasury yield over the same 3-month window, and the change in the Baa-Treasury spread and percent change in commodity prices over the same 3-month window.¹⁴ For the LSAP instrument z_t^{lsap} , I also consider the one-month change in the Chicago Fed’s National Financial Conditions Index, for the reasons discussed below.¹⁵ Finally, I include two lagged values of the left-hand side variables z_t^i , since Miranda-Agrippino and Ricco (2021) found the high-frequency surprises to be serially correlated. The first column of Table 1 reports results for the federal funds rate instrument z_t^{ff} , the second column for the forward guidance instrument z_t^{fg} , and the third column for the LSAP instrument z_t^{lsap} . Standard errors computed using 10,000 bootstrap replications are reported beneath each coefficient estimate.¹⁶

Under the common assumptions of full information and rational expectations (FIRE) in financial markets, the high-frequency instruments z_t^i should be uncorrelated with all of the vari-

¹³Macroeconomic data release surprises are computed in the same way as in Bauer and Swanson (2023a,b), Gürkaynak, Sack, and Swanson (2005b), and Swanson and Williams (2014): as the actual released value of the statistic minus the market’s *ex ante* expectation of that release, as reported by Money Market Services. The unemployment rate surprise, core CPI surprise, and GDP surprise are in percentage points; the nonfarm payrolls surprise is in thousands of workers, divided by 1000 to make the scale more comparable to the other statistics. If there are multiple monetary policy announcements in a given month, I use the most recent surprise that predates the first announcement that month.

¹⁴S&P 500 data are from Yahoo! Finance, the shadow federal funds rate is the update of Wu and Xia (2016) from Cynthia Wu’s website, Treasury yield data are from the updated Gürkaynak, Sack, and Wright (2007) database at the Federal Reserve Board, the Baa-Treasury spread is from the St. Louis Fed FRED database, and the commodity price index is the Bloomberg total commodity price index minus 0.4 times the Bloomberg agricultural commodity price index, as in Bauer and Swanson (2023b). The 3-month changes are computed as the change in the asset price from 13 weeks before to the day before the first monetary policy announcement of the month.

¹⁵The data is the National Financial Conditions Index for credit from the Federal Reserve Bank of Chicago’s website. The data are weekly on Fridays, so the one-month change is computed from four weeks before the Friday prior to the monetary policy announcement to the Friday before the monetary policy announcement.

¹⁶As in Gertler and Karadi (2015) and Bauer and Swanson (2023b), I compute these standard errors using the Wild bootstrap to allow for heteroskedasticity. Results using asymptotic heteroskedasticity-consistent standard errors are very similar.

TABLE 1: REGRESSIONS OF HIGH-FREQUENCY INSTRUMENTS z_t^{ff} , z_t^{fg} , AND z_t^{lsap} ON ECONOMIC AND FINANCIAL NEWS PREDATING THE MONETARY POLICY ANNOUNCEMENTS

	z_t^{ff}	z_t^{fg}	z_t^{lsap}
Macroeconomic News			
Nonfarm payrolls surprise	1.39** (.572)	-0.018 (.034)	0.008 (.009)
Unemployment surprise	-0.31 (.371)	-0.010 (.023)	-0.016 (.010)
GDP surprise	-0.09 (.120)	-0.003 (.007)	0.002 (.002)
Core CPI surprise	0.09 (.489)	-0.035 (.032)	-0.014 (.009)
Financial News			
$\Delta \log$ S&P500 (3m)	0.26 (.995)	0.118** (.049)	0.017 (.035)
Δ shadow fed funds rate (3m)	0.52*** (.200)	-0.022** (.011)	0.002 (.005)
Δ 2-year Treasury (3m)	0.16 (.243)	0.053*** (.014)	-0.005 (.009)
Δ 10-year Treasury (3m)	-0.15 (.195)	-0.025*** (.010)	0.001 (.009)
Δ Baa spread (3m)	-0.33* (.189)	0.020* (.012)	0.007 (.008)
$\Delta \log$ Commodity prices (3m)	0.43 (.959)	0.151*** (.053)	0.019 (.024)
Δ Chicago Fed NFCI (1m)	—	—	-0.057** (.025)
Lagged Monetary Policy Surprises			
z_{t-1}^i	-0.17* (.104)	-0.199** (.078)	-0.104 (.093)
z_{t-2}^i	-0.26*** (.096)	-0.115* (.063)	0.031 (.114)
R^2	0.18	0.12	0.06

Notes: Coefficient estimates ψ from regressions $z_t^i = \delta + \psi'X_{t-} + \zeta_t$. Variables X_{t-} are observed prior to the monetary policy announcements each month: the most recent nonfarm payrolls surprise, unemployment rate surprise, GDP surprise, and core CPI surprise, percent change in S&P 500 from 3 months before to the day before the monetary policy announcement, change in shadow federal funds rate, 2-year Treasury, and 10-year Treasury yield over the same period, change in Baa spread over the same period, change in a commodity price index over the same period, change in the Chicago Fed National Financial Conditions index over the month before the monetary policy announcement, and lagged values of the endogenous variable z_t^i . Bootstrapped standard errors in parentheses. *, **, and *** denote statistical significance at the 10%, 5%, and 1% levels, respectively. Sample: 1988:1–2019:12. See text for details.

ables in Table 1, because those variables pre-date the monetary policy announcements each month. Nevertheless, many of the coefficients in Table 1 are statistically and economically significant, with regression R^2 of up to 18 percent. The coefficients in Table 1 are consistent with the Fed Response to News channel (Bauer and Swanson, 2023a,b): financial markets seem to have underestimated how aggressively the Fed would respond to the incoming data. For example, a negative surprise in nonfarm payrolls or an increase in the Baa credit spread leads the Fed to cut the federal funds rate by more than markets expected (first column), while lower stock or commodity prices lead the Fed to ease forward guidance by more than markets expected (second column). More generally, if the short end of the yield curve (the shadow federal funds rate) has fallen recently, the Fed tends to cut the federal funds rate by more than markets expected, while if the 2-year Treasury yield has fallen recently, the Fed tends to give easier forward guidance than markets expected. The federal funds rate and forward guidance instruments are also negatively serially correlated, so that large surprises tend to be followed by small surprises in the opposite direction for the next two months.

Turning to the LSAP instrument in the third column of Table 1, the macroeconomic and financial news that predict the Fed’s upcoming funds rate and forward guidance announcements generally do not help to predict the Fed’s upcoming LSAP announcements. However, recent changes in credit conditions, as measured by the Chicago Fed’s National Financial Conditions Index, are a strong predictor of those LSAP announcements, with an increase in credit market stress leading the Fed to lower long-term interest rates via LSAPs by more than markets expected (third column of Table 1).

Note that these results cannot be explained by a “Fed Information Effect” (Romer and Romer, 2000; Nakamura and Steinsson, 2018), because all of the variables in Table 1 are publicly observed prior to each monetary policy announcement in the regression.¹⁷ Instead, Bauer and Swanson (2023a,b) present a simple model of the Fed Response to News channel that explains these results based on private sector agents having imperfect information about the Fed’s monetary policy reaction function and (optimally) updating their beliefs about that reaction function over time.

For the VAR analysis in the present paper, the exact reasons for the correlations in Table 1 are not particularly important; what is crucial is that the correlations are present in the data and will tend to bias VAR impulse response function estimates if they are not removed, as discussed

¹⁷Other authors have found similar predictability—see the Introduction, above, and Bauer and Swanson (2023a,b).

in Section 3, above, as shown by Bauer and Swanson (2023a,b), and as I also verify below.

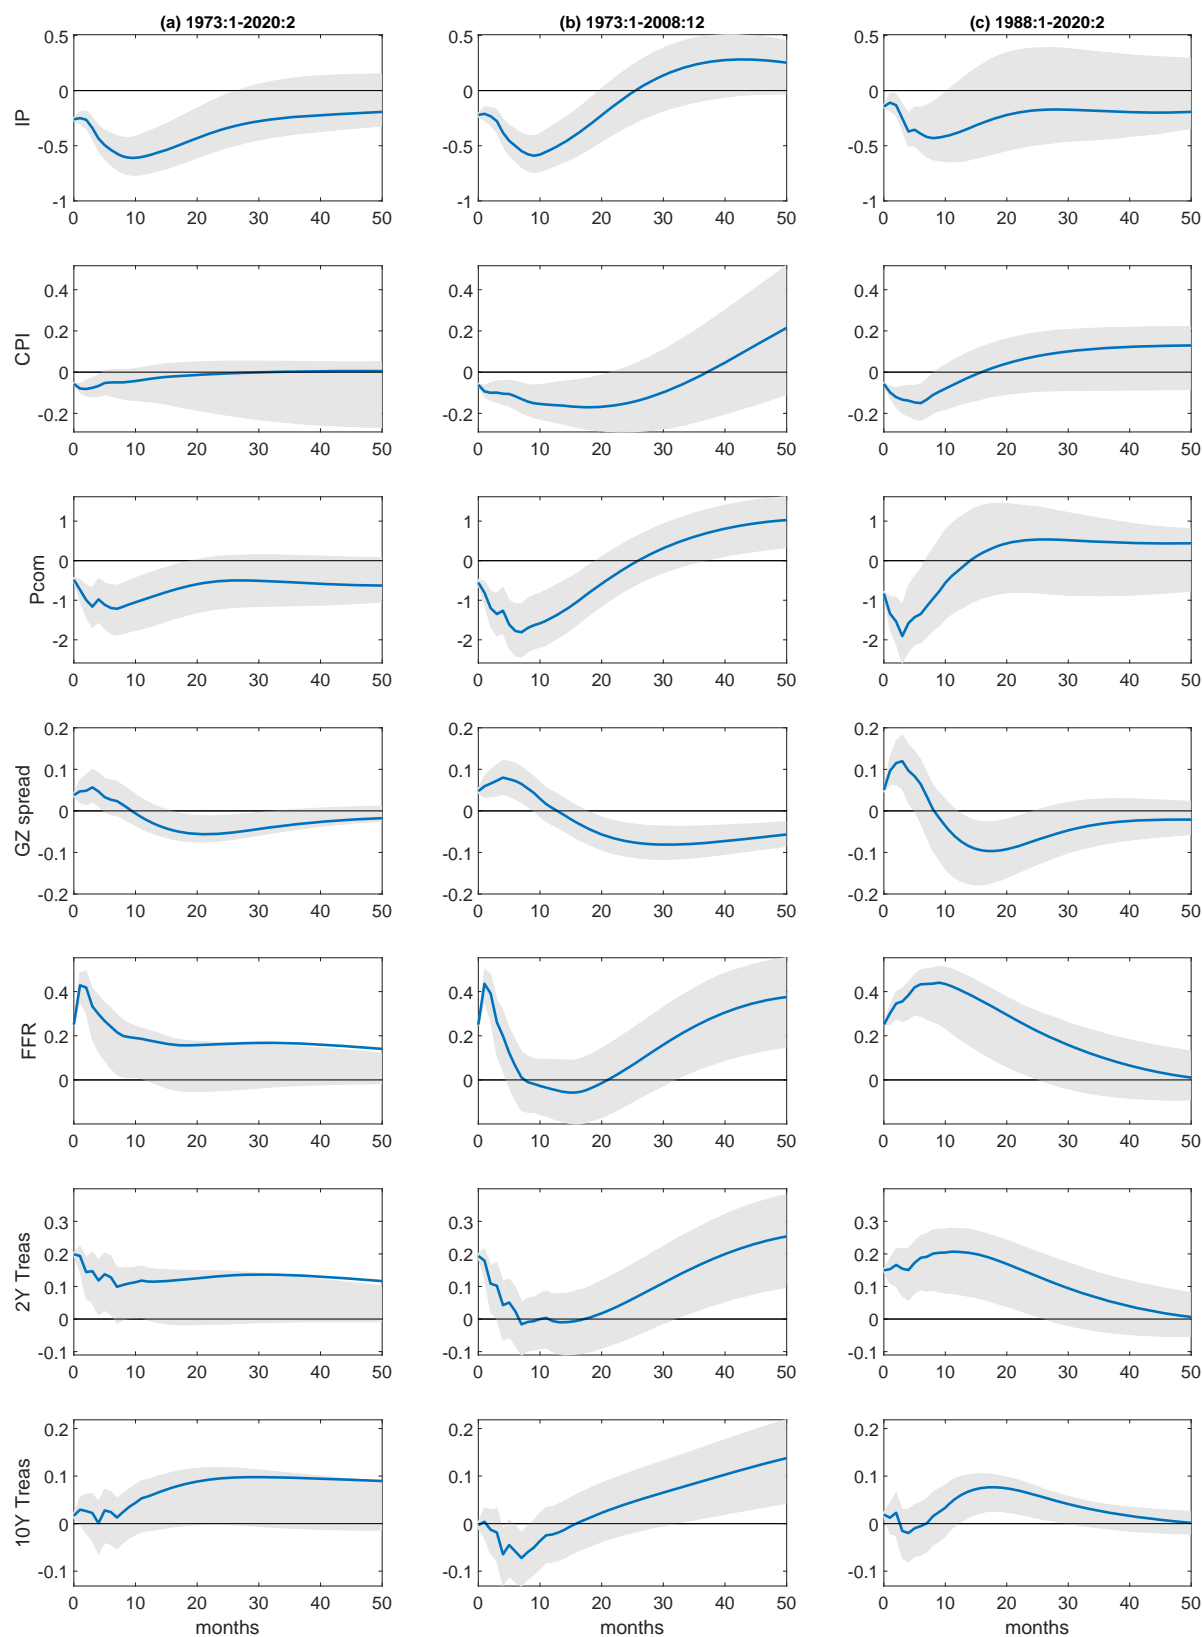
4.2 Economic Effects of Changes in the Federal Funds Rate

Figure 1 reports estimated impulse response functions for the federal funds rate, using the instrument z_t^{ff} . I first use the unadjusted instrument z_t^{ff} , without orthogonalization, for consistency with the standard high-frequency identification approach used by previous authors, such as Stock and Watson (2012, 2018), Gertler and Karadi (2015), and Ramey (2016); later I will orthogonalize the instrument as in Bauer and Swanson (2023b) and equation (9) above and show what difference it makes. Note that these impulse response functions are for a change in the federal funds rate *only* and do not include any additional forward guidance, so they are not strictly comparable to the estimates in Gertler and Karadi (2015), Ramey (2016), and Bauer and Swanson (2023b), who look at the impulse responses to changes in interest rate futures with several months to maturity.

Each column of Figure 1 presents results from estimating the VAR over a different sample: the first column reports results for the full sample, 1973:1–2020:2, the second column for a sample that ends before the U.S. zero lower bound period (which avoids any concerns related to the shadow federal funds rate), and the third column for a sample that begins in 1988:1, which is when my high-frequency instrument data begin. The solid blue line in each panel reports the estimated impulse response function, while the shaded gray region reports 90% standard-error bands, computed using the same methods as in Bauer and Swanson (2023b) and Gertler and Karadi (2015), with 10,000 bootstrap simulations. In each column, the impact effect of the federal funds rate shock on the (shadow) federal funds rate is normalized to 25bp.

Over the full sample, in the first column, a 25bp federal funds rate tightening is followed by an additional increase in the federal funds rate two months later to about 40bp and then a gradual return back toward baseline. The two-year Treasury yield rises about 20bp on impact and gradually returns back toward baseline, while the 10-year Treasury yield does not respond significantly. The Gilchrist-Zakrajsek (2012) credit spread rises a few basis points on impact, increases a bit further over the next few months, and then decreases before returning to baseline. Commodity prices fall about 0.5% in the impact month and decline further, by a little over 1%, before converging back toward baseline. Industrial production falls slightly on impact and declines further to a trough of about 0.6 percent after 9 months, and the CPI responds only slightly. Overall, these impulse responses are consistent with theory and with previous empirical estimates. The first-stage F -statistic for the instrument z_t^{ff} for this sample is 10.5, which is above

FIGURE 1: STRUCTURAL VAR IMPULSE RESPONSES TO A FEDERAL FUNDS RATE SHOCK, DIFFERENT SAMPLE PERIODS



Structural VAR impulse response functions to a 25bp federal funds rate shock, identified using the unadjusted high-frequency instrument z_t^{ff} around FOMC announcements, post-FOMC press conferences, FOMC minutes releases, Fed Chair speeches, and Vice Chair speeches, for three different sample periods: (a) full sample, 1973:1–2020:2, (b) pre-ZLB sample, 1973:1–2008:12; and (c) 1988:1–2020:2, since my high-frequency data begin in 1988. Shaded regions report bootstrapped 90% standard-error bands. See text for details.

the weak instruments threshold value of 10 suggested by Stock and Watson (2012).

The estimated impulse responses for the two shorter samples in the second and third columns of Figure 1 are similar to those in the first column, although the standard errors are larger due to the shorter samples. The first-stage F -statistics for the instrument z_t^{ff} over these samples are also smaller—9.6 and 9.8, respectively—right around the weak instruments threshold suggested by Stock and Watson (2012). Given the similarities across samples and smaller standard errors for the longest sample, I focus on the full 1973:1–2020:2 sample for the remainder of the paper.

I now follow Bauer and Swanson (2023b) and examine the importance of the Fed Response to News bias for these estimates. I orthogonalize the instrument z_t^{ff} with respect to the variables in Table 1, as in equation (9), and take the residuals $z_t^{ff\perp}$. The orthogonalized federal funds rate surprises $z_t^{ff\perp}$ should then better satisfy the instrument exogeneity conditions in equation (4).

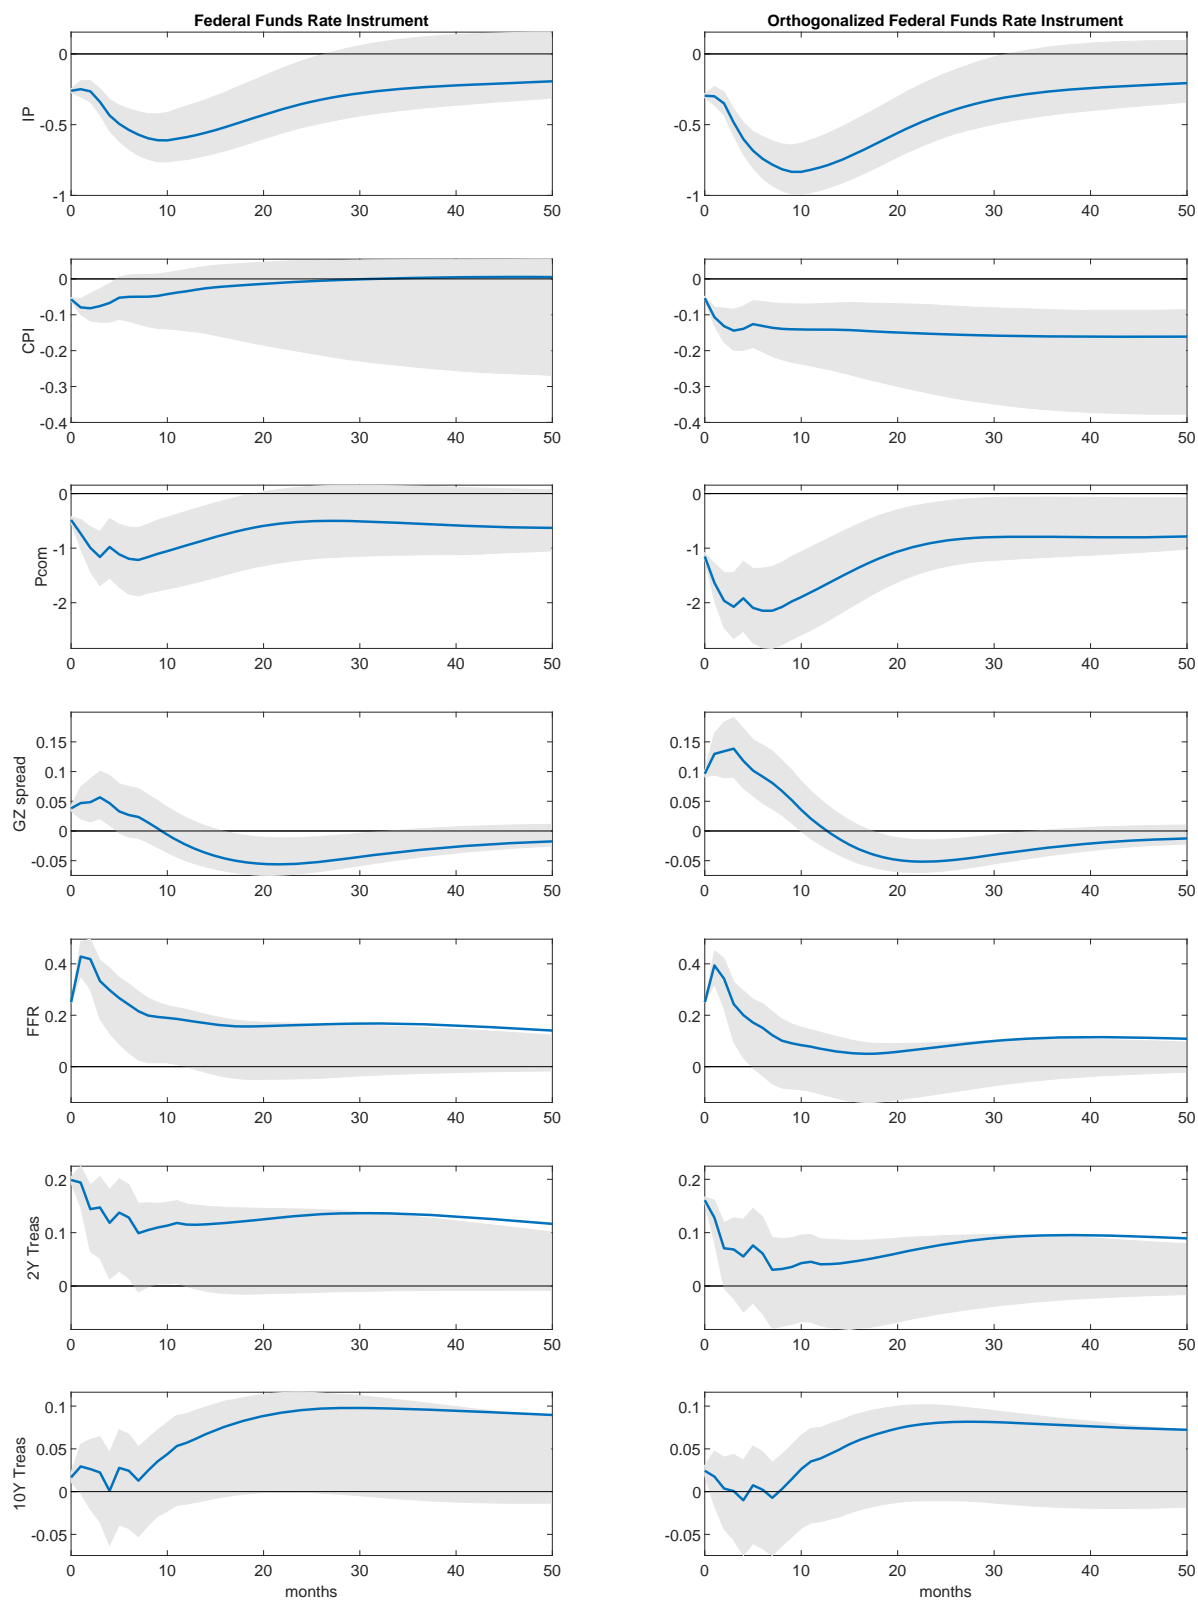
Figure 2 compares the estimated impulse response functions using the unadjusted federal funds rate instrument z_t^{ff} and the orthogonalized instrument $z_t^{ff\perp}$, estimated over the full sample 1973:1–2020:2. The first column of Figure 2 is thus the same as the first column of Figure 1, while the second column presents the new estimates using $z_t^{ff\perp}$.

There are several important points to take away from Figure 2. First, there is evidence of a Fed Response to News bias in the first column: relative to the second column, the impulse response functions in the first column exhibit attenuated responses of commodity prices, the credit spread, the CPI, and output in response to the shock. As in Bauer and Swanson (2023b), the results here suggest that positive macroeconomic and financial news (which tends to be followed by tighter monetary policy) is pushing all of these variables in the opposite direction to true monetary policy shocks, leading to the attenuation bias in the first column.

Second, the impulse response functions in the first column of Figure 2 are attenuated, but not by as much as for the *mps* measure in Bauer and Swanson (2023b). The latter measure of monetary policy surprises is based on interest rate futures out to a horizon of one year and thus includes a substantial role for forward guidance. The less attenuated responses in Figure 2 here thus suggest that changes in the federal funds rate are less contaminated by a Fed Response to News bias than are changes in forward guidance—i.e., financial markets seem to have better understood how the Fed would change the federal funds rate in response to incoming economic news, relative to forward guidance.

Third and finally, orthogonalizing the federal funds rate instrument reduces its power. The first-stage F -statistic for the orthogonalized instrument $z_t^{ff\perp}$ is low: only 5.3 vs. 10.5 for the un-

FIGURE 2: STRUCTURAL VAR IMPULSE RESPONSES TO A FEDERAL FUNDS RATE SHOCK, UNADJUSTED VS. ORTHOGONALIZED INSTRUMENTS



Structural VAR impulse response functions to a 25bp federal funds rate shock, identified in the left column using the unadjusted high-frequency instrument z_t^{ff} around FOMC announcements, post-FOMC press conferences, FOMC minutes releases, Fed Chair speeches, and Vice Chair speeches, and in the right column using the high-frequency changes z_t^{ff} orthogonalized with respect to economic and financial news available prior to the announcements each month. Sample: 1973:1–2020:2. Shaded regions report bootstrapped 90% standard-error bands. See text for details.

adjusted instrument z_t^{ff} . Stock and Watson (2012) suggest that values less than 10 indicate a potential weak instrument problem, which raises concerns about the estimates in the second column of Figure 2. In contrast to Bauer and Swanson (2023b), Chair speeches and other non-FOMC monetary policy announcements do not help here because those non-FOMC announcements do not affect the current federal funds rate, only expectations about its future path. However, the robustness of the results across the two columns of Figure 2 provides some reassurance that the impulse responses in the second column are not spurious results from a weak instrument.

4.3 Economic Effects of Changes in Forward Guidance

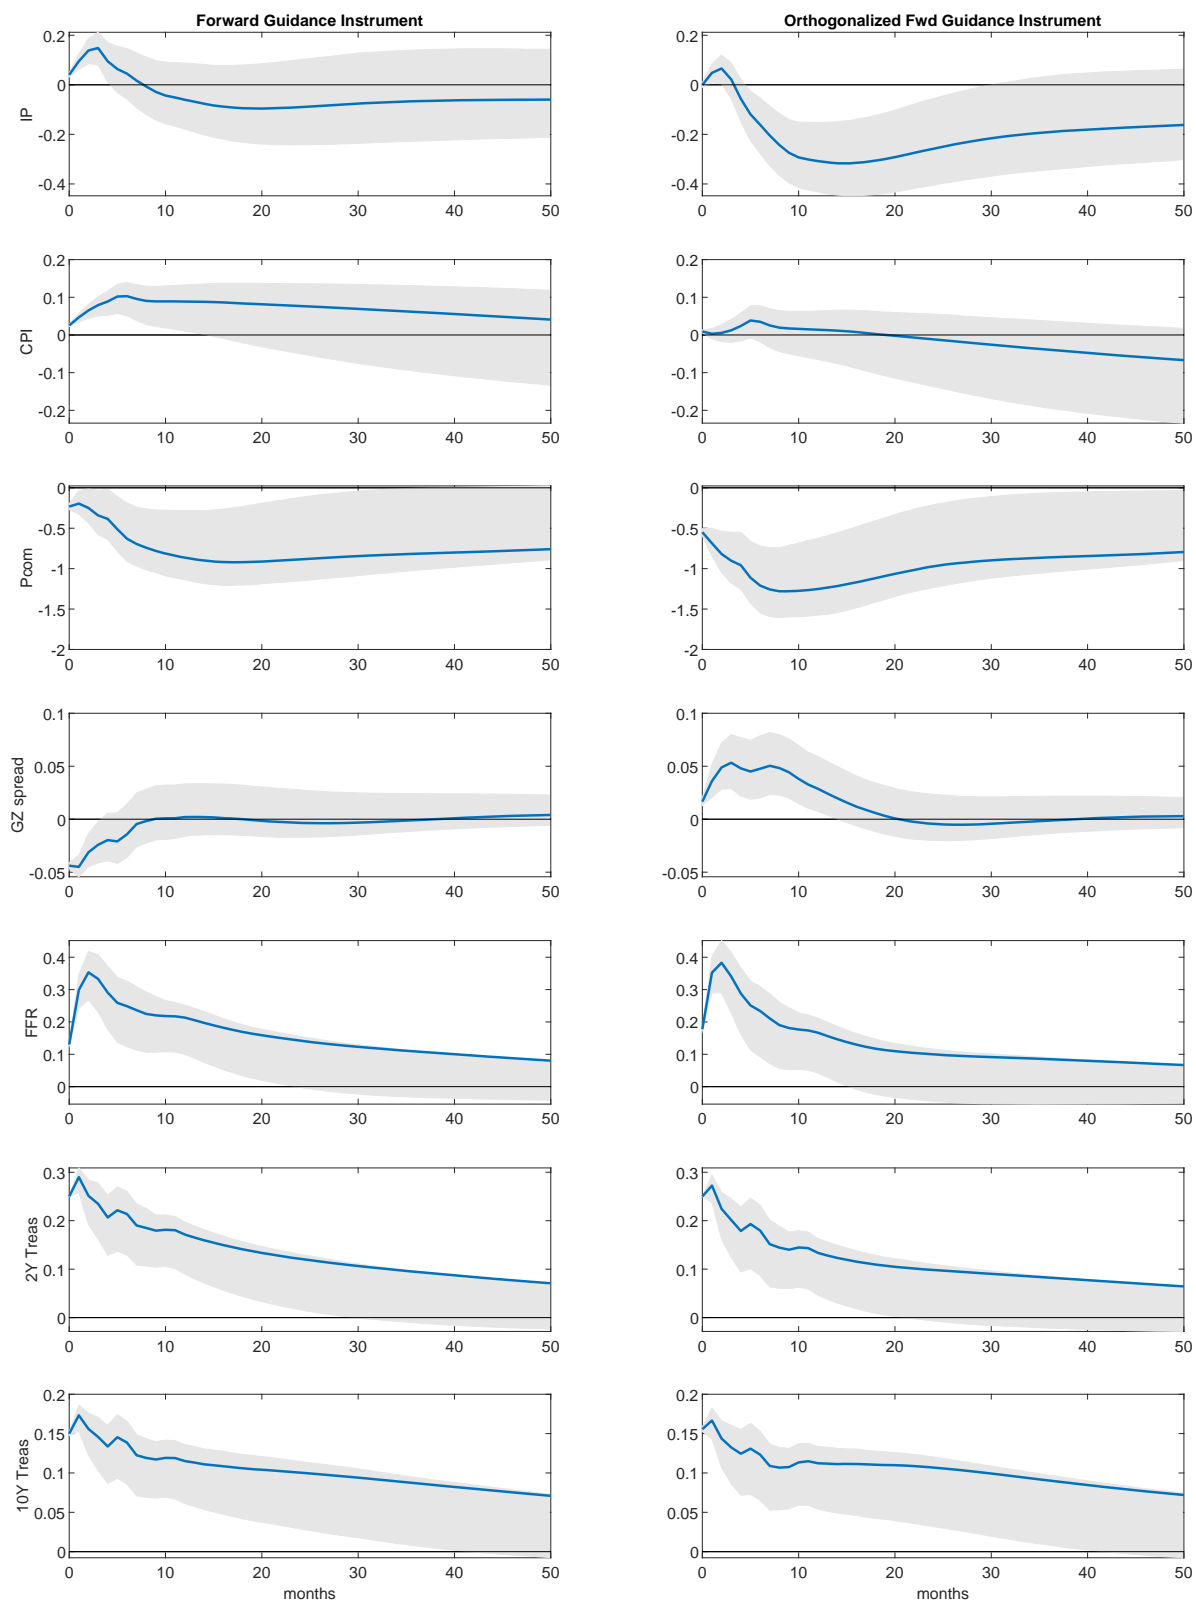
Figure 3 presents results for the effects of forward guidance using the high-frequency forward guidance instrument z_t^{fg} . Just as in Figure 2, the first column of Figure 3 reports results using the unadjusted instrument z_t^{fg} , while the second column reports results using the orthogonalized instrument $z_t^{fg\perp}$.

In Figure 3, the forward guidance instrument is normalized to have a 25bp effect on the two-year Treasury yield in both columns. The first point to note is that the first-stage F -statistics for these regressions are much higher than was the case for the federal funds rate regressions earlier: 39.7 for the first column of Figure 3 and 28.5 for the second column. These large F -statistics suggest that my instrument for forward guidance in the VAR is quite strong. The reason why z_t^{fg} and $z_t^{fg\perp}$ are so much stronger than z_t^{ff} and $z_t^{ff\perp}$ is that the Swanson-Jayawickrema (2023) data contain many non-FOMC announcements, such as Fed Chair speeches, that have very important implications for forward guidance.

The results in the first column of Figure 3 for the unadjusted forward guidance instrument are similar to those in Miranda-Agrippino and Ricco (2023, henceforth MAR). MAR use high-frequency VAR methods that are very similar to those used here, but MAR's high-frequency instrument is just the unadjusted change in forward guidance around FOMC announcements from Swanson (2021). One problem with this approach is that MAR's instrument is much weaker than the one here, with a first-stage F -statistic of just 3.3 (and just 1.2 after orthogonalizing as in equation (9)), which raises serious weak instrument concerns.

In response to a 25bp forward guidance tightening, the estimates in the first column of Figure 3 imply that the two-year Treasury yield converges gradually back to baseline after the shock; the federal funds rate increases about 15bp in the impact month and about 20bp more over the next two months, consistent with the forward guidance, and then converges gradually back to

FIGURE 3: STRUCTURAL VAR IMPULSE RESPONSES TO A FORWARD GUIDANCE SHOCK, UNADJUSTED VS. ORTHOGONALIZED INSTRUMENTS



Structural VAR impulse response functions to a 25bp forward guidance shock, identified in the left column using the unadjusted high-frequency instrument z_t^{fg} around FOMC announcements, post-FOMC press conferences, FOMC minutes releases, Fed Chair speeches, and Vice Chair speeches, and in the right column using the high-frequency changes z_t^{fg} orthogonalized with respect to economic and financial news available prior to the announcements each month. Sample: 1973:1–2020:2. Shaded regions report bootstrapped 90% standard-error bands. See text for details.

baseline; and the 10-year Treasury yield rises about 15bp on impact before slowly converging back to baseline. Commodity prices fall moderately, but the credit spread, CPI, and output all display puzzling behavior: the credit spread declines in response to the shock before returning to baseline, while the CPI and output increase. MAR conclude that their forward guidance instrument is probably contaminated by correlation with past and contemporaneous shocks, which is consistent with a Fed Response to News bias, and my results here also suggest that is likely.

In the second column of Figure 3, these puzzles are greatly reduced or eliminated when I use the orthogonalized forward guidance instrument $z_t^{fg\perp}$. The GZ credit spread increases in response to the tightening, commodity prices respond substantially more negatively, the CPI response is slightly negative after several months, and output responds much more negatively.

These estimates imply a very substantial Fed Response to News bias for forward guidance, much larger than the bias I estimated for the federal funds rate in Figure 2, and consistent with the large bias documented by Bauer and Swanson (2023b) for their one-dimensional monetary policy surprise measure mps , which included a substantial role for forward guidance. Intuitively, in the first column of Figure 2, positive economic and financial news is correlated with increases in the high-frequency instrument z_t^{fg} , and the positive effects of that news on the economy are offsetting the true effects of a forward guidance tightening.

Finally, I compare the estimated impulse response functions in the second columns of Figure 2 and Figure 3. Notice first that the effects of a 25bp federal funds rate change on the economy in Figure 2 are all larger than the effects of a 25bp change in the 2-year Treasury yield due to forward guidance in Figure 3. In fact, the differences between the effects of the federal funds rate and forward guidance are even greater than this in practice, because a one-standard-deviation change in the federal funds rate over my sample is 7.6bp—about one-third as large as the shock in Figure 2—while a one-standard-deviation change in the 2-year Treasury yield due to forward guidance is only 5.5bp—about one-fourth as large as the shock in Figure 3.¹⁸ As a result, I estimate that changes in the federal funds rate had substantially more powerful effects on the economy than changes in forward guidance.

¹⁸I measure the standard deviation of a federal funds rate shock here as the standard deviation of the monthly high-frequency instrument $z_t^{ff\perp}$ in the pre-ZLB period, 1988–2008, excluding months in which there was no change in the federal funds rate. I measure the standard deviation of a forward guidance shock as the standard deviation of the monthly high-frequency instrument $z_t^{fg\perp}$ from 1988–2019 (results for other samples are similar).

4.4 Economic Effects of Changes in LSAPs

Figure 4 presents the estimated effects of changes in the Federal Reserve’s large-scale asset purchases. As in Figures 2 and 3, the first column reports results for the unadjusted high-frequency LSAP instrument, z_t^{lsap} , while the second column reports results using the orthogonalized instrument $z_t^{lsap\perp}$. In both columns, the instrument is normalized to have a 25bp effect on the 10-year Treasury yield.

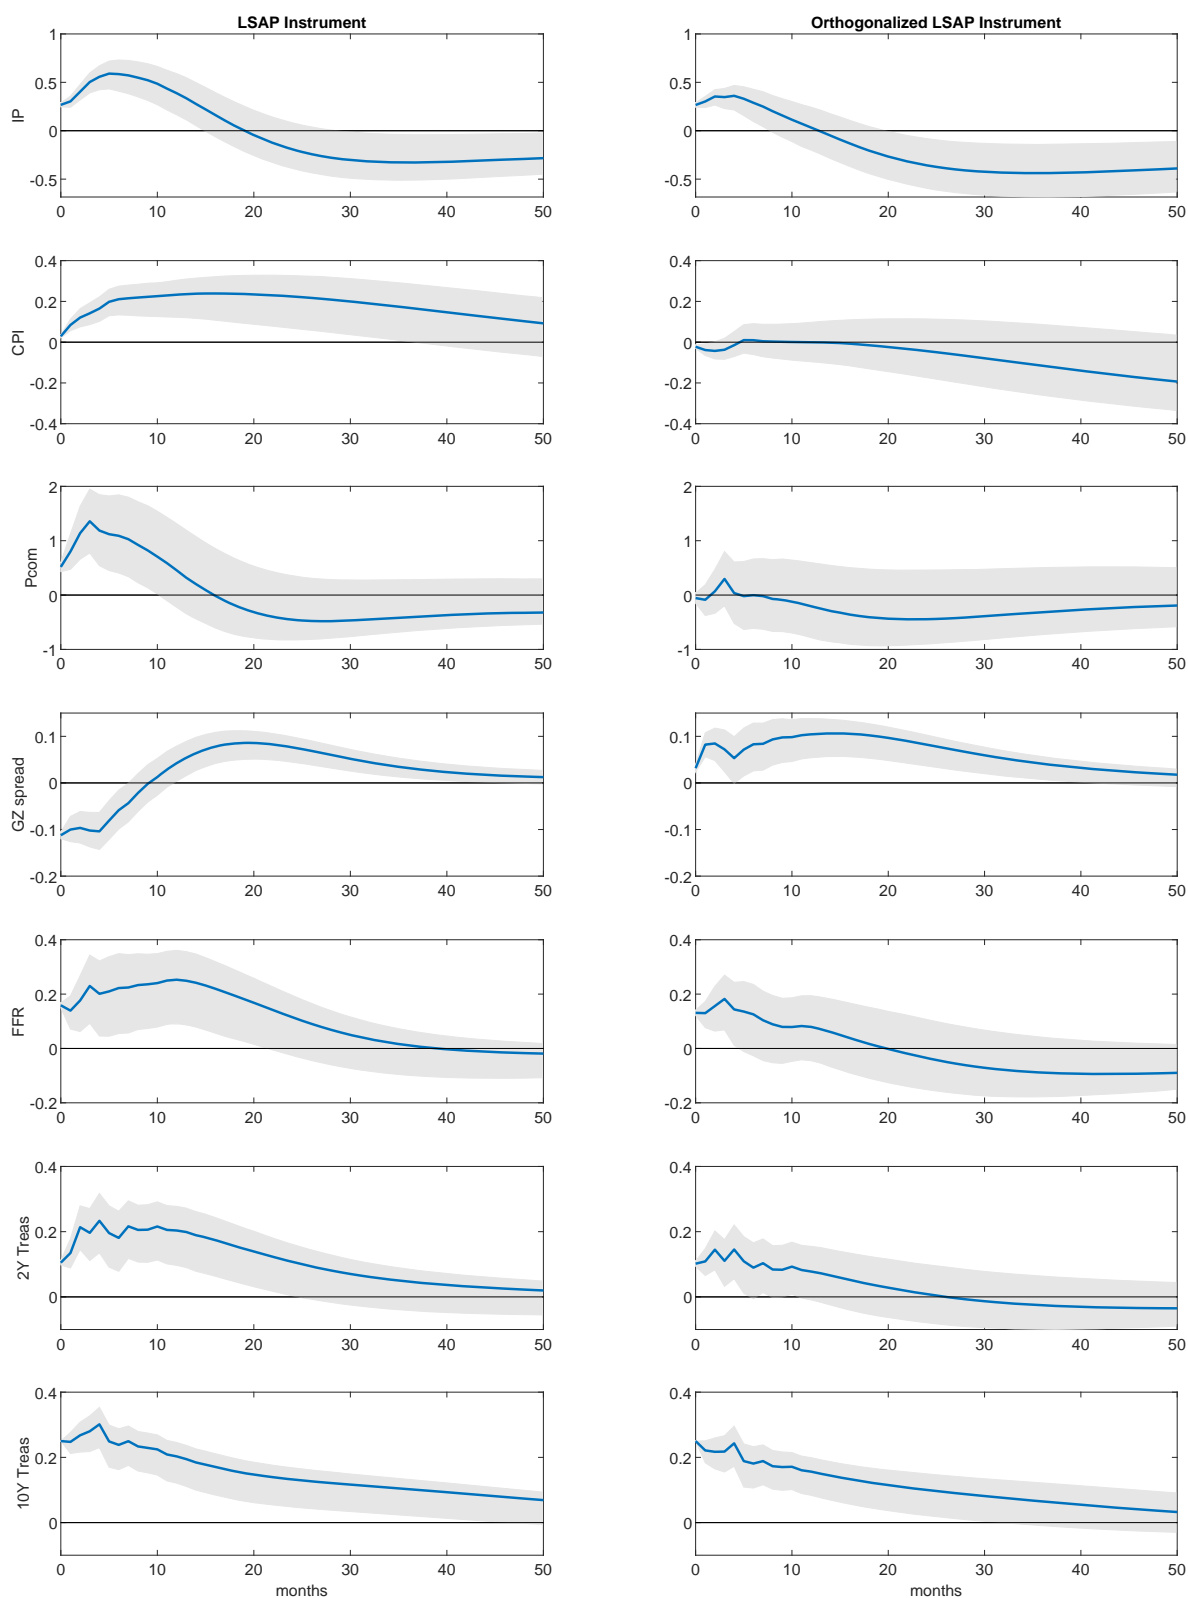
The first-stage F -statistics for these regressions are not as high as for forward guidance. For the unadjusted LSAP instrument z_t^{lsap} , the first-stage F -statistic is 9.2, while for the orthogonalized instrument $z_t^{lsap\perp}$, the statistic is 8.3, both of which are a bit below the weak instruments threshold of 10 suggested by Stock and Watson (2012). The results in Figure 4 are thus considerably less well estimated than those in Figure 3. One of the reasons the LSAP instrument is less powerful than that for forward guidance is that z_t^{lsap} and $z_t^{lsap\perp}$ are only nonzero beginning in January 2009, resulting in a smaller effective sample.

The results in the first column of Figure 4 are again similar to Miranda-Agrippino and Ricco (2023): an increase in the 10-year Treasury yield due to a Fed LSAP causes a puzzling, large, positive response of the CPI and shorter-lived puzzling responses of output, commodity prices, and the GZ credit spread. Miranda-Agrippino and Ricco conclude that, like their forward guidance instrument, their LSAP instrument is probably contaminated by correlation with other past and contemporaneous shocks, consistent with the Bauer-Swanson Fed Response to News story.

The orthogonalized LSAP instrument, used in the second column of Figure 4, largely eliminates these puzzles. In the second column, the credit spread responds positively and significantly to the LSAP shock. Commodity prices and the CPI respond slightly negatively, and output responds negatively after about a year. Intuitively, decreases in credit market stress are correlated with subsequent increases in the 10-year Treasury yield due to an LSAP (see Table 1), and the economic effects of the change in credit market conditions offset the true effects of the monetary policy shock in the first column of Figure 4.

Comparing the second column of Figure 4 to Figures 2 and 3, the economic effects of the federal funds rate in Figure 2 are the largest, followed by the forward guidance effects in Figure 3, and lastly the LSAP effects in Figure 4. In fact, the differences between these policies are even greater than is suggested by Figures 2–4 in practice, because a one-standard-deviation change in the 10-year Treasury yield due to an LSAP is 4.6bp over my sample, only about one-fifth as

FIGURE 4: STRUCTURAL VAR IMPULSE RESPONSES TO AN ASSET PURCHASES SHOCK, UNADJUSTED VS. ORTHOGONALIZED INSTRUMENTS



Structural VAR impulse response functions to a 25bp asset purchases shock, identified in the left column using the unadjusted high-frequency instrument z_t^{lsap} around FOMC announcements, post-FOMC press conferences, FOMC minutes releases, Fed Chair speeches, and Vice Chair speeches, and in the right column using the high-frequency changes z_t^{lsap} orthogonalized with respect to economic and financial news available prior to the announcements each month. Sample: 1973:1–2020:2. Shaded regions report bootstrapped 90% standard-error bands. See text for details.

large as the shock in Figure 4, while one-standard-deviation changes in the federal funds rate and forward guidance are about one-third and one-fourth the sizes of the shocks in Figures 2 and 3, respectively, as discussed above.¹⁹ I thus estimate that the effects of LSAPs on the U.S. economy are the weakest and least robust of the three policies.

Finally, there are a few caveats to bear in mind in Figure 4 and in the comparison of those results to Figures 2 and 3. First, the first-stage F -statistics in Figure 4 are below 10, albeit only slightly, suggesting that the instrument may be somewhat weak and the results should be treated with some caution. Second, the LSAP instrument z_t^{lsap} is only nonzero after 2009, while the federal funds rate instrument z_t^{ff} is essentially zero from 2009–15, implying that the samples over which the LSAP impulse responses and federal funds rate impulse responses are identified are different. Comparing the effectiveness of these two monetary policies thus requires assuming that the behavior of the economy itself is relatively stable across these two samples—an assumption that is consistent with the evidence in Wu and Xia (2016) and Debortoli, Galí, and Gambetti (2020), but should still be acknowledged.²⁰ Third and finally, the orthogonalized instrument $z_t^{lsap\perp}$ substantially reduced the puzzling impulse responses in the first column of Figure 4, but some puzzling behavior in those impulse responses still remains. For example, commodity prices do not respond negatively to the monetary policy tightening, and output responds positively for the first several months before turning negative later. In constructing the orthogonalized instrument $z_t^{lsap\perp}$, I have tried to control for economic and financial news that drives monetary policy, but it is possible that there is additional news that is affecting both monetary policy and the economy that is missing from Table 1. If so, there could be some additional Fed Response to News bias remaining even in the second column of Figure 4. Alternatively, perhaps speeches by the Fed Chair and Vice Chair, press conferences, and minutes releases contain some information about the state of the economy that is unknown to financial markets, so that there is a “Fed Information Effect” (Romer and Romer, 2000; Nakamura and Steinsson, 2018) in Figure 4 that explains the remaining puzzling behavior in the second column. Additional research investigating these issues would be interesting.

¹⁹The standard deviation of the LSAP shock is calculated as the standard deviation of the monthly high-frequency instrument $z_t^{lsap\perp}$ over the sample from 2009–15. Over longer samples the standard deviation of the LSAP shock is even smaller.

²⁰Wu and Xia (2016) and Debortoli, Galí, and Gambetti (2020) test for structural breaks or time-varying parameters in the U.S. economy around 2009 and find that the behavior of the economy is relatively stable.

5. Discussion and Conclusions

In this paper, I separately identify and estimate the effects of changes in the federal funds rate, forward guidance, and large-scale asset purchases (LSAPs) on the U.S. economy. I identify these three monetary policies using high-frequency interest rate changes around major Fed monetary policy announcements, including FOMC announcements, post-FOMC press conferences, FOMC meeting minutes releases, and speeches and testimony by the Federal Reserve Chair and Vice Chair. I use these high-frequency interest rate changes as an external instrumental variable in a structural monetary policy VAR to estimate the effects of changes in those policies on the U.S. economy.

As in Bauer and Swanson (2023a,b), I find that there is a substantial Fed Response to News bias in these VAR estimates. Intuitively, the Fed appears to have responded to incoming economic and financial news by more than financial markets expected, leading to an *ex post* correlation between that news and the high-frequency interest rate changes that I use as an instrumental variable for the VAR estimation. Failing to correct for this correlation leads to impulse response functions that are attenuated or even have a puzzling, opposite sign.

After correcting for the Fed Response to News bias, I find that the effects of all three monetary policies are consistent with the predictions of standard macroeconomic models, with tighter monetary policy leading to higher credit spreads, lower output, and lower prices, perhaps with a delay. I also find that changes in the federal funds rate had the largest effects on the U.S. economy, followed by changes in forward guidance and, lastly, changes in LSAPs. The estimated effects of changes in the federal funds rate are substantially stronger and more robust than those of the other two policies.

My finding that the federal funds rate has larger effects on the economy is somewhat surprising, given that Swanson (2021) estimates that unconventional monetary policies were about equally as effective as changes in the federal funds rate at moving financial market variables such as Treasury yields, corporate bond yields, and exchange rates. However, I estimate here that the federal funds rate has larger effects on commodity prices, and Swanson (2021) finds that the federal funds rate has larger effects on stock prices, so the larger effects on those variables could be part of the reason why inflation and output respond more strongly to changes in the federal funds rate.

My results suggest that central banks should continue to focus on short-term interest rates as their primary monetary policy tool going forward. Forward guidance and LSAPs certainly

have large effects on financial markets (Swanson, 2021) and do seem to have moderate effects on the economy, but the effects of the federal funds rate on the economy appear to be larger and more robust.

My results also suggest several avenues for future research. First, understanding the exact mechanisms by which the federal funds rate seems to have larger effects on the stock market, commodity prices, the CPI, and output are an important priority. Second, there may still be some Fed Response to News bias remaining in my impulse response function estimates; indeed, the presence of an omitted variable driving both Fed policy and the economy could explain why there still seem to be some minor price and output puzzles remaining in the second columns of Figures 3 and 4. Third, although Bauer and Swanson (2023a) presented substantial evidence that the Fed Information Effect for FOMC announcements is small, they did not consider Fed Chair speeches or other non-FOMC announcements in their analysis. It would be interesting to explore whether the case for a Fed Information Effect is stronger for Fed Chair speeches and other non-FOMC announcements than it is for FOMC announcements.

References

- ARUOBA, S. BORAĞAN, AND THOMAS DRECHSEL (2023). “Identifying Monetary Policy Shocks: A Natural Language Approach,” unpublished manuscript, University of Maryland.
- ARIAS, JONAS, JUAN RUBIO-RAMÍREZ, AND DANIEL WAGGONER (2021). “Inference in Bayesian Proxy-SVARs,” *Journal of Econometrics* 225, 88–106.
- BAUER, MICHAEL, AND MIKHAIL CHERNOV (2023). “Interest Rate Skewness and Biased Beliefs,” *Journal of Finance* forthcoming.
- BAUER, MICHAEL, AND ERIC SWANSON (2023a). “An Alternative Explanation for the ‘Fed Information Effect’,” *American Economic Review* 113, 664–700.
- BAUER, MICHAEL, AND ERIC SWANSON (2023b). “A Reassessment of Monetary Policy Surprises and High-Frequency Identification,” *NBER Macroeconomics Annual*, 37, 87–155.
- BERNANKE, BEN, AND KENNETH KUTTNER (2005). “What Explains the Stock Market’s Reaction to Federal Reserve Policy?” *Journal of Finance* 60, 1221–1257.
- CALDARA, DARIO, AND EDWARD HERBST (2019). “Monetary Policy, Real Activity, and Credit Spreads: Evidence from Bayesian Proxy SVARs,” *American Economic Journal: Macroeconomics* 11, 157–192.
- CIESLAK, ANNA (2018). “Short-Rate Expectations and Unexpected Returns in Treasury Bonds,” *Review of Financial Studies* 31, 3265–3306.
- COCHRANE, JOHN, AND MONIKA PIAZZESI (2002). “The Fed and Interest Rates—A High-Frequency Identification,” *American Economic Review* 92, 90–95.
- DEBORTOLI, DAVIDE, JORDI GALÍ, AND LUCA GAMBETTI (2020). “On the Empirical (Ir)Relevance of the Zero Lower Bound Constraint,” *NBER Macroeconomics Annual* 34, 141–170.
- EBERLY, JANICE, JAMES STOCK, AND JONATHAN WRIGHT (2020). “The Federal Reserve’s Current Framework for Monetary Policy: A Review and Assessment,” *International Journal of Central Banking* 16, 5–71.
- FAUST, JON, JOHN ROGERS, ERIC SWANSON, AND JONATHAN WRIGHT (2003). “Identifying the Effects of Monetary Policy Shocks on Exchange Rates Using High-Frequency Data,” *Journal of the European Economic Association* 1, 1031–1057.
- FAUST, JON, ERIC SWANSON, AND JONATHAN WRIGHT (2004). “Identifying VARs Based on High-Frequency Futures Data,” *Journal of Monetary Economics* 51, 1107–1131.
- GAGNON, JOSEPH, MATTHEW RASKIN, JULIE REMACHE, AND BRIAN SACK (2011). “The Financial Market Effects of the Federal Reserve’s Large-Scale Asset Purchases,” *International Journal of Central Banking* 7, 3–43.
- GERTLER, MARK, AND PETER KARADI (2015). “Monetary Policy Surprises, Credit Costs, and Economic Activity,” *American Economic Journal: Macroeconomics* 7, 44–76.
- GILCHRIST, SIMON, VIVIAN YUE, AND EGON ZAKRAJSEK (2019). “US Monetary Policy and International Bond Markets,” *NBER Working Paper* 26012.
- GILCHRIST, SIMON, AND EGON ZAKRAJSEK (2012). “Credit Spreads and Business Cycle Fluctuations,” *American Economic Review* 102, 1692–1720.
- GÜRKAYNAK, REFET, BRIAN SACK, AND ERIC SWANSON (2005a). “Do Actions Speak Louder than Words? The Response of Asset Prices to Monetary Policy Actions and Statements,” *International Journal of Central Banking* 1, 55–93.
- GÜRKAYNAK, REFET, BRIAN SACK, AND ERIC SWANSON (2005b). “The Sensitivity of Long-Term Interest Rates to Economic News: Evidence and Implications for Macroeconomic Models,” *American Economic Review* 95, 425–436.

- GÜRKAYNAK, REFET, BRIAN SACK, AND JONATHAN WRIGHT (2007). “The U.S. Treasury Yield Curve: 1961 to the Present,” *Journal of Monetary Economics* 54, 2291–2304.
- KARNAUKH, NINA, AND PETRA VOKATA (2022). “Growth Forecasts and News about Monetary Policy,” *Journal of Financial Economics* 146, 55–70.
- KIM, KYUNGMIN, THOMAS LAUBACH, AND MIN WEI (2020). “Macroeconomic Effects of Large-Scale Asset Purchases: New Evidence,” *Federal Reserve Board Finance and Economics Discussion Series* 2020–047.
- KUTTNER, KENNETH (2001). “Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market,” *Journal of Monetary Economics* 47, 523–544.
- MIRANDA-AGRIPPINO, SILVIA (2017). “Unsurprising Shocks: Information, Premia, and the Monetary Transmission,” *Staff Working Paper* 626, Bank of England.
- MIRANDA-AGRIPPINO, SILVIA, AND GIOVANNI RICCO (2021). “The Transmission of Monetary Policy Shocks,” *American Economic Journal: Macroeconomics* 13, 74–107.
- MIRANDA-AGRIPPINO, SILVIA, AND GIOVANNI RICCO (2023). “Identification with External Instruments in Structural VARs,” *Journal of Monetary Economics* 135, 1–19.
- NAKAMURA, EMI, AND JON STEINSSON (2018). “High-Frequency Identification of Monetary Non-Neutrality: The Information Effect,” *Quarterly Journal of Economics* 133, 1283–1330.
- PLAGBORG-MØLLER, MIKKEL, AND CHRISTIAN WOLF (2021). “Local Projections and VARs Estimate the Same Impulse Responses,” *Econometrica* 89, 955–980.
- RAMEY, VALERIE (2016). “Macroeconomic Shocks and Their Propagation,” *Handbook of Macroeconomics* 2A, 71–162.
- ROGERS, JOHN, CHIARA SCOTTI, AND JONATHAN WRIGHT (2018). “Unconventional Monetary Policy and International Risk Premia,” *Journal of Money, Credit, and Banking* 50, 1827–1850.
- ROMER, CHRISTINA, AND DAVID ROMER (2000). “Federal Reserve Information and the Behavior of Interest Rates,” *American Economic Review* 90, 429–457.
- ROMER, CHRISTINA, AND DAVID ROMER (2004). “A New Measure of Monetary Shocks: Derivation and Implications,” *American Economic Review* 94, 1055–1084.
- SASTRY, KARTHIK (2021). “Disagreement about Monetary Policy,” unpublished manuscript, MIT.
- STOCK, JAMES, AND MARK WATSON (2012). “Disentangling the Channels of the 2007–09 Recession,” *Brookings Papers on Economic Activity*, Spring, 81–135.
- STOCK, JAMES, AND MARK WATSON (2018). “Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments,” *Economic Journal* 128, 917–948.
- SWANSON, ERIC (2006). “Have Increases in Federal Reserve Transparency Improved Private-Sector Interest Forecasts?” *Journal of Money, Credit, and Banking* 38, 791–819.
- SWANSON, ERIC (2018). “The Federal Reserve Is Not Very Constrained by the Lower Bound on Nominal Interest Rates,” *Brookings Papers on Economic Activity* Fall, 555–572.
- SWANSON, ERIC (2021). “Measuring the Effects of Federal Reserve Forward Guidance and Asset Purchases on Financial Markets,” *Journal of Monetary Economics* 118, 32–53.
- SWANSON, ERIC (2023). “The Importance of Fed Chair Speeches as a Monetary Policy Tool,” *American Economic Association Papers & Proceedings* 113, 394–400.
- SWANSON, ERIC, AND VISHUDDHI JAYAWICKREMA (2023). “Speeches by the Fed Chair Are More Important than FOMC Announcements: An Improved High-Frequency Measure of U.S. Monetary Policy Shocks,” unpublished manuscript, University of California, Irvine.

- SWANSON, ERIC, AND JOHN WILLIAMS (2014). “Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates,” *American Economic Review* 104, 3154–3185.
- WOLF, CHRISTIAN (2020). “SVAR (Mis)identification and the Real Effects of Monetary Policy Shocks,” *American Economic Journal: Macroeconomics* 12, 1–32.
- WRIGHT, JONATHAN (2012). “What Does Monetary Policy Do to Long-Term Interest Rates at the Zero Lower Bound?” *Economic Journal* 122, F447–F466.
- WU, JING CYNTHIA, AND FAN DORA XIA (2016). “Measuring the Macroeconomic Impact of Monetary Policy at the Zero Lower Bound,” *Journal of Money, Credit, and Banking* 48, 253–291.