

What Does Anticipated Monetary Policy Do?

Stefania D'Amico* Thomas B. King*

April 5, 2017

Abstract

Applying sign and zero restrictions to survey forecasts embedded in a VAR, we study the economic effects of news about future monetary policy—the type of shock induced by credible "forward guidance." We find that such policy has large, immediate, and persistent effects on inflation and real activity, that these effects are larger than those of unanticipated monetary policy, and that the economic responses grow larger as the horizon of the guidance moves farther into the future. Our results also suggest that conventional monetary-policy shocks themselves are effective only because they shift interest-rate expectations.

*Federal Reserve Bank of Chicago. Contacts: stefania.damico@chi.frb.org and thomas.king@chi.frb.org. For helpful comments, discussions, and suggestions we thank Jonas Arias, Bob Barsky, Jeff Campbell, Marco Del Negro, Charlie Evans, Jordi Gali, Matteo Iacoviello, Alejandro Justiniano, Thomas Laubach, Stefano Neri, Daniel Rees, Eric Swanson, Harald Uhlig, Min Wei, and conference and seminar participants at the 2015 WEAI meetings, the 2016 AEA meetings, the 2016 International Research Forum on Monetary Policy, the Banque de France, the Bank of Canada, the Federal Reserve Board, the Federal Reserve Bank of Chicago, and the University of Tor Vergata. The views expressed here do not reflect official positions of the Federal Reserve.

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 to stimulate the economy. 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 accommodative 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, there is no conclusive evidence on whether this mechanism is actually present in the data. We address this question and quantify the dynamic effects of monetary-policy expectations and forward guidance on output and inflation.

Just as identifying a conventional monetary-policy shock requires capturing a deviation from the historical policy rule, identifying an *anticipated* monetary-policy shock requires capturing an *anticipated* deviation from the historical policy rule.² The type of shock that interests us in theory—and that is of interest to policymakers hoping to stimulate the economy through forward guidance—is one in which agents come to correctly anticipate policy easing beyond what the expected state of the economy would normally warrant. To identify this shock, we embed survey forecasts of the short-term interest rate and key macroeconomic variables in a structural VAR model and use the following sign restrictions: the expected short rate must move in the opposite direction of expected inflation and expected GDP. This pattern of changes in expectations is unique to anticipated monetary-policy deviations in theoretical models.

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

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

²The 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, most fluctuations in short-rate forecasts will reflect expectations of such systematic policy responses to future economic conditions. Previous evidence, such as Romer and Romer (2000) and Campbell et al. (2012), supports this notion: when people expect looser policy, it is usually because they expect the economy to deteriorate. If we did not control for this expected endogenous policy response, it would obscure the effects we are interested in. Changes in short-rate forecasts alone are not sufficient to isolate the exogenous component of policy, as the distinguishing feature is the corresponding change in expected economic conditions. Our sign restrictions, being contemporaneously imposed on three key forecasted variables (short rate, GDP, and inflation) and excluding all cases in which those move in the same direction, isolate the anticipation of exogenous policy innovations.

To ensure that our shocks can be interpreted as "news" about future monetary policy, our identification also requires that the survey forecasts correctly anticipate, on average, the economic fluctuations resulting from these shocks. While the interpretation of forward guidance as a news shock is not uncommon (e.g., Milani and Treadwell, 2012), the empirical identification of news about the monetary-policy rate is novel. The empirical literature on news shocks mostly identifies news about variables, such as technology (e.g., Beaudry and Portier, 2006; Barsky and Sims, 2011) and non-cyclical components of fiscal policy (e.g., Ramey, 2011; Mountford and Uhlig 2009), that are assumed to be exogenous—an assumption that certainly does not hold for the monetary-policy rate. Consequently, the identifying procedures developed in those papers are not applicable to our problem. Instead, our restrictions identify news by linking changes in survey forecasts with subsequent changes in economic data.

We find that, consistent with an FOMC that usually hews closely to a clearly communicated policy rule, anticipated deviations from that rule are small; nevertheless, they have large, immediate, and persistent effects on both prices and real activity. We estimate that a 10-basis-point decline in expectations for the average short-term rate over the following year—roughly the size of the shock associated with the FOMC's adoption of calendar-based forward guidance in August 2011—raises current output, prices, and hours by about 0.6 percent in the

near term. These responses are larger and occur faster than those produced by an identical path of the short rate that is unanticipated, but the effects of the unanticipated policy largely catch up after a few quarters. To perform this comparison, we identify unanticipated policy innovations using the same restrictions on the survey forecasts that we use when identifying anticipated policy innovations, but we require the actual short rate to decline contemporaneously. This novel identification method allows for the expectations channel to have a role in both types of shock, maximizing their comparability. However, our finding that anticipated policy moves are more effective than unanticipated ones also holds when we use more-standard identification schemes (Christiano et al., 1999; Uhlig, 2005) for the unanticipated policy shocks.

Finally, we construct a new quarterly dataset of very long-range survey forecasts and use those expectations to test how the magnitude of the effect of monetary-policy news depends on its horizon. We find that shocks to long-term policy expectations have larger effects than in the shorter-run cases. For example, based on our median estimates, anticipated policy innovations are roughly twice as powerful at the 11-year horizon as they are at the one-year horizon. The longer-horizon results are also more statistically significant than the shorter-horizon ones.

These findings support the basic mechanisms at work in the standard New Keynesian model. We confirm that news about future monetary-policy easing has immediate stimulative effects, that these are larger than the effects of similarly sized unanticipated policy shocks, and that the economic responses to monetary-policy news grow larger as the period it pertains to moves farther into the future. Although the magnitudes of our estimates—particularly for the very long-horizon news—are smaller than standard models would suggest, they are consistent with extensions, such as McKay et al. (2016), that introduce discounting in the IS curve and thereby dampen the effects of news at long horizons. We conclude, therefore, that credible forward guidance can likely provide significant economic stimulus, just as the theory suggests.

The reason that forward guidance is more effective than conventional policy in the short run is that interest-rate expectations adjust immediately to forward guidance but only sluggishly to conventional policy, and economic outcomes ap-

pear to depend primarily on these expectations. Indeed, we also explore the channels through which conventional policy operates, and we find that when an unanticipated policy shock occurs, it is not the change in the short rate *per se* that matters for economic outcomes, but rather the associated shift in expected short rates. We reach this conclusion by using our estimates to construct hypothetical scenarios in which there is a shock to today's short rate but expectations for future short rates remain unchanged. This type of unanticipated shock in which the expectations channel is shut down has no significant effect on GDP or inflation.

To our knowledge, these findings are the first to assess the macroeconomic effects of monetary-policy expectations and forward guidance 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. Standard versions of New Keynesian models deliver macroeconomic effects that are generally viewed as implausibly large (Del Negro et al., 2012; Carlstrom et al., 2012), and relatively minor variations, such as those explored in Levin et al. (2010) and Werning (2015), can yield impacts of policy expectations that differ markedly from each other. A number of papers have estimated versions of New Keynesian models that incorporate news or forward guidance about monetary policy (Milani and Treadwell, 2012; Del Negro and Schorfheide, 2013; Campbell et al., 2016; Gomes et al., 2017). Though informative, those papers face the usual limitation of DSGE models, which is that the results may hinge upon the particular structural specification. In contrast, our VAR-based approach imposes only a minimal set of restrictions.

In the only previous attempt to estimate the effects of forward guidance in a model-free way, Campbell et al. (2012) found the opposite of what theory predicts: when the Fed signals that lower rates are coming, survey expectations of GDP and inflation decline. They argued that this 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 endogenous policy responses to the economy versus expectations for exogenous policy innovations, a distinction that our methodology is specifically designed to address. Gurkaynak et al. (2005) and Gertler and Karadi (2015) also examined aspects of changes in the anticipated path of interest rates, but neither distinguished expectations for exogenous policy innovations from expectations for endogenous policy responses as we do, nor did they consider differential effects of expectations over different horizons.

In Section 2 of the paper, we motivate our identification strategy within a simple New Keynesian model in which forward guidance is thought of as a news shock, whose qualitative and quantitative implications provide a benchmark for our empirical results. Section 3 shows that our survey-augmented VAR is a valid reduced-form of that type of model and describes the VAR implementation. Section 4 details our structural identification scheme. Section 5 summarizes the baseline results. In Section 6, we use those results to construct scenarios that illustrate the importance of anticipation in monetary policy, including forward-guidance scenarios. Section 7 conducts a battery of robustness checks, and Section 8 concludes the paper.

2 Policy expectations in a New Keynesian model

We begin by considering what happens when we allow for news about future monetary policy in an otherwise standard New Keynesian (NK) model. The purpose of this exercise is threefold. First, it illustrates in a familiar setting the mechanism underlying the impact of anticipated monetary-policy innovations. Second, it demonstrates the qualitative and quantitative responses to such shocks implied by theory, providing hypotheses (such as the forward-guidance puzzle) to be tested in our empirical work. Lastly, it illustrates how the sign restrictions that will be used to identify anticipated policy innovations in our VAR are prescribed by the theory. The model has some overlap with previous studies of forward guidance and monetary-policy news (e.g., Eggertsson and Woodford, 2003; Werning, 2011; Milani and Treadwell, 2012; Campbell et al., 2012, 2016), but we emphasize the particular features that will be of interest for our empirical tests.

2.1 Model Description

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

$$\pi_t = \beta E_t \pi_{t+1} + \kappa y_t \quad (1)$$

$$y_t = E_t y_{t+1} - \frac{1}{\sigma} (i_t - E_t \pi_{t+1} - r^*) \quad (2)$$

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

$$i_t = \phi_y y_t + \phi_\pi \pi_t + v_t \quad (3)$$

where

$$v_t = \rho v_{t-1} + \varepsilon_t \quad (4)$$

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

We depart from the standard treatment only by allowing agents to have some knowledge about the policy innovation ε_t prior to period t . Agents potentially receive correct information (i.e., news) about this innovation every period before it occurs. Let a_t^{t+h} be agents' anticipation of the policy innovation ε_{t+h} as of period t , and let u_t denote the component of ε_t that is unanticipated as of period $t-1$. That is

$$\varepsilon_t = a_{t-1}^t + u_t \quad (5)$$

For any fixed period $T > t$, rational expectations implies that $\{a_t^T\}$ follows the martingale process

$$a_t^T = a_{t-1}^T + \eta_t^T \quad (6)$$

where η_t^T is a serially independent news shock. For the remainder of the paper we will refer to this shock as the "policy-expectations shock." Clearly, the standard

model without news is a special case in which $a_t^{t+h} = 0$ for all t and h .

Given this setup, one can solve forward to obtain a solution for time- t inflation, output, and the short-rate as a function of the current policy stance v_t and an infinite-order moving average of the future policy innovations that are expected as of time t :

$$\pi_t = \psi_{0,\pi}v_t + \sum_{h=1}^{\infty} \psi_{h,\pi}a_t^{t+h} \quad y_t = \psi_{0,y}v_t + \sum_{h=1}^{\infty} \psi_{h,y}a_t^{t+h} \quad i_t = \psi_{0,i}v_t + \sum_{h=1}^{\infty} \psi_{h,i}a_t^{t+h} \quad (7)$$

The effects of unanticipated shocks u_t in this model are standard and are the same regardless of whether the anticipated component exists or not. In particular, Appendix A shows that $\psi_{0,\pi}$ and $\psi_{0,y}$ are negative, meaning that time- t inflation and output move in the opposite direction of an unanticipated monetary-policy shock u_t . Under standard parameterizations, $\psi_{0,i}$ 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.

2.2 Responses to Policy-Expectations Shocks

In the model described above, when agents receive a shock η_t^{t+h} that alters the anticipated monetary-policy innovation h periods ahead, their expectations of what the economic variables will be in period $t+h$ react as follows:

$$\Delta E_t [\pi_{t+h}] = \psi_{0,\pi}\eta_t^{t+h} \quad \Delta E_t [y_{t+h}] = \psi_{0,y}\eta_t^{t+h} \quad \Delta E_t [i_{t+h}] = \psi_{0,i}\eta_t^{t+h} \quad (8)$$

Intuitively, expectations of future inflation, output, and interest rates are affected by expectations of a future monetary-policy change in the same way that current inflation, output, and interest rates are affected by a current unanticipated policy shock. Thus, a policy-expectations shock causes both $E_t [\pi_{t+h}]$ and $E_t [y_{t+h}]$ to move in the *opposite* direction of $E_t [i_{t+h}]$. This observation motivates our sign-based identification scheme in the VAR. Notably, no other shock in standard models of this type can produce this response pattern.³

³For 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.

The empirical question is, given a shock to today's expectations of policy h periods in the future, how does the economy respond today? The predictions of the NK model with respect to this question are given by the multipliers $\psi_{h,\pi}$, $\psi_{h,y}$, and $\psi_{h,i}$ in equation (7), whose closed-form solutions are provided in Appendix A. To get a sense of the magnitudes involved and how the expectational horizon matters, Figure 1 illustrates the immediate responses and the subsequent economic dynamics under a standard calibration. Specifically, 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 $\rho = .5$. We focus on policy-expectations shocks η_t^{t+h} that are sufficient to lower the expected h -period-ahead annualized interest rate by 25 basis points, where $h = 1, \dots, 4$ quarters.

As shown in panel A, inflation rises immediately in response to the policy-expectations shocks, and it rises by more the farther into the future those innovations are expected to occur. For an anticipated monetary easing of 25 basis points 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: the current value of the interest rate rises to offset the stimulative effects of its anticipated future declines. For this calibration, that policy response is large enough to drive the output gap negative in early periods, even though the shock itself is a stimulative one. In general in this economy, depending on the parameter values, output and inflation responses to policy-expectations shocks are ambiguous in the short run.⁴

However, if the path of the economy shown in panel A resulted from deliberate policy, the central bank would find itself in the somewhat bizarre position of mechanically raising rates in response to its own accommodative forward guidance. Since this seems implausible, we consider a second, more realistic scenario: at time t , the central bank announces that it will lower the policy rate by 25 basis points in $t + h$ but also that it will maintain the rate unchanged at its $t - 1$ level until that time. This is equivalent to introducing an unanticipated shock in

A "markup shock," which would appear as an additional stochastic term in equation (1), would generally move inflation and output in opposite directions.

⁴ Indeed, in Appendix A we show that those responses are only guaranteed to be positive if $\phi_\pi < 1/\beta$. Since β is typically close to one, this condition will nearly always be violated if the model satisfies the Taylor principle.

period t and a series of anticipated innovations in periods $t + 1$ through $t + h - 1$ that are sufficient to offset the systematic response of the short rate. As shown in panel B, once the mechanical policy response to expected future easing is shut down, the output gap rises substantially. Given the stabilized nominal short rate, the higher inflation results in a significant reduction in real rates.⁵ Panel C shows that a similar outcome occurs if the central bank promises to lower the short rate by 25 basis points, not just in period $t + h$, but for the entirety of the 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 using our empirical results in Section 6.

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 25-basis-point shock to the actual short rate (not shown) 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 one year ahead. 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, or its assumed policy rule. Rather, as discussed by McKay et al. (2016), 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 several authors 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 so far no empirical work estimating this impact in a model-free way.

Summing up, the stylized NK model shows that policy-expectations shocks, as we have defined them, move expected short-term interest rates in the opposite direction of expected future output and inflation. This will be a key identifying feature of our empirical approach. In terms of predictions, the model shows

⁵At the zero lower bound (ZLB), the situation depicted in panel B arises. From the perspective of the linear model, a short-rate that does not move in response to forward guidance because of the ZLB looks like an accommodative shock. Thus, the results in panel B reproduce those of previous theoretical models that have explicitly taken the ZLB into account (e.g., Eggertsson and Woodford, 2003; Werning, 2011; Del Negro et al., 2013).

that, if the systematic policy response to anticipated future innovations is not shut down, the immediate macroeconomic effects of policy-expectations shocks can go in either direction. Thus, the sign of these effects is an empirical question. However, if the systematic policy response to anticipated future innovations *is* shut down, the model also makes two firm, qualitative predictions about the effects of such shocks. First, policy-expectations shocks have immediate effects on output and inflation that are in the same direction of their expected future effects. Second, the effects of policy-expectations shocks grow larger as the horizon they pertain to moves farther into the future. Our empirical results provide tests of both predictions.

3 Empirical Specification

3.1 Reduced Form

Our empirical strategy identifies policy-expectations shocks as “news” shocks in a VAR that jointly models economic variables and survey forecasts, treating the latter as direct measures of agents’ expectations. Survey forecasts have been previously used to measure expectations in numerous empirical macro/finance applications.⁶ However, their inclusion in a structural VAR is somewhat novel and raises questions about their added value and whether such a VAR is a valid reduced form for a plausible set of economies. We therefore briefly derive the VAR from first principles, based on a general structural model that nests the NK model discussed above.⁷

Consider a linear rational-expectations economy of the form

$$\mathbf{x}_t = \mathbf{A}\mathbf{x}_{t-1} + \mathbf{B}E_t[\mathbf{x}_{t+1}] + \boldsymbol{\varepsilon}_t \tag{9}$$

where \mathbf{x}_t is the state vector of economic variables and $\boldsymbol{\varepsilon}_t$ is an exogenous sto-

⁶These applications include term-structure modeling (e.g., Chun, 2011; Kim and Orphanides, 2012; Piazzesi et al. 2015), improving forecasting models (e.g., Wright, 2013; Frey and Mokinski, 2015), and identification of DSGE models (e.g., Milani, 2011; Del Negro and Schorfheide, 2013).

⁷The only previous study known to us in which surveys are used to identify shocks to expectations in a VAR setting is Leduc and Sill (2013), which takes the VAR reduced form as given.

chastic process that is independent of current and past values of \mathbf{x}_t and has zero unconditional mean. To introduce the possibility that agents may form expectations based on information other than that contained in \mathbf{x}_t , we generalize the conditional expectation of $\boldsymbol{\varepsilon}_t$ itself to be a stochastic process, as we did in the NK model of the previous section. In particular, we assume that agents may receive an unbiased signal in each period about the future value of $\boldsymbol{\varepsilon}_t$, potentially at all horizons $h > 0$. Thus, for each period t , we can write the innovation $\boldsymbol{\varepsilon}_t$ as the sum of an unanticipated mean-zero component \mathbf{u}_t and a component anticipated in the previous period \mathbf{a}_{t-1}^t :

$$\boldsymbol{\varepsilon}_t = \mathbf{a}_{t-1}^t + \mathbf{u}_t. \quad (10)$$

Under the standard (no-news) assumption, $\boldsymbol{\varepsilon}_t = \mathbf{u}_t$, the conditional mean of $\boldsymbol{\varepsilon}_t$ is also always zero, and the economy has a reduced-form VAR representation:

$$\mathbf{x}_t = \boldsymbol{\Phi}\mathbf{x}_{t-1} + \boldsymbol{\Psi}_0\mathbf{u}_t \quad (11)$$

where $\boldsymbol{\Psi}_0\mathbf{u}_t$ is a mean-zero, iid vector of reduced-form errors, assuming that the matrices of structural parameters \mathbf{A} and \mathbf{B} satisfy the conditions for invertibility. The reduced-form parameters $\boldsymbol{\Phi}$ are easily estimated by OLS or Bayesian methods.

In the more general case where the news component may be non-zero, \mathbf{a}_{t-1}^t and \mathbf{u}_t are independent, and, for any fixed period $T > t$, rational expectations imply that \mathbf{a}_t^T must follow a martingale. Thus, the time- t news shocks, $\boldsymbol{\eta}_t^T = \mathbf{a}_t^T - \mathbf{a}_{t-1}^T$, are mean-zero and iid. We note that this way of specifying news shocks mathematically nests the treatment of those shocks in much of the previous theoretical literature on news about productivity (e.g., Beaudry and Portier, 2006; Schmitt-Grohe and Uribe, 2012; Jaimovich and Rebelo, 2009) and on forward guidance in monetary policy (Lasseen and Svensson, 2011; Campbell et al., 2012; Milani and Treadwell, 2012; Del Negro and Schorfheide, 2013).⁸

⁸Lorenzoni (2009), Barsky and Sims (2011), and Blanchard et al. (2013) consider a related class of models in which, rather than forming beliefs about variables that are only realized in the future, agents solve a filtering problem in which they try to learn current productivity from a noisy signal. Although the models in those papers differ slightly from the structure assumed here, the mechanism is similar because the primary economic effects occur through changes in agents' expectations of future values. In particular, when they infer higher productivity, they revise upward their beliefs of future output, causing an immediate increase in demand.

In Appendix B, we show that, in the presence of news shocks, \mathbf{x}_t follows the process

$$\mathbf{x}_t = \mathbf{\Phi}\mathbf{x}_{t-1} + \mathbf{\Psi}_0\mathbf{u}_t + \mathbf{\Psi}_0\mathbf{a}_{t-1}^t + \sum_{h=1}^{\infty} \mathbf{\Psi}_h\mathbf{a}_t^{t+h} \quad (12)$$

where $\mathbf{\Psi}_h$, are the multipliers on future expectations.⁹ The reduced-form parameters $\mathbf{\Phi}$ and disturbance $\mathbf{\Psi}_0\mathbf{u}_t$ are the same as in equation (11). Thus, the presence of news generalizes the no-news model by adding a linear combination of the anticipated state innovations. These terms are necessarily correlated with \mathbf{x}_{t-1} because the \mathbf{a}_t^{t+h} are persistent. Consequently, estimating the VAR on $\{\mathbf{x}_t\}$ alone will produce omitted-variable bias with respect to $\mathbf{\Phi}$, as well as an estimated variance of the reduced-form errors that is too large.

Since we do not observe the \mathbf{a}_t^{t+h} , we cannot estimate (12) directly. However, we can rewrite the system solely in terms of observable variables in one of two ways, according to the following proposition.

Proposition 1 *In the linear rational-expectations economy with news shocks, characterized by the structural equations (9) and (10), the state vector \mathbf{x}_t follows the process*

$$\mathbf{x}_t = \boldsymbol{\theta}_1\mathbf{x}_{t-1} + \boldsymbol{\theta}_2\mathbf{x}_{t-2} + \mathbf{e}_{1,t-1} - \boldsymbol{\theta}_1\mathbf{e}_{2,t-1} + \mathbf{e}_{2,t} \quad (13)$$

where $(\mathbf{e}_{1,t} \ \mathbf{e}_{2,t}) \sim Niid[\mathbf{0}, \boldsymbol{\Sigma}]$, and $\boldsymbol{\theta}_1$, $\boldsymbol{\theta}_2$, and $\boldsymbol{\Sigma}$ are matrices of reduced-form parameters. Equivalently, the joint dynamics of \mathbf{x}_t and its one-period-ahead expectation can be written as the VAR

$$\begin{pmatrix} E_t[\mathbf{x}_{t+1}] \\ \mathbf{x}_t \end{pmatrix} = \begin{pmatrix} \boldsymbol{\theta}_1 & \boldsymbol{\theta}_2 \\ \mathbf{I} & \mathbf{0} \end{pmatrix} \begin{pmatrix} E_{t-1}[\mathbf{x}_t] \\ \mathbf{x}_{t-1} \end{pmatrix} + \begin{pmatrix} \mathbf{e}_{1,t} \\ \mathbf{e}_{2,t} \end{pmatrix} \quad (14)$$

Proof. See Appendix B. ■

Equation (13) shows that \mathbf{x}_t does not have a VARMA representation, let alone the simple VAR(1) dynamics of the no-news economy (11). This observation again highlights the potential for misspecification error in the presence of news if survey data (or other direct measures of expectations) are not included. Indeed, in general, the reduced-form parameters in equation (13) cannot be identified

⁹Note that equation (7) is the NK special case.

from the observations $\{\mathbf{x}_t\}$ alone. Intuitively, if Σ is of full rank, there are innovations in $E_t[\mathbf{x}_{t+1}]$ that are not linear combinations of innovations in \mathbf{x}_t and therefore cannot be recovered only from data on \mathbf{x}_t . The issue is related to the non-invertibility induced by news shocks that is discussed by Sims (2012) and Leeper et al. (2013), although the arguments in those papers have to do with identifying the structural parameters of the system *given* estimates of the reduced form. Here, even the reduced-form parameters themselves are not fully identified.

On the other hand, equation (14) shows that a VAR that incorporates survey data is properly specified, and its reduced-form parameters are both identified and straightforward to estimate. Thus, if news shocks are present, incorporating direct measures of expectations in the VAR is not only admissible but also *necessary* in order to obtain consistent reduced-form estimates.¹⁰ The empirical news literature has long recognized the need to include information from forward-looking variables, such as stock prices (Beaudry and Portier, 2006) and consumer confidence (Barsky and Sims, 2012). One advantage of the survey data is that they directly reveal the magnitudes of the expectations for \mathbf{x}_{t+h} . The ability to measure the reaction of these expectations at the time of a shock allows us to impose that this reaction equals the subsequent change in the actual \mathbf{x}_{t+h} on average, which is a defining feature of news.

3.2 Data and reduced-form estimation details

As our baseline specification we include in \mathbf{x}_t the 3-month Treasury Bill (TBill) rate, log GDP, log CPI, log hours worked, and the growth rate of M2 at a quarterly frequency over the period 1983-2015. Following Christiano et al. (1999) and others, we include GDP and CPI in log levels. We have also run the model in first differences and with a larger \mathbf{x}_t that includes several additional macro variables, with little effect on the results.

As direct measures of expectations, we employ forecasts from widely used macroeconomic surveys. We denote the survey forecast of a variable x_{t+h} as

¹⁰Indeed, in the presence of news, measures of expectations are required even if the researcher has no interest in the effects of news shocks themselves and only wishes to estimate the effects of the unanticipated structural shocks \mathbf{u}_t .

$E_t^S [x_{t+h}]$, where the superscript S on the expectations operator distinguishes it from the model-implied expectation, which we continue to denote $E_t [x_{t+h}]$. Our baseline model uses information from the Blue Chip Survey (BCS), which begin in 1983 Q2 for our series of interest. The Survey of Professional Forecasters (SPF) is used as a robustness check. Each survey reports the respondents' average forecasts of real GDP, CPI, and the 3-month TBill rate, which we use as a proxy for the monetary-policy instrument. The survey data on future GDP and CPI are in log levels, while data on the TBill rate are reported as averages over the period $t + 1$ to $t + h$. Due to idiosyncrasies in the conventions and timing of their reporting, the survey data require some manipulation to be useful in our VAR model, particularly for longer horizons. Our method for obtaining constant-horizon quarterly series and their properties are detailed in Appendix C. We are able to construct constant-horizon quarterly series for survey expectations of key variables at horizons ranging from 3 months to 11 years. Use of these data over a long sample period is novel and provides important information for testing the forward guidance puzzle.

With these data, we estimate a series of models

$$\begin{pmatrix} E_t^S [\mathbf{x}_{t+h}] \\ \mathbf{x}_t \end{pmatrix} = \boldsymbol{\theta}_0 + \boldsymbol{\Theta}(L) \begin{pmatrix} E_t^S [\mathbf{x}_{t+h}] \\ \mathbf{x}_t \end{pmatrix} + \begin{pmatrix} \mathbf{e}_{1,t} \\ \mathbf{e}_{2,t} \end{pmatrix} \quad (15)$$

where L is the lag operator. Note that, while this specification has the form of equation (14), it is more general. In particular, we do not impose the bottom row of $\boldsymbol{\Theta}(L)$ to be $(\mathbf{I} \ \mathbf{0})$ or the intercept term to be zero. This is because, although we impose that rational expectations hold for the specific set of shocks we are interested in, we want to allow for the possibility that the survey data themselves might be biased or that their forecast errors might be predictable in response to other types of shocks that we do not specify.¹¹

We estimate separate models for each survey forecast horizon, 1, 2, 3, and 4 quarters and 6 and 11 years. In each case, we estimate $\boldsymbol{\Theta}(L)$ and $\boldsymbol{\Sigma}$ by Bayesian methods, using a flat normal-Wishart prior and selecting the lag length by BIC. We also ran specifications in which we included multiple surveys within the same

¹¹Numerous studies suggest bias and predictability among consensus survey forecasts. See for example, Davies and Lahiri (1995, 1999) and Coibion and Gorodnichenko (2012, 2015).

model, simultaneously identifying shocks to expectations at different horizons. The results were not appreciably different from those we report below.

4 Shock Identification

Let $\mathbf{\Gamma}$ denote the matrix of multipliers on the structural shocks, such that

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

and $\boldsymbol{\Sigma} = \mathbf{\Gamma}'\mathbf{\Gamma}$. (Appendix B derives the elements of $\mathbf{\Gamma}$ as functions of structural parameters for the linear rational-expectations economy of the previous section.) Although the vectors of news and unanticipated shocks, $\boldsymbol{\eta}_t$ and \mathbf{u}_t , may contain arbitrarily many elements, we single out two: a policy-expectations shock (η_t) and an unanticipated monetary-policy shock (u_t). We denote by Γ_u^x and Γ_η^x the elements of $\mathbf{\Gamma}$ corresponding to the impact of shock u_t and η_t on variable x . Thus, for example, Γ_η^{GDP} is the contemporaneous effect of a policy-expectations shock on GDP, and $\Gamma_u^{E^S[i]}$ is the contemporaneous effect of an unanticipated policy shock on the survey forecast of the TBill rate. In order to identify the necessary elements of $\mathbf{\Gamma}$, we impose a combination of exact and partial identification restrictions, which we now describe.

4.1 Policy-expectations shocks

Our sign restrictions for the policy-expectations shock η_t enforce the following condition: the time- t impact on the expected average TBill rate over periods t to $t+h$ must be in the opposite direction of the time- t impact on expectations of the time- $t+h$ GDP and CPI. This assumption about the contemporaneous impacts of policy-expectations shocks is consistent with the predictions of the NK model discussed earlier and, indeed, with a large class of forward-looking macroeconomic models.¹² To ensure that we are not conflating policy expectations

¹²One 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

with unanticipated policy, we impose that the current TBill rate does not fall in response to an accommodative expectations shock. In other words, we isolate cases in which current policy does not ease but expected future policy does, implying that the short rate can move in the opposite direction of its expectation, which is the behavior predicted by the NK model in Section 2.¹³ In Section 7, we consider a specification in which this rate does not change at all in the period of the policy-expectations shock.

Further, to narrow the admissible set of shocks to those that contain news about future policy innovations, in line with theoretical treatments, we restrict anticipated monetary-policy changes to materialize as expected on average. Specifically, the observed responses of GDP, inflation, and the short rate in the h periods subsequent to the shock are equal, on average, to the time- t change in the survey forecasts of these variables at horizon $t+h$. Structural VARs that do not include actual measures of expectations cannot impose this rationality condition, but in most previous studies of monetary policy the treatment of expectations was not of primary importance as the focus was not on news. Note that this condition does not rule out the possibility that agents' expectations may not always be rational or that they do not respond to noise, since there are additional VAR shocks that we do not identify.

As a normalization, we consider policy-expectations shocks that move expected short rates in the negative direction (i.e., anticipated future monetary-policy easing). Thus, the restrictions to identify policy-expectations shocks can be summarized as follows:

$$\Gamma_{\eta}^{ES[i]} \leq 0, \quad \left\{ \Gamma_{\eta}^{ES[GDP]}, \Gamma_{\eta}^{ES[CPI]} \right\} \geq 0, \quad \Gamma_{\eta}^i \geq 0 \quad (16)$$

$$\Gamma_{\eta}^{ES[GDP]} = E_t \left[\frac{\partial GDP_{t+H}}{\partial \eta_t} \right], \quad \Gamma_{\eta}^{ES[CPI]} = E_t \left[\frac{\partial CPI_{t+H}}{\partial \eta_t} \right], \quad \Gamma_{\eta}^{ES[i]} = \frac{1}{H} \sum_{h=1}^H E_t \left[\frac{\partial i_{t+h}}{\partial \eta_t} \right] \quad (17)$$

would rise in that scenario. Our restrictions ensure that we do not include such situations.

¹³In a different context, Leduc and Sill (2013) show empirically that the current short rate rises when consumers receive positive news about the future.

where H is the horizon of the survey forecast used in the VAR. Note that, consistent with the reporting conventions of the surveys, our restrictions are on the future *levels* of GDP and CPI at the end of the forecast period but on *averages* of the short rate over the forecast period.

4.2 Unanticipated policy shocks

We next turn to the unanticipated monetary-policy innovations u_t . The primary purpose of identifying these shocks is to compare their effects to those of the anticipated innovations. Since survey forecasts should also respond within a quarter to unanticipated policy shocks, in our baseline specification for u_t , we impose on those forecasts the same set of restrictions used for η_t . This allows the expectation channel to be present in both types of shock, maximizing comparability. The only difference is that unanticipated policy shocks involve a contemporaneous decline in the level of the current short rate, whereas such decline can occur only with a lag when the policy innovation is anticipated. Thus, the restrictions are:

$$\Gamma_u^{ES[i]} \leq 0, \quad \left\{ \Gamma_u^{ES[GDP]}, \Gamma_u^{ES[CPI]} \right\} \geq 0, \quad \Gamma_u^i < 0 \quad (18)$$

$$\Gamma_u^{ES[GDP]} = E_t \left[\frac{\partial GDP_{t+H}}{\partial u_t} \right], \quad \Gamma_u^{ES[CPI]} = E_t \left[\frac{\partial CPI_{t+H}}{\partial u_t} \right], \quad \Gamma_u^{ES[i]} = \frac{1}{H} \sum_{h=1}^H E_t \left[\frac{\partial i_{t+h}}{\partial u_t} \right] \quad (19)$$

As with the policy-expectations shock, these restrictions on the unanticipated policy shock are consistent with the predictions of standard NK models (including the one discussed in Section 2), and there are no other shocks within standard versions of those models that result in the same patterns.

There is, of course, a large literature on identifying unanticipated monetary-policy innovations, and for robustness we will consider two of the most commonly used identification schemes (Christiano et al., 1999, and Uhlig, 2005) in Section 7. However, at least for the purposes of this paper, the restrictions in equations (18) and (19) have two advantages. First, they ensure rationality of the shock responses. This is not guaranteed by other identification schemes. Second, they

ensure consistency between the policy-expectations and unanticipated shocks. For example, standard approaches may lead to a price puzzle, but this would be inconsistent with our identifying assumption that anticipated policy easing cause expected inflation to rise.

4.3 Implementation of sign and zero restrictions

The standard method for imposing sign (or other partial-identification) restrictions is to draw a set of candidate factorizations of the estimated covariance matrix Σ , discarding the draws that are inadmissible. The drawing procedure makes use of the result that, if Γ is a factorization of Σ , then so is $\Gamma\mathbf{R}$, where \mathbf{R} is an orthonormal matrix. Thus, random draws are generated by taking an easily computable factor (such as the Cholesky decomposition) and multiplying it by randomly generated orthonormal matrices. Standard procedures for generating these matrices produce draws of Γ that are uniformly distributed under the Haar measure.

Although this procedure has become standard in the sign-restriction literature, Baumeister and Hamilton (2015) point out that draws that are uniform under the Haar measure are not generally unconditionally uniform for any given element of the structural parameters Γ^{-1} . They recommend treating the problem as a fully Bayesian one by specifying a prior for the structural parameters. Since the data do not provide any information about the factorization of Σ , the posterior distribution of Γ^{-1} is proportional to this prior within the region of the parameter space that satisfies the partial restrictions.

For our problem, two difficulties render the Baumeister-Hamilton approach infeasible. First, while it might be possible to use theory to come up with priors for the structural parameters applying to our two shocks of interest, some of the variation in our VAR comes from other shocks that we do not identify and for which we have no structural interpretation. Thus, there is no obvious way to produce priors for all of the structural parameters of our model. Second, we need to impose zero restrictions in addition to sign restrictions, and Baumeister and Hamilton do not provide any method for doing this.

Fortunately, Arias et al. (2016) have developed an algorithm for drawing

uniformly (i.e., with agnostic priors) across either the structural parameters of the model or across the impulse-response functions (IRFs) given an arbitrary set of zero restrictions. To implement the identification in equations (16) - (19), and make sure that the identification *only* comes from this set of restrictions, we follow their procedure.¹⁴ For our baseline results, we use a uniform prior over the structural parameters. We show in Section 7 that imposing a uniform prior over the IRFs produces almost the same results.

To compute IRFs, 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 dynamic effects of a one-standard-deviation shock under each draw. We report the median and 16%-to-84% range across all draws, noting that these distributions reflect both statistical uncertainty about the reduced-form parameters and "model uncertainty" over the possible structural rotations that are consistent with our restrictions for any given set of reduced-form estimates. Our focus on the quantiles of the IRF distribution accounting for uncertainty about both the model parameters and the impact matrix effectively treats the problem as a Bayesian one, as recommended by Baumeister and Hamilton (2015).

5 Results

5.1 Shorter Horizons

We begin with the results for the models that include survey data with a horizon of one year or less. As we will show, these results are very close to one another, suggesting that the economic multipliers on policy-expectations shocks at short horizons are very similar. Then, we compare their effects to those of the unanticipated innovations. The results using longer-horizon surveys differ somewhat, and we defer them to a separate subsection.

5.1.1 The time series of policy-expectations shocks

Figure 2 offers some evidence that the policy-expectations shocks we have identified do indeed correspond to periods in which the path of anticipated monetary

¹⁴We thank Jonas Arias for providing us with Matlab code that greatly assisted in this effort.

policy had reasons to shift.¹⁵ In particular, using each parameter draw from each model characterized by a specific survey horizon, we compute the time series of the shocks and plot the median across all draws for each model. All four models deliver very similar results.

The largest positive policy-expectations shocks occur in 2008 Q4 and 2009 Q1, just as the short-term rate hits the zero lower bound. This is precisely as we would expect if forward-looking agents incorporate the ZLB constraint in their expectations. From the perspective of a linear model, agents' knowledge that, because of the ZLB, monetary policy cannot be as accommodative as the usual rule would prescribe constitutes an expectation for tightening. Another noticeable positive shock occurs in 1994 Q4, as market participants were surprised by the extent of the tightening at the November FOMC, action that the Fed motivated "as necessary to keep inflation contained."¹⁶

The largest negative policy-expectations shock in the sample occurs in 2011 Q3, coincident with Federal Reserve's adopting calendar-based forward guidance, which specified that economic conditions were "likely to warrant exceptionally low levels for the federal funds rate at least through mid-2013." Similarly, in 2012 Q1, another accommodative shock appears as the Committee extended the date of its calendar-based guidance to "late 2014." Finally, the shock identified in 2014 Q2, coincides with the FOMC switching to qualitative guidance about likely deviations from the conventional policy rule by announcing: "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." Indeed, between 2011 Q3 and 2014 Q2—roughly the period during which the FOMC was most explicit in its use of forward guidance—our estimated policy expectations shocks at the one-year horizon average approximately -0.6 standard deviations per quarter. No other consecutive three-year period in our sample

¹⁵The dates in the figure reflect the dates of the identified shock and not necessarily those of the corresponding FOMC statements, as we have taken care to account for the timing of the surveys relative to FOMC meetings. In particular, the BCS data 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.

¹⁶See the December 1994 Bluebook for a detailed description of this episode.

shows this level of anticipated accommodation.¹⁷

There are also sizeable policy-expectations shocks in our sample prior to the ZLB period that correspond to identifiable events. For example, we find large accommodative shocks in early 2008, when the FOMC aggressively cut rates in an intermeeting move and noted the "downside risks" associated with the emerging financial crisis; in 2003 Q3, when it introduced language to signal that "policy accommodation can be maintained for a considerable period;" in 2002 Q3 when it tilted the balance of risks toward the downside mentioning that "the risks are weighted mainly toward conditions that may generate economic weakness;" and in 1998 Q4, when in an intermeeting move it significantly eased policy pointing to "unsettled conditions in financial markets."

Prior to the late 1990s, explicit FOMC communication was more sparse, and thus specific events that might relate to our estimated shocks are harder to find. For example, the large accommodative expectations shocks that occur in early 1984 do not correspond to any obvious Fed communication, but it roughly coincides with the "inflation scare" episode pointed out in Goodfriend (1993), in which a significant rise in expected long-run inflation occurred in the absence of an aggressive funds rate tightening. An advantage of our approach is that it does not require that shocks arise only from explicit forward guidance. Any information that has caused agents to (correctly) change their beliefs about future monetary policy innovations should be identified as policy-expectations shocks in the data.

5.1.2 Impulse-Responses

Policy-expectations shocks Figure 3 presents IRFs for policy-expectations shocks in our baseline VARs. Each row reports the results using survey forecasts at a different horizon (1 to 4 quarters ahead). For the moment, we concentrate on the results using the four-quarter forecast (last row). We estimate that the size of a one-standard-deviation shock to policy expectations at this horizon is about 2 basis points (i.e., the survey-forecasted TBill rate declines by about 2

¹⁷The relatively high volatility of the policy-expectations shocks during the ZLB period raises the possibility that that episode might be driving the results discussed below. In Section 7, we estimate the model excluding this period and obtain similar impulse-responses.

basis points in response to such a shock), with a credibility interval of about 1 to 4 basis points. The small size of these shocks is consistent with a monetary authority that mostly adheres to a consistent rule and rarely uses explicit forward guidance. However, in spite of the small change in short-rate expectations, the expected economic changes are nontrivial: the GDP and CPI survey forecasts for the next four quarters increase by about 0.15 percent, according to the posterior medians.

Our central question is how these changes in expectations affect actual macroeconomic outcomes. The last four columns of IRFs in Figure 3 show the evidence on this question: in response to a one-year policy-expectations shock of average size, GDP and CPI immediately increase by about 0.16 percent, increases that are "statistically significant" in the sense that about 95 percent of the posterior probability mass in both cases lies above zero. The point estimate of hours worked rises by 0.21 percent, though its credibility band is wider and includes zero. All three variables revert to their steady-state values only very slowly. Meanwhile, in the period of the shock, the actual TBill rate rises by about 2 basis points. (Recall that our restrictions guaranteed that this response would be non-negative.) However, it quickly reverses sign, and subsequent changes in the actual TBill rate are closely aligned with those of the expected TBill rate. The response of M2, which is not central to our story, is not significantly different from zero.

Now comparing the results across rows, we note that the responses of most of the variables are very similar, regardless of the horizon of the expectations used in the model. This observation, together with the high correlation among the shocks evident in Figure 2, implies that the economic multipliers on short-horizon policy-expectations shocks are very similar and thus, from now on, we can focus just on one of these shocks.

In interpreting these IRFs, it is important to be careful to avoid confusing assumptions with results. Although our identifying restrictions guarantee that the cumulative responses of actual GDP and CPI after h periods must equal the initial changes in their respective h -period-ahead survey forecasts, nothing predetermines the magnitude of the response at impact, its significance, and its evolution until period h . Our results suggest that the economy responds

immediately to expectations, with nearly all of the eventual change in current variables occurring in the period of the shock. The magnitude of economic responses to the policy-expectations shock could have been small (in principle even zero), but that is not the case. We get a 15- to 20-basis-point increase in GDP, prices, and hours from a 2-basis-point decline in the one-year expected short rate, an economically significant response. Further, in principle, shocks to four-quarter expectations could have generated much different results than shocks to one-quarter expectations. But, we find that there are virtually no differences across expectational horizons.

Finally, and perhaps most importantly, we analyze how the effects of expected changes in monetary policy compare to those of unexpected changes. This comparison, which is what will ultimately allow us to discuss the importance of the expectations channel in conventional monetary policy and the effects of forward guidance, requires estimating the responses to unanticipated policy shocks, which we turn to next.

Unanticipated policy shocks Responses to the unanticipated policy shocks are shown in Figure 4. Recall that these shocks are produced within the same set of models depicted in Figure 3, using identifying restrictions that differ from the policy-expectations shocks only in the contemporaneous response of the current TBill rate, which must now decline at impact. We first note that, looking down rows, the responses are very similar when any shorter-term survey horizon is used, which is in line with what we observed in the case of the policy-expectations shocks.

Focusing on the four-quarter horizon for brevity, our median estimate is that a one-standard-deviation unanticipated policy shock lowers the actual short rate by about 10 basis points.¹⁸ It also moves average short-rate expectations by about 6 basis points on impact, because of the strong persistence of the TBill series. Recall that the policy-expectations shock moved the four-quarter expected short rate by only about 2 basis points on impact. In this sense, the aver-

¹⁸Note that this is considerably smaller than the size of monetary-policy shocks that are typically estimated in VAR models, especially at a quarterly frequency. The reason is the inclusion of the survey data, which help the VAR to predict the one-step-ahead short rate with a relatively high degree of accuracy.

age unanticipated policy shock is "bigger" than the average policy-expectations shock. Despite this, the response of the macroeconomic variables is of similar magnitude. GDP and CPI immediately increase by about 0.15 percent in response to the unanticipated policy shock, and both remain around those levels for the remainder of the projection period. The initial response of hours — a statistically insignificant 0.04 percent — is somewhat weaker than it was in Figure 3. These findings highlight the potency of near-term anticipated monetary policy innovations, as a small policy-expectations shock can generate changes in current GDP and CPI that are at least as large as those produced by a larger unanticipated policy shock. In the robustness section, we will show that this result holds when we use more-standard identification schemes (Christiano et al., 1999; Uhlig, 2005) for the unanticipated policy shocks.

5.2 Longer Horizons

Figure 5 shows the IRFs when we apply our model to the 6- and 11-year survey data. It is worth pointing out that when using longer-horizon expectations, the short-run behavior of the economy is less likely to be constrained by the news restrictions. For example, if the cumulative response of actual GDP has to equal on average the initial change in the 6-year forecast of GDP, then this response has 24 quarters to satisfy this condition. In contrast, when using one-year expectations, the GDP trajectory has 4 quarters to satisfy the constraint.

In response to a one-standard-deviation policy-expectations shock at the six-year horizon, GDP rises by about 0.2 percent, CPI rises by about 0.15 percent, and hours worked rises by 0.12 percent, according to the posterior medians. The responses are persistent and, except for hours, statistically significant. These are similar to the responses observed after a one-standard-deviation near-term policy-expectations shock. However, the change in the one-year average expected TBill rate was 2 basis points, whereas here it is only about 1 basis point.¹⁹ Thus, a given change in policy expectations is more potent when it lasts for 6 years.

This idea is further confirmed by the results for the 11-year surveys. There, the movement in the average expected TBill rate over the survey horizon is again

¹⁹This is not surprising considering the little variation displayed by long-range survey forecasts in Figure A2 of Appendix C.

about 1 basis point. But in response median GDP rises by 0.25 percent, CPI by 0.22 percent, and hours by 0.28 percent, with all three responses statistically significant and persistent. These responses are larger than in the six-year case. We interpret these results to suggest that, at least qualitatively, the predictions of the NK theory are correct—policy expectations at longer horizons have a larger effect on current outcomes than policy expectations at shorter horizons.

As shown in the bottom panel of Figure 5, using the 6- and 11-year survey data, a one-standard-deviation unanticipated policy shock moves the contemporaneous TBill rate by about 10 basis points. Differently from the shorter-horizon estimates, the GDP response becomes statistically significant only in the medium-run, when it rises above 0.1 percent, but IRFs of CPI and hours are insignificant. These findings indicate that also at long horizons anticipated policy moves are more effective than unanticipated ones.

6 The effects of anticipation

In this section, we use the results from the models estimated above to examine the effects of anticipation in monetary policy, which is the question at the core of this study. We address this question from two angles. First, we ask how much of the economic reaction to unanticipated policy shocks is due to the effect that those shocks have on expectations for future interest rates, rather than on the level of today's short rate. Second, we ask whether a given deviation from the historical policy rule has a different effect when it is anticipated than when it is not, which is crucial to understand the marginal effect of credible forward guidance. In both cases, our approach is to simulate a combination of policy-expectations and unanticipated policy shocks that produce the hypothetical short-rate path we wish to consider. This approach is similar in spirit to the fiscal policy scenarios in Mountford and Uhlig (2009).

6.1 Expectations in conventional monetary policy

To understand the importance of the expectation channel for the propagation of unanticipated policy shocks, we construct a hypothetical scenario in which

there is a shock to today's short rate but expectations for future monetary policy remain unchanged. One way to approach this question, using direct estimation rather than simulation, would have been to use additional zero restrictions to identify a policy shock in which the short rate moved contemporaneously but expectations did not:

$$\left\{ \Gamma_u^{ES[i]}, \Gamma_u^{ES[GDP]}, \Gamma_u^{ES[CPI]} \right\} = 0, \quad \Gamma_u^i < 0 \quad (20)$$

However, when we attempted to estimate this model, we could not find any draws of Γ for which these restrictions were met. This corroborates the implausibility of shocks that move the policy rate today without affecting expectations of its future values.²⁰

Alternatively, we can use the results of our baseline model to construct the following hypothetical experiment: Concurrent with an unanticipated shock u_0 , we introduce a policy-expectations shock η_0 that is equal to

$$\eta_0 = -\frac{\Gamma_u^{ES[i]}}{\Gamma_\eta^{ES[i]}} u_0 \quad (21)$$

and thus exactly offsets the effect of u_0 on the expected short rate. This produces an immediate change in the actual short rate but no change in the average expected short rate over the survey horizon.

Figure 6 illustrates the results of this scenario using the model with four-quarter surveys. In panel A, we consider a pair (u_0, η_0) that lowers the contemporaneous TBill rate by 25 basis points such that expectations for the future short rate do not move in the period of the shock. In contrast to the IRFs shown in Figure 4, GDP and CPI do not respond significantly to the unanticipated monetary-policy shocks when the expectation effects are removed. Indeed, the median of the response is negative (though insignificant). This could be due to the fact that, in Panel A, the TBill rate is higher after about a year, and anticipation of that move may depress current output and inflation. We therefore

²⁰This is consistent with Gurkaynak et al. (2005), who show that market-based measures of short-rate expectations at different horizons react very fast and significantly to unanticipated monetary-policy shocks.

introduce a second offsetting expectations shock η_1 to ensure that policy expectations are not only unchanged in the period of the unanticipated policy shock but also in the subsequent period. These results are shown in panel B. The introduction of the second expectations shock is sufficient to keep short-rate expectations essentially unchanged for the entirety of the projection period. And, as a result, there is virtually no impact on the economic variables.

We repeat this exercise using the longer-horizon survey data (not shown). The models with both the 6- and 11-year surveys also produce insignificant results. This exercise illustrates how powerful the expectation channel is: when policy expectations do not change, conventional monetary-policy shocks have no significant effect on current economic outcomes.

6.2 Forward guidance

We now use the results of the models estimated in Section 5 to analyze the marginal effects of forward guidance about the policy-rate path. As we did in the NK model (see panels B and C of Figure 1), we simulate a scenario in which the Fed introduces a policy-expectations shock at time 0, η_0 , that changes the anticipated path of the short rate by a certain amount and an unanticipated policy shock, u_0 , to ensure that the change in the actual short rate is zero at time 0. u_0 effectively shuts down any systematic policy response to the economic expansion induced by the forward-guidance announcement. Thus, if η_0 is the initial policy-expectations shock, the required unanticipated policy shock in period 0 is given by

$$u_0 = -\frac{\Gamma_\eta^i}{\Gamma_u^i}\eta_0. \quad (22)$$

In our exercise, we set the initial change in the expected TBill rate over the next four quarters, $\left(\Gamma_\eta^{ES[i]} - \frac{\Gamma_u^{ES[i]}\Gamma_\eta^i}{\Gamma_u^i}\right)\eta_0$, equal to -10 basis points. For a sense of scale, this is approximately the size of the change in TBill expectations implied by our median estimates in response to the adoption of calendar-based forward guidance in 2011 Q3.²¹

²¹Recall from Figure 2 that, according to the model with the four-quarter surveys, the calendar-based forward guidance was associated with a policy-expectations shock of approximately a 2.5 standard deviations. In that model, our median estimate of the multiplier

We compare the outcome of this experiment to a separate scenario in which the short rate follows an identical path over quarters 0 to 4, but this path is not pre-announced and therefore there is no policy-expectations shock. That is, we simulate a sequence of unanticipated policy shocks $(\tilde{u}_1, \dots, \tilde{u}_4)$ that produces the same short-rate path as in the forward-guidance scenario. (Note that we do not need a \tilde{u}_0 , because the TBill rate does not move in period 0 under forward guidance.) The difference between the IRFs produced by (η_0, u_0) and $(\tilde{u}_1, \dots, \tilde{u}_4)$ tells us the *marginal* effect of forward guidance—that is, the extra effect on the economy that is achieved by credibly announcing a given policy path in advance.

Figure 7 summarizes the outcome, with the forward-guidance scenario shown in blue and the unanticipated scenario shown in red. The top-right panel shows that the realized TBill rate does not move in period 0 but gradually falls over the next four quarters and averages -10 basis points over that period. By construction, the paths are identical in both scenarios over these four quarters (although they are allowed to differ after that time). The top-left panel shows that, in the case of forward guidance, rate expectations fall immediately in period 0, whereas in the unanticipated policy scenario, rate expectations only change slowly, in response to the shocks to the actual policy path. The initial disparity in the TBill expectations in that panel is the key difference between the two scenarios.

The bottom panels show that the forward-guidance announcement causes current GDP and CPI to increase by about 0.6 percent on impact, according to the posterior medians. Furthermore, the distribution of outcomes is skewed to the upside, so that responses of double that size cannot be ruled out. For all three variables, more than 84 percent of the probability mass lies above zero for several years after the shock. Hours rise by about 0.7 percent at impact, although after the first quarter their response is not statistically significant. In the unanticipated policy scenario, the median responses of GDP and CPI are at most half as big as in the forward guidance scenario, and the response of hours is proportionally even smaller. The statistical significance of the responses is also weaker. In addition, the effects take longer to occur: in the forward-guidance scenario, the

$$\left(\Gamma_{\eta}^{E^S[i]} - \frac{\Gamma_u^{E^S[i]} \Gamma_{\eta}^i}{\Gamma_u^i} \right) \text{ is about } -4.$$

economic variables essentially jump to new levels in the period of the shock and remain there, but in the case of unanticipated policy, they take approximately a year to reach their maximum levels. However, we note that the differences between the forward-guidance scenario and the unanticipated policy scenario are not statistically different after the first couple of quarters, in the sense that their posterior credibility intervals overlap. The relatively rapid convergence of the two scenarios in the case of short-horizon guidance is to be expected given that the disparity between the expected policy rate paths dissipates quickly.

We re-run this exercise for each of the survey horizons using the same size shock (as measured by the change in the average expected TBill rate) and compute the median distance between the economic responses to the anticipated and unanticipated policies (i.e., the difference between the blue and red solid lines in Figure 7) at different points in time, from the period of the shock to 20 quarters after the shock hits. The results are tabulated in Table 1. We also report a "significance" test for this difference — that is, the percentage of the posterior probability mass for which the anticipated policy generates a larger response than the unanticipated policy. In all cases, we find that the forward-guidance shocks generate larger responses. For the 6- and 11-year scenarios, the marginal effect of forward guidance is larger than at the one-year horizon, and it is almost always statistically significant. For example, we find that if forward guidance lowered policy expectations for six years, rather than one year, the near-term response of GDP would have been about 1.06 percent and the near-term response of the CPI would have been about 0.82 percent. Further, forward guidance at the 11-year horizon elicits bigger economic effects than at the 6-year horizon, as predicted by the NK models.

7 Robustness and Extensions

In this section, we examine a variety of alternative specifications. We start with analyzing the robustness of the positive and substantial marginal effects of forward guidance to different identification schemes for the unanticipated policy shocks, as this result is at the core of our study. Then, we turn to demonstrate that our novel methodology for the policy-expectations shock is robust to various

modifications of the identifying restrictions and implementation.

7.1 Alternative Identification for Unanticipated Shocks

Because the way we identify unanticipated policy shocks is novel, we first single out that aspect of the estimation. Specifically, we check our results against two standard methods for identifying unanticipated monetary-policy shocks: short-run exclusion restrictions similar to those of Christiano, Eichenbaum, and Evans (CEE) (1999) and sign restrictions similar to those of Uhlig (2005). For the CEE-like shocks, we impose that GDP, CPI and hours can only respond to such shocks with a one-period lag. For the Uhlig-like shocks, we impose that unanticipated policy shocks do not cause CPI or M2 to fall in either the quarter of the shock or the subsequent quarter.²²

The primary purpose of identifying the unanticipated policy shocks was to compare them to the policy-expectations shocks. Therefore, robustness can be demonstrated most clearly by repeating the forward-guidance exercise of Section 6.2. Figure 8 shows the results for the VAR model using one-year survey data, with the outcomes of the CEE-like unanticipated shocks in panel A and those of the Uhlig-like shocks in panel B.

In both cases, the 10-basis-point forward guidance over the next year (depicted in blue) has a very similar effect to what we found in the baseline model, with median GDP, CPI, and hours all rising by approximately half a percentage point on impact and decaying back to zero only very slowly. The statistical significance of the results is also similar to what we found above. In other words, changing the identification strategy for the unanticipated shock has little consequence for the estimated effects of anticipated policy. (See the last two rows of the top panel in Table 2 for more details on the IRFs.)

The red regions show the effects of the same short-rate path when it is produced by unanticipated shocks under the alternative identification schemes. In both cases, unanticipated policy has small and largely insignificant effects. (Under the Uhlig-like scheme, the effect on CPI is positive and significant over the first two quarters by construction, but it is still economically small.) Consequently, we

²²Uhlig (2005) uses bank reserves rather than M2 as the measure of the money supply, but this series is no longer practical to include because it exhibits a large structural break in 2008.

continue to estimate that the marginal effects of forward guidance are substantial in the short run—indeed, a bit bigger than in the baseline model—but become statistically insignificant after a few quarters. When we use longer-horizon survey forecasts (not shown), we find bigger and more persistent marginal effects of forward guidance, again consistent with our baseline results.

The result in panel B that the real variables do not respond significantly to unanticipated policy is consistent with Uhlig’s (2005) original finding. Our estimated response of inflation is slightly weaker than what he found, but the difference would not be enough to meaningfully affect our conclusions. The panel A results are more puzzling, since they seem to contradict Christiano et al.’s (1999) finding that real activity responds significantly to unanticipated policy shocks. A likely reason for the discrepancy is that Christiano et al.’s model did not account for forward-looking information about policy and thus may have conflated the effects of unanticipated policy shocks with those of policy-expectations shocks.

7.2 Other Robustness Checks

Here, for each specification, we calculate the IRFs for the policy-expectations shocks, and report their values at the time of the shock (time 0), one year after the shock ($1y$), and 5 years after the shock ($5y$) in Table 2. To keep the analysis manageable, we focus on the results using only the one-year and six-year survey data, summarized in the top and bottom panels, respectively. In the the first row of each panel, we report the baseline results as a term of comparison.

Flat prior on IRFs. As noted in Section 4.3, our results adopt a flat prior for the structural parameters ($\mathbf{\Gamma}^{-1}$), and here we verify whether imposing a flat prior on the IRFs produce similar results. The results using that approach are reported in the second row of each panel. Overall, using both the one- and six-year survey data, the economic effects of one-standard-deviation policy-expectations shock appear to be quite similar to the baseline. However, it should be noted that, in the top panel, the size of the shock is smaller but the effects for GDP and hours are larger and more significant. Therefore, the results would have been stronger under this assumption.

Restricting the response of the TBill rate. The NK model of Section 2 implies that the short-term interest rate ought to rise in response to a policy-expectations shock, and we allow that response to be non-negative in our baseline model. An alternative is simply to identify shocks for which the current policy rate does not move in the period of the shock. In other words, rather than the sign restriction $\Gamma_{\eta}^i \geq 0$ in equation (16), we impose the zero restriction $\Gamma_{\eta}^i = 0$. The results are shown in the third row of each panel. At the one-year horizon, the economic responses are very similar to the baseline. For the six-year horizon, even if the magnitude of the shock is slightly bigger, some of the effects are a bit smaller than the baseline.

Excluding the ZLB period. Since the effective zero bound on nominal interest rates could be an important source of non-linearity in our sample period, we check 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 guidance has been most actively used as a policy tool. The results of the models that are re-estimated stopping the sample in Q4 2008 are shown in the fourth row of each panel. First of all, it should be noted that the size of one-standard-deviation policy-expectations shocks is the same as the baseline for both the one- and six-year VAR models, and second, that their economic impact is very similar to that of the baseline specification, except that at the six-year horizon is larger and statistically significant beyond the short-run.²³

Alternative survey measure. Our baseline model relied on forecasts from the BCS. We check robustness to SPF forecasts. One way in which those data differ is that they begin in Q3 1981, rather than in Q1 1983, so that this also serves to extend the sample slightly. However, until the 1990s, the SPF does not include observations on our variables of interest at horizons beyond one year, thus the

²³These 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.

results of this exercise appear only in the fifth row of the top panel. In this case, a one-standard-deviation policy-expectations shock is smaller in magnitude but its effects, although mostly quite similar to the baseline, appear to be a bit larger for CPI. This result could be due in part to the extension of the sample back to 1981. As shown in Figure A1 in the Appendix, 1981 and 1982 are particularly important for the inflation expectations process and the early 1980s is when the SPF and BCS forecasts diverge the most.

Even more robustness checks. We also conducted a variety of other robustness checks that are not reported in the table. As noted earlier, we experimented with different combinations of macro variables and different lag lengths in the VAR and ran the model in first differences rather than in levels. We also imposed the sign restriction on the real yield rather than the nominal yield. (The real yield is calculated as the h -quarter survey forecast of the average nominal TBill rate minus the h -quarter forecast of CPI inflation.) Finally, instead of using a normal-Wishart prior on the VAR parameters, we used the Minnesota prior, and we also performed simple maximum-likelihood estimation, treating the parameters as known for the purpose of drawing structural factorizations. Again we found very similar results in all of these cases.

8 Conclusion

We have used time-series data to examine the dynamic effects of news about future monetary policy on the economy. The central identification challenge is teasing out expected deviations from the historical policy rule—that is, anticipated changes in the short-term interest rate that cannot be explained by expectations of future growth and inflation. In the context of explicit monetary-policy communication, the shocks we identify correspond to what Campbell et al. (2012) call “Odyssean” forward guidance: commitments to future accommodation beyond what the expected state of the economy would normally warrant. We accomplish this identification by combining data on interest rate expectations with data on expectations of macro conditions. In particular, we use quarterly survey forecasts, at horizons ranging from three months to 11 years, back to the

early 1980s. We look for shocks that move the survey forecast of the short-term interest rate in the opposite direction of the survey forecasts of GDP and inflation. In a wide class of theoretical models, only anticipated innovations to monetary policy can have this effect.

We find that anticipated policy easing over the near term is associated with immediate and persistent increases in output and inflation. In particular, we estimate that a 10-basis-point decline in the expected short-term interest rate over the following year raises current GDP, prices, and employment by about 0.6 percent, in our baseline model. These are large effects, and we observe that when longer-horizon (6- and 11-year) expectations are used, the effects are larger still. Further, when we compare these economic responses to those elicited by conventional monetary-policy shocks, we find that the marginal effect of forward guidance is positive and significant and that the conventional policy shocks themselves work entirely through the expectation channel.

Overall, these results corroborate the main mechanism underlying New Keynesian models and their qualitative predictions. News about future monetary policy has significant effects on the current economy, and they grow larger as the horizon of the news moves farther into the future. The theoretical literature has debated the magnitude of the impact of forward guidance, and our results provide some quantitative benchmarks for models that study such policies. Our estimates for the effects of forward guidance at long horizons are smaller than those produced by standard NK models, which many researchers consider implausibly large. But our results are consistent with extensions, such as McKay et al. (2016), that introduce discounting in the IS curve and thereby dampen the effects of news at long horizons. We conclude, therefore, that credible forward guidance can likely provide significant economic stimulus, just as the theory suggests.

References

- [1] Arias, Jonas E., Juan F. Rubio-Ramirez, and Daniel F. Waggoner, 2016. "Inference Based on SVARs Identified with Sign and Zero Restrictions: Theory and Applications." Working Paper (August).
- [2] Barakchian, Mahdi S., and Christopher Crowe, 2013. "Monetary policy matters: Evidence from new shocks data." *Journal of Monetary Economics*, 60, 950-966.
- [3] Barsky, Robert B., and Eric. R. Sims, 2011. "News shocks and business cycles." *Journal of Monetary Economics*, 58, 273-289.
- [4] 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.
- [5] 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.
- [6] Baumeister, Christiane and James D. Hamilton, 2015. "Sign Restrictions, Structural Vector Autoregressions, and Useful Prior Information." *Econometrica*, 83(5), 1963-1999.
- [7] Beaudry, Paul, and Frank Portier, 2006. "Stock prices, news and economic fluctuations," *American Economic Review*, 96(4), 1293-1307.
- [8] 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.
- [9] Campbell, Jeffrey R., Chales L. Evans, Jonas D.M. Fisher, and Alejandro Justiniano, 2012. "Macroeconomic Effects of FOMC Forward Guidance." *Brookings Papers on Economic Activity*, Spring.

- [10] Campbell, Jeffrey R., Jonas D.M. Fisher, Alejandro Justiniano, Leonardo Melosi, 2016. "Forward Guidance and Macroeconomic Outcomes Since the Financial Crisis." NBER *Macroeconomics Annual* 2016, Volume 31.
- [11] Carlstrom, Charles T., Fuerst, Timothy. S., and Paustian, Matthias, 2012. "How Inflationary Is an Extended Period of Low Rates?" FRB Cleveland Working Paper 1202.
- [12] Christiano, Lawrence J., Eichenbaum, Martin, and Evans, Charles L., 1999. "Monetary policy shocks: What have we learned and to what end?," Handbook of Macroeconomics, in: J. B. Taylor & M. Woodford (ed.), Handbook of Macroeconomics, edition 1, volume 1, chapter 2, 65-148 Elsevier.
- [13] Chun, Albert L, 2011. "Expectations, bond yields and monetary policy." *Review of Financial Studies*, 24 (1), 208-247.
- [14] Coibion, Olivier and Gorodnichenko, Yuriy, 2012. "What Can Survey Forecasts Tell Us about Information Rigidities?" *Journal of Political Economy* 120(1): 116-59.
- [15] Coibion, Olivier and Gorodnichenko, Yuriy, 2015. "Information rigidity and the expectations formation process: A simple framework and new facts." *American Economic Review*, 105(8): 2644-2678.
- [16] Davies, Antony and Kajal Lahiri, 1995. "A new framework for testing rationality and measuring aggregate shocks using panel data." *Journal of Econometrics* 68: 205-227.
- [17] Davies, Antony and Kajal Lahiri, 1999. Re-examining the Rational Expectations hypothesis using panel data on multiperiod forecasts. In C. Hsiao, et al. (Eds.), *Analysis of Panels and Limited Dependent Variable Models*, Cambridge Univ. Press: 226-254.
- [18] Del Negro, Marco, Marc Giannoni, and Christina Patterson, 2012. "The Forward Guidance Puzzle." Federal Reserve Bank of New York, Staff Report No. 574, May.

- [19] Del Negro, Marco and Schorfheide, Frank, 2013. "DSGE Model-Based Forecasting." *Handbook of Economic Forecasting* 2: 57-140.
- [20] Eggertsson, Gauti B. and Michael Woodford, 2003. "The Zero Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity*, Spring.
- [21] Frey, Christoph and Frieder Mokinski. "Forecasting with bayesian vector autoregressions estimated using professional forecasts." *Journal of Applied Econometrics*, 31(6): 1099-1255.
- [22] Galí, Jordi, 2008. "Monetary Policy Design in the Basic New Keynesian Model," in: *Monetary Policy, Inflation, and the Business Cycle: An Introduction to the New Keynesian Framework*, Princeton University Press.
- [23] Gertler, Mark, and Karadi, Peter, 2015. "Monetary Policy Surprises, Credit Costs and Economic Activity." *American Economic Journal: Macroeconomics*, 7(1): 44-76.
- [24] Gomes, Sandra, Iskrev Nikolay, and Mendicino, Caterina, 2017. "Monetary Policy Shocks: We Got News!" *Journal of Economic Dynamics and Control* 74, 108-128.
- [25] Goodfriend, Marvin, 1993. "Interest rate policy and the inflation scare problem: 1979-1992." *Federal Reserve Bank of Richmond Economic Quarterly* Volume 79/1 Winter.
- [26] 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.
- [27] Jaimovich, Nir and Sergio Rebelo, 2009. "Can news about the future drive the business cycle?" *American Economic Review*, 99(4), 1097-1118.
- [28] 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.

- [29] 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.
- [30] Laseen, Stefan and Svensson, Lars E. O., 2011. "Anticipated Alternative Policy Rate Paths in Policy Simulations." *International Journal of Central Banking* 7(2): 1-35.
- [31] Leduc, Sylvain and Sill, Keith, 2013. "Expectations and Economic Fluctuations: An Analysis Using Survey Data." *Review of Economics and Statistics* 95(4): 1352-67.
- [32] Leeper, Eric M., Sims, Christopher A., and Zha, Tao, 1996. "What Does Monetary Policy Do?" *Brookings Papers on Economic Activity*, vol. 27(2), 1-78.
- [33] Leeper, Eric M., Todd B. Walker, and Shu-Chun Susan Yang, 2013. "Fiscal Foresight and Information Flows." *Econometrica* 81(3): 1115-45.
- [34] 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.
- [35] Lorenzoni, Guido, 2009. "A Theory of Demand Shocks." *American Economic Review* 99(5), 2050-84.
- [36] McKay, Alisdair, Emi Nakamura, and John Steinsson, 2016. "The Power of Forward Guidance Revisited." *American Economic Review*, 106(10), 3133-3158.
- [37] Milani, Fabio, 2011. "Expectation Shocks and Learning as Drivers of the Business Cycle." *The Economic Journal*, 121, 379-401.
- [38] Milani, Fabio, and John Treadwell, 2012. "The Effects of Monetary Policy "News" and "Surprises". *Journal of Money, Credit, and Banking* 44(8): 1667-92.
- [39] Mountford, Andrew and Harald Uhlig, 2009. "What Are the Effects of Fiscal Policy Shocks?" *Journal of Applied Econometrics*, 2009, 24(6), 960-92.

- [40] Piazzesi, Monika, Juliana Salomao, and Martin Schneider, 2015. "Trend and Cycle in Bond Premia." Working Paper, March.
- [41] Ramey, Valery A., 2011. "Identifying Government Spending Shocks: It's All in the Timing." *Quarterly Journal of Economics* 126(1): 1-50.
- [42] Romer, Christina D. and David H. Romer, 2000. "Federal Reserve Information and the Behavior of Interest Rates." *American Economic Review* 90(3): 429-57.
- [43] Romer, Christina D. and David H. Romer. 2004. "A New Measure of Monetary Shocks: Derivation and Implications." *American Economic Review*, 94(4): 1055-1084.
- [44] Schmitt-Grohe, Stephanie, and Martin Uribe, 2012. "What's news in business cycles." *Econometrica*, 80(6), 2733-2764.
- [45] Sims, Eric, 2012. "News, Non-Invertibility, and Structural VARs." University of Notre Dame Working Paper (June).
- [46] Uhlig, Harald, 2005. "What Are the Effects of Monetary Policy on Output? Results from an Agnostic Identification Procedure." *Journal of Monetary Economics* 52 (2), 381-419.
- [47] Werning, Iván, 2011. "Managing a liquidity trap: Monetary and fiscal policy." NBER, Working Paper, No. 17344, August.
- [48] Werning, Iván, 2015. "Incomplete Markets and Aggregate Demand." NBER Working Paper No. 21448.
- [49] Wright, Jonathan H., 2013. "Evaluating real-time VAR forecasts with an informative democratic prior." *Journal of Applied Econometrics*, 2013, 28(5), 1099-1255.

Appendix A. Derivation of Results in the NK Model

Substituting (3) and (4) into (1) and (2) and rearranging gives inflation and output as functions of their own expected future values and of v_t :

$$\pi_t = \Lambda \{E[\pi_{t+1}] [\kappa + \beta (\phi_y + \sigma(1 - \rho))] + \kappa (r^* + \sigma E[y_{t+1}] - v_t)\}$$

$$y_t = \Lambda \{E[\pi_{t+1}] [1 - \beta\phi_\pi] + (1 - \beta\rho) (r^* + \sigma E[y_{t+1}] - v_t)\}$$

where $\Lambda = [\kappa (\phi_\pi - \rho) + (1 - \beta\rho) (\phi_y + \sigma(1 - \rho))]^{-1}$. Thus the effects of unanticipated shocks, which are given by the multipliers on v_t , are

$$\psi_{0,\pi} = -\kappa\Lambda \quad \text{and} \quad \psi_{0,y} = -(1 - \beta\rho)\Lambda \quad (23)$$

Substituting (23) into the policy rule and taking the derivative with respect to v_t also gives

$$\psi_{0,i} = [(1 - \rho) (1 - \beta\rho) \sigma - \kappa\rho] \Lambda \quad (24)$$

(23) and (24) reproduce the results in Galí (2008). Since $\phi_\pi > \rho$, $\psi_{0,\pi}$ and $\psi_{0,y}$ are necessarily negative. For $\psi_{0,i}$ to be positive requires $(1 - \rho) (1 - \beta\rho) \sigma > \kappa\rho$. Sufficient for this is $\sigma > \kappa/(1 - \beta)$, which is satisfied under standard parameterizations.

To find the ψ_h at longer horizons, note that, since v_t is only backward-looking, the multipliers on a_t^{t+1} are simply the derivatives of π_t , y_t , and i_t with respect to the one-period-ahead expectations. Indeed, by iterating forward, we obtain that, at any horizon h , the multipliers are given by the recursion:

$$\begin{pmatrix} \psi_{h,\pi} \\ \psi_{h,y} \end{pmatrix} = \mathbf{R} \begin{pmatrix} \psi_{h-1,\pi} \\ \psi_{h-1,y} \end{pmatrix} \quad (25)$$

where

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

In addition,

$$\psi_{h,i} = \phi_\pi \psi_{h,\pi} + \phi_y \psi_{h,y}.$$

Given the admissible values of the parameters, all elements of \mathbf{R} except the bottom-left are necessarily positive. That element is positive if and only if $\phi_\pi < 1/\beta$. Consequently, for arbitrary values of ϕ_π and β , $\psi_{h,y}$ may take either sign for any $h > 0$, and $\psi_{h,\pi}$ may take either sign for any $h > 1$.

Appendix B. Proof of Proposition 1

We conjecture that (9) has a solution of the form (12). From equation (12), expectations of the economic variables follow the process

$$E_t[\mathbf{x}_{t+1}] = \mathbf{\Phi}\mathbf{x}_t + \sum_{h=0}^{\infty} \mathbf{\Psi}_h \mathbf{a}_t^{t+h+1} \quad (26)$$

Substituting the latter in for the expectation term in (9):

$$(\mathbf{I} - \mathbf{B}\mathbf{\Phi})\mathbf{x}_t = \mathbf{A}\mathbf{x}_{t-1} + \mathbf{B} \left(\mathbf{\Psi}_0 \mathbf{a}_t^{t+1} + \sum_{h=1}^{\infty} \mathbf{\Psi}_h \mathbf{a}_t^{t+h+1} \right) + \boldsymbol{\varepsilon}_t \quad (27)$$

Given invertibility, this can be written in the form of (12), which verifies that such a solution exists, where $\mathbf{\Psi}_h$, is given by the recursion:

$$\begin{aligned} \mathbf{\Psi}_0 &= (\mathbf{I} - \mathbf{B}\mathbf{\Phi})^{-1} \\ \mathbf{\Psi}_h &= (\mathbf{I} - \mathbf{B}\mathbf{\Phi})^{-1} \mathbf{B}\mathbf{\Psi}_{h-1} \quad \text{for } h = 1, \dots \end{aligned}$$

Since $\mathbf{x}_t = \mathbf{\Phi}\mathbf{x}_{t-1} + \mathbf{\Psi}_0 \mathbf{u}_t + \mathbf{\Psi}_0 \mathbf{a}_{t-1}^t + \sum_{h=1}^{\infty} \mathbf{\Psi}_h \mathbf{a}_t^{t+h}$ and given (26), \mathbf{x}_t can be trivially written as the sum of its lagged expectation and the contemporaneous effects of the news shock $\boldsymbol{\eta}_t^{t+h}$ and unanticipated shocks \mathbf{u}_t :

$$\mathbf{x}_t = E_{t-1}[\mathbf{x}_t] + \sum_{h=1}^{\infty} \mathbf{\Psi}_h \boldsymbol{\eta}_t^{t+h} + \mathbf{\Psi}_0 \mathbf{u}_t \quad (28)$$

Substituting this into (26),

$$E_t[\mathbf{x}_{t+1}] = \mathbf{\Phi}E_{t-1}[\mathbf{x}_t] + \mathbf{\Phi}\mathbf{\Psi}_0 \mathbf{u}_t + \sum_{h=0}^{\infty} [(\mathbf{\Phi}\mathbf{\Psi}_{h+1} + \mathbf{\Psi}_h) \boldsymbol{\eta}_t^{t+h+1} + \mathbf{\Psi}_h \mathbf{a}_{t-1}^{t+h+1}] \quad (29)$$

Due to the constant-horizon structure of the observable expectations, the unobserved \mathbf{a}_{t-1}^{t+h+1} still appears in the above equation. As a device for eliminating this term, we note that without loss of generality, we can write the anticipated component as a factor structure of arbitrarily large dimension,²⁴ with factors \mathbf{f}_t and loadings given by a sequence of matrices $(\gamma_1, \gamma_2, \dots)$, such that

$$\mathbf{a}_t^{t+h} = \gamma'_h \mathbf{f}_t$$

for all h , where, necessarily, $\dim \mathbf{f}_t \geq \dim \boldsymbol{\varepsilon}_t$. The fact that \mathbf{a}_t^{t+h} is a martingale implies

$$E_{t-1} [\gamma'_h \mathbf{f}_t] = \gamma'_{h+1} \mathbf{f}_{t-1}$$

Since this is true at all horizons, there must exist matrices $\boldsymbol{\rho}$ and $\boldsymbol{\gamma}$ such that $E_{t-1} [\mathbf{f}_t] = \boldsymbol{\rho} \mathbf{f}_{t-1}$ and $\gamma_h = \boldsymbol{\gamma} \boldsymbol{\rho}^h$ for all h . Assuming that expectations of the far future always move by less than expectations of the near future (that is, for any $\epsilon > 0$, $\exists \delta > 0$ such that $h > \delta$ implies $\det \boldsymbol{\gamma}_h < \epsilon$), the eigenvalues of $\boldsymbol{\rho}$ lie inside the unit circle, and the factors follow a stationary, first-order vector autoregression, $\mathbf{f}_t = \boldsymbol{\rho} \mathbf{f}_{t-1} + \boldsymbol{\eta}_t$.²⁵ Given that $\boldsymbol{\eta}_t^{t+h} = \mathbf{a}_t^{t+h} - \mathbf{a}_{t-1}^{t+h}$, it follows that $\boldsymbol{\eta}_t^{t+h} = \boldsymbol{\gamma}'_h \boldsymbol{\eta}_t$.

Applying the factor structure to equation (12),

$$\begin{aligned} \mathbf{x}_t &= \boldsymbol{\Phi} \mathbf{x}_{t-1} + \left(\sum_{h=1}^{\infty} \boldsymbol{\Psi}_h \boldsymbol{\gamma}_h \right) \mathbf{f}_t + \boldsymbol{\Psi}_0 (\boldsymbol{\gamma}_1 \mathbf{f}_{t-1} + \mathbf{u}_t) \\ &= \boldsymbol{\Phi} \mathbf{x}_{t-1} + \boldsymbol{\gamma} \left(\sum_{h=0}^{\infty} \boldsymbol{\Psi}_h \boldsymbol{\rho}^{h+1} \right) \mathbf{f}_{t-1} + \boldsymbol{\gamma} \left(\sum_{h=1}^{\infty} \boldsymbol{\Psi}_h \boldsymbol{\rho}^h \right) \boldsymbol{\eta}_t + \boldsymbol{\Psi}_0 \mathbf{u}_t \end{aligned} \quad (30)$$

²⁴Our use of the factor structure for expectations generalizes the "path factor" of interest rates computed by Gurkaynak et al. (2005) in the sense that we allow for (and will empirically identify) the loadings of these factors on expectations of macro variables as well as of interest rates.

²⁵Note that the factor \mathbf{f}_t governing expectations about the future is a mean-reverting process, even though the expectation about any particular date \mathbf{a}_t^T is a martingale.

Rearranging gives

$$\mathbf{f}_{t-1} = \left(\gamma \sum_{h=0}^{\infty} \Psi_h \rho^{h+1} \right)^{-1} \left[\mathbf{x}_t - \Phi \mathbf{x}_{t-1} - \gamma \left(\sum_{h=1}^{\infty} \Psi_h \rho^h \right) \boldsymbol{\eta}_t - \Psi_0 \mathbf{u}_t \right] \quad (31)$$

Meanwhile, applying the factor structure to equation (30) gives

$$\begin{aligned} E_t[\mathbf{x}_{t+1}] &= \Phi E_{t-1}[\mathbf{x}_t] + \Phi \Psi_0 \mathbf{u}_t + \left(\gamma \sum_{h=0}^{\infty} (\Phi \Psi_{h+1} + \Psi_h) \rho^{h+1} \right) \boldsymbol{\eta}_t \\ &\quad + \left(\gamma \sum_{h=0}^{\infty} \Psi_h \rho^{h+1} \right) \rho \mathbf{f}_{t-1} \end{aligned}$$

Substituting (31) into this equation,

$$\begin{aligned} E_t[\mathbf{x}_{t+1}] &= \Phi E_{t-1}[\mathbf{x}_t] + \Phi \Psi_0 \mathbf{u}_t + \left(\gamma \sum_{h=0}^{\infty} (\Phi \Psi_{h+1} + \Psi_h) \rho^{h+1} \right) \boldsymbol{\eta}_t \\ &\quad + \rho \left[\mathbf{x}_t - \Phi \mathbf{x}_{t-1} - \gamma \left(\sum_{h=1}^{\infty} \Psi_h \rho^h \right) \boldsymbol{\eta}_t - \Psi_0 \mathbf{u}_t \right] \end{aligned} \quad (32)$$

Using (28) to eliminate \mathbf{x}_t from (32) gives

$$\begin{aligned} E_t[\mathbf{x}_{t+1}] &= (\Phi + \rho) E_{t-1}[\mathbf{x}_t] - \rho \Phi \mathbf{x}_{t-1} \\ &\quad + \left(\gamma \sum_{h=0}^{\infty} (\Phi \Psi_{h+1} + \Psi_h) \rho^{h+1} \right) \boldsymbol{\eta}_t + \Phi \Psi_0 \mathbf{u}_t \end{aligned}$$

$$\text{or} \quad E_t[\mathbf{x}_{t+1}] = \boldsymbol{\theta}_1 E_{t-1}[\mathbf{x}_t] + \boldsymbol{\theta}_2 \mathbf{x}_{t-1} + \boldsymbol{\Gamma}_{11} \boldsymbol{\eta}_t + \boldsymbol{\Gamma}_{12} \mathbf{u}_t \quad (33)$$

where

$$\boldsymbol{\theta}_1 \equiv \Phi + \rho \text{ and } \boldsymbol{\theta}_2 \equiv -\rho \Phi$$

$$\boldsymbol{\Gamma}_{11} \equiv \gamma \sum_{h=0}^{\infty} (\Phi \Psi_{h+1} + \Psi_h) \rho^{h+1} \text{ and } \boldsymbol{\Gamma}_{12} = \Phi \Psi_0$$

and thus, combining (33) and (28), the system has the reduced-form VAR repre-

sentation given in (14), where

$$\mathbf{\Gamma} = \begin{pmatrix} \mathbf{\Gamma}_{11} & \mathbf{\Gamma}_{12} \\ \sum_{h=1}^{\infty} \Psi_h \gamma_h & \mathbf{\Psi}_0 \end{pmatrix} \text{ and } \begin{pmatrix} \mathbf{e}_{1,t} \\ \mathbf{e}_{2,t} \end{pmatrix} = \mathbf{\Gamma} \begin{pmatrix} \boldsymbol{\eta}_t \\ \mathbf{u}_t \end{pmatrix}.$$

Equation (13) then follows from direct substitution.

Appendix C. 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 TBill 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 2015:4, as well as "nowcasts" of the contemporaneous data and "backcasts" of the lagged data. The main issue 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 levels, this is not always useful 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 vintage of 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 through $t + 4$ levels of GDP and CPI that are based on the same indexation as the 2015 data. Finally, for each quarter, we average the $t + 1$ through $t + h$ forecasts of the TBill rate to obtain its average forecasts over the following h quarters.

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 short-run BCS forecasts are reasonably straightforward, and we construct one-year expectations by averaging the forecasts for quarters $t + 1$ through $t + h$ in the last month of each quarter. However, the long-range BCS data 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 2015.

Specifically, since 1983, the BCS has been providing semiannual long-range consensus forecasts over two 5-year periods (1- to 6-year and 6- to 11-year) for various interest rates, including the average 3-month TBill 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 and 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.

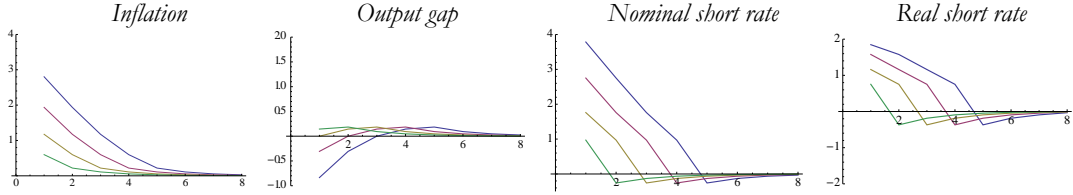
We address both of these issues through interpolation. Specifically, from 1983 to 1996, when the long-range forecasts are available only in March and October, we use the results from the BCEI and linearly interpolate to obtain June, September, and December values. Once the June and December values become directly observable, we interpolate to obtain only the September value. (Interpolation was not necessary for the short horizons, 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 over the first year following the survey, it is possible to compute the 6-year expectations by

taking the weighted average of the one-year and 1-6-year expectations, and the 11-year expectations by taking the weighted average of the one-year, 1-6-year, and 6-11-year expectations, respectively.

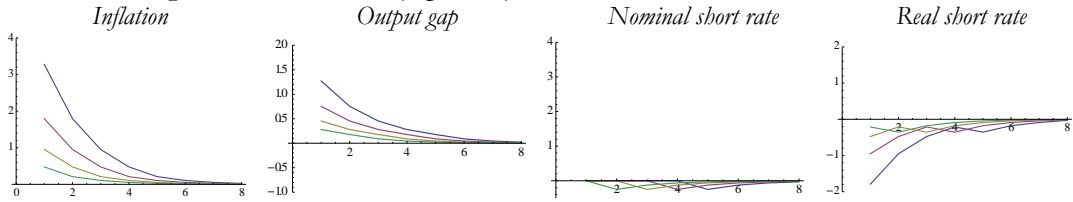
The three panels of Figure A1 plot the resulting time series of the survey-based expectations of the average 3-month TBill rate, CPI inflation, and GDP growth over the next year. The projections of the 3-month TBill rate and CPI inflation are very similar between the two surveys. In the case of GDP growth, on the other hand, the SPF projections are more volatile and, at least through about the year 2000, 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 TBill, 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 1. Monetary-policy news and forward-guidance shocks in the New Keynesian model

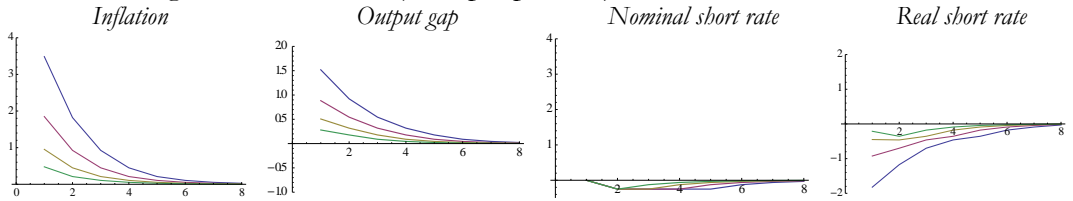
A. News shocks (1 period)



B. Forward-guidance scenario (1 period)

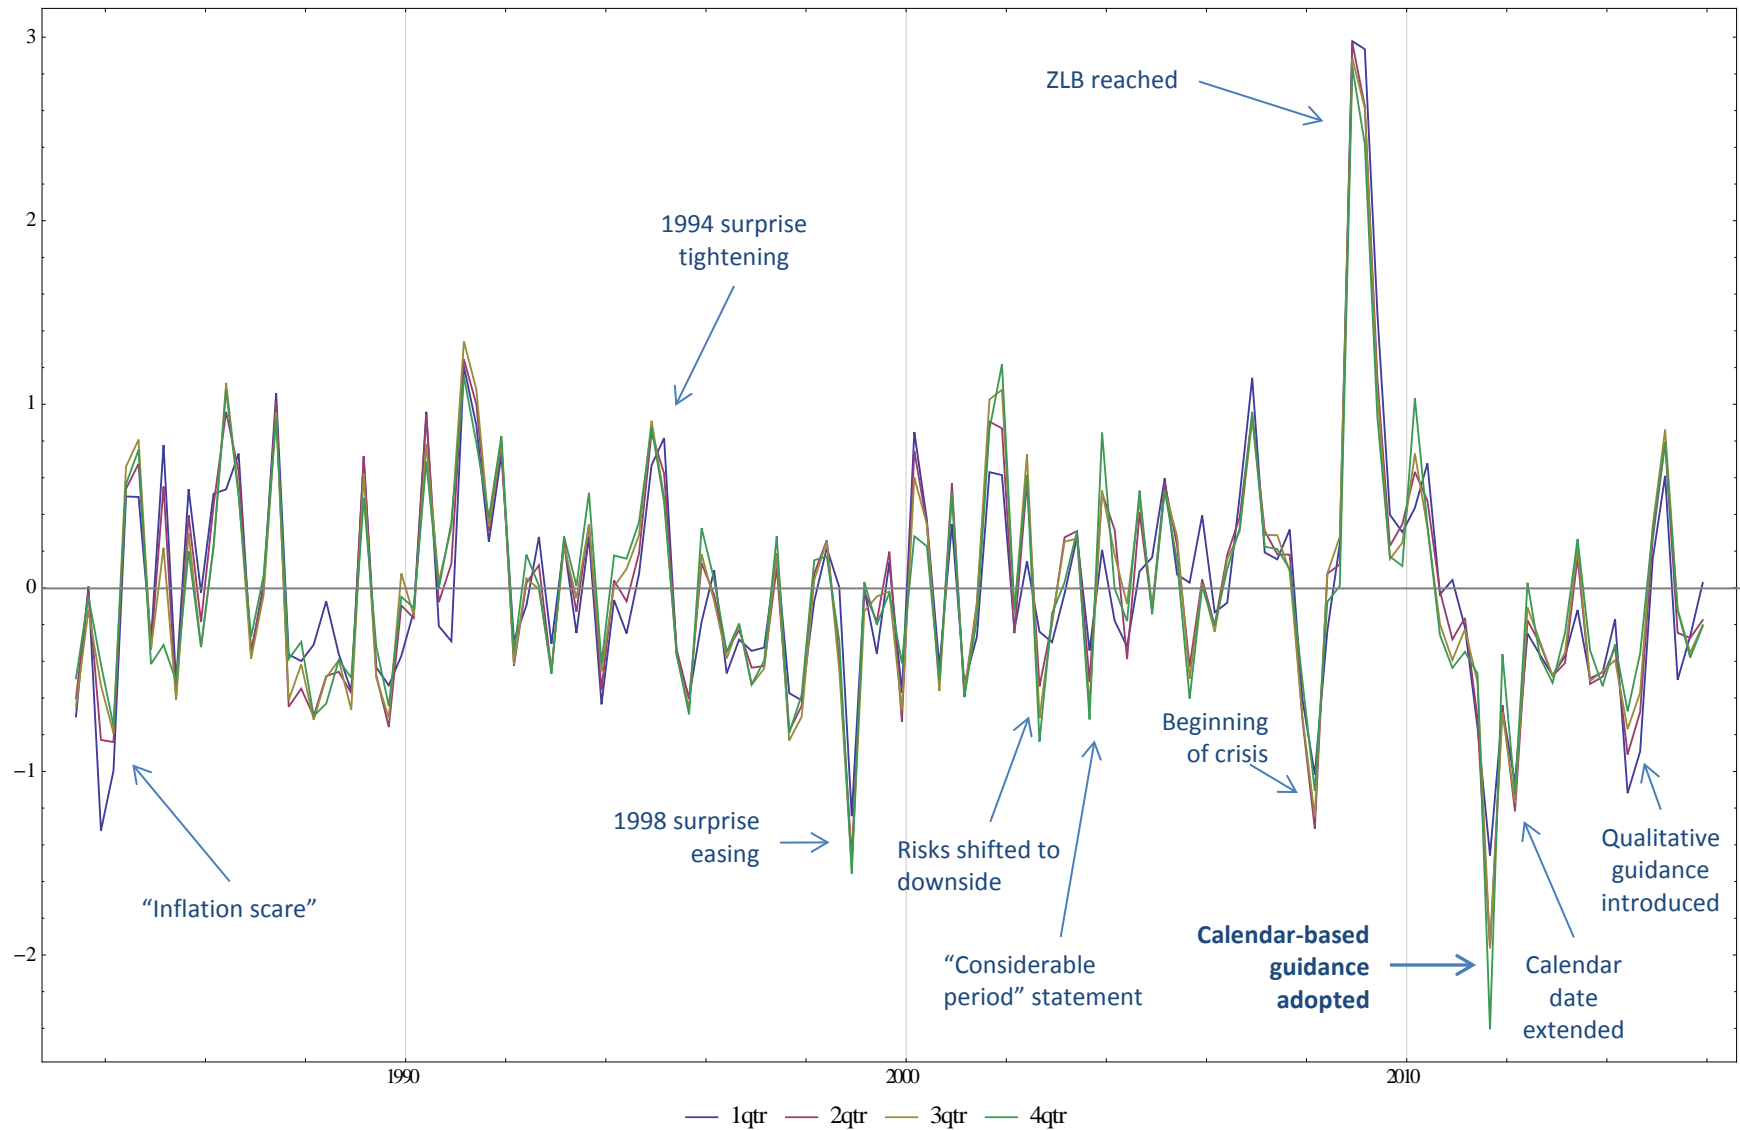


C. Forward-guidance scenario (multiple periods)



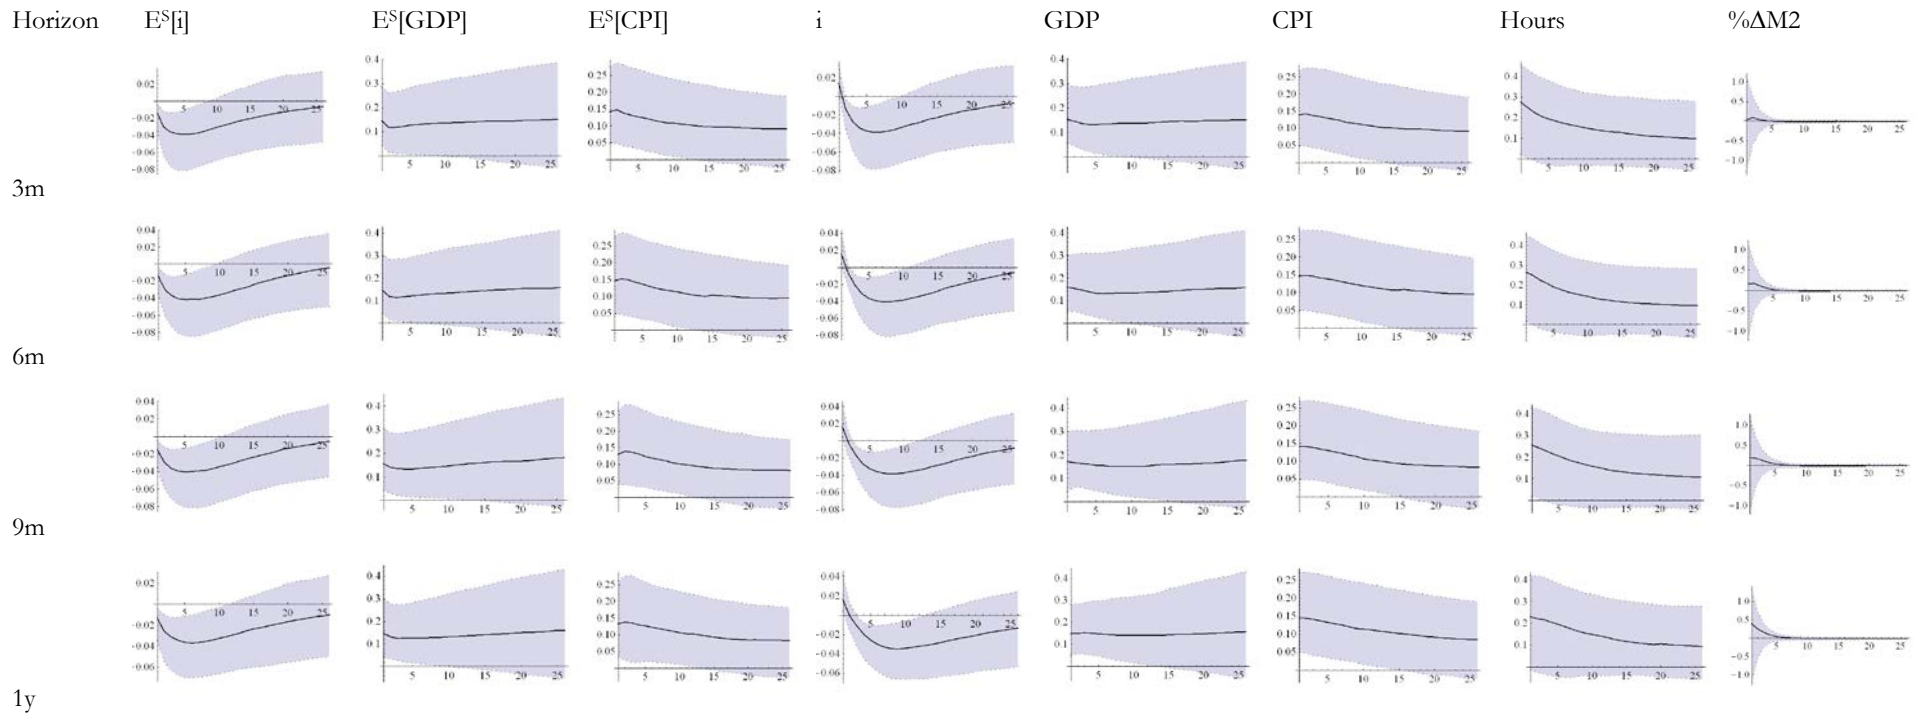
Notes: The figures show the effects of a monetary-policy news shock that is sufficient to lower expectations of the short-term interest rate b periods ahead by 25 basis points, where $b=1$ (green), 2 (yellow), 3 (red), and 4 (blue). In panel A, the central bank hews to its standard policy rule following the shock and therefore raises nominal interest rates in the short term. In panel B, the central bank introduces an unanticipated shock and a series of policy-expectations shocks in order to maintain the nominal short rate at zero in periods t through $t+b-1$. In panel C, the central bank holds the short rate at -0.25 in periods $t+1$ through $t+b$. Responses are shown for the period of the policy-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. Shorter-horizon policy-expectations shocks in the baseline identification



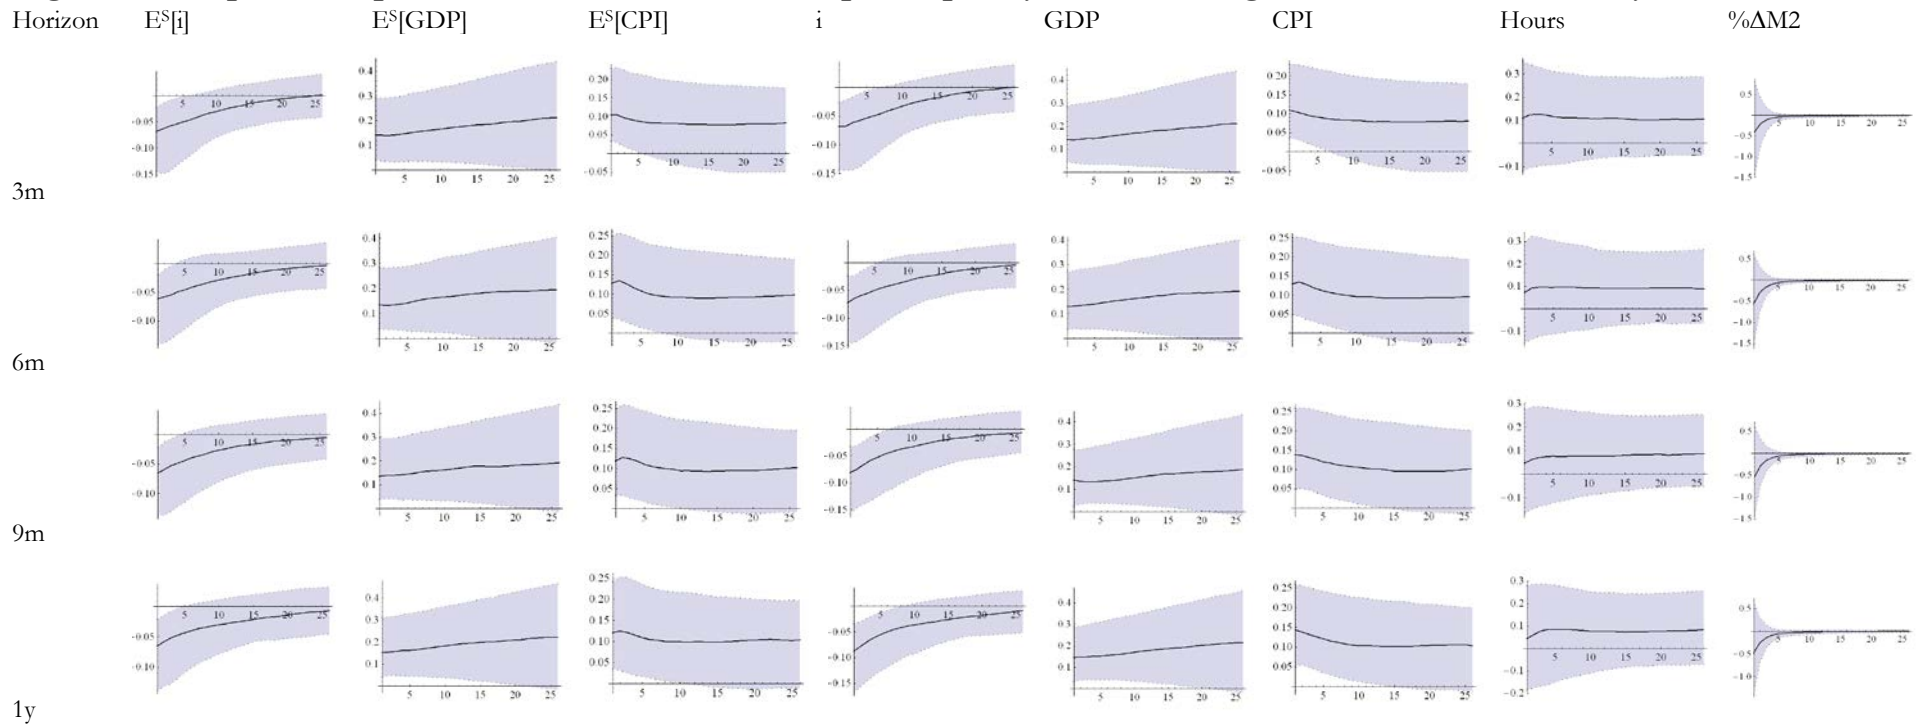
Note: The figure plots the estimated policy-expectations shocks in standard deviation unit (posterior medians) in our baseline models, using survey forecasts 1 quarter to 4 quarters ahead.

Figure 3. Impulse-response functions for policy-expectations shocks using shorter-horizon survey data



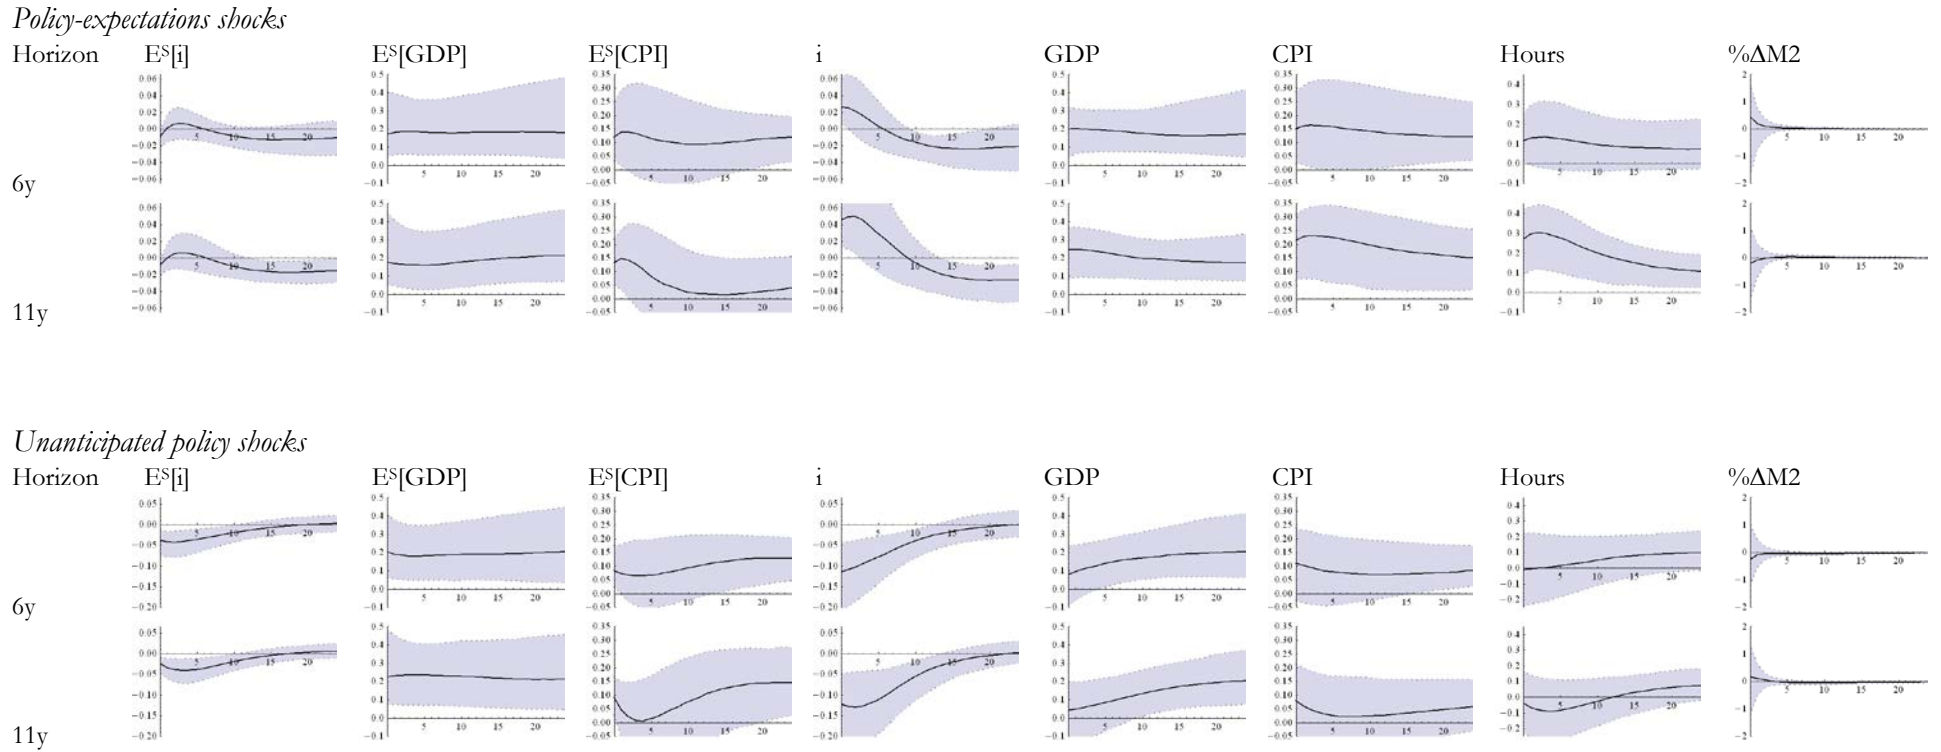
Note: The figure plots impulse-response functions following a one-standard-deviation policy-expectations shock, as defined by the restrictions (16) and (17), for each of the variables included in our baseline model specification. Each row corresponds to a VAR that includes survey forecasts at a different horizon, as indicated in the first column. The solid black lines are posterior medians, while the shading depicts the 16 – 84% posterior credibility regions, taking account of uncertainty about both the reduced-form parameter values and the structural factorization, as discussed in the text. All variables are in percentage points.

Figure 4. Impulse-response functions for unanticipated policy shocks using shorter-horizon survey data



Note: The figure plots impulse-response functions following a one-standard-deviation unanticipated policy shock, as defined by the restrictions (18) and (19), for each of the variables included in our baseline model specification. Each row corresponds to a VAR that includes survey forecasts at a different horizon, as indicated in the first column. The solid black lines are posterior medians, while the shading depicts the 16 – 84% posterior credibility regions, taking account of uncertainty about both the reduced-form parameter values and the structural factorization, as discussed in the text. All variables are in percentage points.

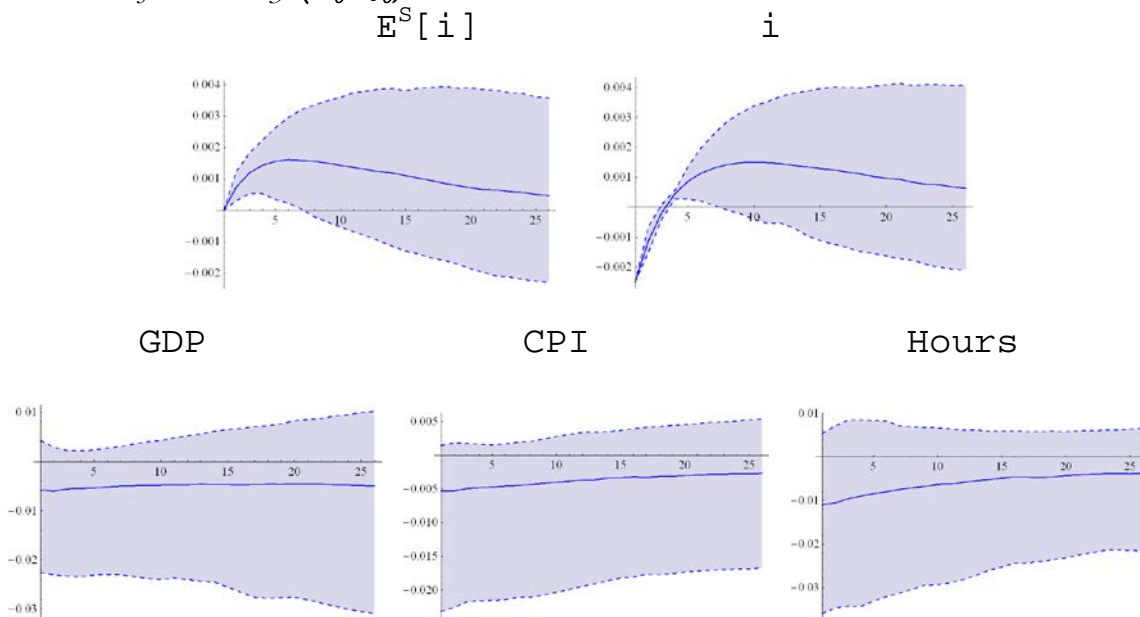
Figure 5. Impulse-response functions using longer-horizon survey data



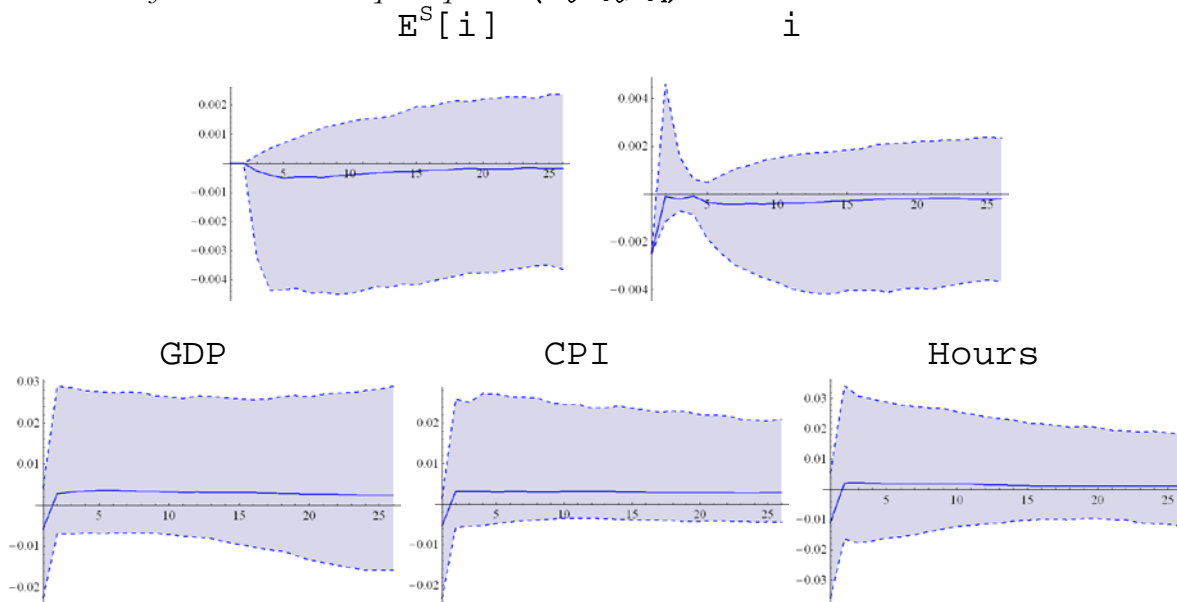
Note: The figure plots impulse-response functions following a one-standard-deviation policy-expectations shock (top panel) and unanticipated policy shock (bottom panel), as defined by the restrictions (16) through (19), for each of the variables included in our baseline model specification. Each row corresponds to a VAR that includes survey forecasts at a different horizon, as indicated in the first column. The solid black lines are posterior medians, while the shading depicts the 16 – 84% posterior credibility regions, taking account of uncertainty about both the reduced-form parameter values and the structural factorization, as discussed in the text. All variables are in percentage points.

Figure 6. Effect of unanticipated policy shock with the expectations channel neutralized.

A. *Period of shock only* ($u_\phi \eta_\phi$)

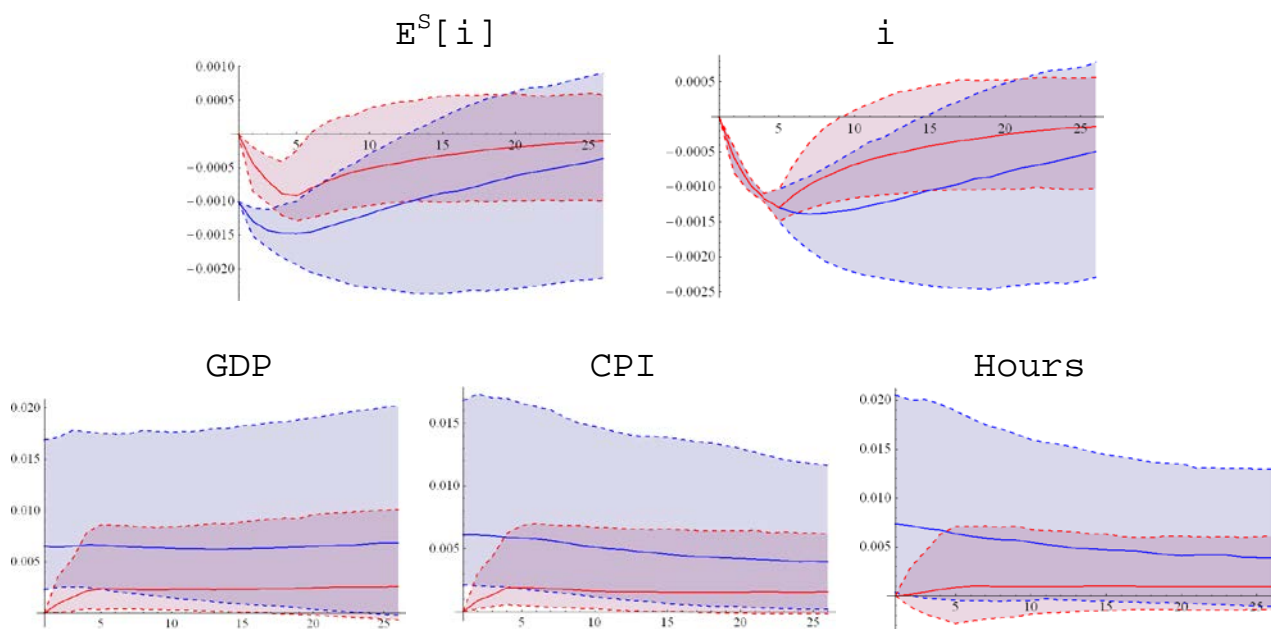


B. *Period of shock and subsequent period* ($u_\phi \eta_\phi \eta_D$)



Note: The figures show the results of simulations, using the results of our baseline VAR model with one-year survey data, in which there is an unanticipated shock to monetary policy and, simultaneously, an offsetting shock to policy expectations that leaves the survey forecast of the TBill rate unchanged, thus effectively shutting down the effect of policy expectations. In panel A, expectations are neutralized in the period of the unanticipated shock only; in panel B, they are also neutralized in the subsequent period. Solid lines are posterior medians, and shading depicts 16 – 84% credibility intervals. All variables are in percentage points.

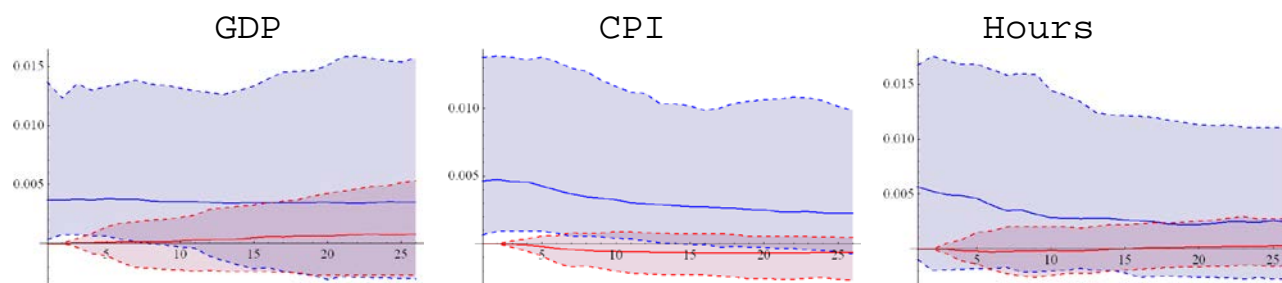
Figure 7. Estimated marginal effect of forward guidance



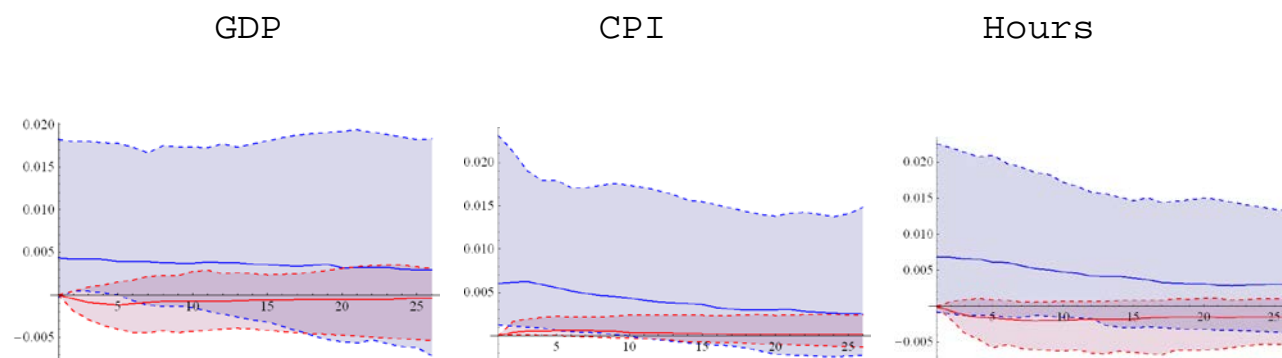
Note: The blue regions in the figure show the results of simulations, using the results of our baseline VAR model with one-year survey data, in which there is a policy-expectations shock and, simultaneously, an offsetting unanticipated-policy shock that leaves the current TBill rate unchanged. The red regions show the effects of a series of unanticipated policy shocks that exactly reproduce the blue path of the TBill rate over the subsequent four quarters. Thus, the difference between the two sets of impulse-response functions can be interpreted as the marginal effects of forward guidance about a given interest-rate path. Solid lines are posterior medians, and shading depicts 16 – 84% credibility intervals. All variables are in percentage points.

Figure 8. Forward guidance with alternative identifications for unanticipated shocks

A: Christiano-Eichenbaum-Evans (1999)



B: Uhlig (2005)



Note: The blue regions in the figures show the results of simulations in which there is a policy-expectations shock and, simultaneously, an offsetting unanticipated-policy shock that leaves the current TBill rate unchanged. The red regions show the effects of a series of unanticipated policy shocks that exactly reproduce the blue path of the TBill rate over the subsequent four quarters. Thus, the difference between the two sets of impulse-response functions can be interpreted as the marginal effects of forward guidance about a given interest-rate path. In panel A, unanticipated policy shocks are identified using short-run ordering restrictions similar to those of Christiano et al. (1999); in panel B, unanticipated policy shocks are identified using sign restrictions similar to those of Uhlig (2005). In both cases, policy-expectations shocks are identified as in our baseline model, and the survey horizon is four quarters. Solid lines are posterior medians, and shading depicts 16 – 84% credibility intervals. All variables are in percentage points.

Table 1. Marginal effects of forward guidance at different horizons

Effect on impact

FG horizon	GDP		CPI		Hours	
	Median	<i>Prob. > 0</i>	Median	<i>Prob. > 0</i>	Median	<i>Prob. > 0</i>
1y	0.65	<i>0.96</i>	0.61	<i>0.99</i>	0.74	<i>0.85</i>
6y	1.07	<i>0.91</i>	0.85	<i>0.92</i>	0.57	<i>0.83</i>
11y	1.30	<i>0.91</i>	1.21	<i>0.93</i>	1.16	<i>0.89</i>

Effect after 2 qtrs

FG horizon	GDP		CPI		Hours	
	Median	<i>Prob. > 0</i>	Median	<i>Prob. > 0</i>	Median	<i>Prob. > 0</i>
1y	0.42	<i>0.87</i>	0.38	<i>0.86</i>	0.60	<i>0.77</i>
6y	1.08	<i>0.92</i>	0.82	<i>0.86</i>	0.62	<i>0.79</i>
11y	1.35	<i>0.93</i>	1.20	<i>0.91</i>	1.24	0.90

Effect after 4 qtrs

FG horizon	GDP		CPI		Hours	
	Median	<i>Prob. > 0</i>	Median	<i>Prob. > 0</i>	Median	<i>Prob. > 0</i>
1y	0.29	<i>0.75</i>	0.27	<i>0.75</i>	0.53	<i>0.72</i>
6y	1.04	<i>0.91</i>	0.77	<i>0.83</i>	0.59	<i>0.75</i>
11y	1.32	<i>0.93</i>	1.11	<i>0.89</i>	1.15	<i>0.89</i>

Effect after 8 qtrs

FG horizon	GDP		CPI		Hours	
	Median	<i>Prob. > 0</i>	Median	<i>Prob. > 0</i>	Median	<i>Prob. > 0</i>
1y	0.27	<i>0.71</i>	0.26	<i>0.73</i>	0.39	<i>0.70</i>
6y	0.92	<i>0.88</i>	0.61	<i>0.78</i>	0.48	<i>0.73</i>
11y	1.23	<i>0.92</i>	1.02	<i>0.87</i>	1.02	<i>0.90</i>

Effect after 20 qtrs

FG horizon	GDP		CPI		Hours	
	Median	<i>Prob. > 0</i>	Median	<i>Prob. > 0</i>	Median	<i>Prob. > 0</i>
1y	0.26	<i>0.65</i>	0.18	<i>0.66</i>	0.24	<i>0.68</i>
6y	0.70	<i>0.82</i>	0.51	<i>0.81</i>	0.39	<i>0.71</i>
11y	1.01	<i>0.90</i>	0.82	<i>0.85</i>	0.89	<i>0.82</i>

Notes: The table shows the estimated effect of a pre-announced exogenous decline of 10 basis points in the short-term interest rate over various horizons, *minus* the effect of the *same* path of policy when it is not announced in advance. (The exercise is the same as that depicted, for the 4-quarter horizon, in Figure 7.) The columns labeled “Prob. > 0” report the posterior probability that the pre-announced policy results in a greater change than the unanticipated policy after the indicated number of quarters. All variables are in percentage points.

Table 2. Effects of policy-expectations shocks under alternative identifying assumptions and implementations.

1-Year Survey Horizon

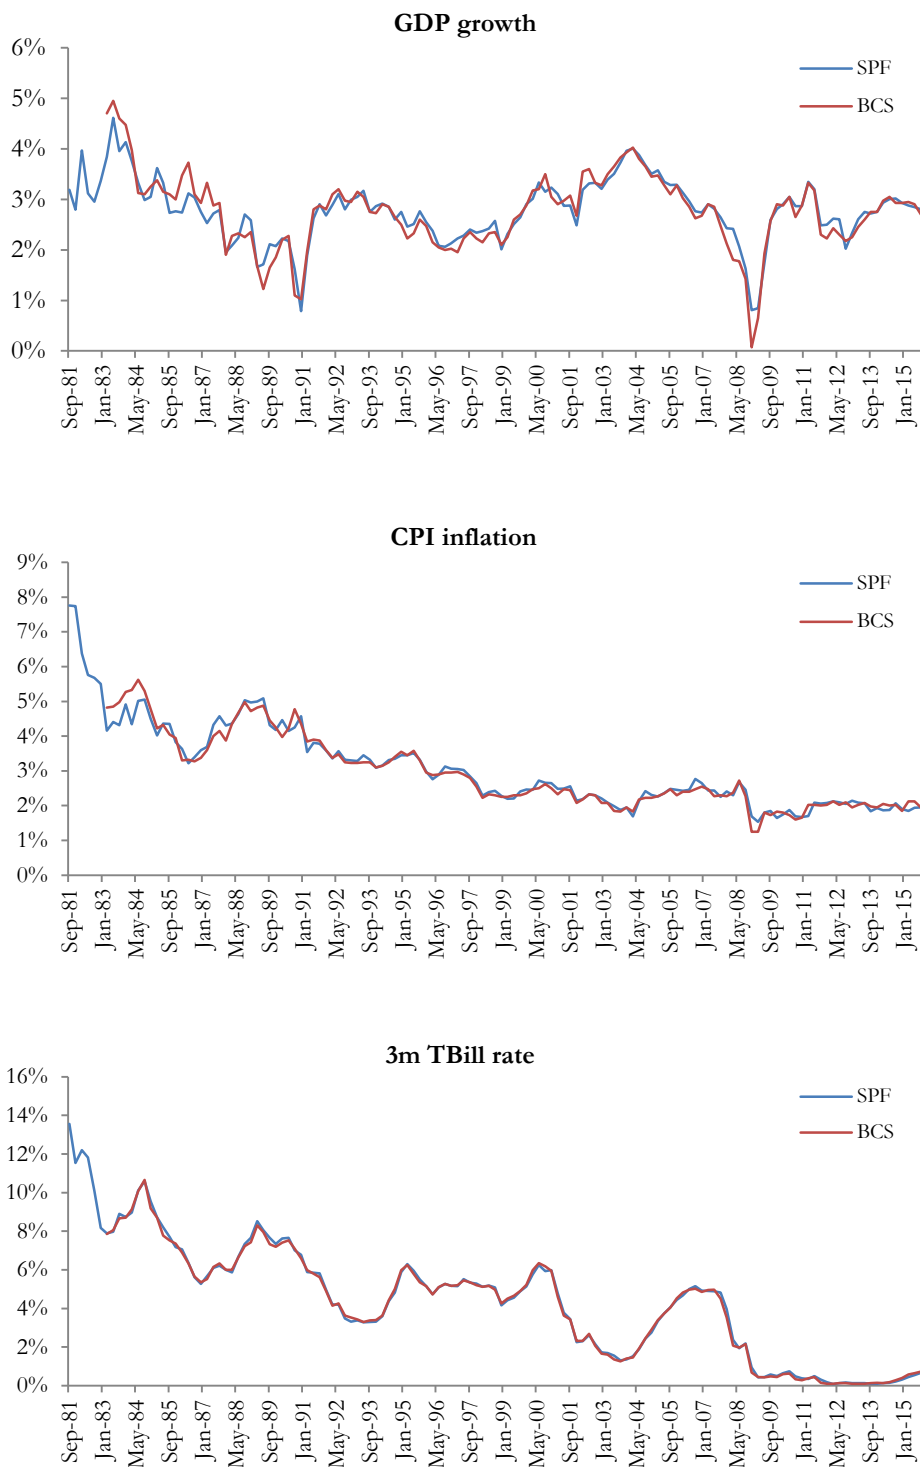
MODEL	Initial Reaction of $E^S[i]$	<u>GDP</u>			<u>CPI</u>			<u>Hours</u>		
		0	1y	5y	0	1y	5y	0	1y	5y
Baseline	-0.02*	0.16*	0.17*	0.18	0.16*	0.15*	0.09	0.21	0.18	0.11
Flat IRF prior	-0.01*	0.21*	0.23*	0.23*	0.14*	0.15*	0.11*	0.27*	0.23*	0.16*
$\Gamma_{\eta}^i = 0$	-0.01*	0.16*	0.17*	0.17	0.14*	0.14*	0.10*	0.18	0.15	0.10
Excluding ZLB	-0.02*	0.17*	0.14*	0.10	0.12*	0.15*	0.12*	0.12	0.05	0.02
SPF	-0.01*	0.19*	0.14*	0.11	0.18*	0.18*	0.14*	0.14	0.08	0.06
CEE	-0.02*	0.18*	0.18*	0.18	0.18*	0.17*	0.11	0.23*	0.19	0.13
Uhlig	-0.02*	0.18*	0.22*	0.25*	0.14*	0.13*	0.10*	0.22*	0.17	0.14

6-Year Survey Horizon

MODEL	Initial Reaction of $E^S[i]$	<u>GDP</u>			<u>CPI</u>			<u>Hours</u>		
		0	1y	5y	0	1y	5y	0	1y	5y
Baseline	-0.01*	0.18*	0.19*	0.21*	0.15*	0.17*	0.14*	0.13*	0.14	0.09
Flat IRF prior	-0.01*	0.19*	0.21*	0.22*	0.18*	0.18*	0.13*	0.24*	0.22*	0.13
$\Gamma_{\eta}^i = 0$	-0.02*	0.14*	0.17*	0.20*	0.14*	0.13*	0.11*	0.16	0.13	0.10
Excluding ZLB	-0.01*	0.16*	0.16*	0.20*	0.16*	0.19*	0.15*	0.19*	0.19*	0.10*
CEE	-0.02*	0.18*	0.19*	0.23*	0.15*	0.13*	0.11*	0.19*	0.14	0.12*
Uhlig	-0.01*	0.22*	0.21*	0.19*	0.20*	0.21*	0.15*	0.26*	0.25*	0.13*

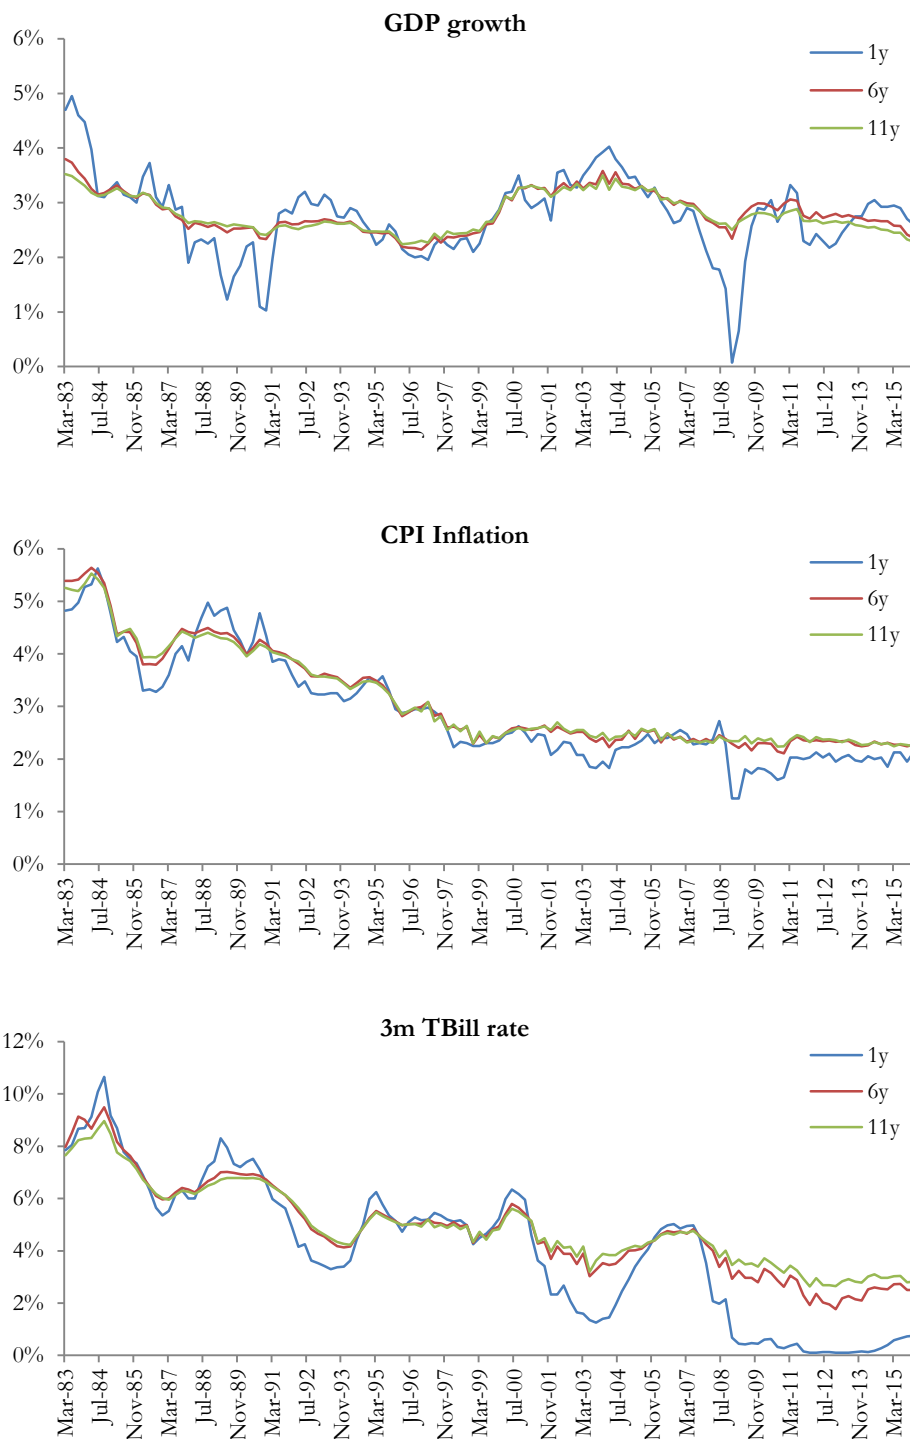
Notes: The tables show the effects of one-standard-deviation policy-expectations shocks on economic variables at different horizons, using several alternative identification schemes and VAR implementations, as described in the text. The top panel reports the results for models using one-year survey forecasts; the bottom panel reports the results for models using 6-year survey forecasts. Reported values are posterior medians. Asterisks indicate that at least 84% of the posterior probability mass has the same sign as the median. All variables are in percentage points.

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



Notes: The figures depict the consensus forecasts from the Blue Chip Survey (BCS) and Survey of Professional Forecasters (SPF) used in our VAR models, at the one-year horizon. Construction of the data is described in the text of Appendix C.

Figure A2. Term structure of BCS forecasts



Notes: The figures depict the long-range forecasts from the Blue Chip Survey used in our VAR models, and compare them to the one-year forecasts. Construction of the data is described in the text of Appendix C.