What Does Anticipated Monetary Policy Do?

Stefania D'Amico[a,1], Thomas B. King[a,1,*]

[a]*Federal Reserve Bank of Chicago, 230 S La Salle St, Chicago, 60604, IL, US*

Abstract

Applying sign restrictions to measures of expectations in a VAR, we quantify the economic effects of imperfectly and perfectly anticipated monetary policy innovations—the type of shocks induced by partially and fully credible forward guidance (FG). We find that fully credible FG one-year ahead has large near-term effects on prices and real activity, but, consistent with the forward guidance puzzle, these effects do not grow much larger as the FG horizon extends into the future. We also estimate that anticipated policy innovations are very noisy. At the one-year horizon, over 70 percent of the monetary-policy signal consists of noise, and the prevalence of noise increases at longer horizons. FG subject to this level of noise would be only partially credible and have significantly smaller macroeconomic effects.

JEL Codes: E32, E52, E58

Keywords: Monetary Policy, Forward Guidance, Sign Restrictions, Survey-augmented VARs.

---

[*]corresponding author

Email addresses: `stefania.damico@chi.frb.org` (Stefania D'Amico), `thomas.king@chi.frb.org` (Thomas B. King)

[1]For helpful comments, discussions, and suggestions we thank anonymous referees, as well as Jonas Arias, Boragan Aruoba, Bob Barsky, Jeff Campbell, Marco Del Negro, Charlie Evans, Jordi Gali, Matteo Iacoviello, Alejandro Justiniano, Thomas Laubach, Stefano Neri, Daniel Rees, Santiago Sordo Palacios, Eric Swanson, Harald Uhlig, Min Wei, Jonathan Wright, and conference and seminar participants at the 2015 WEAI meetings, the 2016 AEA meetings, the 2016 International Research Forum on Monetary Policy, the Banque de France, the Bank of Canada, the Federal Reserve Board, the Federal Reserve Bank of Chicago, and the University of Tor Vergata. The views expressed here do not reflect official positions of the Federal Reserve Bank of Chicago or the Federal Reserve System.

1. Introduction

As central banks have relied increasingly on forward guidance (FG) in recent years, it has become particularly important to understand how beliefs about future monetary policy affect the current state of the economy. A number of studies find that the impact of FG can be rapid and large in New Keynesian (NK) models, but those results may hinge upon the particular structural specification and calibration, and relatively minor variations in assumptions can cause the effects of policy expectations to differ markedly.[2] Thus, the quantitative impact of anticipated monetary policy remains an open question. Using an agnostic approach that relies on a minimal set of restrictions, we empirically quantify the economic effects of imperfectly and perfectly anticipated monetary-policy innovations—the type of shocks induced by partially and fully credible FG. We find that expectations about near-term policy innovations have large contemporaneous effects on prices and real activity. But policy-expectations formation is imbued with a lot of noise, which may complicate policymakers' efforts to shape beliefs and provide stimulus through FG.

Estimating the efficacy of FG requires us to tackle two questions. First, we ask what happens to the economy in reaction to an exogenous change in monetary-policy expectations. Second, we ask how this reaction depends on the information content of policy signals—that is, their *credibility*. Both questions are of interest to policymakers hoping to stimulate the economy through FG. While FG credibility has been the subject of a few theoretical treatments (e.g., Campbell et al., 2019), it has been largely disregarded in the empirical analysis of anticipated monetary-policy shocks. We address both questions by identifying anticipated monetary-policy innovations in a structural VAR and using the resulting impulse-responses to infer the amount of noise contained in typical monetary-policy signals. The results allow us to study the effects of partially and fully credible FG over various horizons, in a manner consistent with NK models.

First, we focus on identifying the anticipated monetary-policy shock—that is, an expected deviation from the historical policy rule.[3] To do so, we embed survey forecasts of the short-term interest rate and key macroeconomic variables in a structural VAR and impose the following restrictions: the expected interest rate must move in the opposite direction of expected inflation, expected GDP, and the current interest rate. This pattern of changes in expectations is *unique* to the anticipation of exogenous monetary-policy innovations in theoretical models. In particular, our sign restrictions, by excluding all cases in which GDP, inflation, and short rate forecasts move in the same direction, eliminate both the "Fed information effect" (Romer and Romer, 2000; Campbell et al., 2012; Nakamura and Steinsson, 2018) and the "Fed response to news" channel (Bauer and Swanson, 2023). Importantly, although the inclusion of survey forecasts in our VAR helps to discipline the model and aids identification, we do not take the surveys as straight reads of agents' beliefs. Rather, our sign restrictions apply to measures of expectations that are informed by both the surveys and the VAR itself.[4]

We find that, consistent with an FOMC that hews closely to its policy rule, anticipated deviations from that rule are small. Nevertheless, these anticipated policy shocks have significant and persistent effects on both prices and real activity. We estimate that a 10-basis-point decline in one-year expectations for the average short-term rate raises prices by 0.3 percent and output and hours worked by about 0.2 percent in the near term. Using these estimates we show that fully credible FG—an announcement about future policy that is fully believed and actually occurs—generates a significant near-term response well above the response caused by an identical policy-rate path that is *not* anticipated. This is because interest-rate expectations adjust immediately to fully credible FG but only sluggishly to conventional policy, and economic outcomes

[2]Campbell et al. (2020) and Lunsford (2020) discuss the recent use of FG in practice. Examples of NK treatments include Krugman (1998), Eggertsson and Woodford (2003), Laseen and Svensson (2011), Werning (2011), Milani and Treadwell, (2012), Del Negro and Schorfheide (2013), and Gomes et al. (2017). Levin et al. (2010) and Werning (2015) show how results in these models can be sensitive to assumptions.

[3]The title of our paper is a homage to Leeper et al. (1996), which was among the first studies to grapple with this identification problem with respect to conventional monetary policy.

[4]VARs containing only backward-looking information are likely to be misspecified and produce biased results (e.g., Sims, 1992; Caldara and Herbst, 2019). The forward-looking information in survey forecasts mitigates this misspecification. Indeed, survey forecasts either by themselves or in combination with VAR forecasts have been found to have excellent properties relative to multiple econometric predictions (e.g., Aiolfi et al., 2010; Ang et al., 2007; Faust and Wright, 2009; Tallman and Zaman, 2020).

appear to depend primarily on these expectations.

Our estimates uncover the responses to policy-expectations shocks without taking a stand on how expectations are formed. In particular, we do not assume that agents have perfect information about future monetary policy. Put differently, anticipated policy shocks are a mix of actual future policy shocks and noise.[5] Hence, we proceed to separate the dynamic effects of these two components using the equivalence result of Chahrour and Jurado (2018), which demonstrates how to recover estimates of noise and future "fundamental" shocks (actual future policy shocks, in our case) from observations on anticipated and unanticipated shocks like those identified in our VAR. By this measure, we find that anticipated policy innovations reflect a large amount of noise. At the one-year horizon, noise shocks have a variance almost 3 times as large as that of actual future policy shocks, implying that agents who learn rationally will initially adjust their expectations only 1/4 of the way toward a policy signal given one year ahead. FG that is partially credible can be modeled as a signal that, though perfectly informative, is believed by agents to contain noise *ex ante*. Thus, when we consider a policy announcement only as informative as the typical signal estimated in our data, the economic effects are significantly damped relative to fully credible FG.

As the horizon of anticipation extends, we find that the prevalence of noise grows, to the point that, beyond four years, signals about future monetary policy seem all but useless. Clearly, FG's credibility is crucial to its efficacy, and modelling FG as a noisy signal can have very important implications for lower-for-longer type of monetary policies (Bernanke et al., 2019). However, even when we consider fully credible FG, we do not find that the economic effects get increasingly larger as the horizon of the guidance extends. Hence, we conclude that the FG puzzle is indeed a puzzle (Del Negro et al., 2012; Carlstrom et al., 2012) and cannot be resolved by appealing to credibility alone.

Two key innovations of our paper are the empirical identification of noise in (exogenous) anticipated monetary-policy shocks and the estimation of the effects of partially and fully credible FG. These novelties allow us to show that ignoring the amount of noise in anticipated monetary-policy changes may produce misleading conclusions about the efficacy of FG. For instance, according to our estimates, a 10-basis-point FG shock could result in only a 3-basis-point change in policy expectations, as the latter do not fully adjust to noisy signals. This calls for caution in using observed changes in policy expectations at face value to draw conclusions about the efficacy of FG, as those two concepts do not map one-to-one.

1.1. Related Empirical Literature

Our paper is about anticipated monetary policy, which has been a subject of empirical investigation since at least Kuttner (2001). In an important paper along these lines, Gürkaynak et al. (2005) differentiated between unexpected changes to the current policy rate (target surprises) and changes in expectations for the future policy rate (path factor) around FOMC meetings. A few papers have exploited their path factor to study FG (e.g., Bundick and Smith, 2016; Swanson, 2021). However, using only the path factor is not sufficient to control for the Fed information effect, since an anticipated trajectory for the policy rate could reflect either endogenous or exogenous policy actions. We address this problem using simultaneous sign restrictions on expected interest rates, inflation, and real activity.

Following the working paper version of the present paper (D'Amico and King, 2015), several other empirical studies have addressed the complication introduced by the Fed information effect (e.g., Cieslak and Schrimpf, 2019; Jarocinski and Karadi, 2020; Miranda-Agrippino and Ricco, 2021; Bu et al., 2021; Andrade and Ferroni, 2021; Acosta, 2023). Our approach avoids a number of potential shortcomings in these studies. First, most papers have not distinguished between unanticipated (current) and anticipated (future) changes to the policy rate, as their measures of monetary policy surprises conflate the two. We want to separate these two types of shocks in order to (1) assess the marginal effects of FG over unanticipated monetary policy, and (2) derive the identification of fundamental and noise shocks. No previous study that

[5]Our framework for thinking about this issue is similar to the "noisy news" of Forni et al. (2017), although they do not analyze monetary policy. Their insight that future data perfectly reveal current structural shocks helps in overcoming the identification issues pointed out in Blanchard et al. (2013). As we show in Section 6, knowing that the dynamic response of the policy rate to noise shocks must be zero in the future is key to our identification of fundamental policy shocks and noise shocks.

has addressed the Fed information effect has examined either of these questions. Second, several papers in this literature use high-frequency market-based measures of policy expectations, which include risk premia. Yet, evidence, such as Bernanke and Kuttner (2005) and Hanson and Stein (2015), indicates that monetary policy actions cause changes in both the quantity and price of risk across securities markets. Papers that ignore this phenomenon may thus introduce significant measurement error.[6] Risk premia are likely to be particularly large at longer horizons, and since we analyze anticipated policy from one to 11 years ahead, it is especially important for us to avoid this contamination. This is one reason that we favor surveys over market measures. Third, while a few other studies have used sign restrictions to purge information effects, they are typically applied to measures of future interest rates and *either* real activity (e.g., stock prices) *or* inflation (e.g., rates on inflation swaps). Such identification schemes may not be sufficient to separate anticipated monetary policy shocks from other types of anticipated shocks, such as those to productivity or markups. In contrast, our sign restrictions on *both* expected future output and inflation are theoretically consistent only with exogenous changes in anticipated monetary policy.

The insights of Wright (2013), Robertson et al (2005), and Doh and Smith (2020) have guided the approach used in this paper to impose a loose form of rational expectations (RE) in the survey-augmented VAR. Wright (2013) finds that using survey forecasts to inform priors about the variable means in BVARs improves forecast accuracy for inflation and short rates, relative to multiple benchmarks. Robertson et al. (2005) use a relative entropy procedure to "tilt" draws from a VAR predictive distribution toward restrictions that have to be met exactly. This allows one to make theoretically coherent predictions when the restrictions are motivated by economic models, as in the case of RE in NK models. We adopt a similar technique, but, like Doh and Smith (2020), we do not impose our restrictions exactly. We use an informative prior over how closely the VAR impulse responses and the agents' beliefs match each other at the horizon of interest. This is equivalent to imposing a loose form of RE in the VAR system. Importantly, since this forecast-consistent prior is applied to our impulse-response functions, it informs both the reduced-form and structural VAR coefficients. As explained in Doh and Smith (2020), this is a great advantage with sign restrictions that identify VAR parameters up to a set: the forecast-consistency restrictions help distinguishing between models that are equally probable from the point of view of the data.

2. Policy Expectations in a New Keynesian Model

To fix ideas, we consider anticipated monetary policy in a standard New Keynesian (NK) model, as this provides a useful framework for decomposing anticipated policy innovations into fundamental and noise shocks that we will apply to our analysis. The model has some overlap with previous studies of FG and monetary-policy news (e.g., Eggertsson and Woodford, 2003; Werning, 2011; Milani and Treadwell, 2012; Campbell et al., 2012, 2019), but we emphasize the particular features that are of interest for our empirical tests. Indeed, all the impulse response functions (IRFs) estimated with our approach correspond to the exercises one would run using this NK model. In this sense, we estimate the effects of imperfectly and perfectly anticipated policy innovations in a manner consistent with partially and fully credible FG in NK models.

2.1. Model Description

Under standard NK assumptions, the equilibrium conditions can be written as follows:

$$\pi_t = \beta \mathrm{E}_t \pi_{t+1} + \kappa y_t \tag{1}$$

$$y_t = \mathrm{E}_t y_{t+1} - \frac{1}{\sigma}\left(i_t - \mathrm{E}_t \pi_{t+1} - r^*\right) \tag{2}$$

[6] Cieslak and Schrimpf (2019) is the one paper in this literature to take risk premia seriously, and they show that changes in risk premia do indeed contaminate measured policy surprises, especially at the zero lower bound (ZLB). Kaminska et al. (2020) argue that risk premia were relevant for interest-rate surprises even before the ZLB period, explaining about one quarter of the variation in the 3-month to the 2-year maturity and almost half of the variation at longer maturities.

where π_t is inflation, y_t is the output gap, E_t is the expectation conditioned on time$-t$ information, i_t is the nominal short-term interest rate, r^* is the natural rate of interest, $0 < \beta < 1$ is the rate of time preference, $\sigma > 0$ is the coefficient of relative risk aversion, and the Phillips Curve slope $\kappa > 0$ is a nonlinear combination of structural parameters. Following Gali (2015, c. 3), assume that the short-term interest rate is set by the central bank according to the rule:

$$i_t = \phi_y y_t + \phi_\pi \pi_t + \xi_t \tag{3}$$

$$\xi_t = \rho \xi_{t-1} + v_t \tag{4}$$

with v_t being a mean-zero random disturbance, and $\phi_\pi > 1$, $\phi_y \geq 0$, and $0 \leq \rho < 1$. As is typical, we assume parameters are such that i_t responds positively to v_t in equilibrium.

We depart from the standard treatment only by allowing agents to have some knowledge about the policy innovation v_t prior to period t, and potentially adjust their expectations of v_t every period before t. Let a_t^{t+h} be agents' anticipation of v_{t+h} as of period t, and let u_t denote the component of v_t that is unanticipated as of period $t-1$. That is,

$$v_t = a_{t-1}^t + u_t. \tag{5}$$

Rationality implies that anticipated policy follows the martingale process:

$$a_t^{t+h} = a_{t-1}^{t+h} + \eta_t^h, \tag{6}$$

where η_t^h is a serially independent shock. For the reminder of the paper we refer to η_t^h as the "anticipated policy shock" at horizon h. This shock is "news" in the sense of Chahrour and Jurado (2018). The standard NK model without news is a special case in which $a_t^{t+h} = 0$ for all t and h.[7] It is straightforward to show that anticipated policy shocks in this model move expected short-term interest rates in the opposite direction of expected future output and inflation. (See Appendix A.) This will be a key identifying feature of our empirical approach.

2.2. Anticipated Policy with Noise

As shown in equation (5), the policy innovation v_t equals the sum of anticipated policy a_{t-1}^t and unanticipated policy u_t. But anticipated shocks can reflect either fundamental information about future policy or noise. In Section 6, we will empirically distinguish these two components. To set the stage for that exercise, we introduce a signal-extraction problem into the model.

Suppose that agents receive an unbiased signal s_t^h about the policy innovation that will occur at time $t+h$. This signal is observed with noise ϵ_t^h so that:

$$s_t^h = v_{t+h} + \epsilon_t^h \tag{7}$$

where ϵ_t^h is iid and uncorrelated with fundamentals. In general, the variance of ϵ_t^h may differ across horizons and we will be able to estimate how it varies. But here, for expositional purposes, we simplify to the case where the noise variance, σ_ϵ^2, is horizon-independent.

We assume that agents know the stochastic process governing v_t and ϵ_t^h and update their beliefs in a Bayesian fashion:

$$E_t[v_{t+h}] = K\left(s_t^h - E_{t-1}[v_{t+h}]\right) + E_{t-1}[v_{t+h}] \tag{8}$$

where K is the Kalman gain

$$K = \frac{\sigma_v^2}{\sigma_v^2 + \sigma_\epsilon^2}. \tag{9}$$

Under the assumption of homoskedastic ϵ_t^h, K does not vary with the horizon of the expectation, so learning is constant-gain. The revision in expectation for v_{t+h} resulting from the signal received in t maps into the

[7] This way of specifying shocks to monetary-policy expectations nests many previous models of FG, including Lasseen and Svensson (2011), Campbell et al. (2012), Milani and Treadwell (2012), and Del Negro and Schorfheide (2013).

anticipated policy shock η_t^h introduced earlier:

$$\eta_t^h = E_t[v_{t+h}] - E_{t-1}[v_{t+h}] = a_t^{t+h} - a_{t-1}^{t+h}. \tag{10}$$

Thus, changes in policy expectations can result from either noise shocks (ϵ_t^h) or correct information about future policy shocks (v_{t+h}). (We refer to the latter as the "fundamental" policy shocks, following the terminology of Chahrour and Jurado, 2018.)

As an example, consider signals that begin to arrive four quarters before the policy innovation might occur. We trace out the average economic effects in two polar cases: either the initial signal reflects only the fundamental shock ($s_0^4 = v_4$) or it turns out to be entirely noise ($s_0^4 = \epsilon_0^4$). In both cases, we assume that agents know that the policy rate will remain unchanged through quarter 3. (This is achieved by offsetting the systematic response of policy in quarters 0 through 3, as described in Appendix A.)

Beginning in time $t = 0$, the top-left panel of Figure 1 shows how expectations about the policy innovation in the fourth quarter, v_4, evolve under the two scenarios. We have calibrated $K = 0.27$, based on our empirical results at the one-year horizon, shown later. (This implies that σ_ϵ^2 is about three times σ_v^2, consistent with the one-year-ahead policy signal being quite noisy.) We normalize to an initial signal that generates an anticipated policy change of $\eta_0^4 = -10$ basis points. That is, $s_0^4 = -0.10/0.27 = -37$ basis points. Since agents cannot initially distinguish between fundamental and noise shocks, the initial change in policy expectations is the same regardless of the information contained in the signal. Subsequently, if the signal turns out to be a fundamental shock (red line), expectations for v_t converge toward -0.37 over time; and, if it is noise (gray line), expectations approach zero over time, as agents learn that there will be no policy shock. In both cases, in $t = 4$ the truth is revealed and there is a discontinuous jump. From the perspective of the agents, this jump is the unanticipated shock u_4.

Since the macroeconomic variables in our model are just linear functions of current-period policy expectations, we can derive the corresponding impulse response functions (IRFs) under anticipated policy innovations that turn out to be either fundamental or noise shocks. The second, third, and fourth panels in Figure 1 show the evolution of the short-term rate, inflation, and output gap under each scenario. For comparison, the dashed line shows the perfect-foresight case, characterized by the same-sized signal but $\sigma_\epsilon = 0$ and $K = 1$. With the relatively uninformative signal that occurs in the noisy case, the initial responses of inflation and output are just 27% of what they would be under perfect foresight. Signals that provide fundamental information but are not fully believed result in IRFs that have qualitatively different *shapes*, not just lower magnitudes, than IRFs under perfect-foresight. The red lines for inflation and output slope downward much less dramatically than the dashed black lines.

These results have important implications for the analysis of FG. If we view the perfect-foresight case as what happens when a central bank communicates fully credible FG—so that agents impound the central bank's signal one-for-one into their beliefs—then it is clear that FG can be quite powerful in NK models. However, the introduction of noise can significantly dampen FG's effects, even when the guidance turns out to be truthful. If the FG signal contains noise, as in the red lines of Figure 1, agents will not fully take it on board, and its macroeconomic effects will be smaller. We call this the case of *partially credible* FG.

Finally, the dotted black lines show the average of the red and gray lines, weighted by the variance of the noise and fundamental shocks, respectively. These are the IRFs that one would estimate in response to an anticipated policy shock $\eta_0^4 = -0.10$, using a large amount of data generated by this model. They are thus comparable to the raw IRFs resulting from our baseline VAR in the next section. In Section 6, we will use (9) to recover the noise and fundamental IRFs from the VAR estimates. Note that, because the model is linear, the dotted line, reflecting the average effects of changes in expectations, is just a rescaling of the dashed line. We will use this fact when we consider the empirical effects of *fully credible* FG.

3. Empirical Specification

3.1. Reduced Form and Expectation Measurement

We begin by identifying anticipated policy innovations in a VAR that jointly models economic variables and survey forecasts, treating the latter as imperfect measures of agents' true expectations. Survey forecasts have been previously used in VARs to purge fiscal variables of their predicted innovations (e.g., Ramey, 2011; Auerbach and Gorodnichenko, 2012), but their inclusion in a structural VAR to identify news about monetary policy is novel. When survey forecasts are introduced into a VAR as direct measures of beliefs, an inconsistency arises between the survey- and VAR-implied measures, and it is not clear which to treat as the "true" expectation. We reduce the measurement errors in expectations by taking linear combinations of the VAR and survey forecasts. The forecast combinations are our proxy of agents' true beliefs.

Specifically, we estimate a series of models characterized by the following reduced form:

$$\begin{pmatrix} \mathrm{E}_t^S\left[\mathbf{x}_{t+h}\right] \\ \mathbf{x}_t \end{pmatrix} = \boldsymbol{\theta}_0 + \boldsymbol{\Theta}(L) \begin{pmatrix} \mathrm{E}_t^S\left[\mathbf{x}_{t+h}\right] \\ \mathbf{x}_t \end{pmatrix} + \begin{pmatrix} \mathbf{e}_{1,t} \\ \mathbf{e}_{2,t} \end{pmatrix} \tag{11}$$

where L is the lag operator, $(\mathbf{e}_{1,t}\ \mathbf{e}_{2,t})$ are iid vectors of mean-zero disturbances with covariance matrix $\boldsymbol{\Sigma}$, and $\boldsymbol{\theta}_0$ and $\boldsymbol{\Theta}(L)$ are matrices of reduced-form parameters. $\mathrm{E}_t^S\left[x_{t+h}\right]$ denotes the survey forecast of a variable x_{t+h}, and the superscript S distinguishes it from the agents' true expectations, which we denote $\mathrm{E}_t\left[x_{t+h}\right]$. We proxy this expectation with a linear combination of survey and VAR forecasts, $\mathrm{E}_t^{VAR}\left[x_{t+h}\right]$:

$$\mathrm{E}_t\left[x_{t+h}\right] = \alpha \mathrm{E}_t^S\left[x_{t+h}\right] + (1-\alpha)\mathrm{E}_t^{VAR}\left[x_{t+h}\right] \tag{12}$$

In our baseline specification, we use $\alpha = 0.5$, and we show robustness to a range of possible values in Appendix E. Agents' expectations under this approach are still not generally the same as the VAR forecasts, so we have to consider the economic implications of the gap between them, an issue we address in Section 3.3 below.

We estimate separate models for each survey forecast horizon of 1, 2, 3, 4, 5, and 11 years in our data. In each case, we estimate $\boldsymbol{\Theta}(L)$ and $\boldsymbol{\Sigma}$ by Bayesian methods, using an uninformative normal-Wishart prior.

3.2. Shock Identification and Implementation

Let $\boldsymbol{\eta}_t$ and $\mathbf{u}_t$ be vectors of anticipated and unanticipated shocks, and let $\boldsymbol{\Gamma}$ denote the matrix of multipliers on the structural shocks, such that

$$\begin{pmatrix} \mathbf{e}_{1,t} \\ \mathbf{e}_{2,t} \end{pmatrix} = \boldsymbol{\Gamma} \begin{pmatrix} \boldsymbol{\eta}_t \\ \mathbf{u}_t \end{pmatrix} \tag{13}$$

and $\boldsymbol{\Sigma} = \boldsymbol{\Gamma}'\boldsymbol{\Gamma}$. Although $\boldsymbol{\eta}_t$ and $\mathbf{u}_t$ may contain arbitrarily many elements, we single out two, corresponding to the theoretical treatment of Section 2: an anticipated monetary-policy shock (η_t^h) and an unanticipated monetary-policy shock (u_t). We denote by Γ_u^x and Γ_η^x the elements of $\boldsymbol{\Gamma}$ corresponding to the impact of shock u_t and η_t^h on variable x.

In order to identify the necessary elements of $\boldsymbol{\Gamma}$, we impose a set of partial-identification restrictions. Specifically, as shown in Table 1, our sign restrictions for the contemporaneous impacts of the anticipated monetary-policy shock η_t^h enforce the following condition: the time-t impact on the expected average TBill rate over periods t to $t+h$ must be in the opposite direction of the time-t impact on the expected levels of GDP and CPI at time-$t+h$. This assumption is consistent with the predictions of NK models, like that in Section 2, for an anticipated policy shock and *only* for this type of shock.[8] To ensure that we are not conflating anticipated with unanticipated policy, we also impose that the current TBill rate moves in the opposite direction of η_t^h, implying that the short rate moves in the opposite direction of its expectation. This is the behavior predicted by the NK model in Section 2.

Since expectations should also respond within a quarter to unanticipated policy shocks, in our baseline specification for u_t, we impose on the expected variables the same set of restrictions used for η_t^h. This allows the expectation channel to be present in both types of shock, maximizing comparability. The only difference is that the current short rate moves in the same direction of u_t. (See Table 1.) These restrictions on the unanticipated policy shock, like the restrictions used to identify our anticipated shocks, are consistent with

[8] For example, our restrictions exclude aggregate-demand shocks at time-t since the Fed *raises* expected future rates in response to exogenous increases in output and inflation.

the predictions of standard NK models, and there are no other shocks within those models that result in the same patterns.

To compute IRFs, we draw jointly from the posterior distribution of the VAR parameters and the set of admissible $\mathbf{\Gamma}$'s, and we simulate the dynamic effects of a one-standard-deviation shock under each draw. We discard all draws that violate the sign restrictions and keep drawing until $50,000$ draws are accepted. We report the median and 16%-to-84% range across all draws. Our focus on the quantiles of the IRF distribution, accounting for both statistical uncertainty about the reduced-form parameters and "model uncertainty" over the possible structural rotations consistent with our restrictions for any given set of reduced-form estimates, effectively treats the problem as a Bayesian one, as recommended by Baumeister and Hamilton (2015). Specifically, we sample factorizations of $\mathbf{\Sigma}$ that are consistent with a uniform prior distribution over the structural parameters $\mathbf{\Gamma}^{-1}$ using algorithms developed by Rubio-Ramirez, Waggoner, and Zha (2010) and Arias, Rubio-Ramirez, and Waggoner (2018).[9]

3.3. Forecast Consistency

Under full-information RE, subjective and physical probabilities are identical, implying that our measures of expectations should always be equal to our VAR forecasts. However, given the evidence reported in Coibion and Gorodnichenko (2012, 2015), it is difficult to defend the assumption of full RE in surveys. Our baseline model thus does not impose strict consistency between $E_t\left[x_{t+h}\right]$ and $E_t^{VAR}\left[x_{t+h}\right]$. However, although we want to allow these objects to differ, we adopt an informative prior that they do not differ very much. In other words, we think that deviations from rationality are generally not too large and that anticipated changes to future monetary policy are close, on average, to realized changes. Our approach is similar to the shrinkage toward surveys adopted by Doh and Smith (2020).[10]

Specifically, the VAR responses of GDP, CPI, and the short rate in the h periods subsequent to the shock are on average in the proximity (controlled by the parameter δ) of the change in these variables anticipated at time-t for the horizon $t+h$. To implement this idea, we adopt the prior distributions:

$$\frac{\Gamma_\eta^{E_t[x_{t+h}]}}{E_t^{VAR}\left[\frac{\partial x_{t+h}}{\partial \eta_t}\right]} \sim N[0,\delta] \tag{14}$$

where h is the horizon of the survey forecast used in the VAR.[11] We adopt a similar set of priors for the unanticipated policy shock. We implement these "loose rationality" priors by importance sampling—that is, by reweighting the posterior VAR parameter draws of $\mathbf{\Theta}$ and $\mathbf{\Gamma}$ by the joint prior distribution in (14) and drawing 50,000 times again. In our baseline specification, we take $\delta = 0.5$. Thus, *a priori*, there is about a 32% probability that agents' expectations differ from the VAR forecasts by more than 50% for each of the three variables. We experiment with a range of values for δ in Appendix E.

The parameters α and δ address related but distinct specification issues. While α concerns a measurement issue (how to reconcile the two measures of expectations in our model), δ concerns an economic issue (to what extent are expectations rational). Setting either $\alpha = 0$ or $\delta = 0$ enforces full RE by ensuring that expectations are always equal to the VAR forecasts. But $\delta = 0$ equates the VAR forecast to the survey forecast, while $\alpha = 0$ places no restriction on their relationship, implying that even in these extreme cases survey information plays different roles. In general, one must decide how to weight the VAR expectations against the survey expectations and how strongly to shrink the model toward rationality.

3.4. Data

As our baseline specification we include in $\mathbf{x}_t$ the 3-month Treasury Bill (TBill) rate, log GDP, log CPI, log hours worked, the 2-year nominal Treasury yield, and stock market returns (measured by the Wilshire

---

[9] We thank Jonas Arias for providing us with Matlab code that greatly assisted in this effort. Our results are also robust to uniform priors over the IRFs or the Haar measure.

[10] The deviation between subjective and VAR-based forecasts can also indicate weak identification. But the shrinkage procedure, as explained in Doh and Smith (2020), also helps to better identify the model.

[11] Note that, consistent with the reporting conventions of the surveys, our restrictions are on the future *levels* of GDP and CPI at the end of the forecast period but on *averages* of the short rate over the forecast period.

5000 Index), all at a quarterly frequency over the period 1983 Q1-2020 Q1. $E_t^S[\mathbf{x}_{t+h}]$ is measured by the Blue Chip Survey (BCS), which begins in 1983 Q1 for our series of interest. Each survey reports the respondents' average forecasts of real GDP, CPI, and the 3-month TBill rate, which we use as a proxy for the monetary-policy instrument. The GDP and CPI forecasts are in levels (which we transform into logs), while data on the TBill rate are reported as averages over the period $t+1$ to $t+h$. The constructed series of expected GDP, CPI, and the 3-month TBill rate are at horizons ranging from one to 11 years. Due to idiosyncrasies in the conventions and timing of their reporting, the survey data at horizons beyond one year require interpolation to obtain quarterly time series of constant-horizon forecasts.[12]

4. Baseline Results

We begin by describing the results for the anticipated monetary-policy shocks in the baseline VAR specification that includes survey data with a one-year horizon, and two lags (as indicated by the AIC). As noted, our baseline model uses $\alpha = 0.5$ (indicating that survey and VAR forecasts are equally weighted in measuring expectations), and $\delta = 0.5$ (a loose prior on rationality).

4.1. The Time Series of Anticipated Policy Shocks

Figure 2 plots the time series of η_t^4, using the median across draws, to offer some evidence that the one-year-ahead anticipated policy shocks we have identified do indeed correspond to periods in which the expected policy-rate path had reasons to shift.[13] The largest positive anticipated policy shocks occur in 2008 Q4, just as the short-term rate hits the ZLB. This is precisely what we would expect if forward-looking agents incorporate the ZLB constraint in their expectations. From the perspective of a linear model, if agents understand that the ZLB constrains the amount of monetary policy accommodation, then they would expect tightening relative to the usual policy rule. This shock is soon followed by two large negative shocks, as in March 2009 the FOMC introduced the guidance that rates would stay exceptionally low for an "extended period." Then, from 2011 to 2014 there is a sequence of negative anticipated shocks as, in 2011 Q3, the Fed adopted calendar-based FG specifying that economic conditions were "likely to warrant exceptionally low levels for the federal funds rate (FFR) at least through mid-2013;" in December 2012, the it introduced threshold-based FG (i.e., the ZLB would bind at least as long as certain thresholds were met for unemployment and inflation); and, in 2014 Q2, it switched to qualitative guidance about likely deviations from the conventional policy rule.[14]

In our sample, there are also sizable anticipated policy shocks prior to the ZLB period that are recognizable. For example, we find large easing shocks in 2004 Q1, when the FOMC indicated that "it can be patient in removing policy accommodation;" in 2002 Q3 when it tilted the balance of risks toward the downside mentioning that "the risks are weighted mainly toward conditions that may generate economic weakness;" and in 2000 Q3, when the FOMC left the FFR unchanged and stated that "the expansion of aggregate demand is moderating" rather than "may be" moderating, solidifying expectations for the end of the tightening cycle. A notable tightening shock occurs in 2001 Q4. As indicated in the December 2001 Bluebook, following the November FOMC meeting, investors revised the expected FFR one- and two-year ahead by about 100 basis points.

From the mid-1980s to the mid-1990s, explicit FOMC communication was more sparse, and thus specific events that might relate to our estimated shocks are harder to find. An advantage of our approach is that it does not require that shocks arise only from explicit FG. Any information that has caused agents to exogenously change their beliefs about future monetary policy should be identified as anticipated policy

[12] We thank Santiago Sordo-Palacios for assistance with this interpolation.

[13] The dates in the figure reflect the dates of the identified shock and not necessarily those of the corresponding FOMC statements, as we have taken care to account for the timing of the surveys relative to FOMC meetings. In particular, the BCS data are gathered in the first week of each month, while the last FOMC meeting of each quarter usually takes place in the second or third week of the month, and therefore would not be reflected in survey responses until the following quarter.

[14] That is, "The Committee currently anticipates that, even after employment and inflation are near mandate-consistent levels, economic conditions may, for some time, warrant keeping the target federal funds rate below levels the Committee views as normal in the longer run."

shocks in the data. For example, the notable anticipated easing shock in 1995 Q2 corresponds with a firming in expectations that "the period of policy tightening might be drawing to a close," as described in the March-24 and May-19 Bluebooks of 1995. The large positive expectations shocks that occur in early 1984 and 1987 do not correspond to any obvious Fed communication, although they roughly coincide with the "inflation scare" episodes pointed out in Goodfriend (1993), during which the Fed took drastic actions to fight inflation.

Appendix B provides further evidence on the plausibility of our shock by comparing it to others proposed in the related literature. We show that, among those considered, only our shock is significantly related to macroeconomic forecast revisions in a manner consistent with exogenous shifts in policy expectations.

4.2. Baseline Impulse-Responses

Figure 3 presents IRFs to one-year-ahead anticipated policy shocks in our baseline VAR.[15] Our measure of agents' expectations derived from the survey and VAR forecasts are reported in the first row. The the second through fourth rows show the variables actually contained in the VAR: the BC survey forecasts, the macroeconomic data, and the financial variables. We estimate that the size of a one-standard-deviation exogenous shock to policy expectations at the one-year horizon is about 4 basis points (as shown by the immediate response of the expected TBill rate to such a shock), with a credible interval of about 1 to 7 basis points. The small size of these shocks is consistent with a monetary authority expected to adhere closely to its rule most of the time. However, despite the small change in short-rate expectations, the anticipated economic changes are nontrivial: one-year expected GDP and CPI increase by about 0.15 percent, according to the posterior medians in the first row. We note that the behavior of our measure of expectations is similar, but not identical, to the behavior of the survey forecasts in the second row, reflecting that surveys only receive partial weight in our measure ($\alpha < 1$).

Our central question is how the changes in expectations affect actual macroeconomic outcomes (third and fourth rows of IRFs). In response to a one-year anticipated policy shock of average size, GDP immediately increases by about 0.05 percent and reaches a peak of about 0.2 percent after 10 quarters, while the CPI immediately increases by about 0.15 percent and reaches a peak of about 0.2 percent within one year. These increases are "statistically significant" in the sense that at least 84 percent of the posterior probability mass for both variables lies above zero. The point estimate of hours worked does not seem to rise at impact and reaches about 0.1 percent after one year, though its credible band is wider and includes zero for the first 4 quarters. It is possible that hours worked increase sluggishly because of the presence of frictions in the labor market. All three variables revert toward their initial values only very slowly. Meanwhile, in the period of the shock, the actual TBill rate rises by about 2 basis points. However, it quickly reverses sign, and subsequent changes in the actual TBill rate are closely aligned with those of the expected TBill rate. Overall, our results suggest that the economy responds quite fast to expectations, with a large share of the total eventual changes occurring within a year from the shock.

We also find that the response of the stock market to an anticipated policy shock is not statistically different from zero. Exogenous monetary-policy shocks of the type we have identified may generally be small relative to other movers of equity prices, including changes in risk and risk aversion. In contrast, the reaction of the 2-year Treasury yield to an anticipated policy shock is quantitatively consistent with the response of the TBill rate over 8 quarters, indicating that the reaction is mainly driven by a change in the expectation component, as opposed to the term premium. (Note that, unlike the survey variables in the VAR system, the 2-year Treasury yield is not subject to a loose-rationality prior.)

Finally, for robustness, we also considered an alternative identification scheme in which a contemporaneous zero restriction, rather than a sign restriction, applies to the short rate in the case of the anticipated policy shock. Appendix F shows that this alternative identification produces results very similar to those obtained under the baseline identification. Further, given that Aruoba et al. (2022) demonstrate that the ZLB affects the propagation of shocks in structural VARs, in Appendix G we show how the baseline results vary when we exclude the ZLB period (2008:Q4-2015:Q4) from our sample. Interestingly, except for hours worked,

[15] The responses of the macroeconomic variables in first differences corresponding to Figures 3 through 6 appear in Appendix C. The IRFs to the unanticipated shock appear in Appendix D.

the IRFs are overall little changed. This might be due to the fact the ZLB is a rarely binding constraint for expected future rates; it may also be that including surveys in the VAR mitigates the misspecification associated with the nonlinearity in policy rates.

5. Fully Credible Forward Guidance

Our baseline estimates above show that exogenous changes in policy expectations can have sizable effects on the economy, suggesting the potential for FG to be a useful policy tool. However, for FG to make sense as a policy, two further criteria must be met. First, it must be true that the central bank has some ability to affect expectations. This is the issue of partial credibility, which we take up in the next section. Second, it must be true that a policy-rate path when perfectly anticipated has economic effects that would not occur if the same path were completely unanticipated. In other words, there must be *marginal* effects of anticipated policy over unanticipated policy. In this section, we ask whether this second criterion holds in the data. This allows us to assess fully-credible FG policies in a manner consistent with their treatment in standard NK models.

5.1. Measuring the Marginal Effects of Fully Credible FG

Using the results of the VARs estimated in Section 4, we simulate a combination of anticipated and unanticipated policy shocks that produce the hypothetical short-rate path we wish to consider.[16] Specifically, suppose that, in quarter 0, the Fed: announces that the policy rate will be 10 basis points lower in quarters 1 through 4, is believed by agents, and follows through on its promise. To implement this scenario, labeled fully-credible FG, we simulate an anticipated policy shock at quarter 0, η_0^4, that lowers the expected short rate by 10 basis points and a sequence of unanticipated policy shocks, $\{u_0, ..., u_4\}$, that ensures that the actual short rate remains unchanged at time 0 and is then 10 bp lower over the following four quarters. (Note that u_0 offsets any endogenous policy response to the economic expansion induced by the FG announcement.)[17] After quarter 4, we introduce no further shocks, and the economy simply follows the VAR dynamics back to the steady state. We do this for each of our 50,000 draws.

We compare the outcome of this experiment to a scenario in which the short rate follows an identical path over quarters 0 to 4, but this path is not pre-announced and therefore there is no anticipated policy shock. That is, we simulate a sequence of unanticipated policy shocks in quarters 1 through 4 that produces the same short-rate path as in the fully-credible FG scenario. (We do not need a shock in quarter 0 because the actual TBill rate is assumed not to move when FG is announced.) The distance between the IRFs of the fully-credible FG and unanticipated policy scenarios tells us the *marginal* effect of full credibility—that is, the extra economic effect that is achieved by unequivocally announcing a given policy path in advance.

Figure 4 summarizes the outcome of the fully-credible FG scenario in red, and of the unanticipated policy scenario in blue. The *marginal* effects of credible forward guidance, beyond an identical unanticipated policy path, are the differences between the red and blue lines, shaded in pink. (The black lines and gray shaded regions represent the case of partial credibility, discussed in the next section.) By construction, the TBill rate paths are identical in both the anticipated and FG scenarios until quarter 4 and then are allowed to differ. In the case of fully-credible FG, the expected TBill rate falls immediately to -10 basis points, whereas in the unanticipated policy scenario, it changes more slowly as shocks to the actual policy-rate path arrive. The initial disparity in the TBill rate expectations is the key difference between the two scenarios, and it is what drives the significant increase in GDP and CPI expectations over the first two quarters of the fully-credible FG scenario.

As a result, in the fully-credible FG scenario, current GDP and CPI increase on impact by about 0.17 and 0.3 percent, respectively, according to the posterior medians. (Credible bands are omitted to magnify the

[16] This approach is similar in spirit to the fiscal policy scenarios in Mountford and Uhlig (2009).

[17] To be precise, the anticipated policy shock η_0^4 is determined by setting the initial change in the expected TBill rate over the next four quarters, $\left(\Gamma_\eta^{E[i]} - \frac{\Gamma_u^{E[i]}\Gamma_\eta^i}{\Gamma_u^i}\right)\eta_0^4$, equal to -10 bp. The required unanticipated policy shock in quarter 0 is then given by $u_0 = -(\Gamma_\eta^i/\Gamma_u^i)\eta_0^4$.

scale of the figure.) GDP peaks at about 0.4 percent after three years, while CPI sees its entire increase of about 0.35 percent within the first year. Hours do not rise at impact, but reach 0.2 percent by quarter 4. In the unanticipated policy scenario, the median responses of GDP and CPI are at most half as big as in the FG scenario and take longer to occur, precisely because the expected TBill rate adjusts more slowly. However, as shown in Table 2, the differences between the fully-credible FG and unanticipated-policy scenarios are not statistically different after the first couple of quarters, in the sense that their credible intervals overlap. The relatively rapid convergence of the two scenarios in the case of one-year guidance is to be expected given that the disparity between the anticipated policy-rate paths dissipates quite quickly. Nevertheless, fully-credible FG not only triggers overall larger macroeconomic responses, but also moves them forward as policy expectations immediately adjust to the announced policy path.

5.2. Fully Credible FG at Longer Horizons

As demonstrated by Carlstrom et al. (2012) and Del Negro et al. (2012), the FG puzzle originates from a theoretical implication typical of standard NK models: each one-period extension of a commitment to stay at the ELB has increasingly larger effects the longer the duration of the commitment is. This seems quite implausible. Understanding how the potency of fully-credible FG varies with the horizon of the guidance also has important policy implications. As argued by many scholars and policymakers, one possible strategy against a higher likelihood of recurrent ELB episodes is to keep rates "lower for longer" (e.g., Bernanke et al., 2019). The basic motivation is that promises by the central bank to keep rates "lower for longer" can help reduce longer-term rates and stimulate the economy today, even if further cuts in the policy rate are not feasible.

Having built constant-horizon expectations of the TBill rate, GDP, and CPI from one to 11 years ahead, using our VAR estimates and related simulations, we can test for the FG puzzle using a shock conceptually consistent with the FG shock in NK models. Specifically, extending the horizon of fully-credible FG (i.e., the horizon of the expectations at which the identifying restrictions are imposed in the VAR), we can assess how the macroeconomic effects depend on that horizon. Table 2 tabulates the marginal effects (i.e., the effects of fully-credible FG minus those of unanticipated policy) of those extensions. In each experiment, at time 0 the Fed commits to reduce in the next period the policy rate to -10 basis points and keep it there for a particular length of time. For each variable, at each horizon, we report the marginal effects at impact and after 2, 4, 8, and 20 quarters. Marginal effects marked by a star are statistically "significant" (i.e., the posterior probability mass for which FG generates a response larger than unanticipated policy is $\geq 84\%$).

Overall, we find that fully-credible FG shocks generate bigger responses than unanticipated policy shocks in the levels of GDP and CPI, but these macroeconomic effects do not get increasingly larger as the horizon of the guidance gets longer. For example, if FG lowers policy expectations for three years rather than one year, the near-term response of GDP and CPI does not change much. And, if policy expectations are lowered for 11 years rather than three years, the short-run effects on GDP and CPI only double. Thus, there is some evidence that the effects of fully-credible FG get bigger with the guidance's horizon, but not nearly as much as NK models predict. Importantly, for FG at all horizons, the marginal effects are economically and statistically significant in the period of the announcement, but they dissipate quite quickly. This is because the divergence in policy expectations under fully-credible FG and unanticipated policy is short lived, and it is the quick decline in the anticipated policy-rate path that brings forward the stimulative effects.

We interpret these results to suggest that the FG puzzle seems to be indeed a puzzle, as it is a prediction of standard NK models not corroborated by the data, according to our methodology.[18] Furthermore, by construction in our exercises above, the initial change in expectations is subsequently realized. Thus, the explanation of the FG puzzle cannot lie solely in the reliability or credibility of the policy signal. Nevertheless, in the next section, we will show that it is crucial to differentiate between the marginal effects of partial and full credibility, as indeed anticipated policy changes are estimated to be very noisy in general, with important consequences for the economic impact of FG.

[18]This conclusion contrasts with that reached in Bundick and Smith (2020).

6. Noisy Anticipated Policy and Partial Credibility

Fully credible FG assumes that policymakers can perfectly shape agents' expectations at any horizon. This is a very strong assumption. In reality, central banks only provide noisy signals about the future. Here, we use our estimates to examine how the economy responds to signals about future monetary policy that are in fact correct but are, at least at first, only partially believed. The IRFs estimated in Section 4.2 represent the average response of the economy to signals containing both future fundamental policy shocks and noise shocks. If partially credible FG can be assumed to contain the same amount of noise as the policy signals underlying our data, then an anticipated policy shock that turns out to be a fundamental shock can be used to study the case of partial credibility.

6.1. Estimation of Fundamental and Noise Shocks

To recover from our estimates the IRFs for a fundamental shock (v_{t+h}) that is anticipated in advance, we note that the estimated IRF over each period j, for any variable x, to an anticipated policy shock (η_t^h) is a weighted average of the IRF to a future fundamental shock and the IRF to a noise shock (ϵ_t^h):

$$IRF_{\eta^h}(x_j) = w_v IRF_v(x_j) + (1 - w_v) IRF_{\epsilon^h}(x_j), \tag{15}$$

where the weight w_v on the fundamental IRF is equal to the Kalman gain in the signal-extraction problem of Section 2.2. To see this, note that the response of the TBill rate to a noise shock must equal zero after $j = h$ periods, by definition. (If the signal is noise, the future policy innovation does not materialize, and hence the TBill rate in period $j = h$ does not change.) Therefore, in this case, equation (15) reduces to

$$IRF_{\eta^h}(TBill_h) = w_v IRF_v(TBill_h) \tag{16}$$

Since v_t is itself the period-t innovation in the TBill rate, it is trivially the case that $IRF_v(TBill_h) = v_h$. And, by RE, it must be true that $IRF_{\eta^h}(TBill_h) = \eta_0^h$. Thus, we have:

$$w_v = \frac{\eta_0^h}{v_h} = K_h \tag{17}$$

where the second equality follows from equation (8) when $s_{t-h}^h = v_t$ and, generalizing Section 2.2, we now allow the gain to vary across horizons.

The equivalence result of Chahrour and Jurado (2018) allows one to recover K_h given estimates of $\sigma_{\eta^h}^2$ and σ_u^2, which are the variances of the shocks that one would identify in a VAR. Our problem is complicated slightly by the fact that our observations are on time-series *averages* of the TBill rate over various horizons. In Appendix H, we show that, under the simplifying assumption that η_t^j is homoskedastic over each period j leading up to h, the relevant gain is:

$$\bar{K}_h = \frac{\sigma_{\eta^h}^2}{\left(\sigma_u^2 + \frac{h+1}{2}\sigma_{\eta^h}^2\right)} \tag{18}$$

where the "bar" denotes that the expectation is about an h-period average.

Finally, once we have estimated $\bar{K}_h$, we can derive estimates of the variances of the fundamental and noise shocks as

$$\sigma_v^2 = \frac{\sigma_{\eta^h}^2}{h\bar{K}_h} \tag{19}$$

$$\sigma_{\epsilon^h}^2 = h\sigma_v^2\left(\frac{1}{\bar{K}_h} - 1\right) \tag{20}$$

Since $\text{var}(\bar{\eta}^h) = \bar{K}_h \text{var}(v)$, these calculations simply rearrange equation (9), adjusting for the fact that the variance of the anticipated policy shock we estimate is an h-period average. Having these estimates will allow us to assess the relative importance of v_t and ϵ_t^h at different horizons.

To find the IRFs to the noise shocks, first we simulate an anticipated policy shock that changes the average one-year expected TBill rate by -10 basis points, followed by a series of equal-sized unanticipated policy shocks that cause the path of the actual TBill rate to average zero over the following year.[19] Second, for each draw, we use equation (18) to compute $\bar{K}_h$. The median estimate is $\bar{K}_h = 0.27$, implying that, in our baseline specification with one-year-horizon expectations, noise shocks have a variance about 3 times as large as that of future fundamental policy shocks. Consequently, we infer that noise shocks receive 73% of the weight in the IRFs to η^h in (15).

Using this calculation, Figure 5 shows the differential response to fundamental (black) and noise (red) shocks. In each panel, the two IRFs average to the single IRF shown in each panel of Figure 3, using a weight of 73% for the noise IRFs and 27% for the fundamental IRFs. It is striking that, while the IRFs to the noise shocks go quickly to zero, the IRFs to fundamental policy shocks adjust gradually to new levels, with this adjustment being very persistent. Hence, only fundamental policy shocks generate long-run changes in CPI and GDP. However, the noise shocks do trigger short-lived macroeconomic fluctuations, which are hard to see because of the large scale of the fundamental shocks' effects. Specifically, in the first quarter, CPI increases by 0.3 percent, implying a rise in inflation of 1.2 percent in annualized terms; while, the level of GDP increases by 0.14 percent. Those changes are driven by the initial decline in the expected TBill rate and 2-year Treasury yield, which take a couple of quarters to revert to zero, catching up to the unchanged actual TBill rate. In other words, the noise-induced economic fluctuations are due to interest-rate expectational errors as it takes about two periods for agents to correct their interest-rate expectations.

Overall, our decomposition of the baseline IRFs into their fundamental and noise components indicates that the noise accounts for a significant share of policy signals and, as such, it notably dampens the macroeconomic impact of anticipation. The response to fundamental shocks in Figure 5 would have been much larger if they were *known* to be fundamental from the beginning.

6.2. The Marginal Effect of Partial and Full Credibility

If partially credible FG can be assumed to contain the same amount of noise as the policy signals underlying our data, then an anticipated policy shock that turns out to be a fundamental shock can be used to study the case of partial credibility. To better understand the effects of credibility, the black lines in Figure 4 show the median responses, in the VAR estimated with one-year expectations, to a fundamental shock of -10 basis points that is conveyed with noise ("partially credible" scenario). As above, we introduce additional unanticipated shocks to keep the realized policy path at -10 basis points for one year following the "announcement." The differences between the IRFs in the unanticipated-policy and partially-credible scenarios, indicated by the gray-shaded portion of the pink regions, inform us about the marginal effect of anticipation under the (more realistic) assumption that agents engage in signal extraction. The differences between the IRFs in the fully-credible and partially-credible FG scenarios (red versus black lines) inform us about the additional benefits of full credibility.

While in the fully-credible scenario the expected TBill rate declines at impact by 10 basis points as agents fully believe the announcement, in the partially-credible scenario the expected TBill rate declines at impact by about 3 basis points, as $\bar{K}_h$—the gain determining how much agents trust the noisy signal—equals 0.27 in the VAR estimated with one-year expectations. Nevertheless, the discrepancy in policy expectations between the partially-credible and unanticipated-policy scenarios is sufficient to quickly raise inflation by 0.3 percent more in the partially-credible scenario; on the other hand, that discrepancy does not make much of a difference for the real variables in the short run. GDP and hours do start displaying a larger response to the partially-credible policy shock after about 5 quarters, increasing by about 0.05 percent more than in the case of unanticipated policy shocks. Still, the marginal effects of FG under full credibility are much larger than under partial credibility, as noise notably dampens FG's economic impact.

Table 3 summarizes how the informativeness of the policy signal varies with its horizon. In particular, in the first three columns, we report the standard deviations of the anticipated policy shock, the future fundamental policy shock, and the noise shock, calculated using equations (19) and (20). The calculations

[19] Although there are, of course, an infinite number of possible 4-quarter paths that average to zero, a path that involves the same-sized shock in each period has the highest probability density.

shown in each row employ the results from separate VARs in which the horizon of the expectations is extended from 1 to 11 years. The last two columns report the standard deviation of the policy signal and the Kalman gain estimated for each of the expectational horizons.

The estimated standard deviation of the fundamental shock remain almost unchanged as the expectation horizon extends, as we would expect if we are measuring it correctly. Meanwhile, the standard deviation of the noise shock becomes larger with the horizon, doubling in magnitude from the 1- to 3-year horizon, for example. As a result, K_h becomes smaller as the horizon extends. Therefore the ability of FG with this degree of credibility to affect the economy becomes quite limited. Indeed, the calculations suggest that, under partial credibility, FG trying to shape policy expectations beyond the 4-year horizon would be perceived to be so noisy as to almost defeat the purpose of communicating in advance future policies.

These findings have important implications for lower-for-longer type of policies, as partial credibility can severely limit the efficacy of such policies by reducing the central bank's ability to influence agents' expectations. This conclusion is somewhat in line with what Bernanke et al. (2019) find using simulations from the FRB/US model of the Federal Reserve. Differently from those authors, we have estimated the amount of noise in the policy signal and analyzed its evolution at various horizons. The large amount of noise that we find leads us to conclude that it is challenging to augment stimulus at the ZLB using lower-for-longer policies, unless the Fed's credibility enables it to convey unusually informative signals.

7. Conclusion

Applying sign restrictions to measures of expectations in a VAR, we have identified imperfectly and perfectly anticipated monetary policy innovations. In the case of imperfect anticipation, we have distinguished between anticipated policy changes that turn out to reflect only future fundamental policy and only noise. This has allowed us to estimate the economic response to policy-expectations shocks similar to those induced by FG, and the extent to which FG might affect policy expectations. Both types of estimates are crucial to evaluate the efficacy of FG in a world with imperfect information and in a manner consistent with NK models.

We find that anticipated policy easing over the near term is associated with significant and persistent increases in prices and output: a 10-basis-point decline in one-year expectations for the average short-term rate raises current prices by 0.3 percent, and current output and employment by about 0.2 percent. Consequently, we find that the marginal effects of credible FG relative to unanticipated policy are immediate and big, but they do not grow increasingly larger as the guidance's horizon is extended. Thus, we confirm the existence of the FG puzzle, and—since these results hold even when the changes in expectations are equal to the subsequent fundamental shocks to policy—the puzzle cannot be explained by appealing to credibility. A further complication, however, is that anticipated monetary policy seems to be very noisy: at the one-year horizon we estimate that over 70 percent of policy signals consist of noise, and as the horizon of anticipation extends the prevalence of noise increases. FG with this amount of noise is unlikely to be effective because it is unlikely to be believed. Hence, partial credibility can be problematic for lower-for-longer type of policies.

References

Acosta, M., 2023. The Perceived Causes of Monetary Policy Surprises. Working paper, February .

Aiolfi, M., Capistran, C., Timmermann, A., 2010. Forecast Combinations. In Oxford Handbook of Economic Forecasting Oxford University Press, Oxford.

Andrade, P., Ferroni, F., 2021. Delphic and Odyssean Monetary Policy Shocks: Evidence from the Euro Area. J. Monet. Econ. 117, 816-832.

Ang, A., Bekaert, G., Wei, M., 2007. Do Macro Variables, Asset Markets, or Surveys Forecast Inflation Better? J. Monet. Econ. 54, 1163-1212.

Arias, J., Rubio-Ramirez, J.F., Waggoner, D.F., 2018. Inference Based on SVARs Identified with Sign and Zero Restrictions: Theory and Applications. Econometrica 86, 685-720.

Aruoba, S.B., Mlikota, M., Schorfheide, F., Villalvazo, S., 2022. Svars with occasionally-binding constraints. Journal of Econometrics 231, 477-499.

Auerbach, A.J., Gorodnichenko, Y., 2012. Measuring the output responses to fiscal policy. Am. Econ. J. Macroecon. 4, 1-27.

Bauer, M., Swanson, E.T., 2023. An alternative explanation for the 'Fed information effect.'. Amer. Econ. Rev. 113, 664-700.

Baumeister, C., Hamilton, J.D., 2015. Sign restrictions, structural vector autoregressions, and useful prior information. Econometrica 83, 1963-1999.

Bernanke, B.S., Kiley, M.T., Roberts, J.M., 2019. Monetary policy strategies for a low-rate environment, in: AEA Papers and Proceedings, pp. 421-26.

Bernanke, B.S., Kuttner, K.N., 2005. What explains the stock market's reaction to Federal Reserve policy? J. Finance 60, 1221-1257.

Blanchard, O.J., L'Huillier, J.P., Lorenzoni, G., 2013. News, noise, and fluctuations: An empirical exploration. Am. Econ. Rev. 103, 3045-70.

Bu, C., Rogers, J., Wu, W., 2021. A unified measure of fed monetary policy shocks. J. Monet. Econ. 118, 331-49.

Bundick, B., Smith, A.L., 2020a. The dynamic effects of forward guidance shocks. Rev. Econ. Stat. 102, 946-965.

Bundick, B., Smith, A.L., 2020b. Should We Be Puzzled by Forward Guidance? FRB of Kansas City.

Caldara, D., Herbst, E., 2019. Monetary policy, real activity, and credit spreads: Evidence from Bayesian proxy SVARs. Am. Econ. J. Macroecon. 11, 157-92.

Campbell, J.R., Ferroni, F., Fisher, J.D., Melosi, L., 2019. The limits of forward guidance. J. Monet. Econ. 108, 118-134.

Campbell, J.R., King, T.B., Orlik, A., Zarutskie, R., 2020. Issues regarding the use of the policy rate tool. FEDS Working Paper 2020-70 .

Campell, J.R., Evans, C.L., Fisher, J.D., Justiniano, A., 2012. Macroeconomic effects of FOMC forward guidance. Brookings Papers on Economic Activity (Spring) , 1-80.

Carlstrom, C.T., Fuerst, T.S., Paustian, M., 2012. How inflationary is an extended period of low interest rates? FRB of Cleveland Working Paper 1202 .

Chahrour, R., Jurado, K., 2018. News or Noise? The Missing Link. Am. Econ. Rev. 108, 1702-36.

Cieslak, A., Schrimpf, A., 2019. Non-monetary news in central bank communication. J. Int. Econ. 118, 293-315.

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

Coibion, O., Gorodnichenko, Y., 2015. Information rigidity and the expectations formation process: A simple framework and new facts. Am. Econ. Rev. 105, 2644-78.

D'Amico, S., King, T.B., 2015. What does anticipated monetary policy do? FRB of Chicago Working Paper No. 2015-10 .

Del Negro, M., Giannoni, M.P., Patterson, C., forthcoming. The Forward Guidance Puzzle. J. Pol. Econ. Macro. .

Del Negro, M., Schorfheide, F., 2013. DSGE Model-Based Forecasting, in: Handbook of Economic Forecasting. Elsevier. volume 2, pp. 57-140.

Doh, T., Smith, A.L., 2022. A new approach to integrating expectations into VAR models. J. Monet. Econ. .

Eggertsson, G.B., et al., 2003. Zero Bound on Interest Rates and Optimal Monetary Policy. Brookings Papers on Economic Activity (Spring) 2003, 139-233.

Faust, J., Wright, J.H., 2009. Comparing Greenbook and reduced form forecasts using a large realtime dataset. J. Business & Economic Statistics 27, 468-479.

Fernández-Villaverde, J., Rubio-Ramírez, J.F., Sargent, T.J., Watson, M.W., 2007. ABCs (and Ds) of understanding VARs. Am. Econ. Rev. 97, 1021-1026.

Forni, M., Gambetti, L., Lippi, M., Sala, L., 2017. Noisy news in business cycles. Am. Econ. J. Macroecon. 9, 122-52.

Galí, J., 2015. Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework and Its Applications. Princeton: Princeton University Press.

Gomes, S., Iskrev, N., Mendicino, C., 2017. Monetary Policy Shocks: We Got News! J. of Economic Dynamics and Control 74, 108-128.

Goodfriend, M., 1993. Interest rate policy and the inflation scare problem: 1979-1992. FRB Richmond Economic Quarterly 79, 1-23.

Gürkaynak, R.S., Sack, B.P., Swanson, E.T., 2004. Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements. Int. J. Cent. Bank. 1, 55-93.

Hanson, S.G., Stein, J.C., 2015. Monetary policy and long-term real rates. J. of Financial Economics 115, 429-448.

Jarociński, M., Karadi, P., 2020. Deconstructing monetary policy surprisesThe role of information shocks. Am. Econ. J. Macroecon. 12, 1-43.

Kaminska, I., Mumtaz, H., Šustek, R., 2021. Monetary policy surprises and their transmission through term premia and expected interest rates. J. Monet. Econ. 124, 48-65.

Krugman, P.R., Dominquez, K.M., Rogoff, K., 1998. It's Baaack: Japan's Slump and the Return of the Liquidity Trap. Brookings Papers on Economic Activity 29, 137-206.

Kuttner, K.N., 2001. Monetary Policy Surprises and Interest Rates: Evidence from the Fed Funds Futures Market. J. Monet. Econ. 47, 523-544.

Laséen, S., Svensson, L., 2011. Anticipated Alternative Policy Rate Paths in Policy Simulations. Int. J. Cent. Bank. 7, 1-35.

Leduc, S., Sill, K., 2013. Expectations and Economic Fluctuations: An Analysis Using Survey Data. Rev. Econ. Stat. 95, 1352-1367.

Leeper, E.M., Sims, C.A., Zha, T., Hall, R.E., Bernanke, B.S., 1996. What Does Monetary Policy Do? Brookings Papers on Economic Activity 27, 1-78.

Leeper, E.M., Walker, T.B., Yang, S.C.S., 2013. Fiscal Foresight and Information Flows. Econometrica 81, 1115-1145.

Levin, A.T., López-Salido, D., Nelson, E., Yun, T., 2010. Limitations on the Effectiveness of Forward Guidance at the Zero Lower Bound. Int. J. Cent. Bank. 6, 143-189.

Lunsford, K.G., 2020. Policy Language and Information Effects in the Early Days of Federal Reserve Forward Guidance. Am. Econ. Rev. 110, 2899-2934.

Milani, F., Treadwell, J., 2012. The Effects of Monetary Policy "News" and "Surprises". J. Money, Credit, and Banking 44, 1667-1692.

Miranda-Agrippino, S., Ricco, G., 2021. The Transmission of Monetary Policy Shocks. Am. Econ. J. Macroecon. 13, 74-107.

Mountford, A., Uhlig, H., 2009. What are the Effects of Fiscal Policy Shocks? J. of Applied Econometrics 24, 960-992.

Nakamura, E., Steinsson, J., 2018. High-Frequency Identification of Monetary Non-Neutrality: The Information Effect. Q. J. Econ. 133, 1283-1330.

Ramey, V.A., 2011. Identifying government spending shocks: It's all in the timing. The Quarterly J. of Econ. 126, 1–50.

Robertson, J.C., Tallman, E.W., Whiteman, C.H., 2005. Forecasting Using Relative Entropy. J. Money, Credit, and Banking 37, 383-401.

Romer, C.D., Romer, D.H., 2000. Federal Reserve Information and the Behavior of Interest Rates. Am. Econ. Rev. 90, 429-457.

Rubio-Ramirez, J.F., Waggoner, D.F., Zha, T., 2010. Structural Vector Autoregressions: Theory of Identification and Algorithms for Inference. Rev. Econ. Stud. 77, 665-696.

Sims, C.A., 1992. Interpreting the Macroeconomic Time Series Facts: The Effects of Monetary Policy. European Economic Rev. 36, 975-1000.

Swanson, E.T., 2021. Measuring the effects of federal reserve forward guidance and asset purchases on financial markets. J. Monet. Econ. 118, 32-53.

Tallman, E.W., Zaman, S., 2020. Combining survey long-run forecasts and nowcasts with BVAR forecasts using relative entropy. Int. J. of Forecasting 36, 373-398.

Werning, I., 2011. Managing a liquidity trap: Monetary and fiscal policy. Technical Report. NBER Working Paper, No. 17344.

Werning, I., 2015. Incomplete Markets and Aggregate Demand. Technical Report. NBER Working Paper, No. 21448.

Wright, J.H., 2013. Evaluating real-time VAR forecasts with an informative democratic prior. J. of Applied Econometrics 28, 762–776.

8. Figures and Tables

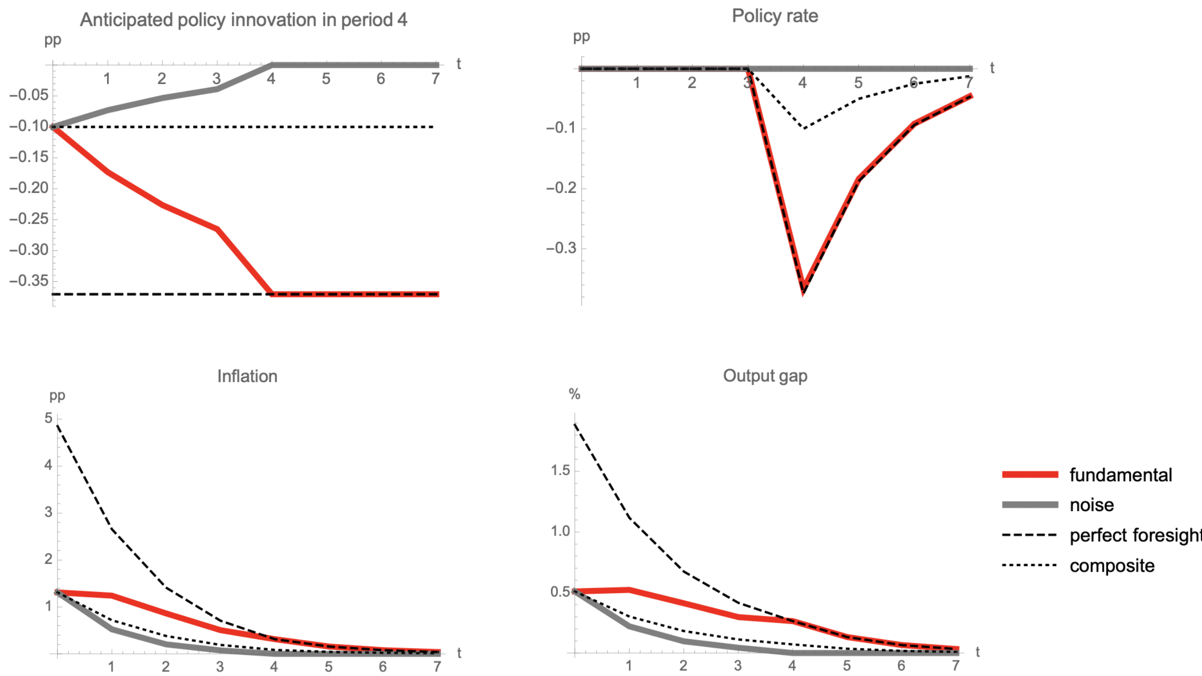

Figure 1: Responses to Anticipated Monetary Policy with Noise in NK Models

Notes: This figure depicts responses to a noisy signal about future monetary policy that turns out to be entirely fundamental (red lines) and entirely noise (gray lines). In each case, the shock is scaled to result in a change of -10 basis points in expectations at period $t = 0$ for the policy rate in time $t = 4$. The dotted black lines show the weighted average of the fundamental and noise shocks, which is the impulse-response that would be estimated on a large dataset containing both types of shocks. For comparison, the dashed black lines show the case where the fundamental shock (about 37 points) is anticipated with perfect foresight.

	E[TBill]	E[GDP]	E[CPI]	TBill
MP Shock				
Anticipated η^h	$-$	$+$	$+$	$+$
Unanticipated u	$-$	$+$	$+$	$-$

Table 1: Identification Restrictions

Notes: Sign restrictions imposed in the VAR to identify the anticipated (η^h) and unanticipated (u) monetary policy shocks.

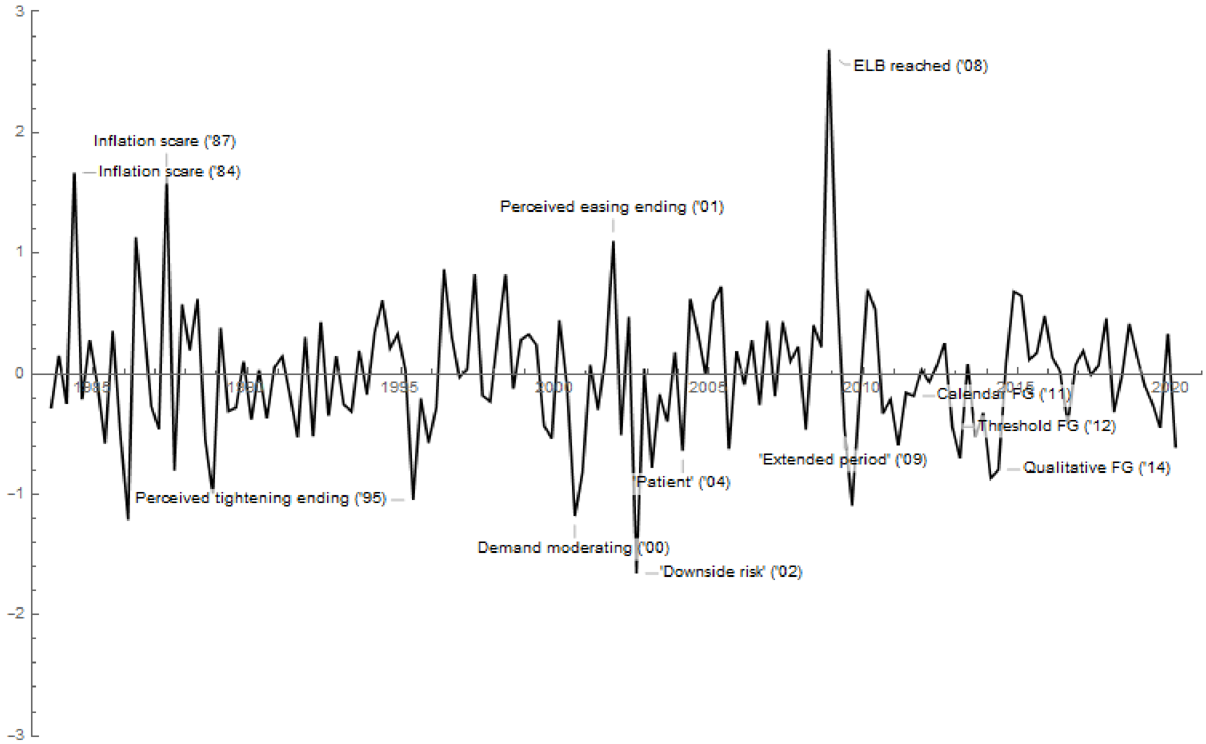

Figure 2: Estimated One-Year-Ahead Anticipated Policy Shocks

Notes: This figure depicts the time series of the median of the one-year-ahead anticipated monetary-policy shocks obtained from the baseline VAR identified with sign restrictions. The unit is one standard deviation, with positive shocks reflecting policy tightening. Some of the largest positive and negative shocks are marked by the year (in parentheses) and the key words characterizing the FG statements or the Bluebook summaries, that is, reports from the staff of the Board of Governors describing inter-meeting changes ahead of each FOMC.

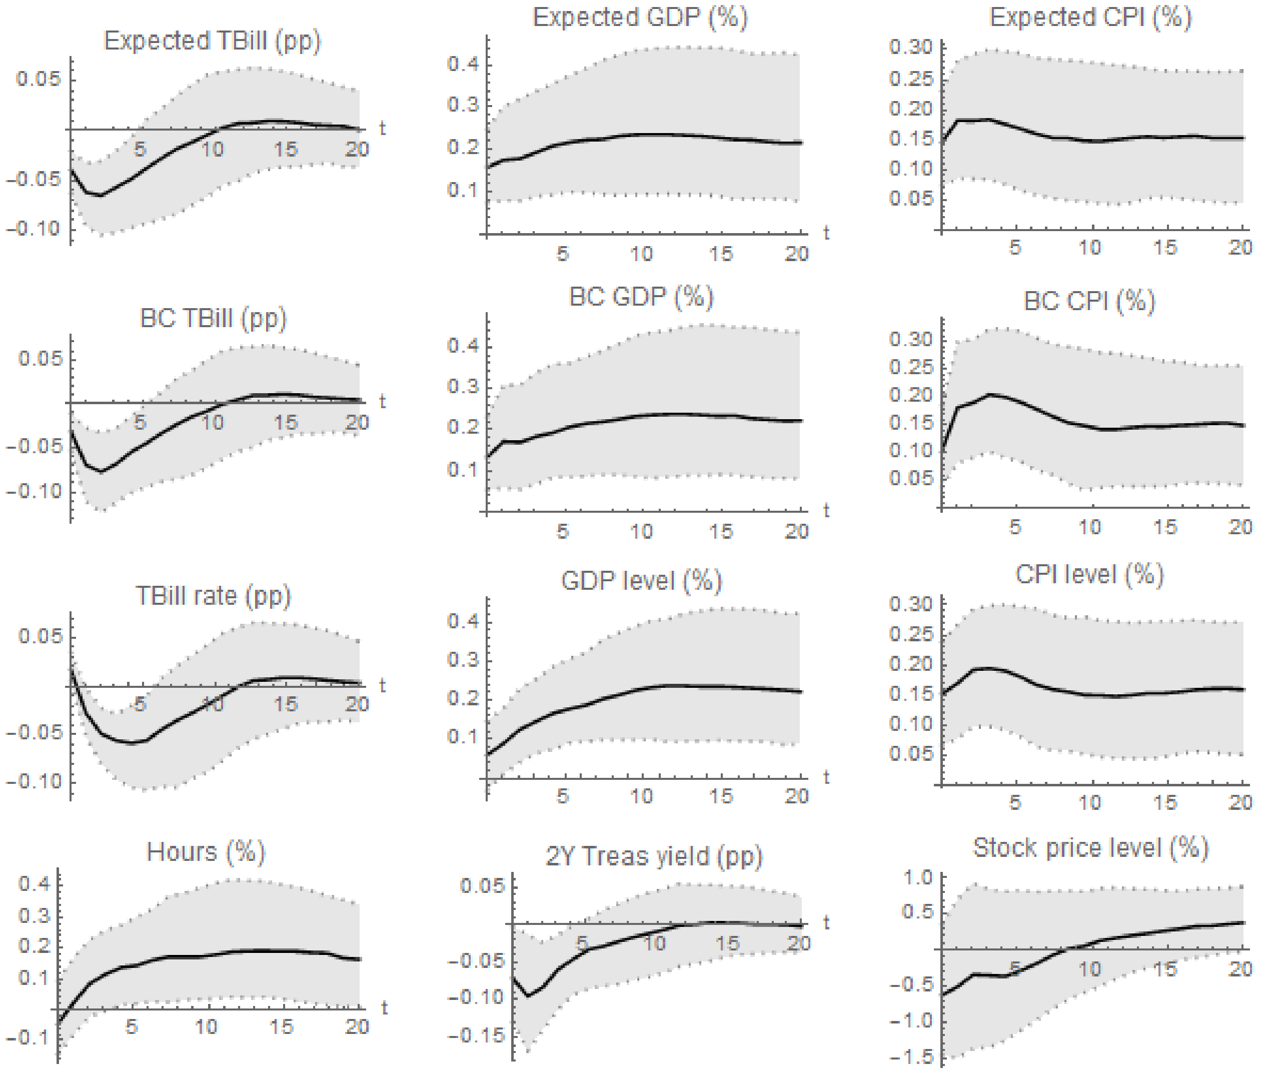

Figure 3: Baseline Impulse Response Functions to a One-Standard-Deviation One-Year-Ahead Anticipated Policy Shock

Notes: The panel reports the IRFs to our baseline one-year-ahead anticipated monetary-policy shock. Specifically, the first row shows the IRFs of the expected variables derived from the combination of the survey and VAR forecasts, the second row shows the IRFs of the BC survey forecasts, the third row shows the IRFs of the actual macroeconomic variables, and the fourth row shows the IRFs of hours worked together with the financial variables. The 68% credible interval is shaded in grey.

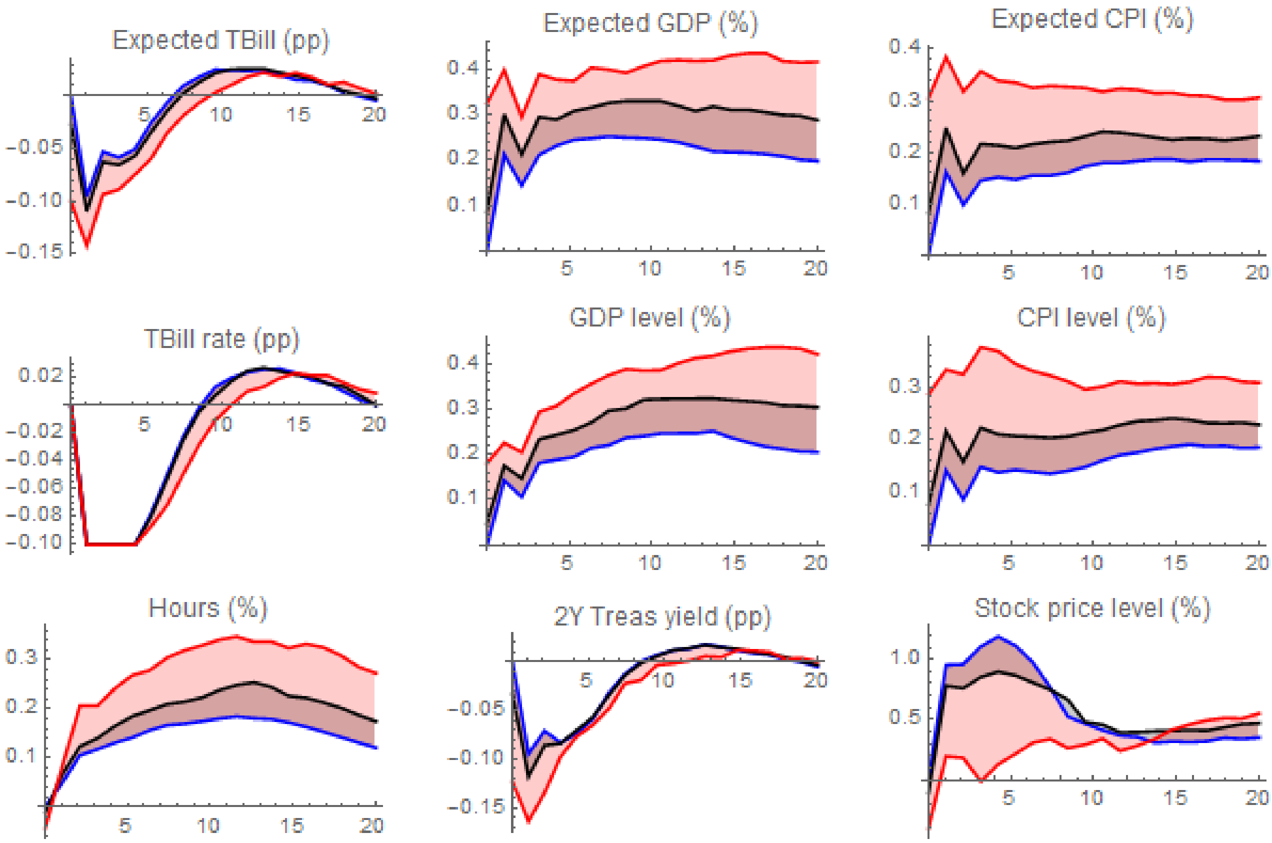

Figure 4: Fully Credible FG (Red) vs. Partially Credible FG (Black) vs. Unanticipated Monetary Policy (Blue)

Notes: The panel shows in red the IRFs of each variable to a perfectly anticipated policy shock that lowers the policy rate by 10 basis points in quarters 1 through 4, i.e., a case of fully credible FG; in blue, the IRFs to an identical policy-rate path that is completely unexpected, i.e., the unanticipated policy scenario. The difference between the red and blue IRFs, the entire pink shaded region, is the marginal effect of fully credible FG. The black lines and gray-shaded portion of the pink regions show the analogous information for the scenario of *partially* credible FG, discussed in Section 6. That is, the difference between the black and blue IRFs is the marginal effect of partially credible FG.

Expectations Horizon	1Y	2Y	3Y	4Y	5Y	11Y
E[TBill]						
Current quarter	-0.10*	-0.10*	-0.10*	-0.10*	-0.10*	-0.10*
2 quarters ahead	-0.05*	-0.05*	-0.05*	-0.08*	-0.07*	-0.08*
4 quarters ahead	-0.03	-0.02	-0.02	-0.06*	-0.04*	-0.07
8 quarters ahead	-0.04	0.00	0.00	-0.03	-0.02	-0.07
20 quarters ahead	0.01	0.00	0.00	-0.01	-0.01	-0.02
log GPD						
Current quarter	0.14*	0.17*	0.13	0.18	0.15	0.32*
2 quarters ahead	0.12	0.09	0.01	0.10	0.17	0.31
4 quarters ahead	0.12	0.07	-0.04	0.09	0.16	0.33
8 quarters ahead	0.19	0.09	-0.02	0.04	0.04	0.30
20 quarters ahead	0.23	0.14	0.03	0.16	-0.01	0.22
GDP growth (ann.)						
Current quarter	0.57*	0.69*	0.51	0.72	0.62	1.27*
2 quarters ahead	0.21	0.03	0.07	0.19	0.11	0.20
4 quarters ahead	0.04	0.00	-0.07	-0.01	-0.04	0.11
8 quarters ahead	0.04	0.04	0.01	-0.02	-0.06	0.06
20 quarters ahead	-0.01	0.00	0.00	0.02	0.00	0.00
CPI						
Current quarter	0.30*	0.30*	0.30*	0.42*	0.43*	0.60*
2 quarters ahead	0.23*	0.21	0.33*	0.52*	0.43*	0.56
4 quarters ahead	0.19	0.14	0.27*	0.44*	0.39	0.64
8 quarters ahead	0.12	0.11	0.11	0.28*	0.17	0.68
20 quarters ahead	0.11	0.08	0.10	0.13	0.04	0.26
Inflation (ann.)						
Current quarter	1.21*	1.18*	1.20*	1.69*	1.72*	2.38*
2 quarters ahead	0.10	-0.12	0.09	0.18	0.15	0.25
4 quarters ahead	-0.03	-0.11	-0.12	-0.10	-0.08	0.00
8 quarters ahead	-0.05	-0.06	-0.08	-0.10	-0.14	0.00
20 quarters ahead	0.00	0.00	-0.01	-0.02	-0.01	-0.03

Table 2: Marginal Effects of Fully-Credible FG at Different Horizons for GDP and CPI

Notes: This table tabulates the marginal effects (i.e., the effects of fully-credible FG minus those of the unanticipated policy) of extensions of fully-credible FG, from one to 11 years. A star indicates that the marginal effect is "significant" (i.e., the percentage of the posterior probability mass for which FG generates a larger response than the unanticipated policy is ≥ 84).

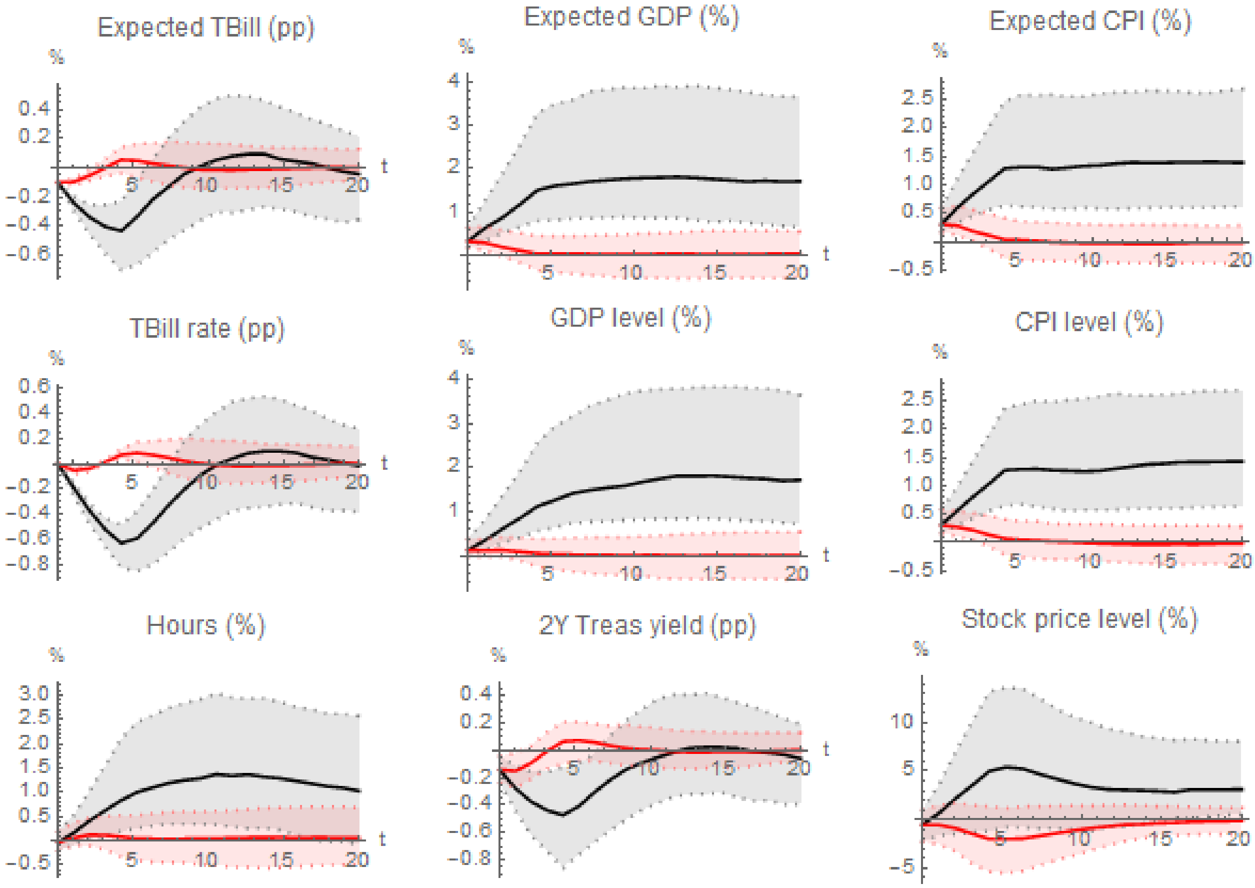

Figure 5: Fundamental Shocks (Black) vs. Noise Shocks (Red)

Notes: The panel shows the IRFs of each variable to anticipated policy shocks that turn out to be future fundamental policy shocks, in black, and to those that turn out to be noise shocks, in red. In each panel, the two IRFs average to the single IRF in each panel of Figure 3, using a weight of 73% for the noise IRF and 27% for the fundamental IRF.

	Anticipated (σ_{η^h})	Fundamental (σ_v)	Noise (σ_{ϵ^h})	Signal (σ_{s^h})	Kalman Gain (K_h)
Horizon					
1-year	0.04	0.09	0.15	0.17	0.27
	(0.02, 0.06)	(0.05, 0.14)	(0.08, 0.24)	(0.11, 0.27)	(0.11, 0.37)
2-year	0.05	0.10	0.23	0.24	0.21
	(0.02, 0.07)	(0.06, 0.15)	(0.12, 0.31)	(0.16, 0.33)	(0.16, 0.22)
3-year	0.05	0.09	0.31	0.32	0.15
	(0.03, 0.07)	(0.05, 0.18)	(0.14, 0.46)	(0.19, 0.46)	(0.13, 0.15)
4-year	0.04	0.09	0.30	0.33	0.12
	(0.02, 0.06)	(0.04, 0.16)	(0.13, 0.50)	(0.18, 0.51)	(0.11, 0.12)
5-year	0.04	0.09	0.38	0.38	0.09
	(0.02, 0.07)	(0.03, 0.15)	(0.15, 0.74)	(0.17, 0.75)	(0.09, 0.10)
11-year	0.02	0.08	0.36	0.37	0.05
	(0.01, 0.04)	(0.02, 0.13)	(0.24, 0.78)	(0.26, 0.79)	(0.05, 0.05)

Table 3: Standard Deviations of Policy Shocks, the Policy Signal, and the Kalman Gain

Notes: This table reports the standard deviations of the: (imperfectly) anticipated policy shock, future fundamental policy shock, noise shock, policy signal, and Kalman gain estimated with our methodology. Reported in parentheses, the credible bands.