ISSN 1936-5330

kc FED Research Working Papers

The Dynamic Effects of Forward Guidance Shocks

Brent Bundick and A. Lee Smith

January 2016; Revised June 2019
RWP 16-02
https://doi.org/10.18651/RWP2016-02

Federal Reserve Bank of Kansas City

The Dynamic Effects of Forward Guidance Shocks*

Brent Bundick† A. Lee Smith‡

First Version: January 2016
Current Version: June 2019

Abstract

We examine the macroeconomic effects of forward guidance shocks at the zero lower bound. Empirically, we identify forward guidance shocks using unexpected changes in futures contracts around monetary policy announcements. We then embed these policy shocks in a vector autoregression to trace out their macroeconomic implications. Forward guidance shocks that lower expected future policy rates lead to moderate increases in economic activity and inflation. After examining forward guidance shocks in the data, we show that a standard model of nominal price rigidity can reproduce our empirical findings. To estimate our theoretical model, we generate a model-implied futures curve which closely links our model with the data. Our results suggest no disconnect between the empirical effects of forward guidance shocks around policy announcements and the predictions from a standard theoretical model.

JEL Classification: E32, E52
Keywords: Monetary Policy Shocks, Forward Guidance, Zero Lower Bound, Impulse Response Matching

*We thank Susanto Basu, Troy Davig, John Keating, Christopher Otrok, Nick Sly, and Stephen Terry for helpful discussions. We thank the editor and three anonymous referees for helpful feedback. We also thank our formal discussants, James D. Hamilton, Silvia Miranda-Agrippino, and Benjamin Born, for useful comments. We have benefited from comments made by participants at various seminars and conferences. Thealexa Becker, Trenton Herriford, and Emily Pollard provided excellent research assistance and we thank CADRE for computational support. The views expressed herein are solely those of the authors and do not necessarily reflect the views of the Federal Reserve Bank of Kansas City or the Federal Reserve System.
Available: http://dx.doi.org/10.18651/RWP2016-02.
Authors' final version. Article accepted for publication in *The Review of Economics and Statistics*.

†Federal Reserve Bank of Kansas City. Email: brent.bundick@kc.frb.org

‡Federal Reserve Bank of Kansas City. Email: andrew.smith@kc.frb.org

1 Introduction

In December 2008, the Federal Open Market Committee (FOMC) lowered the federal funds rate to its effective lower bound. With economic conditions continuing to deteriorate and its conventional policy tool unavailable, the Federal Reserve announced its intention to keep future policy rates exceptionally low "for some time." Such communication about the future path of policy, known as forward guidance, became a fixture of U.S. monetary policy.

However, recent theoretical and empirical works are divided on the macroeconomic effects of forward guidance. In standard models with nominal price rigidities, Eggertsson and Woodford (2003) show that lowering the expected path of policy rates can effectively stimulate economic activity and increase inflation. However, Del Negro, Giannoni and Patterson (2015), McKay, Nakamura and Steinsson (2016), and others argue that these theoretical models overstate the expansionary effects of forward guidance. In contrast, empirical work by Campbell et al. (2012) and Nakamura and Steinsson (2018) argues that communicating lower expected rates may signal bad news about the state of the economy. Through this macroeconomic news effect, these papers suggest that lowering expected policy rates may cause a contraction in expected economic activity.

We examine this apparent disconnect between the empirical evidence and theoretical predictions of macroeconomic models. First, we study the empirical effects of forward guidance shocks at the zero lower bound (ZLB). We identify forward guidance shocks in the data using high-frequency changes in futures contracts around FOMC announcements. To trace out the dynamic effects of these policy changes on macroeconomic aggregates, we embed our identified forward guidance shocks in a standard vector autoregression (VAR). We find forward guidance shocks that lower expected future policy rates result in a persistent economic expansion. Following a one standard deviation forward guidance shock, which lowers 8-quarter ahead futures rate by about 6 basis points, output increases by about 15 basis points and prices are about 5 basis points higher at their peak responses. Similar to conventional policy shocks, we find that forward guidance shocks explain only a small fraction of overall business-cycle fluctuations. Our findings are robust to alternative ordering schemes in the VAR, different measures of economic activity and prices, and alternative measures of expected future interest rates. We also document similar macroeconomic effects when we estimate our empirical model prior to the onset of the zero lower bound.

After identifying forward guidance shocks in the data, we examine their effects in a stan-

dard model of nominal price rigidity. Using a nonlinear solution method, we estimate a standard New-Keynesian model with a zero lower bound constraint. We model a forward guidance shock as an exogenous innovation to the central bank's desired policy rate at the zero lower bound. When desired rates are less than zero, shocks that reduce the desired rate act like an exogenous extension of the zero lower bound episode. This exogenous extension of the zero lower bound lowers expected future policy rates, which we link with our identified forward guidance shock in the data. To appropriately calibrate our forward guidance shock process, we generate a model-implied counterpart to the futures contracts used in our empirical results. Using impulse response matching, we choose the parameters of our nonlinear model such that a forward guidance shock in the model generates the same movements in futures rates that we observe in the data.

Our theoretical model can reproduce the macroeconomic effects of forward guidance shocks we find in the data. In the model, an exogenous decline in expected future policy rates generates movements in economic activity and prices similar in shape and magnitude to our empirical responses. The key features of our model are a reasonable degree of nominal price rigidity, habits in household consumption, investment adjustment costs, and variable capital utilization. Our results suggest that dynamic equilibrium models along the lines of Christiano, Eichenbaum and Evans (2005) remain useful in examining the effects of monetary policy shocks both at and away from the zero lower bound.

We find no disconnect between the empirical effects of forward guidance shocks around FOMC announcements and the predictions from a standard theoretical model. Our findings contrast with Del Negro, Giannoni and Patterson (2015), who argue that standard models with nominal rigidities overestimate the expansionary effects of forward guidance. Our alternative conclusion emerges from the size of the forward guidance shock we estimate. A typical expansionary forward guidance shock around a monetary policy announcement lowers 8-quarter ahead futures rates by about 6 basis points. This shock extends the zero lower bound duration by only one month in our model. Del Negro, Giannoni and Patterson (2015), however, simulate a much longer one-year exogenous extension of the zero lower bound period. In our high-frequency identification of policy shocks around FOMC announcements, we do not observe forward guidance shocks of that size. Thus, our results suggest that standard models work well in analyzing the size of forward guidance shocks we observe in the data around FOMC announcements. However, our findings cannot speak to the plausibility of the model's predictions for substantially larger shocks.

2 Forward Guidance Shocks in the Data

We use a two-step procedure to examine the macroeconomic effects of forward guidance shocks in the data. First, we identify forward guidance shocks associated with regularly-scheduled FOMC meetings using high-frequency changes in interest rate futures. Then, we embed these policy shocks into a Bayesian VAR to trace out their dynamic effects on macroeconomic aggregates. In our baseline results, we focus on the effects of forward guidance shocks during the zero lower bound period (December 2008 - December 2015). After presenting our baseline empirical results, we then also examine the effects of forward guidance shocks prior to the onset of the zero lower bound in Section 2.6.

2.1 High-Frequency Futures Data

We use a combination of federal funds and eurodollar futures contracts to measure unexpected changes in forward guidance. For each regularly-scheduled FOMC meeting from 1994–2015, we compute the daily change in the current month or 3-month ahead federal funds futures rates and the 2-8 quarter ahead eurodollar futures rates.[1] Since any expected changes in policy should be reflected in futures prices ahead of the meeting, the change in futures prices on the day of the meeting provides a measure of the unexpected portion of the policy announcement. Following Gurkaynak, Sack and Swanson (2005), we extract a target and path factor that together summarize almost all of the variation in these futures rates around policy announcements. In this paper, we focus on the path factor which captures unexpected changes to the future path of policy rates that are unrelated to changes in the current policy rate.[2] We scale the path factor so that it moves the 8-quarter ahead eurodollar futures rate one-for-one around FOMC meetings.

The path factor displays significant variation both prior to and during the zero lower bound period and these fluctuations line up with key changes in FOMC forward guidance. Figure 1 plots our forward guidance shock series from 1994–2015 and annotates the dates associated with the some of the largest fluctuations. During the zero lower bound period, we observe large declines in the expected path of rates when the FOMC announced its intention to keep future policy rates exceptionally low for "some time" in December 2008, when "some time" was replaced by "an extended period" in March 2009, and after the

[1] The payoffs of these contracts depend the effective federal funds rate for federal funds futures and the 3-month LIBOR on dollar denominated deposits for eurodollar futures.

[2] In the appendix, we provide additional details on the construction of our target and path factors and provide a comparison with the Gurkaynak, Sack and Swanson (2005) path factor for the overlapping sample.

change to date-based guidance in August 2011. Figure 1 also illustrates that the magnitude of the movements in the path factor is similar before and during the zero lower bound period, which suggests that the FOMC used forward guidance in both periods. Moreover, several of the largest path factor observations before December 2008 coincide with the observations in Table 4 of Gurkaynak, Sack and Swanson (2005) who carefully document the use of forward guidance by the FOMC since the early 1990s. Thus, we leverage this pre-zero lower bound sample in our empirical analysis in two ways. First, we use this earlier sample to elicit priors for our VAR parameters at the zero lower bound. Second, we estimate the macroeconomic effects of forward guidance before the onset of the zero lower bound which, unlike the zero lower bound sample, was not accompanied by simultaneous large-scale asset purchases.

2.2 Baseline Empirical Model

To trace out the macroeconomic effects of a forward guidance shock, we embed the cumulative sum of the path factor in a structural VAR.[3] We estimate our baseline empirical model at a monthly frequency using several indicators of real economic activity, a measure of aggregate prices, and a government bond yield. Specifically, we include a monthly measure of real GDP, a proxy for real equipment investment, capacity utilization, the GDP deflator, the path factor, and the 2-year Treasury yield. The choice to include the 2-year Treasury yield follows Gertler and Karadi (2015) and others who argue that the FOMC's forward guidance operates with roughly a two-year horizon. We use the Macroeconomic Advisers monthly GDP series and its corresponding price deflator to measure aggregate real activity and prices. We proxy equipment investment at a monthly frequency with deflated core capital goods shipments, a series the Bureau of Economic Analysis uses to calculate the official quarterly investment data. Appendix B contains details on the data construction. GDP, the GDP deflator, and investment enter the VAR in natural log form.[4]

Following much of the previous VAR literature studying the effects of conventional mon-

[3]This approach follows Romer and Romer (2004) and Barakchian and Crowe (2013) who point out that including the cumulative sum of unexpected interest rate changes in a VAR is most consistent with the many VAR models which include the federal funds rate in levels. Following these papers, we assign a value of zero to months in which there is no FOMC meeting before cumulatively summing the path factor series.

[4]We could instead follow Romer and Romer (2004) and estimate the coefficients of interest from the vector moving average (VMA) more directly by regressing macroeconomic aggregates on lags of the path factor. However, this approach implies that we lose a number of initial data points. Since our sample is already limited, we use a VAR approach which is asymptotically equivalent assuming that the VMA has a VAR representation.

etary policy shocks, we order our forward guidance shock measure after real activity and the price level but before the 2-year Treasury yield using a recursive identification scheme. This ordering assumes that macroeconomic conditions adjust slowly to changes in expected policy rates but financial markets may respond immediately. At a monthly frequency, the assumption that a monetary policy announcement today does not affect real activity or prices within the period seems plausible. However, our results are not sensitive to this ordering. In Section 2.5, we show that our results are unchanged if we order the path factor first.

We estimate and conduct statistical inference on the VAR from a Bayesian perspective. Our primary interest is to examine the effects of forward guidance shocks during the zero lower bound period. However, Figure 1 illustrates that changes in FOMC forward guidance around policy announcements also occurred prior to the onset of the zero lower bound. Therefore, we use the data from the pre-zero lower bound period as a pre-sample to form our priors for the VAR parameters during the zero lower bound period. We refer to this as our empirical Bayes prior.[5] In Section 2.5, however, we show that we find similar results if we instead use a standard Minnesota prior or an uninformative prior which centers the VAR coefficients at the OLS estimates during the zero lower bound period. Using standard information selection criteria, we include three lags in the VAR.[6]

2.3 Empirical Impulse Responses

We now turn to our key empirical question: What are the macroeconomic effects of forward guidance shocks? Figure 2 plots the estimated impulse responses to an identified forward guidance shock along with 90% probability intervals. A one standard deviation forward guidance shock lowers the path factor by about 6 basis points. Since we normalized our path factor to move one-for-one with the 8-quarter eurodollar future, this decline in the path factor implies that two-year ahead futures rates also decline by about 6 basis points.

A forward guidance shock induces a gradual expansion in economic activity accompanied by increases in investment and capacity utilization and some modest inflationary pressures. Per our ordering assumption, economic activity and prices remain unchanged at impact. In the following months, output rises in a hump-shaped pattern and remains elevated for the

[5]More specifically, for the pre-zero lower bound period, we assume a non-informative natural conjugate Gaussian-Inverse Wishart prior such that the posterior distribution of the VAR parameters is based on the OLS quantities. Then, we use this posterior distribution to define our prior for the VAR parameters at the zero lower bound period.

[6]In the appendix, we illustrate that our results are robust to including twelve lags in the VAR.

next four years. At its peak response, output increases by almost 15 basis points. Investment and capacity utilization also exhibit hump-shaped responses and remain elevated for several years. However, they rise by more than output and reach their peaks more quickly. Prices move up slowly over the horizon of the impulse response and level out after three or four years. Notably, the increase in economic activity and prices persists well beyond the time that expected future rates remain depressed. For instance, the 2-year Treasury yield initially declines with expected future rates but starts to overshoot after just two years. However, output and prices remain elevated throughout the impulse response horizon. Thus, we find that forward guidance shocks share many of the same empirical features of conventional policy shocks as identified by Christiano, Eichenbaum and Evans (2005) and others.

These estimated responses contrast with the work of Campbell et al. (2012) and Nakamura and Steinsson (2018), which finds that an unexpected decline in future policy rates *lowers* forecasts of real activity and inflation. In the appendix, we show that meaningful differences in our policy shock series relative to these two papers help to reconcile our differing conclusions. However, we prefer our path factor shock series for two reasons. First, prior to the zero lower bound, our estimated path factor shock moves closely with the path factor developed in Gurkaynak, Sack and Swanson (2005) and thus is consistent with the seminal work in this area. Second, changes in our path factor during the zero lower bound are consistent with the narrative of FOMC communications. Around some key policy announcements, the policy shocks series of Nakamura and Steinsson (2018) and Campbell et al. (2012) either show little variation or have signs that differ from market commentary at that time.

2.4 Forecast Error Variance Decompositions

Similar to conventional policy shocks, forward guidance shocks only account for a small fraction of business-cycle fluctuations. At the two-year horizon, we find that forward guidance shocks explain less than 10 percent of the total unexpected fluctuations in output. For comparison, we also estimate variance decompositions for conventional monetary policy shocks based on the VAR models of Romer and Romer (2004) and Christiano, Eichenbaum and Evans (2005). These models find that conventional policy shocks explain 25 and 41 percent (respectively) of the unexpected fluctuations in output over the two-year horizon. While forward guidance shocks account for a smaller fraction of the variation in output at business-cycle frequencies compared to conventional monetary policy shocks, these differences do not appear to be statistically meaningful.

2.5 Robustness of Empirical Results

Our baseline results show that a decline in the expected path of the policy rate leads to a persistent expansion in economic activity, a gradual rise in prices, and lower Treasury yields. Before asking whether a standard theoretical model can match these estimated responses, we explore the robustness of our empirical findings. Figures 3 and 4 show the impulse responses to a forward guidance shock from several alternative specifications of the VAR model. To keep these figures easily readable, we omit the probability intervals in the main text. In the appendix, however, we present the probability intervals that accompany the point estimates for each alternative empirical specification. Based on the analysis in this section, we conclude that the qualitative and quantitative features of our estimated responses to a forward guidance shock are robust across several dimensions of our VAR model.

The estimated responses to a forward guidance shock do not meaningfully change when we alter the ordering of the variables in the VAR or use different macroeconomic indicators. Figure 3 shows the impulse responses when we order our policy surprise series first in our recursive VAR. This ordering interprets the policy surprises as predetermined with respect to macroeconomic aggregates and allows economic activity and prices to respond immediately to the forward guidance shock. When the path factor is ordered first, the point estimates of the responses of all the variables are almost identical to our baseline VAR model. Figure 3 also shows the impulse responses if we measure economic activity and prices using industrial production and the consumer price index, the same variables Gertler and Karadi (2015) use in their study of monetary policy shocks. Using these alternative indicators, the peak responses of output and prices occur a bit earlier and the response of output is a bit larger than our baseline model. However, when we account for the uncertainty surrounding the estimated responses, neither difference appears to be significant.

We find similar macroeconomic effects if we use a Minnesota prior rather than the empirical Bayes prior we employ in our baseline specification. This prior is standard in the VAR literature and balances the need to capture the rich dynamics in the data with the concern of over-fitting the data by adding too many lags. Figure 3 illustrates that the responses of real variables and prices are remarkably similar under a Minnesota prior with 13 lags. The only notable difference is the response of the 2-year Treasury yield which exhibits a shallower decline and subsequently over-shoots by more in later months. However, this quantitative difference with our baseline specification is not outside the range of 90% probability intervals for the two empirical models. Figure 3 also shows the estimated responses if we further

restrict the VAR coefficients such that we treat the high-frequency policy surprises as exogenous. In particular, we adjust the Minnesota prior so that only own lags have non-zero coefficients in the path factor equation. The dynamics of the path factor following a forward guidance shock in this alternative specification are similar to those in Figure 2; however, the path factor and 2-year Treasury yield no longer overshoot. The responses of macroeconomic aggregates are also similar compared to the baseline impulse responses.[7]

The estimated effects of forward guidance shocks do not depend on the use of informative priors. Figure 4 shows the impulse responses if we center the VAR parameters at the OLS estimates over the zero lower bound period rather than using an empirical Bayes or Minnesota prior.[8] Using only data beyond December 2008 to inform the VAR coefficients, we observe slightly larger responses of investment, capacity utilization, and prices as well as a more persistent response of overall output to a forward guidance shock. Yet, the hallmark features we previously documented remain, including the hump-shaped responses of real variables, the gradual rise in prices, and the persistence of these effects beyond the reduction in rates.

Up to this point, we have relied on the path factor to measure forward guidance surprises that are orthogonal to unexpected changes in the current federal funds rate. However, focusing only on the zero lower bound period, we can also identify forward guidance shocks by examining fluctuations in the raw interest rate futures around policy announcements since the target range for the federal funds rate remained unchanged during this period. Figure 4 shows that we find similar macroeconomic effects to our path factor during the zero lower bound if we instead use 4-, 8-, or 12-quarter ahead eurodollar futures to measure forward guidance shocks. This finding confirms that the path factor is able to synthesize the behavior of many futures rates around FOMC meetings. The stability of our results when we measure forward guidance shocks using the 12-quarter ahead eurodollar rate is particularly reassuring for two reasons. First, this horizon of rates lies beyond the range of futures we use to construct our path factor. Second, Swanson and Williams (2014) suggest that this rate was less constrained during the zero lower bound period.

[7] In the appendix, we compare our approach to the Proxy VAR approach in Gertler and Karadi (2015) and document similar effects on output following a forward guidance shock.

[8] Given the short sample, we include only one lag for the VARs estimated with uninformative priors.

2.6 Quantitative Easing & Forward Guidance Before 2009

The previous section provides evidence that forward guidance announcements that lower the expected path of rates at the zero lower bound lead to a modest but statistically significant expansion of economic activity. During this period, however, the FOMC also conducted several rounds of large-scale asset purchases known as quantitative easing. Similar to forward guidance, the stated goal of these asset purchases was to ease financial conditions and promote economic activity. Announcements regarding these asset-purchase programs often appeared alongside changes in the FOMC's forward guidance. Thus, one may be concerned that some of the macroeconomic effects we attribute to forward guidance actually emanate from large-scale asset-purchases.

If asset purchases solely operate by signaling the path of future short-term interest rates, then the simultaneous quantitative easing announcements would not affect our results. For example, Krishnamurthy and Vissing-Jorgensen (2011), Woodford (2012), Bauer and Rudebusch (2014), and Bhattarai, Eggertsson and Gafarov (2015) argue that asset purchases acted as a commitment device to reinforce the FOMC's guidance about future policy rates. However, if asset purchases also operate through a portfolio-rebalancing channel, then the simultaneous quantitative easing announcements could bias our estimates of the macroeconomic effects of forward guidance. To address this concern, we estimate three additional empirical specifications which are shown in Figure 5.[9]

The first two robustness checks continue to focus on the use of forward guidance during the zero lower bound period. First, we simply omit observations from the path factor series that coincided with the announcement of a new asset purchase program.[10] Figure 5 illustrates that dropping these observations results in impulse responses that are similar to our baseline estimates. Next, we include one-year ahead Blue Chip forecasts for the short-term interest rates in place of the 2-year Treasury yield. This specification helps determine if our forward guidance measure is capturing revisions to the expected path of interest rates versus changes in risk premia. If our results are solely driven by a portfolio-rebalancing channel, then survey-based expectations of future short-term interest rates would likely be unchanged or even rise following the policy announcement as the portfolio-rebalancing channel would raise output and inflation independent of the FOMC's forward guidance. However, Figure 5 shows that expectations of short-term interest rates *fall* after a forward guidance announce-

[9] We use an uninformative prior for all three of these robustness checks.

[10] Specifically, we drop March 18, 2009, November 3, 2010, September 21, 2011, and September 13, 2012, which correspond to the expansion of QE1 and beginning of QE2, QE3, and the MEP.

ment. Moreover, the decline is of similar magnitude to the movement in the path factor, which suggests that the path factor is primarily capturing revisions to the expected path of future interest rates and not portfolio-rebalancing effects from large-scale asset purchases.

Finally, we examine the macroeconomic effects of forward guidance announcements prior to the zero lower bound, a time when the FOMC made numerous announcements about the future path of policy rates but did not engage in large-scale asset purchases. Therefore, we can use this earlier sample period to trace out the macroeconomic effects of a forward guidance shock without worrying about the simultaneous use of quantitative easing. Figure 5 plots the impulse responses to a one standard deviation path factor shock over the 1994–2008 and 2009–2015 sample periods (estimated separately). We find that forward guidance shocks produce similar macroeconomic effects in both samples which suggests that quantitative easing is not driving our empirical estimates during the zero lower bound period.

Taken together, these results provide further evidence that forward guidance shocks that lower expected future policy rates lead to hump-shaped increases in real variables, a gradual rise in prices, and that these effects persist beyond the reduction in expected rates. Moreover, the presence of quantitative easing does not appear to be biasing these findings. We now take these VAR results as stylized evidence on the effects of forward guidance in the data and ask whether a standard model of nominal price rigidity can reproduce these dynamics. Overall, our findings suggest no disconnect between the empirical effects of forward guidance shocks around FOMC announcements and the predictions from a standard theoretical model.

3 A Theoretical Model of Nominal Price Rigidity

This section outlines the dynamic stochastic general equilibrium model we use to analyze forward guidance shocks. The model shares features with the models of Ireland (2011) and Christiano, Eichenbaum and Evans (2005). Our model features optimizing households and firms and a central bank that systematically adjusts the nominal interest rate to offset shocks to the economy but is constrained by the zero lower bound. We allow for sticky prices using the staggered price-adjustment specification of Calvo (1983). The model considers shocks to household preferences and the central bank's desired policy rate. To appropriately calibrate our forward guidance shock process, we generate a model-implied counterpart to the futures contracts from our empirical results. Following Rotemberg and Woodford (1997) and Christiano, Eichenbaum and Evans (2005), we assume that consumption, investment, and pricing decisions are made prior to the realization of both shocks. This timing assumption ensures

that the impact responses of macroeconomic aggregates in the model following a forward guidance shock are consistent with the recursive identification scheme from our baseline VAR model. Appendix C provides details regarding the model's equilibrium conditions.

3.1 Households

The representative household maximizes lifetime expected utility over streams of consumption C_t and leisure $1 - N_t$. The household derives utility from consumption relative to a habit level H_t. The household receives income from the intermediate goods-producing firm in the form of wages W_t for each unit of labor N_t supplied and through lump-sum dividends D_t. The household has access to zero net-supply one-period nominal B_t and real B_t^R bonds. Nominal bonds pay one dollar and are purchased at a discounted price $1/R_t$, where R_t is the one-period gross nominal interest rate. Real bonds return one unit of consumption and have a purchase price $1/R_t^R$, where R_t^R denotes the one-period gross real interest rate.

The representative household maximizes lifetime utility by choosing $C_{t+s}, N_{t+s}, B_{t+s+1}$, and B_{t+s+1}^R, for all $s = 0, 1, 2, \ldots$ by solving the following problem:

$$\max \mathrm{E}_{t-1} \sum_{s=0}^{\infty} a_{t+s}\beta^s \left(\log\left(C_{t+s} - bH_{t+s}\right) - \xi \frac{N_{t+s}^{1+\eta}}{1+\eta} \right)$$

subject to the intertemporal household budget constraint each period,

$$C_t + \frac{1}{R_t}\frac{B_{t+1}}{P_t} + \frac{1}{R_t^R} B_{t+1}^R \leq \frac{W_t}{P_t} N_t + \frac{B_t}{P_t} + \frac{D_t}{P_t} + B_t^R.$$

λ_t denotes the Lagrange multiplier on the household budget constraint. Consumption habits are external to the household and linked to last period's aggregate consumption $H_t = C_{t-1}$.

The discount factor of the household β is subject to shocks via the stochastic process a_t. We interpret these fluctuations as demand shocks since an increase in a_t induces households to consume more and work less today for no technological reason. We use these shocks to simulate a large decline in household demand which generates a zero lower bound episode. The stochastic process for these fluctuations is as follows:

$$a_t = (1 - \rho_a)\, a + \rho_a a_{t-1} + \sigma^a \varepsilon_t^a, \tag{1}$$

where ε_t^a is an independent and standard normal random variable.

3.2 Final Goods Producers

The representative final goods producer uses $Y_t(i)$ units of each intermediate good produced by the intermediate goods-producing firm $i \in [0, 1]$. The market for final goods is per-

fectly competitive, and thus the final goods-producing firm earns zero profits in equilibrium. Following Ireland (2011), the aggregate price index P_t can be written as follows:

$$P_t = \left[\int_0^1 P_t(i)^{1-\theta_\mu} di\right]^{\frac{1}{1-\theta_\mu}}$$

3.3 Intermediate Goods Producers

Each intermediate goods-producing firm i rents labor N_{it} from the household to produce intermediate good Y_{it}, which is sold in a monopolistically competitive market. Each period, producers can reoptimize their nominal price P_{it} with a constant probability $1-\omega$. Firms that cannot reset their price index it to a weighted combination of past and steady-state inflation. Intermediate-goods firms own their capital stock K_{it} and face a convex cost governed by κ when changing their level of investment I_{it}. Firms also choose the rate of utilization of their installed physical capital U_{it} which affects its depreciation rate. The intermediate goods firms all have access to the same constant returns-to-scale production function. A production subsidy $\Psi = \theta/(\theta-1)$ ensures that the steady state of the model is efficient.

We determine the optimal decisions of the intermediate goods-producing firm in two steps. First, firms determine the minimal cost method to meet the current level of demand for their product. Thus, each firm solves the following cost minimization problem:

$$\min \mathrm{E}_{t-1} \sum_{s=0}^{\infty} \left(\beta^s \frac{\lambda_{t+s}}{\lambda_t}\right) \left(\frac{W_{t+s}}{P_{t+s}} N_{i,t+s} + I_{i,t+s}\right)$$

subject to the production function,

$$Y_{it} \leq \big(K_{it} U_{it}\big)^{\alpha} \big(N_{it}\big)^{1-\alpha}$$

and its capital accumulation equation,

$$K_{it+1} = \left(1 - \delta\Big(U_{it}\Big)\right) K_{it} + \left(1 - \frac{\kappa}{2}\Big(\frac{I_{it}}{I_{it-1}} - 1\Big)^2\right) I_{it}.$$

We assume depreciation depends on utilization via the following functional form:

$$\delta\Big(U_{it}\Big) = \delta + \delta_1\Big(U_{it} - U\Big) + \left(\frac{\delta_2}{2}\right)\Big(U_{it} - U\Big)^2.$$

Ξ_t denotes the marginal cost of producing one additional unit of intermediate good i and q_t is the price of a marginal unit of installed capital. After solving its cost minimization problem,

firms that can reoptimize choose their optimal price to maximize their lifetime discounted real profits. Their profit maximization problem is as follows:

$$\max \mathrm{E}_{t-1} \sum_{s=0}^{\infty} \omega^s \beta^s \frac{\lambda_{t+s}}{\lambda_t} \left(\Psi \, \Pi^{s(1-\chi)} \Pi^{\chi}_{t-1,t-1+s} \, \frac{P_{it}}{P_{t+s}} \, Y_{it+s} - \Xi_{t+s} \, Y_{it+s} \right)$$

subject to the following demand curve,

$$Y_{it+s} = \left[\Pi^{s(1-\chi)} \Pi^{\chi}_{t-1,t-1+s} \, \frac{P_{i,t}}{P_{t+s}} \right]^{-\theta} Y_{t+s}.$$

The inflation rate between periods t and $t+s$ is defined as follows:

$$\Pi_{t,t+s} = \begin{cases} 1 & s = 0 \\ \frac{P_{t+1}}{P_t} \times \frac{P_{t+2}}{P_{t+1}} \times \cdots \times \frac{P_{t+s}}{P_{t+s-1}} & s = 1, 2, \ldots \end{cases}$$

The parameter χ controls the degree of indexation to lagged inflation.

3.4 Equilibrium

In the symmetric equilibrium, all intermediate goods firms face the same marginal costs and hence choose to employ the same amount of labor, capital, and utilization rate. All firms that can change their nominal price choose the same optimal price P_t^*. We denote the gross one-period inflation rate as $\Pi_t = P_t / P_{t-1}$. Under the assumption of Calvo (1983) pricing frictions, the aggregate price index P_t evolves as follows:

$$P_t^{1-\theta} = \theta \Big(\Pi^{1-\chi} \, \Pi^{\chi}_{t-1} \Big)^{1-\theta} \big(P_{t-1} \big)^{1-\theta} + \Big(1 - \theta \Big) \big(P_t^* \big)^{1-\theta}$$

3.5 Monetary Policy

We assume the monetary authority sets the one-period net nominal interest rate $r_t = \log(R_t)$. Due to the zero lower bound on nominal interest rates, the central bank cannot lower its nominal policy rate below zero. In the spirit of Reifschneider and Williams (2000), we assume the monetary authority sets its policy rate according to the following history-dependent rule subject to the zero lower bound:

$$r_t^d = \phi_r r_{t-1}^d + \Big(1 - \phi_r \Big) \Big(r + \phi_\pi \big(\pi_t - \pi \big) + \phi_y \, y_t \Big) + \nu_t \tag{2}$$

$$\nu_t = \rho_\nu \nu_{t-1} + \sigma^\nu \varepsilon_t^\nu \tag{3}$$

$$r_t = \max \Big(0, r_t^d \Big) \tag{4}$$

where r_t^d is the desired policy rate of the monetary authority and r_t is the actual policy rate subject to the zero lower bound. π_t denotes the log of the one-period gross inflation rate Π_t and y_t is the gap between the log of current and log of steady-state output. Finally, ν_t is an autocorrelated monetary policy shock. Away from the zero lower bound, this policy rule acts like a Taylor-type policy rule with interest rate smoothing. Also, an exogenous ε_t^ν shock away from the zero lower bound acts like a conventional monetary policy shock.

When the economy encounters the zero lower bound, however, this history-dependent rule lowers the future path of policy to help offset the previous higher-than-desired nominal rates caused by the lower bound constraint. Households fully internalize this future conduct of policy. When desired rates are less than zero, an exogenous shock to the desired rate ε_t^ν acts like an exogenous extension of the zero lower bound episode. This exogenous extension of the zero lower bound lowers future expected policy rates but leaves current policy rates unchanged, which we link with our identified forward guidance shock in the data. We believe this modeling framework closely aligns with our empirical measure of forward guidance as the path factor is constructed to be orthogonal to changes in the current policy rate.[11]

3.6 Generating Model-Implied Futures Contracts

We want to ensure our simulated forward guidance shock in the model is consistent with the typical forward guidance shock we identify in the data. Since we measure forward guidance shocks empirically using fluctuations in futures rates, we generate a model counterpart to the eurodollar futures contracts we examine in the data.

We denote the price of a n-month ahead eurodollar futures contract at time t by f_t^n. The payoff on this futures contract equals one minus the current annualized 3-month LIBOR rate in the contract expiration month. For the 1-month ahead contract in our model, this payoff concept equals $1 - 12 * \frac{1}{3} * (r_{t+1} + r_{t+2} + r_{t+3})$, where r_{t+n} is the policy rate of the central bank in month $t+n$. Therefore, we calculate the price of the one-month ahead zero net-supply futures contract by including the following equilibrium condition in our model:

$$f_t^1 = \mathrm{E}_t \left\{ 1 - 12 * \frac{1}{3} * \Big(r_{t+1} + r_{t+2} + r_{t+3} \Big) \right\}. \tag{5}$$

[11] Our forward guidance shock specification differs from the works of Del Negro, Giannoni and Patterson (2015) and others, which use anticipated "news" shocks about future monetary policy to model forward guidance shocks. In the appendix, we show that we can achieve identical macroeconomic effects from either our specification or a news-shock approach. We prefer our specification because it is parsimonious and it allows us to estimate our model.

For contracts of maturity longer that one month, we can determine the equilibrium futures prices as follows:

$$f_t^n = \mathrm{E}_t \left\{ 1 - 12 * \frac{1}{3} * \left(r_{t+n} + r_{t+n+1} + r_{t+n+2} \right) \right\}. \tag{6}$$

Note that the structure of the futures contracts implies that an n-month contract at time t becomes an $n-1$ contract at time $t+1$. Therefore, we can also conveniently write the futures prices for maturities longer than one month recursively:

$$f_t^n = \mathrm{E}_t \left\{ f_{t+1}^{n-1} \right\}. \tag{7}$$

For a given horizon, we can determine the futures-implied interest rate by computing one minus the contract price.[12] These model counterparts allow us to determine the appropriate-sized forward guidance shock to simulate in the model.

Since we also examine the effects of forward guidance shock on 2-year Treasury yields in the data, we evaluate the model's predictions for bond yields. Therefore, following Rudebusch and Swanson (2012), we compute the price of a n-month default-free zero-coupon bond that pays one dollar at maturity using the following equilibrium condition:

$$p_t^n = \mathrm{E}_t \left\{ \left(\beta \frac{\lambda_{t+1}}{\lambda_t} \frac{1}{\Pi_{t+1}} \right) p_{t+1}^{n-1} \right\}, \tag{8}$$

where the term in parenthesis is the household's nominal stochastic discount factor and $p_t^0 = 1$. The continuously-compounded yield to maturity on the n-period bond is as follows:

$$y_t^n = -\frac{1}{n} \log p_t^n. \tag{9}$$

To be consistent with the timing assumptions in our structural VAR, we assume that futures rates and bond yields can change in the same period as the forward guidance shock but output and prices are fixed at impact.

3.7 Solution Method

We solve our model using the OccBin toolkit developed by Guerrieri and Iacoviello (2015). This piecewise linear approximation solution method allows us to model the occasionally-

[12] Payoffs on futures contracts are not discounted using the household's stochastic discount factor since, in reality, investors in futures contracts post collateral (which also earns a return). Thus, there is no opportunity cost of funds associated with futures positions and it is not necessary to discount the payoffs until maturity.

binding zero lower bound and solve for the model-implied futures prices in only a few seconds, which permits us to estimate several key model parameters.[13]

3.8 Estimation Strategy

Our primary interest is evaluating the model's ability to reproduce the empirical impulse responses of a forward guidance shock from Section 2. Therefore, we estimate our model using impulse response matching. To compute the impulse responses in our model, we generate two time paths for the economy. In the first time path, we simulate a large negative demand shock which causes the zero lower bound to bind for an extended period. In the second time path, we simulate the same large negative demand shock but also simulate a negative shock to the desired policy rate in Equation 2. We assume that the economy is hit by no further shocks and compute the percent difference between the two time paths as the impulse response to an expansionary forward guidance shock at the zero lower bound. Since the economy is at the zero lower bound, this reduction to the desired rate acts like an exogenous extension of the zero lower bound period.

Our estimation strategy chooses model parameters such that the model's impulse responses come as close as possible to the empirical VAR responses. To implement this strategy, we follow Rotemberg and Woodford (1997) and Christiano, Eichenbaum and Evans (2005) who choose the size of a conventional monetary policy shock such that the movements in their model-implied policy indicator are consistent with the impulse responses from an identified vector autoregression. Since the focus of this paper is on forward guidance shocks during the zero lower bound period, however, we discipline the model using expectations of future policy rates. In particular, our estimation procedure picks the size and persistence of the forward guidance shock process which enables the model to generate the same movement in 8-quarter ahead futures rates that we observe in the VAR. We find that linking the model and data counterparts is crucial in order to evaluate the model's fit. In particular, if we were to leave the movements in model-implied expected future interest rates unconstrained, then it is unclear what size forward guidance shock to simulate in the model.

Following much of the previous literature, we partition the model parameters into two groups. The first group is composed of β, Π, η, ξ, θ, ϕ_π, ϕ_y, ρ_a, σ^a. We calibrate these pa-

[13] Using a fully nonlinear, but simplified, version of our model with a policy function iteration method. We find that the Guerrieri and Iacoviello (2015) toolkit provides a good approximation dynamics of the full nonlinear economy after a forward guidance shock.

rameters using steady-state relationships or results from previous studies. Since the model shares features with the model of Ireland (2011), we calibrate some of our parameters to match his values or estimates. To match our VAR evidence, we calibrate the model to a monthly frequency. We set ξ to normalize output Y to equal one at the deterministic steady state. We choose standard values for the monetary policy reaction to inflation and output ($\phi_\pi = 1.5$, $\phi_y = 0.1$). Our monthly calibrations of β and Π imply a steady state annualized real interest rate of two percent and a two percent annualized inflation target.

We estimate the second set of model parameters which consists of the household habit parameter b, the probability that a firm can not reoptimize its price ω, the degree of lagged inflation indexation χ, the degree of smoothing in the monetary policy rule ϕ_r, the degree of investment adjustment costs κ, the elasticity of the return on capital with respect to capacity utilization $\sigma_\delta = \delta_2/\delta_1$, and the forward guidance shock parameters ρ_ν and σ^ν. In addition, we also estimate the size of the initial negative demand shock ε_0^a which takes the economy to the zero lower bound prior to the forward guidance shock. We collect these parameters into a vector $\gamma = (b, \omega, \chi, \phi_r, \kappa, \sigma_\delta, \rho_\nu, \sigma_\nu, \varepsilon_0^a)$.

Using a Bayesian impulse response matching estimator, we estimate these key model parameters by finding the values which maximize the posterior distribution. Let $\hat{\psi}$ denote the impulse response functions for the 6 variables in our empirical VAR stacked into a single vector with $(6 \times 48 = 288)$ rows and let the diagonal matrix V^{-1} denote a measure of the precision of the empirical impulse responses.[14] Then, let $\psi(\gamma)$ denote the theoretical model's corresponding counterpart to $\hat{\psi}$. Following Christiano, Trabandt and Walentin (2010), we can write the approximate likelihood function as follows:

$$L(\hat{\psi} \mid \gamma, V) = (2\pi)^{-\frac{N}{2}} \mid V \mid^{-\frac{1}{2}} \exp\left[-0.5(\hat{\psi} - \psi(\gamma))' V^{-1} (\hat{\psi} - \psi(\gamma))\right].$$

Let $p(\gamma)$ denote the joint prior density over γ. According to Bayes rule,

$$f(\gamma \mid \hat{\psi}, V) \propto L(\hat{\psi} \mid \gamma, V) p(\gamma), \tag{10}$$

where $f(\gamma \mid \hat{\psi}, V)$ is the posterior density over γ. Our estimator solves the following problem:

$$\max_{\gamma} f(\gamma \mid \hat{\psi}, V). \tag{11}$$

[14] The diagonal of V^{-1} contains one over the squared difference between the 95th and 5th percentile of the empirical probability interval. Omitting off-diagonal terms from V helps make our estimator more transparent as it attempts to place the model's impulse responses inside the empirical probability intervals.

3.9 Priors Over Parameters

For our priors, we use a Beta distribution for parameters that lie between 0 and 1 and a Gamma distribution for parameters which are positive but unbounded. For the household habit parameter b, degree of indexation χ, and the persistence of the forward guidance shock ρ_ν, we center the prior mode at 0.5 with a standard deviation of 0.25. For the Calvo parameter ω, we tightly center our prior mode at 0.93 which is consistent Nakamura and Steinsson (2008)'s evidence that prices remain fixed for about one year on average. We center our prior mode over ϕ_r at 0.95 which is consistent with a large literature arguing that historical Federal Reserve policy features a high degree of inertia. However, we set a loose prior on this parameter since, as we discussed in Section 3.5, its interpretation changes when the economy encounters the zero lower bound.

For the investment adjustment cost parameter κ and elasticity of capital utilization σ_δ, we center our prior at the quarterly estimates of Christiano, Eichenbaum and Evans (2005). However, since our model is calibrated to a monthly frequency, we set loose priors over these parameters to reflect our uncertainty over the exact time-aggregation function. Our prior for the size of the forward guidance shock σ_ν is similarly uninformative. We restrict the initial aggregate demand shock ε_0^a to be negative in order to simulate a decline in aggregate demand that takes the economy to the zero lower bound prior to the forward guidance shock.

4 Estimated Responses to a Forward Guidance Shock

We now analyze the effects of a forward guidance shock in our estimated model and show that the model can reproduce our empirical evidence. Figure 2 plots both the empirical and model-implied impulse responses to a forward guidance shock. At impact, the forward guidance shock causes the model-implied 8-quarter eurodollar rate to decline by about six basis points, which is consistent with our empirical findings. Output, investment, and capacity utilization in the model all rise after the shock in hump-shaped patterns similar to their VAR counterparts. The model also replicates the gradual increase in prices we observe in the data. The peak response of output in the model is quantitatively similar to our empirical results, although the model-implied peak occurs slightly earlier than in the data. As in the data, the 2-year bond yield falls as the forward guidance shock lowers the expected path of short-term interest rates in the economy. Importantly, all of the model's impulse responses fall within the VAR's probability intervals, which suggests the predictions from a standard model of monetary policy are generally in line with the empirical effects of a forward guidance shock.

To provide additional intuition for our results, Figure 6 shows the impulse responses for consumption, additional futures contracts, and real interest rates. Prior to the forward guidance shock, the estimated negative aggregate demand shock implies that the economy is expected to be at the zero lower bound for 20 months. The estimated forward guidance shock then extends the zero lower bound duration by one month. Since households expect the zero lower bound to persist for total of 21 months, 12-month ahead futures rates do not move immediately after the forward guidance shock. However, the 24-month ahead contract falls by several basis points as expected future nominal policy rates decline. The combination of the forward guidance shock, nominal price and investment rigidities, and the zero lower bound produces a hump-shaped response of real interest rates. At impact, current nominal policy rates are fixed at zero and expected inflation rises very slightly due to the nominal rigidity in price setting. Thus, real interest rates only fall modestly while the economy remains at the zero lower bound. However, real rates fall sharply once the economy exits the zero lower bound and the monetary authority can lower its current nominal policy rate. This time path for real interest rates, in addition to habits in consumption, causes a very gradual increase in consumption, which peaks about one year after the forward guidance shock.

4.1 Role of the Initial Demand Shock

While many features of our model are standard, simulating a forward guidance shock at the zero lower bound requires us to estimate the initial conditions in the economy prior to the forward guidance shock. In Figure 6, we illustrate how our estimate of the initial aggregate demand shock affects our main results.

Disciplining the model using futures contracts helps the estimation procedure determine the appropriate zero lower bound episode to simulate in the model. In our baseline results, we find that a total zero lower bound episode of 21 months allows the model to match the data. For comparison, we simulate a larger initial shock to the economy such that the zero lower bound persists for significantly longer.[15] Figure 6 plots the responses under the longer 36-month zero lower bound duration and our baseline 21-month scenario. If we simulate too large of an initial demand shock, the 24-month ahead futures rate fails to move at impact and displays a somewhat hump-shaped pattern. This time path is clearly inconsistent with the empirical evidence from Figure 2 where futures rates fall at impact and rise monotonically.

[15] Our estimated zero lower bound duration of a little less than two years is consistent with the *ex ante* views of professional forecasters as detailed in Figure 4 of Swanson and Williams (2014).

Thus, appropriately choosing the initial demand shock ensures that the model can generate movements in futures rates similar to what we observe in the data.

4.2 Parameter Estimates

The model requires a mix of nominal as well as real rigidities to match the VAR impulse responses. The top panel of Table 2 shows that our estimated degree of nominal rigidity ω implies that prices remain fixed for about 7 quarters, on average. While prices in our model are more persistent than the micro-level estimates of Nakamura and Steinsson (2008), our results are consistent with the findings of Gali and Gertler (1999) and Del Negro, Giannoni and Schorfheide (2015). We find essentially no role for lagged indexation of prices with $\chi = 0.04$, which likely reflects a decline in the persistence of inflation over time.

In addition to a moderate degree of nominal rigidity, three real rigidities help the model reproduce the empirical evidence. Our estimate of consumption habits b is higher than the estimate of Christiano, Eichenbaum and Evans (2005), as one might expect when moving from a quarterly to a monthly frequency. As in Christiano, Eichenbaum and Evans (2005), our estimate of the capacity utilization adjustment cost parameter is very small and not significantly different from zero. Since $1/\sigma_\delta$ governs the elasticity of capacity utilization with respect to the return on capital, our estimate of σ_δ implies a large response of utilization to a given movement in capital returns, which is consistent with our VAR evidence. Turning to investment, we find a much larger monthly investment adjustment cost parameter κ than the quarterly estimates of Christiano, Eichenbaum and Evans (2005), which suggests that, at a monthly frequency, firms make more incremental adjustments in their capital stock.

We estimate a significant degree of desired-rate smoothing in the central bank's policy rule. However, our estimate of $\phi_r = 0.94$ doesn't significantly differ from its prior mode which suggests that ϕ_r may not be well-identified by our impulse response matching procedure. In Appendix G, we explore alternative priors for ϕ_r and consistently find point estimates of ϕ_r which are very near to the prior mode but imply no significant change in the model's fit of the data. This result is not too surprising since we are only informing our estimation procedure with information on monetary policy shocks. Coibion and Gorodnichenko (2012) show that the degree of endogenous interest-rate smoothing is likely better informed by the policy response to non-monetary shocks. However, these additional results show that the overall fit of our model does not rely on a particular assumption about the amount of history

dependence in the central bank's policy rule.[16]

4.3 Quarterly Model Estimates

Parameter estimates from our monthly-frequency model are difficult to compare with the previous literature estimating the effects of conventional monetary policy shocks at a quarterly-frequency. Thus, to facilitate a quantitative comparison of our estimated parameters with those in Christiano, Eichenbaum and Evans (2005), we take the parameter estimates from their Table 2, row 5 (unconditional indexation) and the associated standard errors they estimate to form *priors* for the parameters in a quarterly version of our model.[17] Then, we estimate the posterior distribution of the parameters by minimizing the distance between the impulse responses in the data and the model. Figure 7 shows the resulting impulse responses, which are well within the posterior intervals of the VAR impulse responses for all variables and all horizons. In addition, the model's responses are very close to the point estimates from the VAR for most variables.

What parameters deliver this close fit? The bottom panel of Table 2 shows the the posterior modes and standard deviations of the parameters. All of the estimated parameters are near the Christiano, Eichenbaum and Evans (2005) estimates except for the Calvo pricing friction parameter for which we estimate $\omega = 0.88$ compared to the Christiano, Eichenbaum and Evans (2005) estimate of $\omega = 0.72$. However, our estimate of ω is within a one standard error range of the Christiano, Eichenbaum and Evans (2005) estimate.

Other than a change in the frequency of price adjustment, the same model and parameters that can account well for the dynamics of conventional monetary policy shocks prior to the zero lower bound can also account for the dynamic effects of forward guidance shocks at the zero lower bound. The increase in the Calvo parameter that we find necessary to explain these dynamics relative to the estimated value in Christiano, Eichenbaum and Evans (2005) could represent either the absence of wage rigidity or the absence of a working-capital friction in our model. Christiano, Eichenbaum and Evans (2005) show that removing either of these frictions increases the average duration of prices in their model. Specifically, they find $\omega = 1$ when they allow for flexible wages and $\omega = 0.89$ when they assume that firms do not need to borrow their wage bill in advance. Moreover, our estimate of $\omega = 0.88$ is well within the range reported in the literature, including recent work of Del Negro, Giannoni and Patterson

[16] In the appendix, we also discuss the role of the persistence of the forward guidance shock process.

[17] We choose to work with this set of estimates since we find little need for indexation in our sample.

(2015) and Del Negro, Giannoni and Schorfheide (2015).

4.4 Model-Implied Responses to Specific Forward Guidance Shocks

Our results suggest that a standard model with nominal price rigidities can account for the effects of an average forward guidance shock as estimated from a VAR. However, our linear VAR only allows us to trace out the effects from a typical forward guidance shock in the data which moves futures rates by about 6 basis points. Figure 1 shows that some policy announcements generated significantly larger movements in the path factor. For example, on August 9, 2011, 8-quarter ahead eurodollar futures rates declined by 28 basis points when the FOMC announced that rates were likely to remain low until "mid-2013." This size of movement in futures rates equates to a 4.8 standard deviation forward guidance shock in our model. In order to further evaluate the predictions of our theoretical model, we now examine whether the model generates reasonable quantitative responses following this significantly larger forward guidance shock.

In response to the August 2011 announcement, we find that the model predicts a much larger expansion of economic activity than a typical forward guidance shock. Figure 8 illustrates the model-implied responses to the forward guidance shock on August 9, 2011. To generate these responses, we increase the size of the forward guidance shock such that the 2-year ahead futures rates in the model decline by 28 basis points, the same movement in futures rates we observe in the data. As expected, this substantially larger forward guidance shock generates a much larger increase in economic activity and prices. At its peak response, output rises by about 0.5% which occurs a little over one year after the shock.

Are the model's predicted effects from the August 2011 announcement reasonable? Unlike the responses to a typical forward guidance shock we identify from our VAR, we have no clear data counterpart to compare with the model's responses for this specific shock. Alternatively, to gain some sense of reasonable magnitudes, we can compare the model-implied movement in output with the estimated effects of other unexpected monetary policy interventions. In the VAR model of Christiano, Eichenbaum and Evans (2005), a typical expansionary monetary policy shock raises output by just over 0.5% after about one year. In addition, the Romer and Romer (2004) estimates imply that a one standard deviation expansionary policy shock increases output by around 0.6% at its peak.[18] Therefore, the model's response to one of

[18]We estimate this output response by simulating a one standard deviation shock in Romer and Romer's (2004) three-variable VAR using their shock series and their monthly measures of output and prices.

the largest forward guidance shocks we observe around FOMC meetings seems reasonable in the context of other work measuring the potency of monetary policy.

4.5 Additional Model Results

In the appendix, we provide further robustness checks on the model estimation. Our baseline model assumes that monetary policy responds to deviations of output from its steady-state level. However, Coibion and Gorodnichenko (2011) provide evidence that policymakers respond more to output growth. We find that the model can fit the VAR evidence similarly well if we replace output in the policy rule with output growth. Also, a simplified model with a fixed capital stock, a common benchmark in the literature, can reproduce the empirical impulse responses to a forward guidance shock from a smaller VAR. In addition, we illustrate how our estimated degree of investment adjustment costs facilitate the model's ability to fit the empirical responses of investment and futures rates. We also use simulated data from our structural model to show that our empirical strategy generally works well in recovering the true forward guidance shocks and the associated impulse responses.

5 Discussion

Our empirical results and our conclusions regarding the ability of standard theoretical models to match these results are both at odds with a growing literature on the effects of forward guidance. One strand of this literature emerges from the work of Del Negro, Giannoni and Patterson (2015), which argues that the output response to a forward guidance shock is implausibly large resulting in a "forward guidance puzzle." A second strand of this literature suggests that forward guidance announcements may contain two pieces of news: news about the future state of the economy as well as news about future interest rates. Work by Campbell et al. (2012) and Nakamura and Steinsson (2018) argues that the macroeconomic news effect dominates resulting in an "event-study activity puzzle." In this section, we relate our findings to this literature and provide some explanations as to why we reach different conclusions.

5.1 The Forward Guidance Puzzle

Our findings contrast with the work by Del Negro, Giannoni and Patterson (2015) which argues that standard models with nominal rigidities overestimate the expansionary effects of forward guidance. Our alternative conclusion emerges from the size of the forward guidance

shock we estimate. In both our empirical evidence and theoretical model, a typical expansionary forward guidance shock in a one-day window around a policy announcement lowers 8-quarter ahead futures rates by about six basis points. This shock extends the zero lower bound duration by one month in our model, which produces modest increases in output and inflation that are consistent with our empirical evidence.

Del Negro, Giannoni and Patterson (2015) simulate a much larger forward guidance shock. Motivated by the FOMC's extension of its date-based guidance from "late 2014" to "mid 2015" in September 2012, they simulate an exogenous one-year extension of the zero lower bound period which results in a very large expansion in economic activity. These authors argue this increase in activity is implausibly large and denote their finding the "Forward Guidance Puzzle." However, in our estimated model, a one-year exogenous extension requires a highly unlikely 25+ standard deviation shock. In Figure 9, we reproduce the key forward guidance experiment in Del Negro, Giannoni and Patterson (2015) by simulating a forward guidance shock large enough to exogenously extend the zero lower bound duration by one year. Similar to their findings, we observe an extremely large increase in economic activity compared to our baseline results.

Are the model-implied responses to a one-year exogenous extension plausible or do they illustrate a fundamental flaw in the model? Unfortunately, we do not believe our results can fully answer this question. In our high-frequency identification of policy shocks around FOMC announcements, we do not observe forward guidance shocks of that size in the data. For example, while the change in the *language* of the FOMC statements between the August and September 2012 meetings suggests roughly a one-year extension, 8-quarter ahead eurodollar futures rates only fell by 6 basis points in the one-day window around the September 2012 FOMC meeting. In our model, this shock implies a one-month extension of the zero lower bound episode and Figure 9 shows that this shock has only modest effects on the economy. Therefore, our results suggest that standard models work well in analyzing the typical forward guidance shocks we observe around FOMC announcements. However, our work cannot speak to the plausibility of the model's predictions for substantially larger shocks that may occur between FOMC meetings or all possible experiments that might be of interest to macroeconomic modelers or policymakers.

5.2 Macroeconomic News in Forward Guidance Announcements

Nakamura and Steinsson (2018), building on the work of Campbell et al. (2012), argue

that, in addition to providing news about future interest rates (an Odyssean component), monetary policy announcements contain a significant amount of macroeconomic news (a Delphic component). If the Delphic component dominants, then an unexpected decline in policy rates could *lower* forecasts of economic activity. Our estimates of the effects of forward guidance should be interpreted as an on-average, net-effect of FOMC communication. The VAR evidence we find seems to suggest that, among these two channels, the Odyssean effect dominates, on average. While there may be Delphic components to some FOMC announcements, we find that those influences are typically more than offset by expansionary Odyssean effects. To illustrate this idea, Figure 10 shows the impulse responses if we include the five-year inflation-protected Treasury yield in our VAR model, a key variable of interest in Nakamura and Steinsson (2018). These authors argue that a lower path of the policy rate signals a lower natural real rate of interest, which causes long-term real interest rates to fall and lowers expectations for output growth. In contrast, we find that forward guidance shocks lower financial market measures of long-term real interest rates and *raise* actual output, which is consistent with the predictions of a standard model of monetary policy without a macroeconomic news channel.

6 Conclusion

This paper reconciles empirical evidence on the effects of forward guidance shocks with the predictions from a standard model of nominal rigidity. Our analysis suggests no disconnect between the empirical effects of forward guidance shocks around policy announcements and the predictions from a standard model of monetary policy. Moreover, we find that the estimated parameters that govern the model's key frictions lead us to a model economy that does not appear all that different from the model of conventional policy shocks of Christiano, Eichenbaum and Evans (2005). These findings suggest that the same models economists use to study the effects of conventional monetary policy shocks remain useful in studying the effects of forward guidance shocks at the zero lower bound.

References

Barakchian, S. Mahdi, and Christopher Crowe. 2013. "Monetary Policy Matters: Evidence from New Shocks Data." *Journal of Monetary Economics*, 60(8): 950 – 966.

Bauer, Michael D., and Glenn D. Rudebusch. 2014. "The Signaling Channel for Federal Reserve Bond Purchases." *International Journal of Central Banking*, 10(3): 233–289.

Bhattarai, Saroj, Gauti B. Eggertsson, and Bulat Gafarov. 2015. "Time Consistency and the Duration of Government Debt: A Signalling Theory of Quantitative Easing." Working Paper.

Calvo, Guillermo A. 1983. "Staggered Prices in a Utility-Maximizing Framework." *Journal of Monetary Economics*, 12(3): 383–398.

Campbell, Jeffrey R., Charles L. Evans, Jonas D.M. Fisher, and Alejandro Justiniano. 2012. "Macroeconomic Effects of Federal Reserve Forward Guidance." *Brookings Papers on Economic Activity*, 44(1): 1–80.

Christiano, Lawrence J., Martin Eichenbaum, and Charles L. Evans. 2005. "Nominal Rigidities and the Dynamic Effects of a Shock to Monetary Policy." *Journal of Political Economy*, 113(1): 1–45.

Christiano, Lawrence J., Mathias Trabandt, and Karl Walentin. 2010. "DSGE Models for Monetary Policy Analysis." *Handbook of Monetary Economics*, 3: 285–367.

Coibion, Olivier, and Yuriy Gorodnichenko. 2011. "Monetary Policy, Trend Inflation, and the Great Moderation: An Alternative Interpretation." *American Economic Review*, 101(1): 341–70.

Coibion, Olivier, and Yuriy Gorodnichenko. 2012. "Why are target interest rate changes so persistent?" *American Economic Journal: Macroeconomics*, 4(4): 126–162.

Del Negro, Marco, Marc Giannoni, and Christina Patterson. 2015. "The Forward Guidance Puzzle." Working Paper.

Del Negro, Marco, Marc P. Giannoni, and Frank Schorfheide. 2015. "Inflation in the Great Recession and New Keynesian Models." *American Economic Journal: Macroeconomics*, 7(1): 168–96.

Eggertsson, Gauti B., and Michael Woodford. 2003. "The Zero Lower Bound on Interest Rates and Optimal Monetary Policy." *Brookings Papers on Economic Activity*, 1: 139–211.

Gali, Jordi, and Mark Gertler. 1999. "Inflation Dynamics: A Structural Econometric Analysis ." *Journal of Monetary Economics*, 44(2): 195 – 222.

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

Guerrieri, Luca, and Matteo Iacoviello. 2015. "OccBin: A Toolkit for Solving Dynamic Models with Occasionally Binding Constraints Easily." *Journal of Monetary Economics*, 70: 22–38.

Gurkaynak, Refet S, Brian Sack, and Eric Swanson. 2005. "Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements." *International Journal of Central Banking*, 1(1): 55–93.

Ireland, Peter N. 2011. "A New Keynesian Perspective on the Great Recession." *Journal of Money, Credit, and Banking*, 43(1): 31–54.

Krishnamurthy, Arvind, and Annette Vissing-Jorgensen. 2011. "The Effects of Quantitative Easing on Interest Rates: Channels and Implications for Policy." *Brookings Papers on Economic Activity*, 43(2): 215–287.

McKay, Alisdair, Emi Nakamura, and Jón Steinsson. 2016. "The Power of Forward Guidance Revisited." *American Economic Review*, 106(10): 3133–3158.

Nakamura, Emi, and Jón Steinsson. 2008. "Five Facts about Prices: A Reevaluation of Menu Cost Models." *Quarterly Journal of Economics*, 123(4): 1415–1464.

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

Reifschneider, David, and John C. Williams. 2000. "Three Lessons for Monetary Policy in a Low-Inflation Era." *Journal of Money, Credit, and Banking*, 32(4): 936–966.

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

Rotemberg, Julio J., and Michael Woodford. 1997. "An Optimization-Based Econometric Framework for the Evaluation of Monetary Policy." *NBER Macroeconomics Annual*, 12: 297–346.

Rudebusch, Glenn D., and Eric T. Swanson. 2012. "The Bond Premium in a DSGE Model with Long-Run Real and Nominal Risks." *American Economic Journal: Macroeconomics*, 4(1): 105–143.

Swanson, Eric T., and John C. Williams. 2014. "Measuring the Effect of the Zero Lower Bound on Medium- and Longer-Term Interest Rates." *American Economic Review*, 104(10): 3154–85.

Woodford, Michael. 2012. "Methods of Policy Accommodation at the Interest-Rate Lower Bound." Federal Reserve Bank of Kansas City.

Table 1: Calibrated Parameters in Baseline Model

Parameter	Description	Calibrated Value
β	Household Discount Factor	0.9983
Π	Steady State Inflation Rate	$1.02^{\frac{1}{12}}$
δ_0	Steady State Depreciation	0.1 / 12
δ_1	First-Order Utilization Parameter	$1/\beta - 1 + \delta_0$
η	Inverse Frisch Labor Supply Elasticity	0.5
ξ	Utility Function Constant	58.43
θ	Elasticity of Substitution Intermediate Goods	6.0
α	Capital Share in Production Function	0.33
ϕ_π	Central Bank Response to Inflation	1.5
ϕ_y	Central Bank Response to Output	0.1
ρ_a	Preference Shock Persistence	0.95
σ^a	Std. Dev. of Preference Shock	0.005

Table 2: Estimated Parameters

Baseline Monthly-Frequency Model

		Prior			Posterior	
Parameter	Description	Distribution	Mode	Std. Dev.	Mode	Std. Dev.
b	Habit Persistence	Beta	0.50	0.25	0.8898	0.0146
ω	Calvo Probability	Beta	0.93	0.01	0.9558	0.0015
χ	Degree of Lagged Indexation	Beta	0.50	0.25	0.0358	0.0299
ϕ_r	Policy Rate Smoothing	Beta	0.95	0.25	0.9442	0.0021
κ	Investment Adjustment	Gamma	2.48	60.0	36.7530	2.3256
σ_δ	Capacity Utilization Curvature	Gamma	0.01	60.0	0.0003	0.0002
ρ_ν	Policy Shock Persistence	Beta	0.50	0.25	0.8358	0.0055
σ_ν	Std. Dev. of Policy Shock (APR)	Gamma	0.25	12	0.0850	0.0035

Quarterly-Frequency Model

		Prior			Posterior	
Parameter	Description	Distribution	Mode	Std. Dev.	Mode	Std. Dev.
b	Habit Persistence	Beta	0.63	0.05	0.6134	0.0403
ω	Calvo Probability	Beta	0.72	0.16	0.8858	0.0022
ϕ_r	Policy Rate Smoothing	Beta	0.80	1×10^{-6}	0.8000	3×10^{-6}
κ	Investment Adjustment	Gamma	1.92	0.35	2.3355	0.1896
σ_δ	Capacity Utilization Curvature	Gamma	0.01	1×10^{-6}	0.01	3×10^{-7}
ρ_ν	Policy Shock Persistence	Beta	0.50	0.25	0.0059	0.0218
σ_ν	Std. Dev. of Policy Shock (APR)	Gamma	0.25	4	0.7180	0.0569

Notes: The top panel shows the estimated parameters from our monthly-frequency baseline model. The bottom panel contains the estimated parameters from a quarterly-frequency version of our model. The calibrated parameters in the quarterly model are identical with the exception of β, Π, and δ_0 which are transformed to their implied quarterly values. The priors for the estimated parameters in the bottom panel are based on the estimates of Christiano, Eichenbaum and Evans (2005). In particular, we use the estimates from their model with no indexation (Rows 5 and 6 from their Table 2). The interest-rate smoothing is calibrated in a version of their model so we set a very tight prior at their calibrated value. The policy shock persistence and standard deviation are not estimated in their paper, so we set diffuse priors for these parameters. See Sections 3.9, 4.1, and 4.2 for more details.

Figure 1: Forward Guidance Shock Series

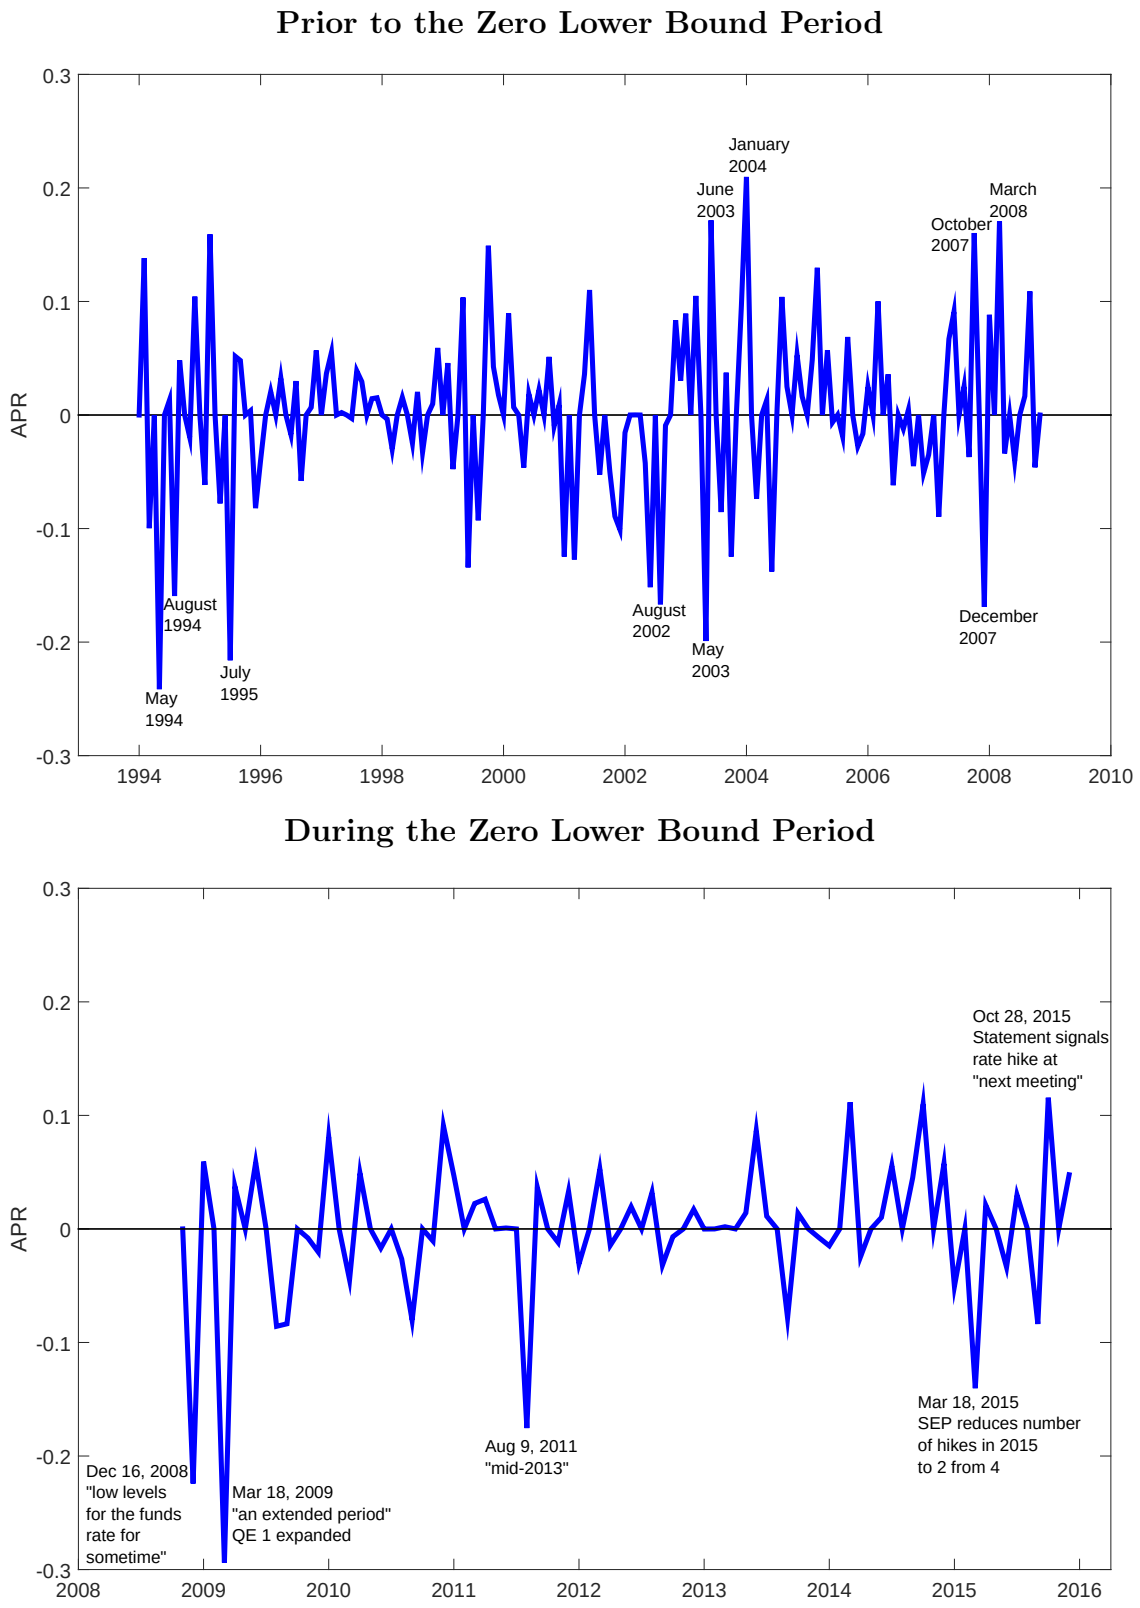

Note: The annotated dates align with some of the largest observations (in absolute value) during each sample period. See Appendix A for additional details on these announcements.

Figure 2: Empirical & Model-Implied Impulse Responses to Forward Guidance Shock

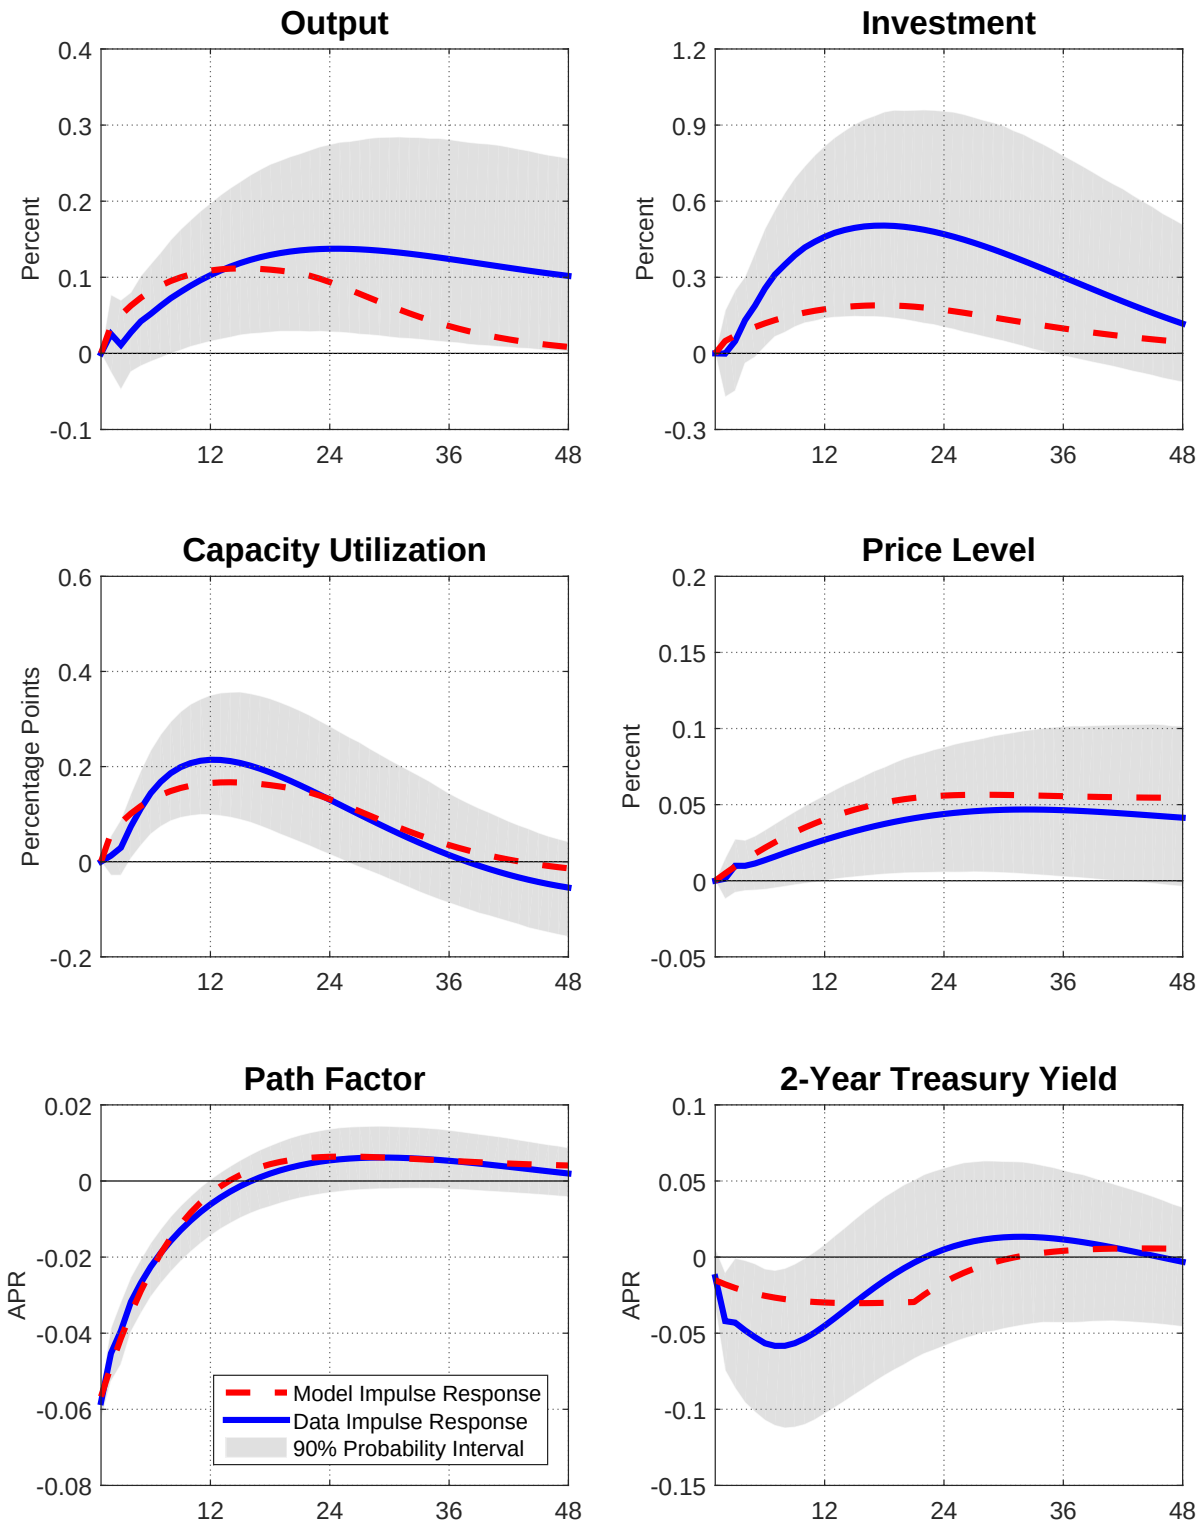

Note: The solid blue lines denote the empirical point estimate to a one standard deviation shock and the shaded areas denote the 90% probability interval of the posterior distribution. The red dashed line denotes the model-implied impulse response.

Figure 3: Empirical Impulse Responses Using Alternative Orderings, Indicators, and Priors

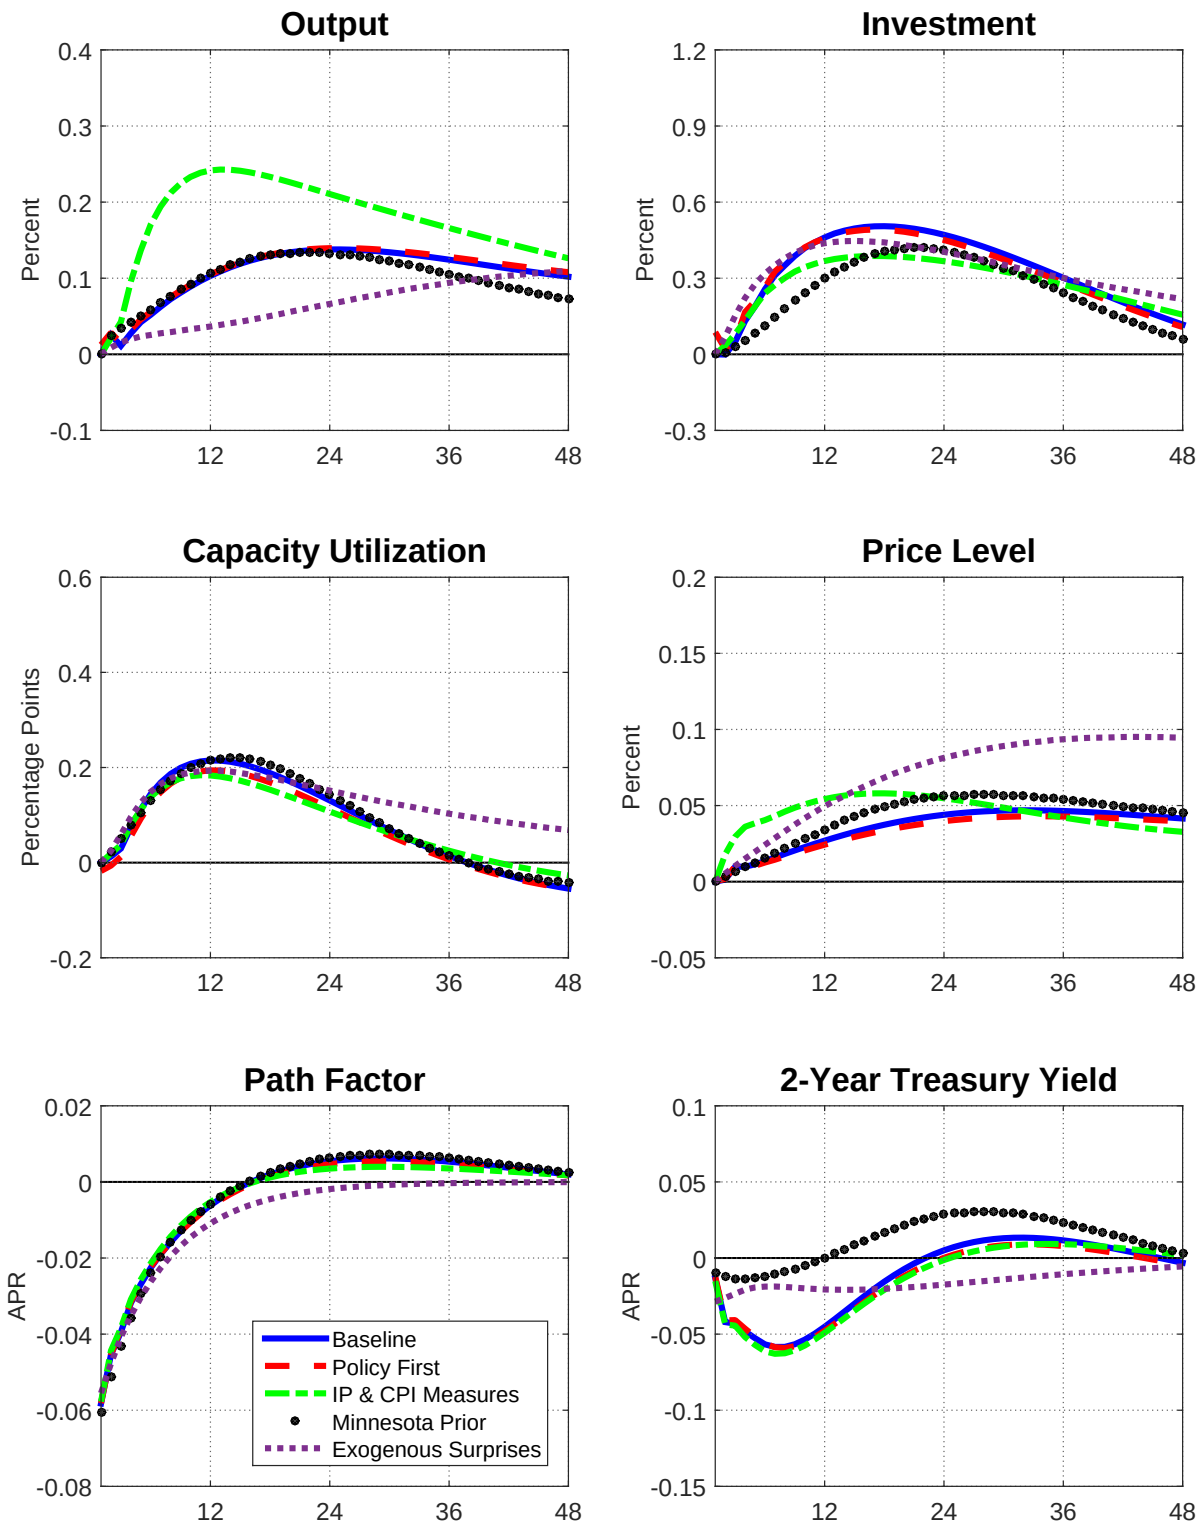

Note: Each line denotes the point estimates to a one standard deviation forward guidance shock from a different empirical specification.

Figure 4: Empirical Impulse Responses Using Different Policy Measures

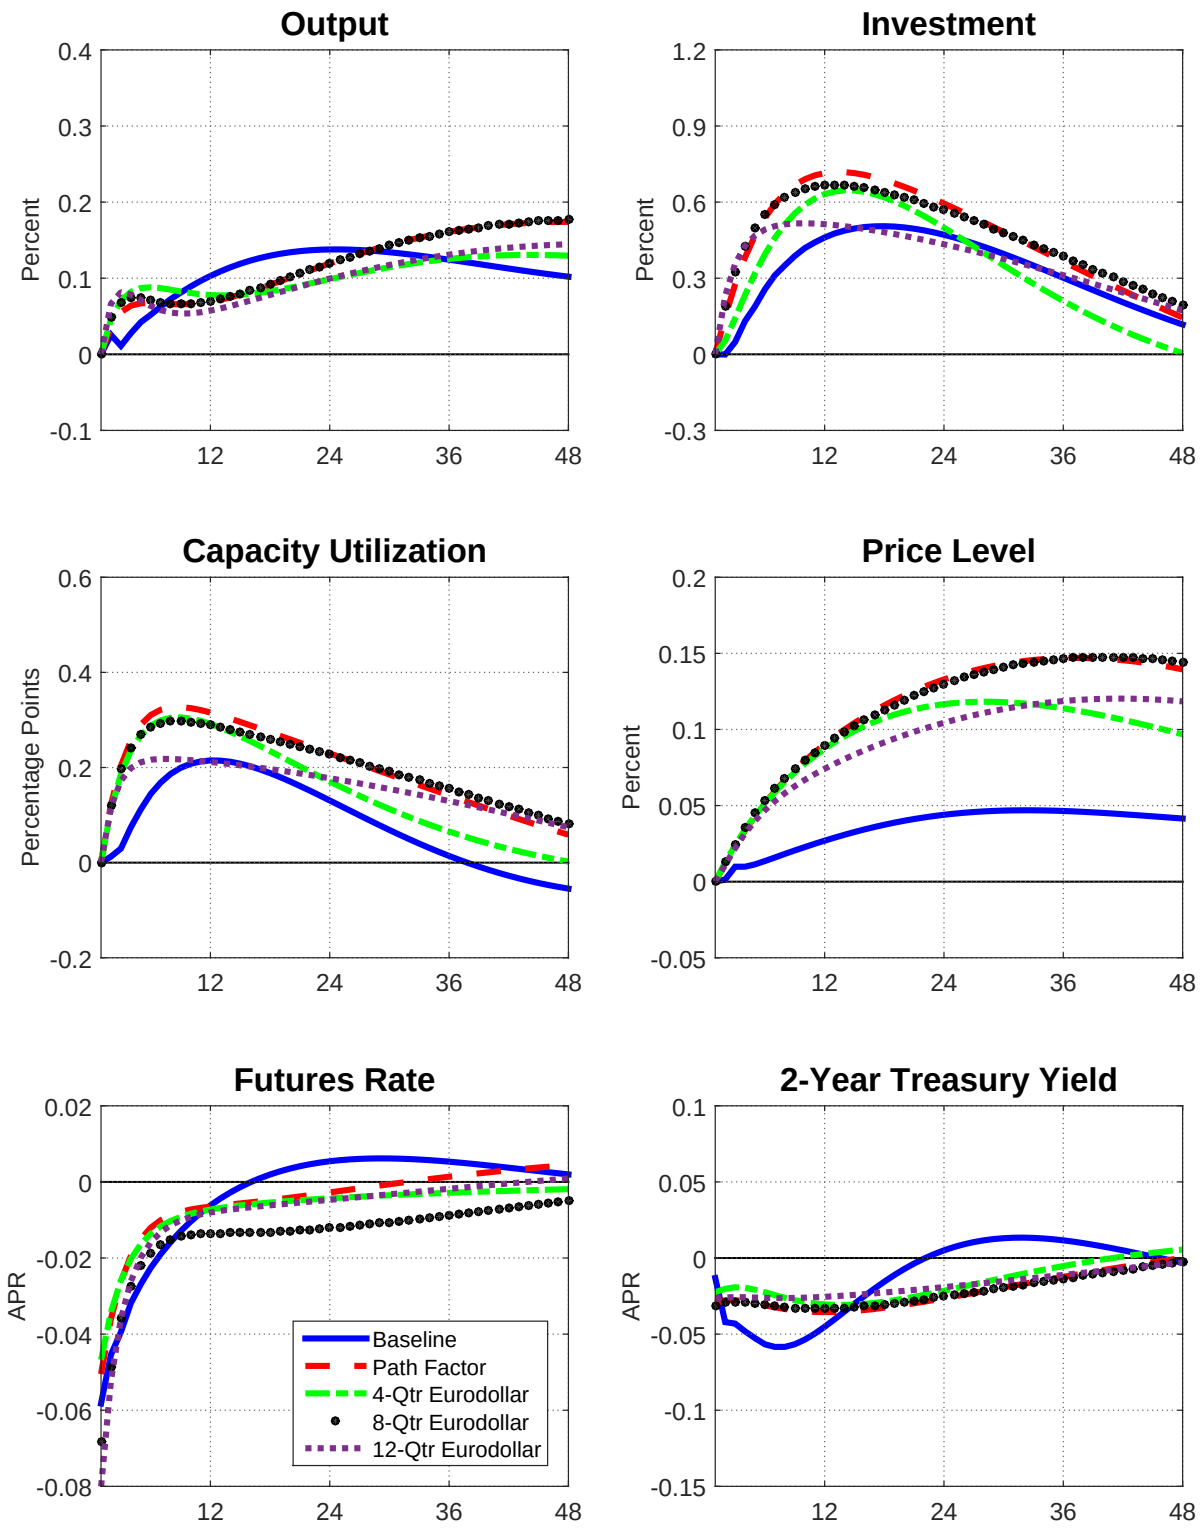

Note: Each line denotes the point estimates to a one standard deviation forward guidance shock from a different empirical specification.

Figure 5: Empirical Impulse Responses: Robustness to QE Announcements

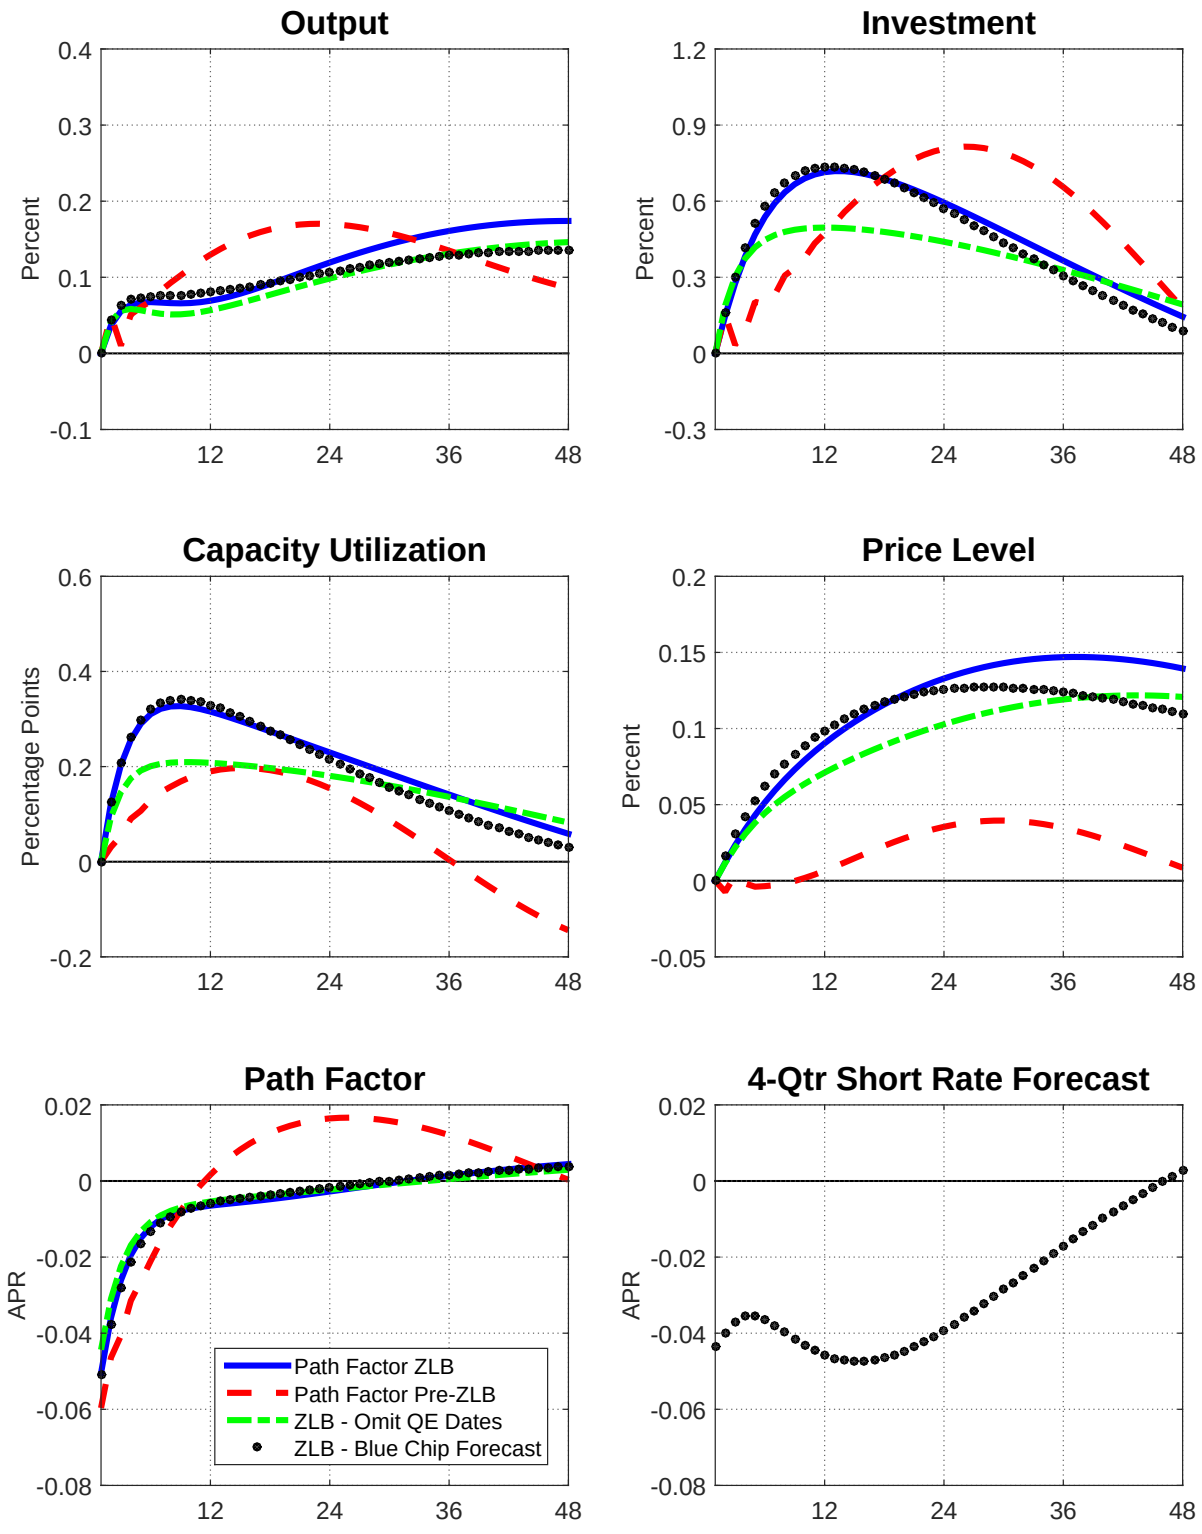

Note: Each line denotes the point estimates to a one standard deviation forward guidance shock from a different empirical specification. The 2-year Treasury yield is included in all models except the "ZLB - Blue Chip Forecast" model.

Figure 6: Role of the Initial Demand Shock

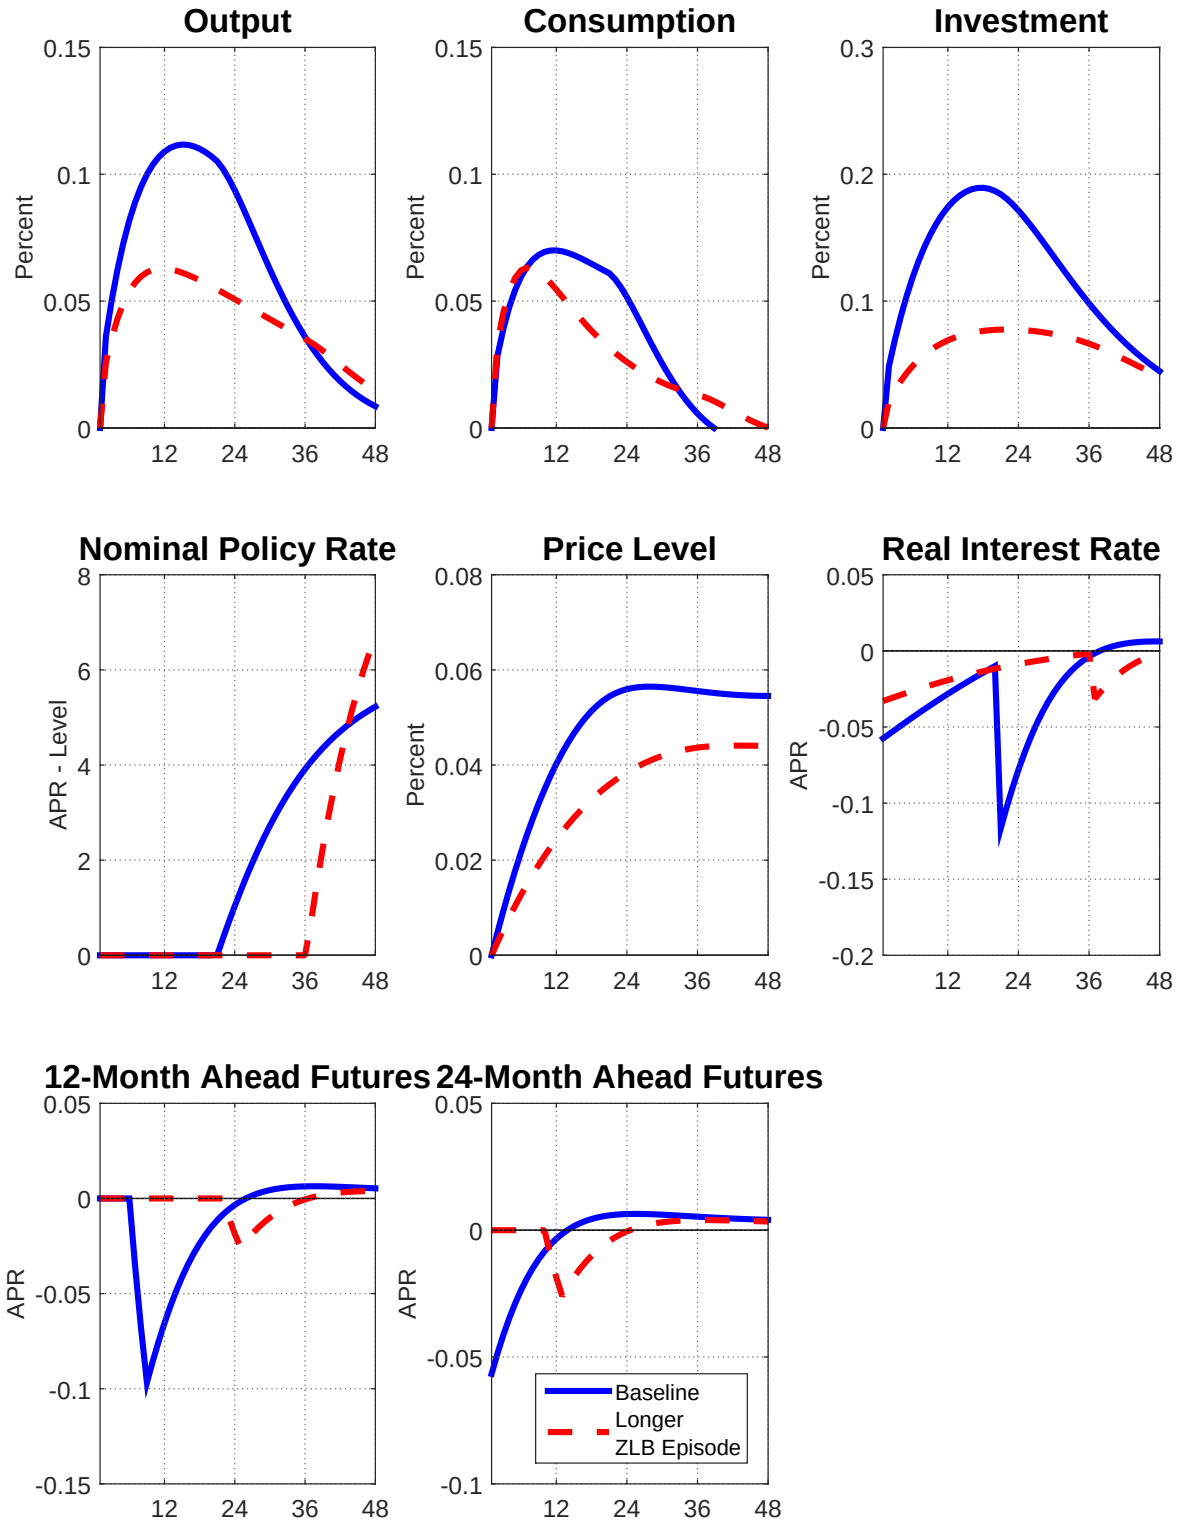

Note: The plot of the nominal interest rate reflects its level after the forward guidance shock.

Figure 7: Implied Impulse Responses to Forward Guidance Shock in a Quarterly Model

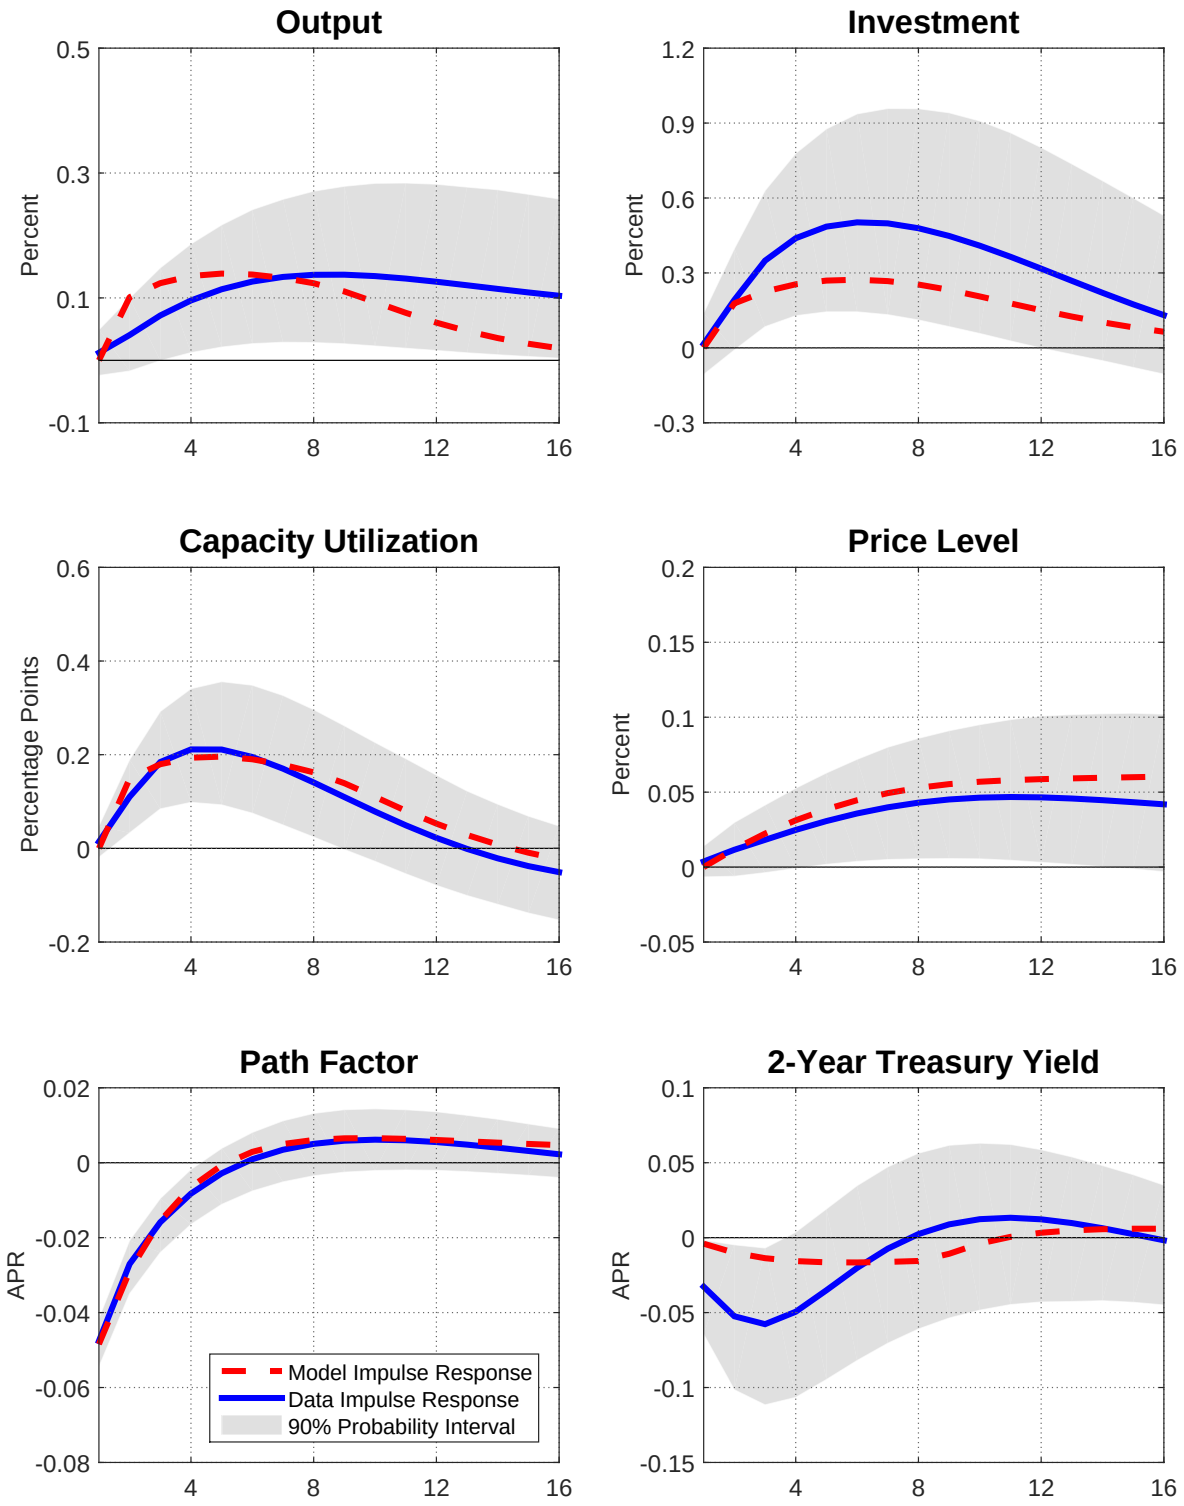

Note: The solid blue lines denote the empirical point estimate to a one standard deviation shock and the shaded areas denote the 90% probability interval of the posterior distribution. The red dashed line denotes the model-implied impulse response. The quarterly empirical impulse responses are three month averages of the monthly empirical impulse responses in Figure 2.

Figure 8: Model-Implied Impulse Responses to August 9, 2011 Announcement

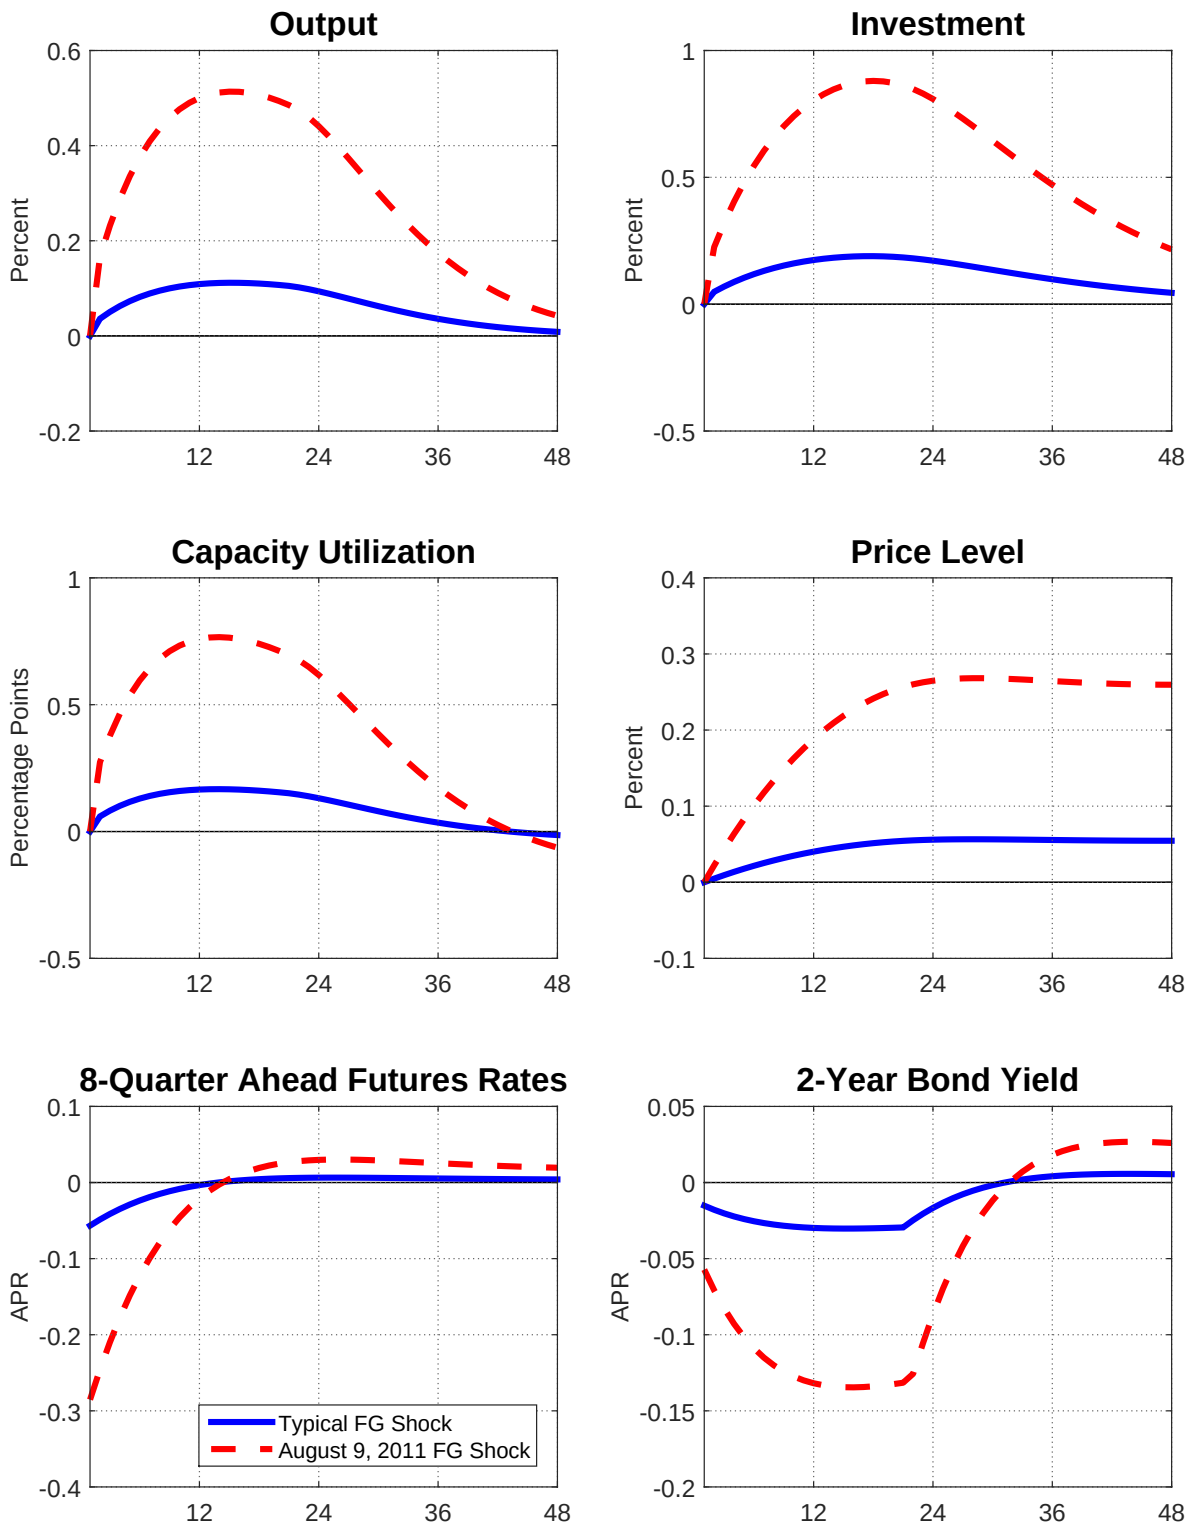

Note: We choose the size forward guidance shock in the model such that the model-implied 2-year Eurodollar rates by the same amount as we observe in the data around the August 9, 2011 FOMC announcement.

Figure 9: Model-Implied Impulse Responses to Exogenous Zero Lower Bound Extensions

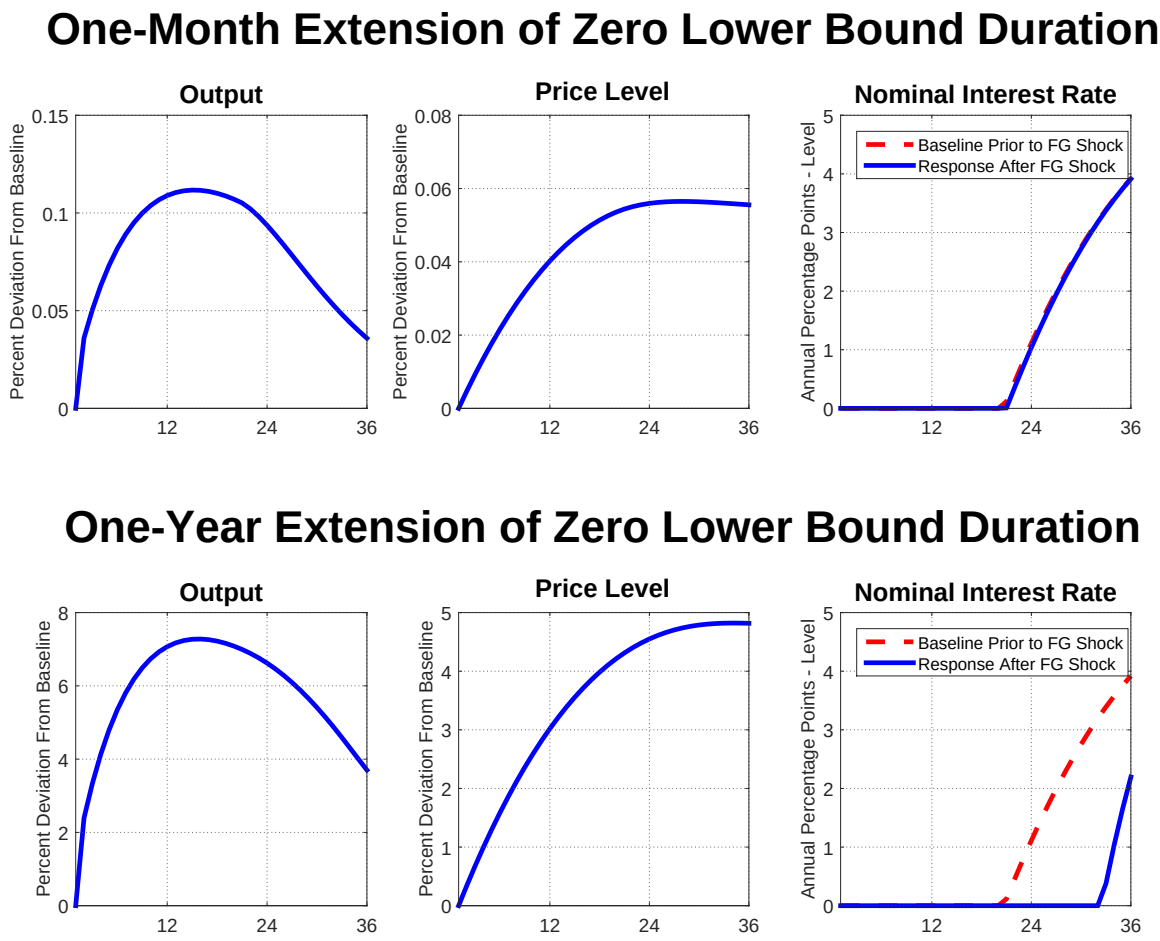

Note: The plot of the nominal interest rate reflects its level after the forward guidance shock. The forward guidance shock in the top panel is chosen to match the 6 basis point decline in the 8-quarter ahead futures rates that occurred after the September 13, 2012 FOMC meeting. The shock in the bottom panel is chosen to generate a one-year extension of the zero lower bound.

Figure 10: Empirical Impulse Responses Including Real Treasury Yields

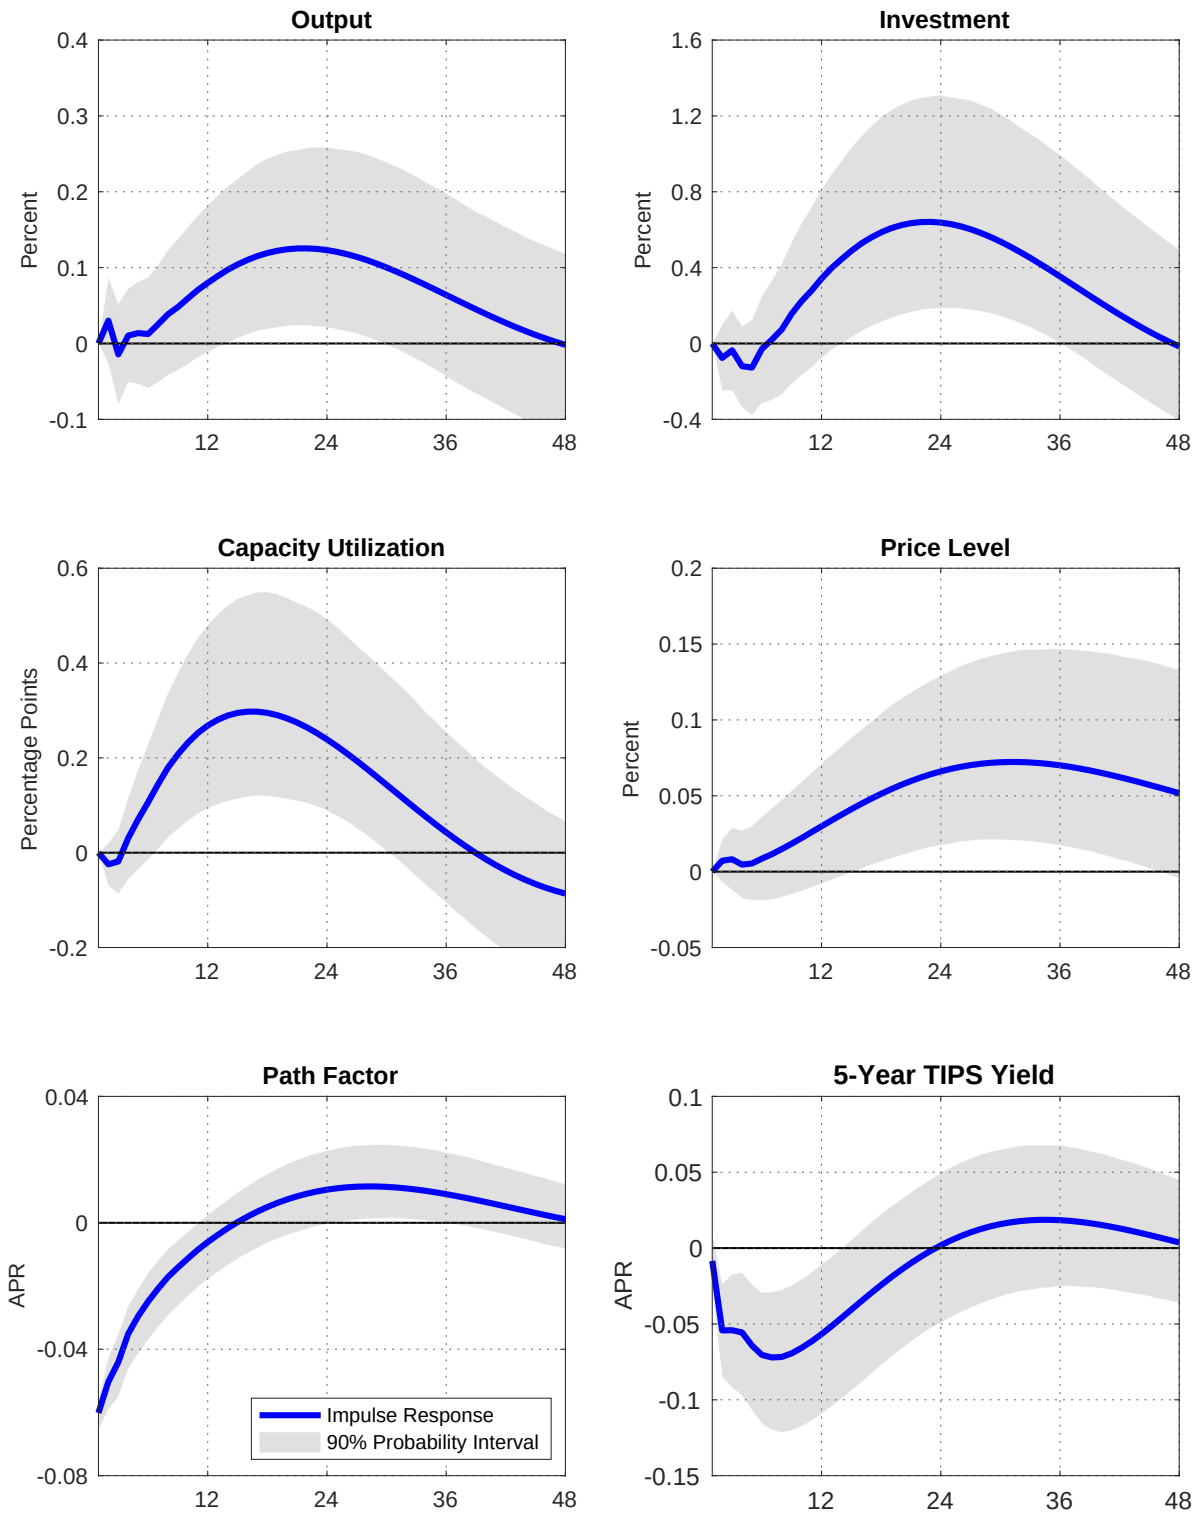

Note: The solid blue lines denote the empirical point estimate to a one standard deviation shock and the shaded areas denote the 90% probability interval of the posterior distribution. Due to the availability of Treasury-Inflation Protected Securities (TIPS) data, our sample for the empirical Bayes prior begins in January 1999.