What Does Anticipated Monetary Policy Do?

Stefania D'Amico Thomas B. King*

March 12, 2016

Abstract

We identify expected future innovations to monetary policy by applying sign restrictions to survey forecasts embedded in a VAR. Expectations for exogenous policy easing that materialize over the subsequent year—similar to those generated by credible "forward guidance"—have immediate stimulative effects on output, inflation, and employment. The effects are larger than those produced by a similar shift in the policy path that is unanticipated. The results are consistent with the mechanism underlying forward guidance in New Keynesian models but suggest that those models overstate the persistence of the inflation response and the effectiveness of forward guidance at longer horizons.

*Federal Reserve Bank of Chicago. Contact: stefania.damico@chi.frb.org, thomas.king@chi.frb.org. Portions of this article originally circulated as part of our companion paper, "Policy Expectations, Term Premia, and Macroeconomic Performance." We thank Jonas Arias, Bob Barsky, Jeff Campbell, Matteo Iacoviello, Alejandro Justiniano, Thomas Laubach, Daniel Rees, Eric Swanson, Min Wei, and conference and seminar participants at the 2015 WEAI meetings, the 2016 AEA meetings, the Banque de France, the Federal Reserve Board, the Federal Reserve Bank of Chicago, and the University of Tor Vergata for helpful comments and suggestions. The views expressed here do not reflect official positions of the Federal Reserve.

1 Introduction

This paper uses time-series data to examine how beliefs about future monetary policy affect the current state of the economy. This question has received particular attention as policymakers have increasingly used "forward guidance" to shape interest rate expectations as a means of providing stimulus. In the U.S., for example, the Federal Open Market Committee has in recent years provided calendar dates through which it expected its target rate to be held near zero, economic thresholds that would warrant an unusually accomodative policy stance, and qualitative signals about potential deviations from its conventional policy rule. In principle, a commitment to future policy accommodation should have immediate stimulative effects on the economy, and in some theoretical models those effects can be quite large.[1] Yet, while this idea has been widely discussed and appears to have already influenced policymakers' thinking, there is no conclusive evidence on whether the mechanism it postulates is actually present in the data. We address this question and quantify the dynamic effects of policy expectations and forward guidance on output, inflation, and employment, providing a benchmark for calibrating and assessing models of forward guidance as well as their policy implications.

Just as identifying the effects of conventional monetary policy requires capturing a deviation from the historical policy rule, identifying the effects of *anticipated* monetary policy requires capturing an *anticipated* deviation from the historical policy rule.[2] The type of shock that matters in theory—and that is of interest to policymakers hoping to stimulate the economy through forward guidance—is one in which agents come to expect policy easing beyond what the expected state of the economy would normally warrant. To identify these shocks, we exploit survey forecasts of key macroeconomic variables and of the short-term interest rate, and we embed those forecasts in a structural VAR model that separates the sources of their innovations. In particular, we use sign restrictions to isolate the anticipation of an *exogenous* innovation to future monetary policy: the expected short rate must move in the *opposite* direction of expected inflation and expected GDP. In standard theoretical models, this directional pattern of changes in expectations is unique to anticipated innovations to monetary policy.

Importantly, an anticipated *endogenous* response of policy to the expected state of the economy moves short-rate expectations in the *same* direction as

[1] For example, see Krugman (1998), Eggertsson and Woodford (2003), Laseen and Svensson (2011), and Werning (2011).

[2] 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 policy.

expected inflation or expected GDP. If policymakers usually act in a clear and consistent way, and if agents understand the policy rule, then most fluctuations in short-rate forecasts will indeed reflect expectations of systematic policy responses to future economic conditions, not expectations of exogenous policy innovations. Previous evidence, such as Romer and Romer (2004) and Campbell et al. (2012), is consistent with this notion: when people expect looser policy, it is usually because they expect the economy to deteriorate. Consequently, if we did not control for expectations of systematic policy responses, they would obscure the expectations of exogenous policy innovations. Information about short-rate forecasts alone is not sufficient to tell these two types of expectations apart, as the distinguishing feature is the corresponding change in forecasted economic conditions. Our sign restrictions are contemporaneously imposed on survey forecasts of three key variables and, by excluding all cases in which these variables move in the same direction, can isolate the anticipation of exogenous monetary policy innovations.

We find that when survey respondents anticipate an exogenous policy easing over the subsequent year, the result is an immediate and persistent increase in both prices and real activity, just as theory predicts. Specifically, a decrease of 25 basis points in expectations for the average short-term rate over the next year results in a short-run increase of about 1 percent in GDP, employment, and prices. These responses occur much faster than those following a conventional monetary-policy shock, which we identify in the same VAR using standard restrictions. After about two years, the shock to policy expectations has about the same effect on output as a conventional policy shock of the same size and a cumulative effect on inflation that is 2 to 3 times as large. We find that shocks to expectations over long horizons (six and eleven years) have effects in the same direction as those to one-year expectations, but they are smaller, less persistent, and not always statistically significant. Importantly, we do not take a stand on whether the anticipated monetary-policy innovations we identify reflect actual information that agents have about future policy ("news") or simply random deviations from full-information rationality ("noise"). Indeed, we show that our results are quantitatively similar regardless of which interpretation one adopts.

We then use these results to simulate a policy scenario consistent with the way that credible forward guidance is modeled in theoretical treatments. In particular, we analyze how the economy responds when there is an initial shock to the expected path of the policy rate, followed by a sequence of shocks to the actual short rate that make the initial change in expectations materialize as anticiapted. We find that this correctly anticipated reduction in the policy-rate path over a one-year horizon results in a significantly larger and more-rapid stimulus than the same reduction does when it is not anticipated. Again, this difference is smaller when we use longer-horizon expectations to conduct this exercise.

To our knowledge, these results are the first to assess the effects of policy expectations and forward guidance on the macroeconomy in a way that is consistent with theoretical treatments but does not impose any particular structural model. They fill a gap in the literature because, even among those who find the theoretical case for forward guidance compelling, its quantitative importance has remained an open question. Different variations of New Keynesian models, such as those explored in Levin et al. (2010), McKay et al. (2014), and Werning (2015), can yield substantially different impacts of policy expectations, and standard versions of such models deliver macroeconomic effects that are generally viewed as implausibly large (Carlstrom et al., 2012; Del Negro et al., 2012). In the only previous attempt to estimate these effects in a model-free way, Campbell et al. (2012) found the opposite of what theory predicts: when the Federal Reserve signals that lower rates are coming, survey expectations of GDP and inflation decline. They argued that this result likely reflected agents interpreting accommodative signals by the Fed as conveying negative information about the prospects for the economy ("Delphic" forward guidance), rather than as commitments to future stimulative deviations from the historical policy rule ("Odyssean" forward guidance). This is an example of the distinction made above between expectations for systematic responses to the economy versus expectations for exogenous policy innovations, a distinction that our methodology is specifically designed to account for.

Our results suggest that, at least in some dimensions, the effects predicted by simple theoretical models are not unrealistic. Using a simple New Keynesian model under a standard calibration we show that an anticipated 25-basis-point cut in the short rate over the coming year raises today's output gap and annual inflation rate by about 1.5 and 3.5 percentage points, respectively—responses that are in the ballpark of our VAR-based estimates. But in other respects our results are considerably at odds with the predictions of standard theory. First, we find that the effect of forward guidance on inflation is short lived, essentially taking the form of a one-time jump in the price level. In contrast, most theoretical models obtain an inflation response that is persistent and often even growing larger over time; consequently, they predict a much larger cumulative effect on prices than we find in the data. Second, New Keynesian models typically predict that forward guidance should have a larger effect on current outcomes when it pertains to horizons that are farther in the future (see, e.g., Carlstrom et al., 2012). Our results suggest the opposite: the effect of changes in policy expectations is smaller when the forecasts involve more-distant horizons. In this sense, we quantify the forward-guidance puzzle discussed in Del Negro et al. (2012).[3]

Beyond the above-mentioned papers, two other studies of monetary policy expectations bear mentioning in relation to our analysis. First, Gurkaynak et al. (2005) estimated that the "path factor" of interest rates (reflecting changes in short-rate expectations orthogonal to the current policy surprise) accounted for a large share of asset prices fluctuations. While their analysis has been widely cited as supporting the practical importance of policy expectations, it does not address the effects of those expectations on the wider economy. Second, Gertler and Karadi (2014) studied monetary policy shocks in a VAR, using one- and two-year Treasury rates as their policy indicators and instrumenting those rates with changes in federal funds futures around FOMC announcements. Their approach captures the idea that policy may operate not just through the current policy rate but through its expected path over the next couple of years. However, their measures of policy shocks conflate changes in the expected path of the policy rate with changes in its current level (as well as with changes in near-term term premia), and thus cannot be used to study the effects of conventional and forward-guidance policies independently. Furthermore, neither Gurkaynak et al. (2005) nor Gertler and Karadi (2014) distinguished expectations for exogenous policy innovations from expectations for endogenous policy responses to future economic conditions as we do, nor did they consider differential effects of expectations over different horizons.

In Section 2 of the paper, we begin by justifying the reduced-form of our survey-augmented VAR, discussing possible interpretations of innovations to survey forecasts, and motivating our identification strategy within a simple New Keynesian model. We also use that model to provide a quantitative theoretical benchmark for our empirical results. In Section 3, we discuss the implementation details of the VAR. Section 4 summarizes the baseline results and discusses what we can infer about the causality of anticipated monetary-policy innovations. In Section 5, we use our results to construct the forward-guidance scenarios. Section 6 conducts a battery of robustness checks, and Section 7 concludes the paper.

[3]On the other hand, our results could be consistent with modifications of the standard model, such as Milani and Treadwell (2012) and McKay et al. (2015), that give forward guidance its most powerful effects at relatively short horizons.

2 Modeling Shocks to Expectations

Our empirical specification relies on the incorporation of survey forecasts in a VAR. While previous studies have used survey data in empirical macro models,[4] our approach is perhaps novel in its use of such data to identify exogenous innovations to monetary-policy expectations. We therefore begin by discussing the circumstances under which the survey-augmented VAR is an admissible representation of the dynamics of the economy and how we ought to interpret innovations in the survey forecasts within this model. In particular, in Section 2.1 we show that the survey-augmented VAR is a valid reduced form for a general class of linear models, including those where agents beliefs may not correspond to the true probability distribution of the economy. In Section 2.2, we discuss how deviations of the survey forcasts from the VAR forecasts may reflect either "news" about future policy not captured by the time-t state or "noise" due to limited information or irrationality, and we discuss the implications for inference about causality. Finally, in Section 2.3, we illustrate how the sign restrictions that will be used to identify structural shocks in our VAR are implied by a standard New Keynesian model and how such a model predicts the economy ought to behave subsequent to an expectations shock, giving us a theoretical benchmark for the empirical results to follow.

2.1 Economic dynamics with subjective expectations

In order to talk meaningfully about "anticipation of innovations," it must be the case that expectations can contain some exogenous component that is not related to the contemporaneous observable fundamentals. In other words, we have to allow for the possibility that subjective beliefs might deviate from statistical beliefs inferred from current and historical data. The literature contains a variety of structures in which agents form expectations in a way that differs from the full-information rational-expectaitons benchmark. For example, agents' information sets may be *larger* than those of the econometrician because they receive informative signals—"news"—about changes in future fundamentals that are not directly reflected in the current state vector (Barsky and Sims, 2012; Leeper et al., 2013). Alternatively, agents may use *less* information than statistical beliefs would imply because information is ei-

[4]For example, see Barsky and Sims (2012), Milani and Rajbhandari (2012), Leduc and Sill (2013), and Barsky et al. (2014), Miyamoto and Nguyen (2014) in the macro literature and Chun (2011), Kim and Orphanides (2012), Orphanides and Wei (2012), and Piazessi et al. (2015) in the term-structure literature.

ther sticky (Mankiw and Reis, 2002) or difficult to extract (Lucas, 1972; Sims, 2003). Finally, agents may not adhere to strict rationality. For example, in Bullard et al. (2008), agents' forecasting models include a "judgment" term that is tacked on the objective forecast, and in Milani (2011) a similar term is motivated by agents' differing degrees of pessimism and optimism.

To encompass these cases (and following, for example, Piazessi et al., 2015), we introduce the subjective expectations operator E_t^S to stand for the first moment of the probability distribution perceived by agents in the economy, conditional on the information they possess at time t. This distribution may or may not be the same as the one that a rational agent with knowledge of both the time-t state and the model parameters would compute. However, as we discuss below, we do not assume that differences between subjective and statistical expectations necessarily reflect irrationality—an important possibility is that they reflect information about the future that cannot be directly inferred from the time-t state. Indeed, forward guidance may be thought of as "news" of this sort: agents anticipate changes in future monetary policy as a result of explicit central bank communication aimed at signaling upcoming deviations from the historical policy rule.[5]

While we take the subjective expectations themselves as primitive objects, we assume that beliefs are consistent in the sense that E_t^S obeys the law of iterated expectations ($E_t^S\left[E_{t+h}^S[y_{t+j}]\right] = E_t^S[y_{t+j}]$ for any stochastic process $\{y_t\}$ and $j > h > 0$), and we also note that E_t^S has the usual linearity property of expectation operators ($E_t^S\left[a + by_{t+h}\right] = a + bE_t^S\left[y_{t+h}\right]$ for constants a and b).

We consider economies of the form

$$x_t = Ax_{t-1} + BE_t^S\left[x_{t+1}\right] + \varepsilon_t \tag{1}$$

where x_t is the state vector, A and B are parameter matrices, and ε_t is a vector of mean-zero, iid fundamental innovations. By the law of iterated expectations, the random variable $E_t^S\left[x_\tau\right]$ follows a martingale under the subjective measure, for any fixed period $\tau > t$. Consequently, we can write

$$E_t^S\left[x_{t+1}\right] = E_{t-1}^S\left[x_{t+1}\right] + F\varepsilon_t + \eta_t \tag{2}$$

where F is an arbitrary parameter matrix and η_t is a random vector representing time-t shocks to expectations of the $t+1$ state. By the properties of a

[5]See for example the following sentence in 2014 and 2015 FOMC statements: "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."

martingale, η_t is independent of $E^S_{t-1}[x_{t+1}]$ under the subjective measure, and we assume that this is also true under the statistical measure (which is tantamount to an assumption that forecast revisions are not predictable). Without loss of generality, we assume that η_t is also uncorrelated with ε_t.

Equation (2) alone does not allow us to examine the dynamic properties of the economy, because it provides only a law of motion for the expectation at a particular *calendar date*—it does not determine how expectations about different periods in the future are linked. To complete the picture, we allow subjective expectations of the state at a *fixed horizon* to follow the general linear process

$$E^S_t[x_{t+2}] = Cx_{t-1} + DE^S_{t-1}[x_{t+1}] + G\eta_t + H\varepsilon_t \tag{3}$$

where C, D, G, and H are arbitrary coefficient matrices. We note that the usual specification of rational expectations conditional on the time-t state is a special case in which $D = 0$ and $\mathrm{var}_t[\eta_t] = 0$, and F, C, and H are reduced-form combinations of A and B.

By direct substitution, (1), (2), and (3) jointly imply that the variables x_t and $E^S_t[x_{t+2}]$ follow a first-order vector autoregression. However, in general, estimation of that system by OLS will not correctly identify the structural shocks because it is possible that ε_t is correlated with the regressor $E^S_{t-1}[x_{t+1}]$ in equation (3). In particular, that correlation will be positive if, in a period prior to t, agents receive an informative signal about the value of ε_t that then influences their forecast in periods subsequent to t. We therefore define the *unanticipated* component, e_t, of the structural innovation by

$$\varepsilon_t = \alpha E^S_{t-1}[x_{t+1}] + e_t \tag{4}$$

where α is the OLS coefficient ensuring that e_t is uncorrelated with $E^S_{t-1}[x_{t+1}]$.

The system then reduces to the following VAR:

$$\begin{pmatrix} x_t \\ E^S_t[x_{t+2}] \end{pmatrix} = \mathbf{\Theta} \begin{pmatrix} x_{t-1} \\ E^S_{t-1}[x_{t+1}] \end{pmatrix} + \mathbf{\Gamma} \begin{pmatrix} e_t \\ \eta_t \end{pmatrix} \tag{5}$$

where

$$\mathbf{\Theta} = \begin{pmatrix} A & B(F\alpha + I) + \alpha \\ C & D + H\alpha \end{pmatrix} \qquad \mathbf{\Gamma} = \begin{pmatrix} BF + I & B \\ H & G \end{pmatrix}$$

Clearly, the representation in (5) can be extended to cases in which (1) involves constant terms and expectations and lags beyond the first order.

If we have direct measures of agents' expectations, we can estimate $\mathbf{\Theta}$ by OLS and obtain unbiased estimates of the reduced-form dynamics of the economy, and we can estimate $\mathbf{\Gamma}$ from the reduced-form VAR error-covariance matrix by applying a suitable set of identifying restrictions. This will be our approach in Section 3. Note that the upper-right element of $\mathbf{\Gamma}$ directly identifies B. From equation (1), this parameter determines the effect that expectations of the future state have on the current state, which is our primary object of interest.

2.2 Interpreting expectations shocks

Although the future fundamental shocks ε_{t+h} are unpredictable from a statistical point of view, they may not be so under the subjective measure. Agents might, rightly or wrongly, believe that they have information about future structural innovations. Indeed, so long as agents recognize the linear structure of the economy (and believe that the value of B is such that the invertibility condition is satisfied), they will perceive the state to have a reduced-form representation

$$x_t = \Phi^S x_{t-1} + \sum_{h=1}^{\infty} \psi_h^S E_t^S [\varepsilon_{t+h}] + \Sigma^S \varepsilon_t \tag{6}$$

where Φ^S, Σ^S, and ψ_h^S $(h = 1, ..., \infty)$ are parameters characterizing subjective beliefs. Together with (2) and the observation that $E_t^S [\varepsilon_{t+h}]$ must follow a martingale under the subjective measure for any fixed period $\tau > t$, equation (6) implies

$$\eta_t = \sum_{h=1}^{\infty} \psi_h^S \Delta E_t^S [\varepsilon_{t+h}] \tag{7}$$

where $\Delta E_t^S [\varepsilon_{t+h}]$ is the change in subjective expectations of ε_{t+h} between periods $t-1$ and t. Thus, the shock η_t to expectations of the level of next period's state is a weighted sum of changes in expectations of structural shocks over all future periods. Put differently, when agents revise their beliefs about the future state, they only revise their expectations about the future stochastic disturbances, not about the deterministic dynamics.

Changes in subjective expectations of future shocks $\Delta E_t^S [\varepsilon_{t+h}]$ may occur for two reasons:

1. *News.* Agents receive correct information (i.e., news) about the shock in period $t + h$ before it occurs. This specification is similar to how forward guidance about monetary policy has been modeled in Laseen and Svensson (2011), Campbell et al. (2012), and Del Negro et al. (2013).

2. *Noise.* Agents' beliefs are subject to random fluctuations that do not correspond to any past, current, or future fundamental shocks. This type of information flow is consistent with models of sentiment or judgment. It is also consistent with models in which, for whatever reason, agents pay attention to a signal that has no economic content.

Clearly, agents must interpret their η_t purely as news. Furthermore, news and noise about ε_{t+h} are observationally equivalent to agents and to econometricians in periods t through $t+h-1$. In other words, since the expectations in equations (6) and (7) only involve the term $E_t^S[\varepsilon_{t+h}]$ (and not its decomposition into news and noise), both types of signals have the same effect on x_t. Only in period $t+h$ does a difference between news and noise emerge, because in that period the shock ε_{t+h} is realized in equation (1) and agents infer how correct the signal was.

As discussed in Blanchard et al. (2013), this observational equivalence means that we will not be able to tell whether expectations shocks represent news or noise or a combination of both in the data. The VAR only identifies the composite term η_t, and there is no way to estimate its decomposition. This issue matters for the attribution of causality. If expectations shocks are all noise, then any fluctuations in the economy that predictably follow such shocks must be caused by anticipation. But if expectations shocks embed a news component, then some of the fluctuations that follow such shocks reflect the effects of the structural innovations themselves, not just their anticipation.

Fortunately, in our model, we will also have an estimate of how the structural innovations in question—exogenous deviations from the monetary-policy rule—affect the economy, as they represent the conventional monetary-policy shocks. Consequently, infering the causality of expectations shocks in the case of news is a straightforward matter of taking the impulse-response to a shock η_t and subtracting the response to a shock e_{t+h}, scaled to be the same size. Thus, while we cannot tell how much of the movements in expectations reflect news versus noise, we can use our results to compute the component of the impulse-response functions (IRFs) that are caused by expectations under an assumption that shocks to expectations are entirely news or noise. We discuss this procedure further in Section 4.2.

A time-t news shock at has the same effects on the state as a time-t noise shock in periods t through $t+h-1$ and, after period $t+h-1$, differs from the effect of a noise shock by an amount equal to the effect of an unanticipated structural shock of the same magnitude. This observation can also be exploited to study cases of forward guidance, which, as we stressed before, is a type of news. In particular, we can construct forward-guidance scenarios by a combination of our identified η_t and conventional policy shocks e_t. A situation in which agents receive credible forward guidance at time t is equivalent to a situation in which their time-t short-rate expectations change exogenously and then subsequent shocks to the actual short rate occur such that the initial change in expectations turn out to be correct. We will use this technique in Section 5 to model forward guidance based on our VAR estimates.

2.3 Expectations shocks in a New Keynesian model

We now apply the above framework to consider what happens when we allow for exogenous fluctuations in subjective expectations of future monetary policy in an otherwise standard New Keynesian (NK) model. The purpose of this exercise is fourfold. First, it illustrates in a familiar setting the mechanism underlying the impact of expectations shocks. Second, it demonstrates the qualitative and quantitative responses to such shocks implied by NK models, providing hypotheses (such as the so-called forward-guidance puzzle) to be tested in our empirical work. Third, it illustrates how the sign restrictions that will be used to identify monetary-policy expectations shocks in our VAR are specified by the theory. Lastly, it shows how forward guidance can be modeled using expectations shocks, setting up the scenarios we will construct with our empirical results.

We borrow the basics of the model from Gali (2008, c. 3). Specifically, under standard NK assumptions, the equilibrium conditions can be written as follows:

$$\pi_t = \beta E_t^S \pi_{t+1} + \kappa y_t \tag{8}$$

$$y_t = E_t^S y_{t+1} - \frac{1}{\sigma}\left(i_t - E_t^S \pi_{t+1} - r^*\right) \tag{9}$$

where y_t is the output gap, 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 $\kappa > 0$ is a nonlinear combination of structural parameters. In addition, assume that the short-term interest rate is set by the central bank according to the rule

$$i_t = \phi_y y_t + \phi_\pi \pi_t + \upsilon_t \tag{10}$$

where

$$\upsilon_t = \rho \upsilon_{t-1} + \varepsilon_t \tag{11}$$

with ε_t being a mean-zero iid shock, and $\phi_\pi \geq 0$, $\phi_y \geq 0$, and $0 \leq \rho \leq 1$.

We depart from the standard treatment only by assuming that expectations are formed under the subjective measure as described in the previous section. However, to keep the exposition simple, we assume that C, D, and F are all equal to zero. Furthermore, we assume that agents have knowledge of the true structural parameters, so that the only departure from the usual case is that they receive iid shocks to their expectations $\Delta E_t^S [\varepsilon_{t+h}]$ for some h. In the special case in which the variance of these shocks is zero, the model reduces to the standard NK model.[6]

To see how monetary-policy expectations shocks affect the current output gap and inflation in this model, note that we can solve forward to obtain a solution for time-t inflation and output as a function of the current policy stance υ_t and an infinite-order moving average of expected future shocks:

$$\pi_t = \psi_{0,\pi}\upsilon_t + \sum_{h=1}^{\infty} \psi_{h,\pi} E_t^S \varepsilon_{t+h}$$

$$y_t = \psi_{0,y}\upsilon_t + \sum_{h=1}^{\infty} \psi_{h,y} E_t^S \varepsilon_{t+h}$$

This is the equivalent of equation (6) for this model, since we have assumed that the subjective multipliers ψ_h^S are equal to their statistical counterparts ψ_h. Gali (2008) shows that the effect of an (unanticipated) monetary-policy shock, which is given by the multipliers on υ_t, is

$$\psi_{0,\pi} = -\kappa\Lambda \tag{12}$$

$$\psi_{0,y} = -\left(1 - \beta\rho\right)\Lambda \tag{13}$$

$$\psi_{0,i} = \left[\left(1 - \rho\right)\left(1 - \beta\rho\right) - \kappa\rho\right]\Lambda \tag{14}$$

where $\Lambda = \left[\kappa\left(\phi_\pi - \rho\right) + \left(1 - \beta\rho\right)\left(\phi_y + \sigma(1 - \rho)\right)\right]^{-1}$. Since ψ_π^0 and ψ_y^0 are necessarily negative for admissible values of the structural parameters, time-t inflation and output move in the opposite direction of the monetary-policy shock ε_t. Under standard parameterizations, ψ_i^0 is positive, implying that the nominal short rate moves in the same direction as the shock, and we will assume for the remainder of the discussion that this is the case.

[6] This specification is similar in some respects to models of beliefs shocks, as, for example, in Schmitt-Grohe and Uribe (2012) and Lorenzoni (2009). Milani and Treadwell (2012) also present a related DSGE analysis in which monetary policy shocks may be anticipated by agents, although the details of their model differ from those here.

Of more interest for our purposes are the multipliers for shocks to expectations about horizon $h > 0$. When the economy is hit by shock to monetary-policy expectations one period ahead, $\Delta E_t^S [\varepsilon_{t+1}]$, expectations of the state variables initially react as:

$$\begin{aligned}\Delta E_t^S \pi_{t+1} &= \psi_{0,\pi} \Delta E_t^S [\varepsilon_{t+1}] \\ \Delta E_t^S y_{t+1} &= \psi_{0,y} \Delta E_t^S [\varepsilon_{t+1}] \\ \Delta E_t^S i_{t+1} &= \psi_{0,i} \Delta E_t^S [\varepsilon_{t+1}]\end{aligned}$$

Intuitively, expectations of a future monetary-policy change have on expectations of future inflation, output, and interest rates the same impact that current policy shocks have on current inflation, output, and interest rates. That is, a shock to future policy expectations causes both $E_t^S \pi_{t+1}$ and $E_t^S y_{t+1}$ to move in the *opposite* direction of $E_t^S i_{t+1}$. This observation motivates our sign-based identification scheme. Notably, no other shock in standard models of this type can produce this response pattern. For example, a shock to expectations about future technology, which would enter through r^*, would generally move the short rate in the same direction as expected output and inflation. A "markup shock," which would appear as an additional stochastic term in equation (8), would generally move inflation and output in opposite directions.

One can show that the multipliers on expectations of ε_{t+h} at longer horizons are given by the recursion:

$$\begin{pmatrix} \psi_{h,\pi} \\ \psi_{h,y} \end{pmatrix} = R \begin{pmatrix} \psi_{h-1,\pi} \\ \psi_{h-1,y} \end{pmatrix} \tag{15}$$

where

$$R = \begin{pmatrix} \kappa + \beta\left(\phi_y + \sigma(1-\rho)\right) & \kappa\sigma \\ 1 - \beta\phi_\pi & \sigma\left(1 - \beta\rho\right) \end{pmatrix} \Lambda$$

To get a sense of the magnitudes associated with these shocks and how the responses vary across the expectational horizon, Figure 1 illustrates how the state of the economy responds in a typical calibration to expectations shocks that are 1 to 4 periods ahead. Taking periods to be quarterly (and again following Gali, 2008), let $\sigma = 1$, $\beta = .99$, $\kappa = .15$, $\phi_y = .125$, $\phi_\pi = 1.5$, and $\rho = .5$. The impulses are shocks $\Delta E_t^S [\varepsilon_{t+h}]$ that are sufficient to lower the expected h-period-ahead annualized interest rate by 25 basis points, where $h = 1, ..., 4$.

We first consider the case in which $\Delta E_t^S [\varepsilon_{t+h}]$ consists purely of noise. That is, once period h arrives, there is no actual policy innovation. As shown in panel A of the figure, inflation rises immediately in response to the expectations shocks, and it rises by more the farther in the future the innovation to the short rate is expected to occur. For an anticipated monetary easing one year ahead (the blue line) current quarterly inflation rises by about 3 percent (at an annual rate). This effect is somewhat damped because of the systematic response of policy, reflected in higher nominal and real short-term rates. For this calibration, the policy response is large enough to drive the output gap negative in early periods, even though the expectations shock itself is a stimulative one. Further, since no fundamental shocks actually materialize, and since the subjective expectational errors are assumed not to be persistent, the state always returns to zero once period h has passed.

Panel B illustrates the case in which $\Delta E_t^S [\varepsilon_{t+h}]$ consists of news. As noted in the previous section, the response of the economy to news shocks is the same as its response to noise shocks up until period $t+h$. After that time, the economy receives the additional stimulus that it would have received from an unanticipated monetary policy shock of 25 basis points at time $t+h$. This case is one that could, in principle, result from credible forward guidance—agents are made to believe that a policy shock will occur in the future, and then it actually does. However, if the path of the economy shown in this panel resulted from deliberate policy, the central bank's behavior would be somewhat bizarre. As in panel A, the systematic response of monetary policy causes nominal and real rates to rise in the near term when a future easing is anticipated; thus, the central bank mechanically finds itself raising rates in response to its own forward guidance. While there is nothing logically inconsistent about this outcome, it seems unrealistic that a central bank would announce an unusually accommodative future policy only to offset part of that accommodation now by adhering to its usual rule.

We therefore consider a more realistic forward-guidance scenario, in which the central bank maintains nominal rates at their time-$t-1$ levels until period $t+h$, at which point it adopts the pre-announced change. To achieve this outcome, it must introduce additional monetary policy shocks in each period $t, ..., t+h-1$ in order to offset its own systematic response to the economy and keep the short rate at its initial level. We assume that these additional policy shocks are also anticipated at time t by agents. (I.e., agents correctly believe that the path of the short rate will be unchanged until period $t+h$.) Panel C shows the impacts. Without the short-term policy offset to the expected future easing, the output gap now rises substantially. Given the stabilized nominal short rate, the higher inflation results in a significant downward movement in real rates.[7]

Finally, in panel D, we show what happens if the central bank promises to lower the short rate by 25 basis points, not just in period $t + h$, but for the entire period $t + 1$ through $t + h$. This is closer to what central banks have done in practice, and it essentially mirrors the forward-guidance experiments we will conduct with our empirical results in Section 5. Perhaps surprisingly, the responses are only slightly greater than in panel C. The reason is that, in the forward-guidance scenario, most of the effect derives from the large short-rate shock that is needed in the first period to keep current rates from rising. The size of that shock primarily depends on the timing and magnitude of the change in expectations that is farthest in the future, which is the same in both panels C and D.

The magnitudes of the responses in Figure 1 are large when compared to those of conventional monetary policy shocks within the same model. For example, an unanticipated shock to the actual short rate has an initial impact of only +0.2 percentage points on both the inflation rate and the output gap—an order of magnitude smaller than when the same shock is anticipated to occur a year in the future. These are manifestations of the "forward guidance puzzle" pointed out by del Negro et al. (2013). They are not specific to the structure of the simple model here, its calibration (within reason), or its assumed policy rule. Rather, as discussed by McKay et al. (2015), they result from the large influence of future interest rates on the path of the output gap and the way that path compounds into inflation via the NK Phillips Curve. While some authors, such as Kiley (2014b) and McKay et al. (2015), have proposed modifications to the basic NK structure that can reduce the effect of forward guidance, it is unclear what a reasonable result from such models should be, since there is no empirical work estimating this impact in a model-free way. Our results below provide some benchmarks along these lines.

[7]The output gap is much higher in panel C than in panel B, but inflation is only slightly higher. The reason is subtle. Since the short-rate shocks that occur in t through $t + h - 1$ in panel C are persistent, less of an expectations shock is needed to achieve a given value of $E_t^S v_{t+h}$ than was needed in panel B. For inflation, the effect of the short-rate and expectations shocks is about the same, so trading off one for the other makes little difference. But for the output gap, negative shocks to the short rate have a positive effect, whereas negative shocks to the *expected* short rate mostly have a negative effect (as shown in panels A and B). Consequently, the shift in the source of the shocks provides a large positive boost.

3 The Survey-Augmented VAR

3.1 Specification and data

We now apply the insights of the previous section to identify shocks to policy expectations in the data. For this purpose, it will be useful to divide the state vector x_t into macroeconomic variables that we will assume cannot respond contemporaneously to conventional, unanticipated monetary-policy shocks ($\mathbf{x}_{1t}$), those that potentially can ($\mathbf{x}_{2t}$), and the short rate itself (i_t). We can stack the data in the vector $\mathbf{X}_t = \begin{pmatrix} \mathbf{x}_{1t} & i_t & \mathbf{x}_{2t} & E_t^S[\mathbf{x}_{t+h}] \end{pmatrix}'$, where $E_t^S[x_{t+h}]$ is the subjective measures of expectations about a subset of the macro data at horizon h. Also let ϵ_t be a corresponding stacked vector of the reduced-form errors with covariance matrix $\mathbf{\Sigma} = \mathbf{\Gamma}\mathbf{\Gamma}'$. We rewrite the system (5) as

$$\mathbf{X}_t = \mathbf{\Theta}(L)\mathbf{X}_t + \epsilon_t. \tag{16}$$

where $\mathbf{\Theta}(L)$ is a lag polynomial of order L.

Our baseline sample uses quarterly data from 1983:3 to 2015:2. To measure agents' subjective expectations, we use survey data from the Blue Chip Survey (BCS) as our baseline, and we use the Survey of Professional Forecasters (SPF) used as a robustness check.[8] Each survey reports the respondents' average forecasts of GDP growth, CPI inflation, and the three-month Treasury bill (3-month T-Bill) rate, which we use as a proxy for the monetary-policy instrument. Due to idiosyncrasies in the conventions and timing of their reporting, the survey data from both sources require some manipulation to be useful in our VAR model. Our method for obtaining constant-horizon quarterly series from these data and the properties of the series are described in detail in the appendix.

Apart from the inclusion of the survey data, the specification of our baseline VAR model is similar to others in the literature. Specifically, we build loosely on Christiano, Eichenbaum, and Evans (2005) in our choice of macro variables. In x_{1t} we include log GDP, log CPI, and log labor productivity.[9] In x_{2t} we include log real profits and the M2 growth rate. (To conserve degrees of freedom, we omit a few variables—consumption, investment, and wages—that were of specific interest to Christiano et al. (2005) but are nearly collinear with the other variables in the VAR.) We also include a long-term Treasury yield in x_{2t}, with a maturity corresponding to the horizon of the survey expectations

[8] We obtained the SPF data from the Federal Reserve Bank of Philadelphia's website.

[9] We will also use the results to discuss employment, which of course can be calculated as the difference between output and productivity.

used in each model specification, in order to capture any effects of expected short rates that operate through longer-term spot rates.

Overall, we found that, although the basic sign and magnitude patterns of our results were consistent across specifications, the choice of lag length mattered in some cases for the confidence bounds around the results. In our baseline model, we use the BIC to select lag length. The robustness checks in Section 6 consider some alternatives to this specification, including different choices of regressors and lag length and estimation using Bayesian methods.

3.2 Identification of structural shocks

We will distinguish three types of structural shocks, conventional (unanticipated) monetary-policy shocks (e_t), policy-expectations shocks (η_t), and a vector of other shocks that we leave unspecified (*other*). The contemporaneous coefficients $\mathbf{\Gamma}$ can be partitioned conformably with $\mathbf{X}_t$:

$$\mathbf{\Gamma} = \begin{pmatrix} \mathbf{\Gamma}_e^{x_1} & \mathbf{\Gamma}_\eta^{x_1} & \mathbf{\Gamma}_{other}^{x_1} \\ \Gamma_e^i & \Gamma_\eta^i & \mathbf{\Gamma}_{other}^i \\ \mathbf{\Gamma}_e^{x_2} & \mathbf{\Gamma}_\eta^{x_2} & \mathbf{\Gamma}_{other}^{x_2} \\ \mathbf{\Gamma}_e^{E^S[x]} & \mathbf{\Gamma}_\eta^{E^S[x]} & \mathbf{\Gamma}_{other}^{E^S[x]} \end{pmatrix} \tag{17}$$

where, for example, $\mathbf{\Gamma}_e^{x_1}$ is a vector containing the response of the block $\mathbf{x}_{1t}$ to the standard monetary-policy shock and Γ_η^i is a scalar representing the contemporaneous response of the policy rate to the policy-expectations shock. (Elements of the $\mathbf{\Gamma}$ matrix highlighted in bold represent vectors, while the others are scalars.)

In order to identify the elements of $\mathbf{\Gamma}$, we impose a combination of exact and partial identification restrictions. As noted, we assume that $\mathbf{x}_{1t}$ does not respond contemporaneously to standard monetary-policy shocks, providing us with a set of short-run exclusion restrictions, as are common in the monetary-policy VAR literature. The exact restrictions then amount to:[10]

$$\mathbf{\Gamma}_e^{x_1} = \mathbf{0} \tag{18}$$

The use of exclusion restrictions to identify the unanticipated monetary-policy shocks introduces an assymetry into our results, since our anticipated policy shocks will be identified with sign restrictions. In the robustness section,

[10]Note that we do not require the other half of the usual short-run restriction: policy can respond contemporaneously to *all* of the macro variables and expectations.

we consider an alternative identification scheme that follows Uhlig (2005) in identifying the unanticipated shocks using sign restrictions as well.

Expectations shocks are identified by drawing from the space of possible $\mathbf{\Gamma}$ matrices that satisfy the exclusion restrictions just described and discarding all draws that do not satisfy sign restrictions on the contemporaneous impacts on the surveys. Those sign restrictions—intended to capture changes in expectations about future monetary policy—enforce the following condition: the time-t impact on the survey forecast of the average T-Bill rate over periods t to $t+h$ must be in the opposite direction of the impact on the survey forecast of the time-$t+h$ GDP and price level. To ensure that expectations shocks are isolated from conventional policy shocks, we also impose that the contemporaneous T-Bill rate does not move in the same direction as the forecasted T-Bill rate in response to an expectations shock.[11] These assumptions about the contemporaneous impacts of expectations shocks are consistent with the predictions of the NK model discussed earlier and, indeed, with a large class of forward-looking macroeconomic models.[12] As a normalization, we consider expectations shocks that move anticipated short rates in the negative direction (i.e., expectations for future policy easing). Thus, the partial restrictions amount to:

$$\begin{aligned} \Gamma_{\eta^h}^{E^S[i]} &< 0 \\ \left\{\Gamma_{\eta}^{E^S[\pi]}, \Gamma_{\eta}^{E^S[y]}, \Gamma_{\eta}^{i}\right\} &> 0 \end{aligned} \tag{19}$$

Our identification assumptions are summarized in Table 1.

To implement this identification scheme, we follow the procedure of Arias et al. (2014), who show how to draw uniformly (under the Haar measure) from the possible set of $\mathbf{\Gamma}$ matrices that satisfy a given set of zero restrictions. To compute impulse-response functions, we draw jointly $10,000$ times from the posterior distribution of the VAR parameters and the set of admissible $\mathbf{\Gamma}$'s, and we simulate the effects of a one-standard-deviation shock under each draw

[11] A positive response of the current short rate to expectations of a declining future short rate could arise if, as we hypothesize, consumers react to looser expected policy by increasing current output and prices and the Fed tightens current policy in response. Leduc and Sill (2013) find empirical support for this hypothesis. We also ran the model restricting this response to be zero, with no appreciable difference in the results.

[12] One might be concerned that we are actually picking up the effects of time-t aggregate-demand shocks: perhaps output and inflation rise today, and persistence causes expectations for their values tomorrow to rise as well, rather than the other way around. But, since the Fed raises rates in response to exogenous increases in output and inflation, short-rate expectations would rise in that scenario. The first condition in (19) ensures that we do not include such situations among the expectations shocks that we identify.

over 44 subsequent quarters. (This matches the horizon of our longest-range surveys.) Following other studies using partial identification, we focus on the pointwise medians across all of the draws.

4 Results

4.1 Baseline

Table 2 displays the contemporaneous impacts of the policy-expectations shocks on the survey expectations themselves. By construction, these shocks have negative effects on expected T-Bill rates and positive effects on expected GDP and inflation. However, nothing guarantees *a priori* the magnitude or significance of these effects. The table shows that exogenous shocks to expectations of the policy rate are modest, resulting, on impact, in an average change of the one-year expected T-Bill rate of just 3 basis points. Such a small magnitude is to be expected in a world in which agents typically do not have much reason to anticipate deviations from the policy rule. In contrast, the total conditional standard deviation of the one-year expected T-Bill rate (i.e., the standard error in that equation of the VAR) is 40 basis points. Consequently, of the total variance in the expected short rate in our sample, less than 1% represents expectations for exogenous deviations from the policy reaction function. As we noted in the introduction, and consistent with the results of Campbell et al. (2012) and others, the overwhelming majority of innovations to policy expectations appear to reflect expectations for systematic responses to the future state of the economy. That is why it is crucial to sweep out those effects.

As also shown in the table, a one-standard-deviation shock to one-year expectations for exogenous policy is associated with a contemporaneous increase in the four-quarter-ahead level of GDP of 0.14 percent and an increase in one-year expected inflation of 0.09 percent. As we will show, these changes in expectations are consistent with what actually happens to the economy subsequent to the shock.

For longer-horizon expectations, the size of the impact of expectations shocks on the expected T-Bill rate is about the same, that is, about -3 basis points. (Again, this is small as a fraction of the total conditional standard deviation of that expectation.) However, since this is the average effect over a much longer time period, it represents a larger change in beliefs than the shock to the same rate over a one-year period and, presumably, should impart a greater amount of stimulus. Yet, while the effects of the 6-year and 11-year policy-expectations shocks on the projected levels of GDP and the CPI are monotonically increasing across maturities, the implied expected average growth rates are decreasing. A policy-expectations shock that lowers the expected average short rate over the next 11 years by 3 basis points only raises expected average annual GDP growth and inflation over the same period by 0.03 percent. A likely explanation for this finding is that agents do not believe that monetary policy innovations can have sustained effects on output and inflation over such long horizons. A policy that results in a 3-basis-point reduction in short-term rates for a period of eleven years in a row may be interpreted as a structural change in the policy rule or the steady state of the economy, rather than as monetary stimulus.

Table 3 offers some evidence that the policy-expectations shocks that we have identified do indeed correspond to periods in which agents' expectations for future monetary policy may have shifted. In particular, we list all of the quarters in which expectations shocks larger than one standard deviation occurred between 1999, when the FOMC began including meaningful forward-looking commentaries in its statements, and 2008, when the effective lower bound was reached.[13] In nearly all of these periods, we can match our identified shocks to obvious changes to the wording of the statement that point toward future policy moves in the same direction of the shocks. Some of these events, such as the introduction of the "patient" language in 2004, are also those that are identified by Gurkaynak et al. (2005) as being particularly potent episodes of forward guidance.[14]

Panel A of Figure 2 shows the responses of key variables to each of the two shocks, standard monetary-policy shocks (e_t) and policy-expectations shocks (η_t), in our baseline model using the one-year BCS forecasts. The IRFs are represented as medians (black line), interquartile ranges (red region), and interdecile ranges (blue regions) across draws of the reduced-form coefficients, and the shocks considered are all one-standard-deviation in an accommodative direction (lower interest rates).

The estimated responses to conventional monetary-policy shocks are fairly standard—for example, they are similar to those found in Christiano et al.

[13]We exclude the ZLB period from this table, even though several clear forward-guidance events do show up in our model as large expectations shocks, because nearly every quarter during this period contains *some* FOMC communication that could be interpreted as dovish.

[14]In constructing this table, we have taken care to account for the timing of the surveys relative to FOMC meetings. In particular, the BCS data for each quarter is typically gathered in the first week of the last month of each quarter, while the last FOMC meeting of the quarter takes place a couple of weeks later and therefore would not be reflected in survey responses until the following quarter. The dates in the table reflect the dates of the identified shock, not of the corresponding FOMC statement.

(2005). Output rises slowly but persistently in the quarters following the shock and peaks at about 0.2 percent after three years, while there is a marginally significant and sluggish response of inflation following an initial modest price puzzle. Interestingly, while these are the standard findings over samples that include the 1970s, exclusion restrictions do not typically deliver these results over the post-1982 period used in this study, likely due to the monetary policy reaction function becoming more forward-looking (see Barakchian and Crowe, 2013). Our use of the survey data in the VAR seems to help with some of the problems related to the existing identification schemes. However, note that the price puzzle remains even though we include the forward-looking survey data on inflation among our regressors.

The expectations shock has a large immediate impact on output, with the response reaching a peak of about 0.1 percent within the first year, but it decays relatively quickly, stabilizing after about three years. The response of the price level to this shock reaches its peak of 0.1 percent at impact, and it reverts only very slowly. The effects of these shocks on employment are also quite large, especially in the short run. (This can be seen by comparing the IRFs for GDP and labor productivity.) These findings indicate that, particularly at short horizons, shocks to expectations can be quite powerful. Note that the expectations shocks do not exhibit a price puzzle.

As noted earlier, a one-standard-deviation expectations shock generates a 3-basis-point change in the one-year survey forecast of the short rate on impact. In contrast, conventional policy shocks move one-year short-rate expectations and the contemporaneous short rate itself by 11 basis points on average. Thus, although the responses depicted in Figure 2 of GDP and inflation to the two shocks are of similar magnitude, the basis-point size of the expectations shock needed to generate these responses is much smaller in the case of the expectations shock. To see this comparison more precisely, Table 4 displays the size of conventional policy shocks required to equal the macroeconomic effect of a −25-basis-point one-year expectations shock. Over the short run, the expectations shock is five times more powerful, as the magnitude of the conventional policy shock would have to be almost 125 basis points to achieve the same cumulative effect on GDP and over 200 basis points to have the same effect on hours. In the longer-run, the expectations shock is about twice as effective as the conventional policy shock for the cumulative effect on GDP and CPI. The point estimate of its effect on hours is smaller, although (not shown in the table) the differences at these horizons are not statistically significant.

Panels B and C of Figure 2 show the effects using longer-horizon surveys. The responses to policy-expectations shocks are similar in magnitude although they are shorter-lived, becoming statistically insignificant after two to three years. The responses to the standard monetary-policy shocks are similar to those in Panel A. When the model is estimated with these longer maturities, the expectations shocks are less important relative to monetary-policy shocks, although they still have some substantial effects, particularly for near-term inflation and employment.

4.2 Causality of expectations shocks: news versus noise

As discussed in Section 2.2, the IRFs shown above describe what happens to the economy in the aftermath of a policy-expectations shock, but they do not necessarily imply that the expectations shock *causes* the entire responses. In the period of the shock itself, this is not an issue—any movement must be due to the effects of changes in expectations, because nothing else has had time to happen yet. However, some of the changes subsequent to the period of the shock could have taken place even if they had not been anticipated. Therefore, in order to isolate the causal effect of the expectations shocks beyond the first period, we need to purge the IRFs of the response that would have materialized anyway.

To do this, we exploit the distinction made earlier between news (anticipated policy changes that actually occur) and noise (anticipated policy changes that do not occur). We consider two extreme cases in which our identified shocks consist either entirely of news or entirely of noise. If the the expectations shocks that we have identified entirely reflect noise, then the IRFs depicted in Figure 2 only reflect causality—since the shocks do not embed any fundamental changes, any response of the economy must arise only from the effects of the shift in beliefs. On the other hand, if the identified shocks entirely reflect news, then the IRFs depicted in Figure 2 are equal to the effect of the change in beliefs *plus* the effect of the subsequent policy change itself. Thus, to isolate the part of the IRF that reflects the causality of expectations in that case, we must subtract the effect of a conventional policy shock of the anticipated size. Fortunately, we also have an estimate of exactly this object.

One complication is that, because of the way the survey data are reported, our measures of subjective expectations are averages over several periods. Thus, there are generally multiple possible expected short-rate paths that would be consistent with any initial expectations shock, and consequently the appropriate series of conventional policy shocks to use in the above-mentioned subtraction is not uniquely determined. (For example, a decline of 25 basis points in the expected one-year average rate could be consistent with a path of short rates that is 25 basis points lower for the entire year or with a path that is unchanged over the first six months and 50 basis points lower over the last six months.) For the purposes of this exercise, we assume shocks that would be sufficient to generate a constant short rate at the anticipated average level over the forecast period. Reasonable variations on this choice make little quantitative difference.

The "all noise" and "all news" cases are obviously extreme. In reality, the expectations shocks that we have identified likely reflect a combination of both noise and news. Intuitively, assuming convexity, as in the linear model of Section 2.1, the true effects of expectations shocks must lie somewhere between these two polar cases that we can actually compute. Consequently, if these two cases are empirically close to each other, we will have a fairly precise idea of the causality of expectations shocks, and the distinction between news and noise will not be particularly relevant.

Figure 3.A shows the upper (blue) bounds for the case of fully noise and the lower causality bounds (red) for the case of fully news for our model using one-year expectations, given an expectations shock of -25 basis points. The difference between news and noise affects our interpretation of the GDP IRF by at most 0.4 percentage points, at medium horizons. At short horizons, and at all horizons for inflation, there is virtually no difference. This is not surprising since we already knew from the estimated responses plotted in panel A of Figure 2 that unanticipated monetary-policy shocks (e_t) have modest effects relative to the policy-expectations shocks (η_t) and therefore, in the case of news, the cumulative impact that needs to be removed is fairly small. Figure 3.B shows the same comparison using the six-year expectations. (The 11-year results are very similar to the 6-year and are omitted from here on.) Again, the bounds are generally quite close to each other quantitatively and are never statistically different. Thus, we conclude that our procedure identifies the effects of policy-expectations shocks—in the causal sense—within a fairly tight range.

5 Modeling forward rate guidance

In this section, we use the results of our model to consider the effects of credible, Odyssean forward guidance—an initial shock to expectations of future short-term rates that is followed by deviations from the policy rule that are sufficient to make the expectations materialize. This method of constructing policy scenarios as combinations of different fundamental structural shocks over several quarters is also employed, for example, by Mountford and Uhlig (2012) in the context of fiscal policy.

Apart from some timing differences, the forward-guidance scenarios we consider here are conceptually the same as those we analyzed in the theoretical model of Section 2.3 (depicted in panels C and D of Figure 1). In particular, we simulate a shock to h-period expectations in period 0, followed by conventional monetary-policy shocks in periods 1 through h that cause the initial change in short-rate expectations to be exactly correct. As above, we assume that the Fed maintains the short rate constant at the level it has announced over the entire period covered by the forecast.

The blue lines in Figure 4, panel A show the effect of a -25-basis-point forward-guidance shock over a one-year horizon, together with $10-90$ percent credible regions. Again, in the period of impact, the responses are the same as those estimated for the expectations shock and pictured in Figure 2. However, in that figure, the actual short rate tightens following the expectations shock. It is precisely this outcome that we neutralize here. To do so, we must impose that the Fed hits the economy with multiple stimulative shocks to the actual short rate following the shift in the expected rate, just as we did in the theoretical model of Section 2.3. As a consequence, the responses of GDP and inflation are generally higher than those that would be generated by a -25-basis-point expectations shock alone. Indeed, the forward guidance raises GDP in the short run by over 1 percent and the price level by nearly 1 percent.

The initial responses of output and inflation to forward guidance are quantitatively similar to those produced by the NK model. (See Figure 1, panel D.) There, the output gap rose by about 1.5 percentage points in response to a one-year forward-guidance shock of -25 basis points, and inflation rose by about 3.5 percent. The estimated increase in the price level is essentially a one-time jump, which, since the data are quarterly, translates into an annualized inflation rate of about 3.7 percent. More importantly, however, while the empirical price level is little changed after the first quarter (and even declines a bit), in the theoretical model inflation continues to be positive for several quarters. Thus, while the response on impact is similar, the theoretical model predicts a much higher *cumulative* change in inflation. The output gap in the theoretical model is much less persistent than our empirical response, but this is not surprising because the theoretical model that we considered did not have much of a built-in persistence mechanism, which, for example, could be introduced through habit formation.

To assess whether forward guidance is itself effective, we need to compare it to what would have happened if the Fed had pursued the same short-rate policy without announcing it in advance. The red lines in Figure 4 plot a simlulation of what would have occurred if the Fed had followed the same short-rate path that it maintained under the forward-guidance scenario but with no preceding change in expectations. (Note that this implies a different series of conventional policy shocks than were used to construct the blue line.) Looking at the difference between the red and the blue lines, the marginal effects of 25-basis-point forward guidance on both GDP and prices are as much as 1 percentage point in the short- to medium-run and about 0.5 percent after ten years, with the differences being statistically significant for this entire period. The marginal effect on hours is also about 1 percent in the short run, although it decays somewhat faster. To the best of our knowledge, the importance of this expectations channel in the data has not previously been documented.

Panel B shows the same type of forward-guidance scenario and comparison for the model based on the six-year survey data. In this case, the experiment considered is the rather heroic one of the Fed announcing a credible 25-basis-point reduction in the short rate for the next six years and then following through on that promise. Despite the much longer horizon of this forward rate guidance, the effect relative to adopting the same short-rate path without pre-announcing it is still about 1 percent on GDP and about 0.5 percent on the price level, with neither statistically significant after two to three years. The marginal effect on hours is similar to that in panel A. These results contrast with the New Keynesian simulations depicted in Figure 1, where the impact of forward guidance was stronger at longer horizons. The discrepancy we document between the model and the data in this respect can be viewed as quantifying the ”forward-guidance puzzle.”

6 Robustness Checks

In this section, we conduct a battery of robustness checks, whose results are summarized in Table 5. Since our most significant results were for one-year expectations, we focus on that horizon. The results using the 6- and 11-year surveys are not shown but are also generally robust. As a summary measure, for each specification the table reports the estimated effect of the forward guidance scenario, relative to the same path of short rates when it is not announced in advance—i.e., the difference between the blue and red lines in Figure 4A. We report the median magnitude of this difference across parameter draws after both one and five years. Asterisks indicate statistical “significance” at the 5% level—that is, whether at least 95% of the parameter draws result in a positive marginal effect of forward guidance.

Since the effective zero bound on nominal interest rates could be an important source of non-linearity in our sample period, our first robustness check aims at verifying the stability of our results to the exclusion of the ZLB period. The pre-ZLB sample is also important to consider in order to show that our results are not just driven by the relatively short period in which forward rate guidance has been most actively used as a policy tool. As shown in the second row of Table 5, when we re-estimate the VAR using only the pre-ZLB period, the estimated impact of forward guidance on inflation and output is almost unchanged, and the effect on hours is only modestly smaller. These results are in contrast to a number of other VAR studies that exhibit apparent structural breaks at the ZLB (e.g., Baumeister and Benati, 2013). A likely explanation for the robustness of our results in this dimension is that direct measures of expectations help with the stability of the reduced-form parameters in the presence of nonlinearities, because those expectations are not required to be linear functions of the data, even though the model itself is linear.

Second, we estimated alternative specifications of the VAR. In particular, we increased the number of lags, we used forecasts from the SPF forecasts instead of the BCS, and we included both the one- and six-year horizon BCS forecasts within the same VAR. As shown in rows 3, 4, and 5 of Table 6, none of these changes altered the results significantly; although, in some cases the statistical significance is weaker at the five-year horizon. When we use the SPF data, overall, the impact of forward guidance seems to be a bit larger across all three key variables. This result could be due in part to the extension of the sample back to 1981.

Third, we also considered different identification schemes for both the expectations shock and the conventional policy shock. In the first case, we imposed the sign restriction on the real yield rather than the nominal yield. (The real yield is calculated as the one-year survey forecast of the average 3-month TBill minus the one-year forecast of CPI inflation.) In the second case, instead of using standard timing restrictions, we followed Uhlig (2005) by imposing that policy shocks lower the short rate and raise the CPI for at least five quarters. As shown in rows 6 and 7, these alternative identifications left the results very little changed.

Finally, instead of using a flat prior on the VAR parameters we used the Minnesota prior (row 8), and again we found very similar results, except that for CPI the cumulative impact of forward guidance becomes a bit larger in magnitude at both horizons.

7 Conclusion

In this paper, we used a survey-augmented VAR with sign restrictions to identify the effects of anticipated monetary policy on the macroeconomy. We found that, at a one-year horizon, accommodative monetary policy expectations shocks lead to large and rapid increases in both GDP and inflation. We argued that most of this response is likely causal. At longer horizons, the effects of these shocks are smaller and not always significant. These results indicate that forward-guidance policies can potentially be quite effective and that they are likely to have the greatest impact when targeted at shorter horizons.

Our results both support and challenge the conclusions of standard New Keynesian models that have been used to argue for forward-guidance policy. On the one hand, we do show that the anticipation channel exist in the data and that, consequently, forward-guidance can have large and immediate effects. On the other hand, those effects are smaller than most theoretical models would predict, particularly with respect to inflation. They also decrease with the horizon of the guidance, in contrast to New Keynesian predictions. Modifications like those proposed in McKay et al. (2015) may help to bring the New Keynesian models into closer alignment with what our results suggest.

Finally, while we think our results are informative for the debate about monetary-policy tools at the zero-lower bound, some caution is warranted in interpreting them in that context. Although expectations for short-term rates are clearly shaped in part by FOMC communications, forward rate guidance has also not historically been a prominent policy instrument. Moreover, FOMC communications about future short-term rates may not be viewed as credible, particularly if they are expressed as Committee forecasts rather than as commitments, and, in this case, they may even have perverse effects on expectations as suggested by Campbell et al. (2012). Thus, while our findings show that forward guidance can be a powerful policy tool under the right conditions, a variety of institutional impediments may dampen its efficacy in practice.

References

Arias, J. E., Rubio-Ramirez, J. F., and Waggoner, D. F., 2014. "Inference Based on SVARs Identified with Sign and Zero Restrictions: Theory and Applications." Federal Reserve Board, FEDS working paper 1100 (April).

Barakchian, Mahdi S., and Christopher Crowe, 2013. "Monetary policy matters: Evidence from new shocks data." *Journal of Monetary Economics*, 60, 950-966.

Barsky, Robert B., Susanto Basu, and Keyoung Lee, 2014. "Whither News Shocks?". 2014 *NBER Macroeconomics Annual*.

Barsky, Robert B., and Eric. R. Sims, 2012. "Information, Animal Spirits, and the Meaning of Innovations in Consumer Confidence," *American Economic Review*, 102(4), 1343-1377.

Baumeister Christiane and Luca Benati, 2013. "Unconventional Monetary Policy and the Great Recession: Estimating the Macroeconomic Effects of a Spread Compression at the Zero Lower Bound." *International Journal of Central Banking, Vol. 9 No. 2, 165-212.*

Beaudry, P., and F. Portier, 2006. "Stock prices, news and economic fluctuations," *American Economic Review*, 96(4), 1293-1307.

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

Bullard, James, George W. Evans, and Seppo Honkapohja, 2008. "Monetary Policy, Judgment, and Near-Rational Exuberance," *American Economic Review*, 98(3), 1163-77.

Campbell, Jeffrey R., Chalres L. Evans, Jonas D.M. Fisher, and Alejandro Justiniano, 2012. "Macroeconomic Effects of FOMC Forward Guidance." *Brookings Papers on Economic Activity*, Spring.

Carlstrom, C. T., Fuerst, T. S., and Paustian, M., 2012. "How Inflationary Is an Extended Period of Low Rates?" FRB Cleveland Working Paper 12-02.

Chun, A.L, 2011. "Expectations, bond yields and monetary policy." *The Review of Financial Studies*, 24 (1), 208-247.

Del Negro, Marco, Marc Giannoni, and Christina Patterson, 2012. "The Forward Guidance Puzzle." Federal Reserve Bank of New York, Staff Report No. 574, May.

Eggertsson, Gauti B. and Michael Woodford, 2003. "The Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity*, Spring.

Gertler. M., and Karadi, P., 2014. "Monetary Policy Surprises, Credit Costs and Economic Activity." NBER working paper 20224 (June).

Gurkaynak Refet S, Brian Sack, and Eric Swanson, 2005. "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary." *International Journal of Central Banking*, Vol. 1 No. 1, 55-93, May.

Gurkaynak Refet S, Brian Sack, and Jonathan Wright, 2007. "The U.S. Treasury yield curve: 1961 to the present." *Journal of Monetary Economics* 54(8): 2291-304.

Kiley, Michael T., 2014. "Policy Paradoxes in the New Keynesian Model," Finance and Economics Discussion Series 2014-29. Board of Governors of the Federal Reserve System.

Kim, Don and Orphanides, Athanasios, 2012. Term Structure Estimation with Survey Data on Interest Rate Forecasts. *Journal of Financial and Quantitative Analysis* 47(1): 241-72.

Krugman, Paul R., 1998. "It's Baaack: Japan's Slump and the Return of the Liquidity Trap." *Brookings Papers on Economic Activity*, 29, 137-206.

Laseen, S. and Svensson, L. E. O., 2011. "Anticipated Alternative Policy Rate Paths in Policy Simulations." *International Journal of Central Banking* 7(2): 1-35.

Leduc, S. and Sill, K., 2011. "Expectations and Economic Fluctuations: An Analysis Using Survey Data." *Review of Economics and Statistics* 95(4): 1352-67.

Leeper, E. M., Sims, C. A., and Zha, T., 1996. "What Does Monetary Policy Do?" *Brookings Papers on Economic Activity*, vol. 27(2), 1-78.

Leeper, Eric M., Todd B. Walker, and Shu-Chun Susan Yang, 2013. "Fiscal Foresight and Information Flows." *Econometrica* 81(3): 1115-45.

Levin, Andrew, David López-Salido, Edward Nelson, and Tack Yun, 2010. "Limitations on the Effectiveness of Forward Guidance at the Zero Lower Bound." *International Journal of Central Banking*, vol. 6(1),143-189, March.

Lorenzoni, Guido, 2009. "A Theory of Demand Shocks." *American Economic Review* 99(5): 2050-84.

Lucas, Robert E., 1972. "Expectations and the Neutrality of Money." *Journal of Economic Theory* 4(2), 103-124.

Mankiw, N. Gregory and Ricardo Reis, 2002. "Sticky Information Versus Sticky Prices: A Proposal to Replace the New Keynesian Phillips Curve." *Quarterly Journal of Economics* 117(4), 1295-1328.

McKay, Alisdair, Emi Nakamura, and John Steinsson, 2015. "The Power of Forward Guidance Revisited." NBER working paper 20882, January.

Milani, Fabio, 2011. "Expectation Shocks and Learning as Drivers of the Business Cycle." *The Economic Journal*, 121, 379-401.

Milani, Fabio, and John Treadwell, 2012. "The Effects of Monetary Policy "News" and "Surprises". *Journal of Money, Credit, and Banking* 44(8): 1667-92.

Mountford, Andrew and Harald Uhlig, H., 2009. "What Are the Effects of Fiscal Policy Shocks?" *Journal of Applied Econometrics*, 2009, 24(6), 960-92.

Orphanides, Athanasios and Wei, Min, 2012. "Evolving macroeconomic perceptions and the term structure of interest rates," *Journal of Economic Dynamics and Control*, vol. 36(2), 239-254.

Piazzesi, Monika, Juliana Salomao, and Martin Schneider, 2015. "Trend and Cycle in Bond Premia." Working Paper, March.

Ramey, V. A., 2011. "Identifying Government Spending Shocks: It's all in the Timing." *The Quarterly Journal of Economics*, 126(1), 1-50.

Romer, Christina D., and David H. Romer, 2010. "The Macroeconomic Effects of Tax Changes: Estimates Based on a New Measure of Fiscal Shocks." *American Economic Review* 100 (June), 763–801.

Rudebusch, Glenn. D., and John C. Williams, 2006. "Revealing the secrets of the temple: the value of publishing central bank interest rate projections." NBER working paper 12638, October.

Schmitt-Grohe, S., and M. Uribe, 2012. "What's news in business cycles." *Econometrica*, 80(6), 2733-2764.

Sims, Christopher A., 2003. "Implications of Rational Inattention," *Journal of Monetary Economics* 50(3), 665-690.

Uhlig, H., 2005. "What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure." *Journal of Monetary Economics* 52, 381-419.

Werning, Ivan, 2011. "Managing a liquidity trap: Monetary and fiscal policy." NBER, working paper, No. 17344, August.

Table 1. Baseline identification restrictions on contemporaneous impact of shocks

Block / Variable		Shock: Policy expectations (η)	Shock: Conventional policy (e)
$x1$	GDP	?	0
	CPI	?	0
	Labor productivity	?	0
$x2$	M2 growth	?	?
	Corporate profits	?	?
	Longer-term yield	?	?
i	3m T-Bill rate	+	–
$E^s[x_{t+h}]$	Survey GDP	+	?
	Survey CPI	+	?
	Survey 3m T-Bill	–	?

Notes: The table shows the restrictions imposed to identify structural shocks in the baseline VAR.

Table 2. Response of survey forecasts to expectations shocks

Change in expectations of…	Shock to 1-year expectations	Shock to 6-year expectations	Shock to 11-year expectations
3m T Bill	-0.03%	-0.03%	-0.02%
Log GDP	0.0014	0.0018	0.0028
Implied expected annual growth	*0.14%*	*0.03%*	*0.03%*
Log CPI	0.0009	0.0016	0.0022
Implied expected annual inflation	*0.09%*	*0.03%*	*0.02%*
Implied change in growth following 25 bp exp. shock	*1.22%*	*0.26%*	*0.26%*
Implied change in inflation following 25 bp exp. shock	*0.79%*	*0.23%*	*0.20%*

Note: Based on VARs using Blue Chip Survey under the baseline identification. Table shows responses to a one-standard-deviation shock to policy expectations at different horizons. Responses are reported for survey-based expectations at the same horizon as the shock in the period when the shock occurs.

Table 3. Largest expectations shocks identified in the VAR, 1999 - 2008

Std. Dev.	**Date**	**FOMC Events**
		Expected-Easing Shocks
-2.5	2000Q3	"Expansion of aggregate demand may be moderating"
-2.1	2001Q3	Sept. 11 attacks
-1.1	2001Q1	Balance of risks shifted to downside; easing cycle begins
-1.1	2006Q1	"Some further policy firming may be needed" (rather than likely)
-1.1	2002Q3	Balance of risks shifted to downside
-1.0	2004Q1	"Committee believes it can be patient…"
-1.0	2006Q3	Removal of phrase "some further policy firming may yet be needed"; "Economic growth has moderated"
-1.0	2008Q1	"Economic growth is slowing… Recent developments... have increased the uncertainty surrounding the outlook"; 75 bp intermeeting cut and downside risks
		Expected-Tightening Shocks
1.9	2005Q4	"Committee judges that some further policy firming is likely" (removed "measured pace" language)
1.5	2001Q4	
1.5	2004Q3	Started tightening cycle
1.5	2005Q2	"Pressures on inflation have picked up in recent months", changed balance of risks from "roughly equal" to "should be kept roughly equal" with "appropriate monetary policy"
1.4	2006Q4	
1.3	2002Q2	"Economy is expanding at a significant pace," downside balance of risks removed
1.2	2000Q2	50bp tightening. "The Committee is concerned that this disparity in the growth of demand and potential supply will continue, which could foster inflationary imbalances."
1.2	2007Q1	"Committee's predominant concern remains the risk that inflation will fail to moderate."

Notes: The table shows the quarters in which the largest expectations shocks occurred during the period 1999 – 2008, as identified by the VAR model. The right-hand column reports changes in the FOMC statement whose timing corresponded to the timing of those shocks.

Table 4. Size of conventional policy shocks required to equal the effect of a -25 bp one-year expectations shock

Horizon	Equal cumulative effect on GDP	Equal cumulative effect on CPI	Equal cumulative effect on hours
1Y	-123	--	-210
2Y	-46	--	-61
4Y	-18	-81	-4
8 Y	-40	-54	-6

Notes: The table shows the size of an exogenous shock to the short rate (a "conventional" policy shock) that would be necessary to equal the effect of a -25bp shock to one-year expectations for the short rate on each of the indicated macro variables. One- and two-year effects on the CPI are omitted due to the negative estimated sign of the conventional policy shock at those horizons.

Table 5. Marginal effects of one-year forward guidance under alternative specifications

	GDP		CPI		Hours	
	1Y	5Y	1Y	5Y	1Y	5Y
Baseline	1.2%*	0.9%*	0.9%*	0.7%*	1.3%*	0.6%*
Pre-ZLB period	1.2%*	1.0%*	0.9%*	0.9%*	0.8%*	0.3%
More lags	1.3%*	1.2%	0.9%*	0.7%	1.4%*	0.6%
SPF instead of BCS (begins 1981)	1.6%*	1.1%*	1.4%*	1.1%*	1.5%*	0.9%*
1Y and 6Y surveys both included	1.0%*	0.7%	0.9%*	0.7%*	0.7%*	0.1%
Expectations shocks use sign restriction on real yield	1.3%*	1.0%*	0.9%*	0.8%*	1.3%*	0.6%
Policy shocks identified by sign restrictions	1.0%*	0.7%	1.0%*	0.7%	1.0%*	0.4%
Minnesota prior	1.0%*	1.2%	1.4%*	1.1%*	1.4%*	0.8%

Notes: Under various modeling alternatives, the table reports simulations of a "forward guidance" policy that commits to maintain the short-term interest rate 25 basis points below the value prescribed by the policy rule for the next four quarters, relative to a situation in which the central bank adopts this same path for the short rate without announcing it in advance. The columns show the difference between the effects of these two policies on macroeconomic variables one and five years after the forward-guidance announcement. See Section 5 of the text for details on the construction of the scenario in the baseline case. GDP and CPI are reported in terms of cumulative log levels. Hours are reported as forward (non-cumulative) log levels. Numbers in the table are medians across 10,000 parameter-vector draws. Asterisks indicate that the marginal effects of forward guidance are positive for at least 95% of the draws.

Figure 1. Expectations and Forward-Guidance Shocks in the New Keynesian Model

A. Noise shocks (1 period)

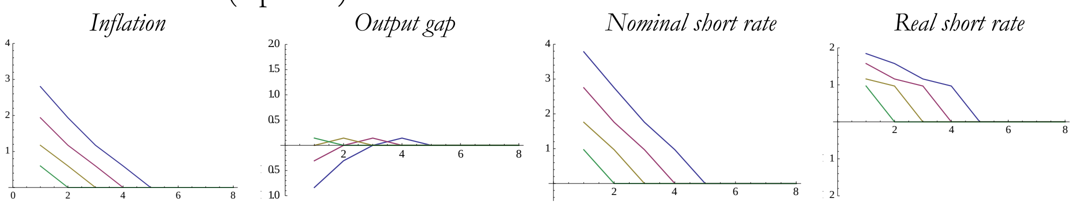

B. News shocks (1 period)

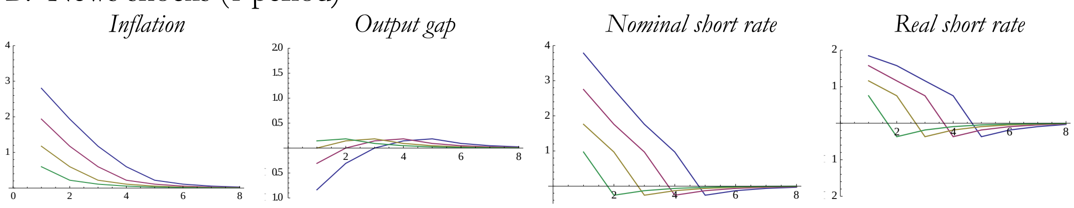

C. Forward-guidance scenario (1 period)

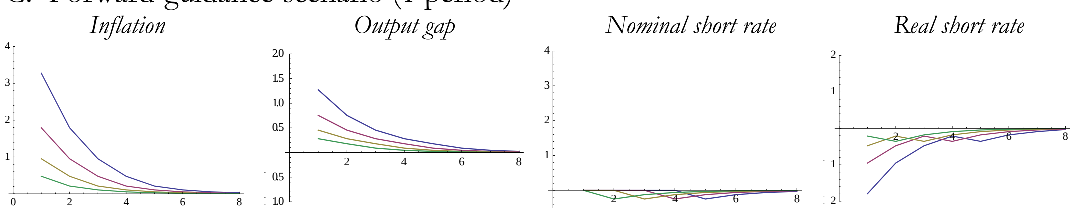

D. Forward-guidance scenario (multiple periods)

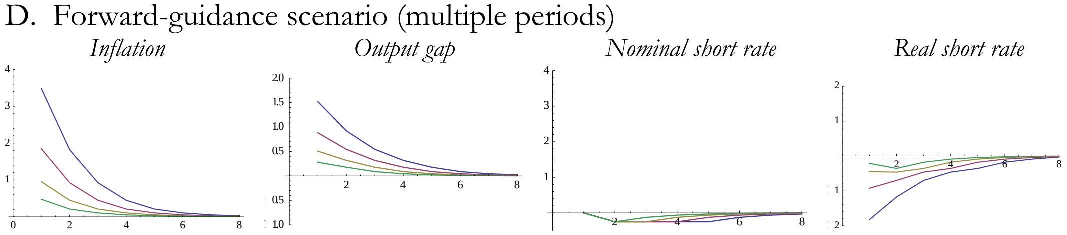

Notes: The figures show the effects of a shock to subjective expectations for monetary policy that is sufficient to lower expectations of the short-term interest rate h periods ahead by 25 basis points, where h= 1 (green), 2 (yellow), 3 (red), and 4 (blue). In panel A, we the anticipated shock to policy does not materialize. In panel B, the anticipated shock does materialize and the central bank takes no other actions. In panel C, the anticipated shock materializes and in addition central bank introduces anticipated policy shocks in order to maintain the nominal short rate at zero in periods t through t+h-1. In panel D, the central bank holds the short rate at -.25 in periods t+1 through t+h. Responses are shown for the period of the expectations shock and seven subsequent periods. Calibration of the model is as described in the text. All variables are in percentage points. Inflation and interest rates are expressed as annual rates.

Figure 2. Impulse-response functions

A. Using 1-year expectations

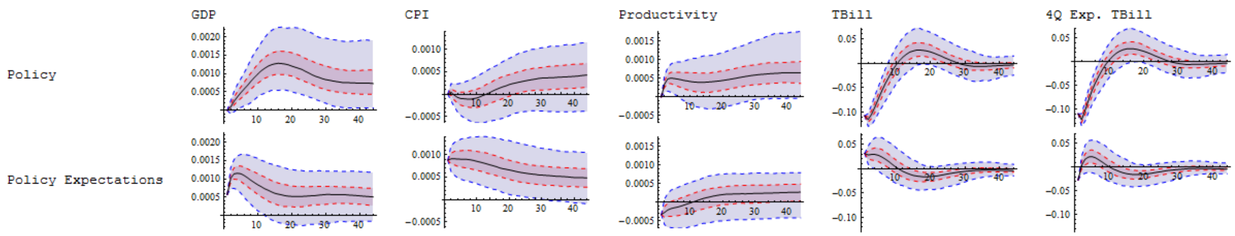

B. Using 6-year expectations

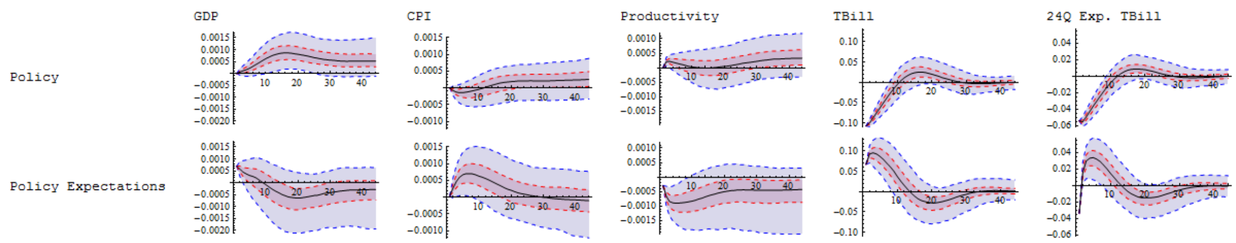

C. Using 11-year expectations

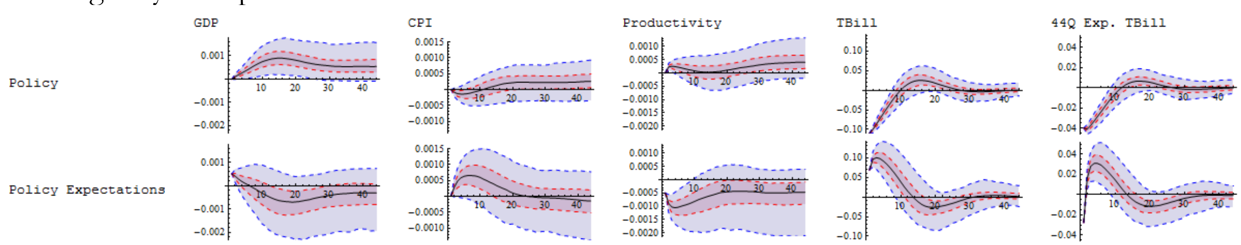

Figure 3. Causality bounds on 25-basis-point expectations shocks

A. Using 1-year expectations

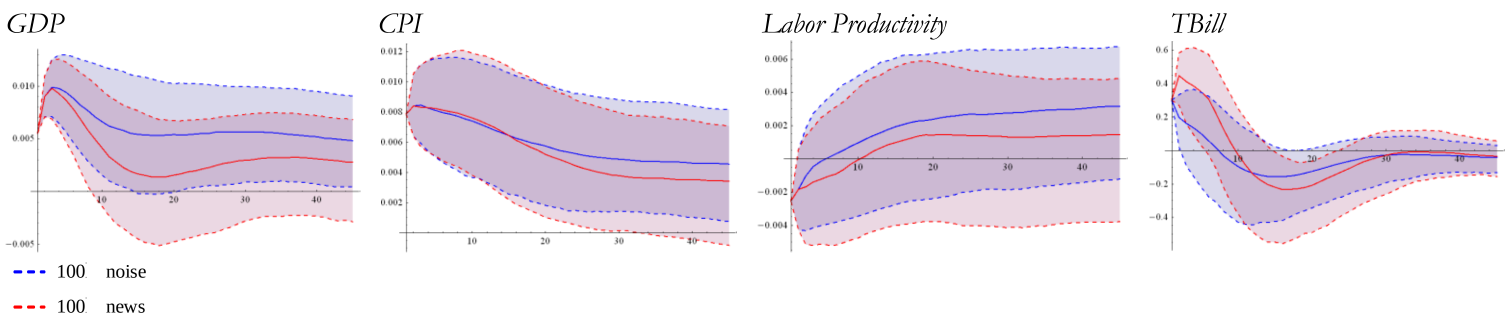

B. Using 6-year expectations

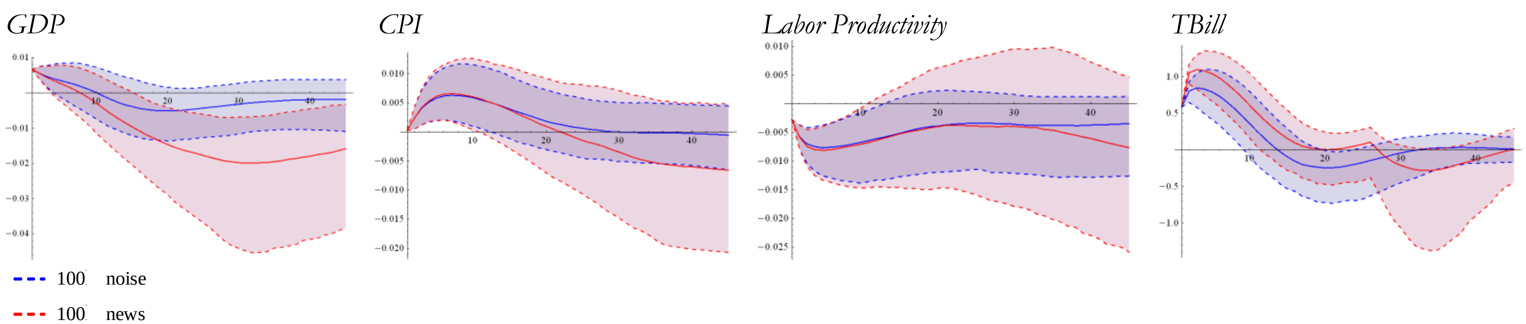

Figure 4. The effects of "forward guidance" scenario

A. One-year forward guidance

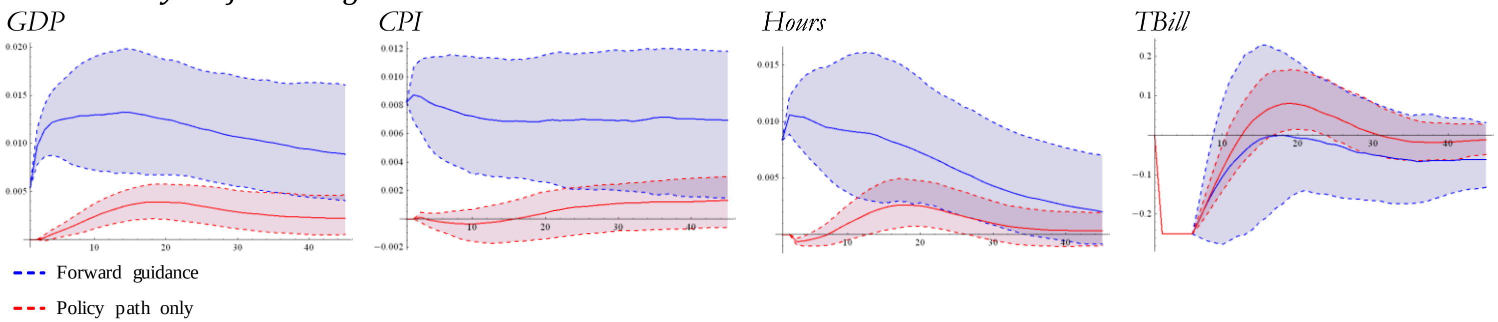

B. Six-year forward guidance

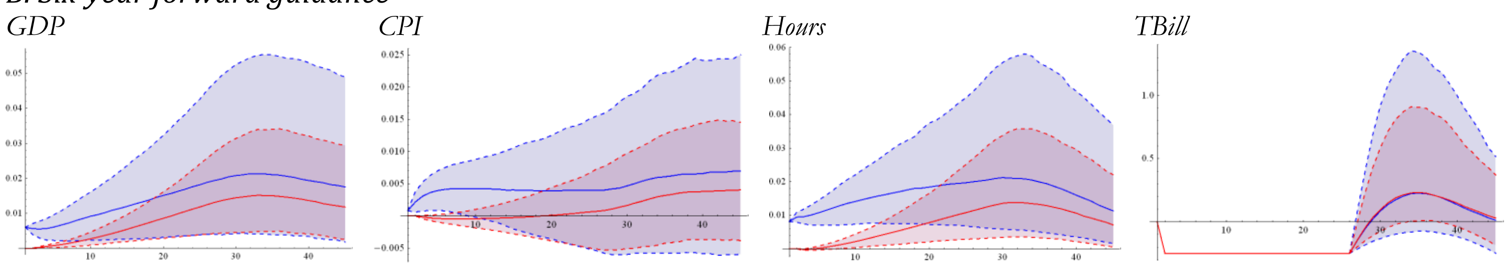

Appendix: Treatment of the Survey Data

We employ survey data from two sources: the Blue Chip Survey (BCS) and the Survey of Professional Forecasters (SPF). The principal advantages of the SPF data are that they begin in 1981 (the year when the three-month Treasury bill rate forecast becomes available) and are reported at a consistent quarterly frequency. However, the longest available forecasting horizon in these data is one year ahead. The BCS data, by contrast, include forecasts of up to 11 years in the future, but they do not begin until 1983 and for some forecasting horizons are reported only twice a year at a slightly irregular interval.

The SPF in quarter t asks respondents for their forecasts in quarters t-1 through t+4. We thus have one-year forecasts reported quarterly from 1981:4 through 2014:3, as well as "nowcasts" of the contemporaneous data and "backcasts" of the lagged data. The main issue we face with these data is transforming the reported forecast growth rates into levels, which we require for our VAR. Although the SPF does ask for GDP and CPI forecasts in terms of levels, this is not always useful to the researcher ex post because re-benchmarking introduces discrete breaks in the series. To obtain consistent series we assume that the average survey backcast of quarter t-1 is correct in the sense that any difference between this value and the revised value we observe in the most-recent data is due entirely to rebenchmarking and does not reflect any fundamental change in agents' beliefs about the economy. By then applying the reported SPF growth rates for the subsequent five quarters, we obtain a forecast for the t+4 levels of GDP and CPI that are based on the same indexation as the 2014 data. Finally, for each quarter, we average the t+1 through t+4 forecasts of the T-Bill rate to obtain forecasts of the average T-Bill rate over the following year.

The same difficulty with benchmarking applies to the BCS, but there we face the added complication that we do not have a backcast for t-1. Therefore, to index the level in the BCS data, we assume that BCS respondents have the same estimate of the quarter-t data level as the SPF respondents. (This is likely a reasonable assumption, given that, as shown below, the SPF and BCS data are generally quite similar in other respects.)

Apart from this, the BCS data on one-year expectations are reasonably straightforward, and we construct one-year expectations by averaging the forecasts for quarters t+1 through t+4 in the last month of each quarter. However, the BCS data also include forecasts at longer horizons, and these involve complications related to the timing and scope of their reporting. To obtain as much consistency as possible from this information, we build a new dataset of long-term expectations from the BCS at a quarterly frequency from 1983 to 2014.

Specifically, since 1983, the BCS has been providing semiannual long-range (2- to 6-year and 7- to 11-year) consensus forecasts for various interest rates, including the 3-month T-Bill rate, as well as real GDP, GDP deflator, and CPI. These long-range consensus forecasts were originally provided every March and October in both the Blue Chip Economic Indicators (BCEI) and the Blue Chip Financial Forecasts (BCFF). Starting in 1996, the BCFF switched to providing these long-range projections in June and December, while the BCEI continued reporting them in March in October. We thus have observations of long-term expectations of our main variables of interest twice per year prior to 1996 and four times per year after that time. These inconsistent frequencies and the fact that the observations are not equally spaced across the year mean that we cannot use these data directly in the VAR.[1]

We address both of these issues through interpolation. Specifically, from 1983 to 1996, when the long-range forecasts are available only in March and October, we use the results from the BCEI and linearly interpolate to obtain June, September, and December values. Once the June and December values become directly observable, we interpolate to obtain only the September value. Interpolation was not necessary for the one-year horizon, because those are available on a monthly basis from the BCEI. Once we have adjusted the timing in this way and computed survey expectations for the average values of variables in the first year following the survey, it is possible to compute medium-term expectations---that is, the expected average value over the next 6 years---by taking the weighted average of the one-year and 2-6-year expectations, and long-term expectations---that is, the expected average value over the next 11 years---by taking the weighted average of the one-year, 2-6-year, and 7-11-year expectations, respectively.

The three panels of Figure A1 plot the resulting time series of the survey-based expectations of the average 3-month T-Bill rate, CPI inflation, and GDP growth over the next year. The projections of the 3-month T-Bill rate and CPI inflation are very similar between the two surveys. In the case of GDP growth, on the other hand, the SPF projections are more volatile and, at least through about the year 2000,

[1] In the pre-1996 part of the sample, the availability of long-range forecasts in BCFF for the same months of BCEI allowed us to compare projections for the variables in common across the two Blue Chip surveys. We found that differences in forecasted values were very small, indicating that it was not inappropriate to splice together the results from both surveys from 1996 to 2014.

more pessimistic than the BCS projections. Figure A2 illustrates the properties of the term structure of BCS forecasts for the same set of variables (3-month T-Bill, CPI inflation, and GDP growth) by plotting their time series at the one-, six-, and 11-year horizon. Shorter-term expectations (blue lines) display much more variation than longer-term expectations, and there is very little difference between 6- and 11-year projections (red and green lines, respectively). These results are consistent with the stylized fact that it is difficult to forecast economic variables far in the future.

Figure A1. Comparison of SPF and BCS one-year forecasts

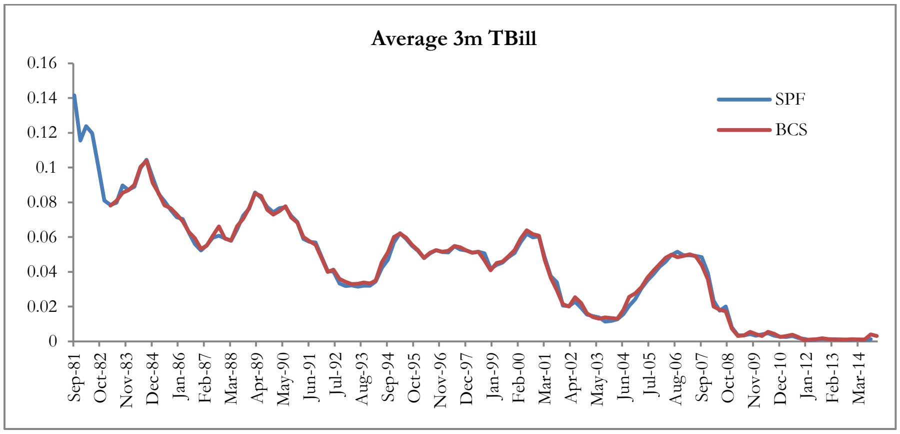

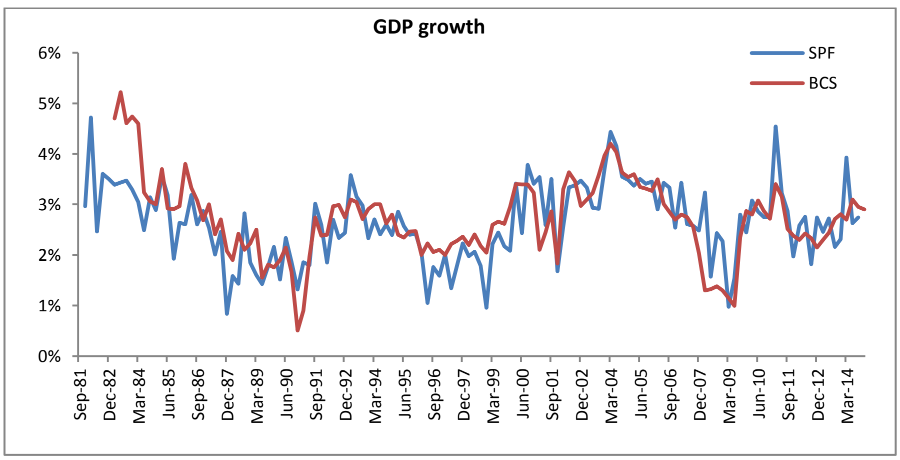

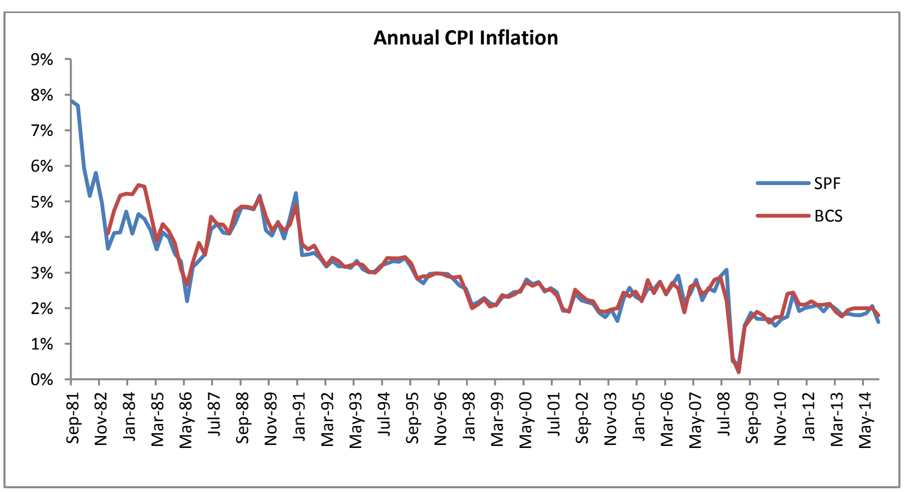

Figure A2. Term structure of BCS forecasts