ISSN 1936-5330

kcFED Research Working Papers

Should We Be Puzzled by Forward Guidance?

Brent Bundick and A. Lee Smith

April 2020
RWP 20-01
http://doi.org/10.18651/RWP2020-01

Federal Reserve Bank of Kansas City

Should We Be Puzzled by Forward Guidance?*

Brent Bundick[†] A. Lee Smith[‡]

May 2020

Abstract

Although a growing literature argues output is too sensitive to future interest rates in standard macroeconomic models, little empirical evidence has been put forth to evaluate this claim. In this paper, we use a range of vector autoregression models to answer the central question of how much output responds to changes in interest rate expectations following a monetary policy shock. Despite distinct identification strategies and sample periods, we find surprising agreement regarding this elasticity across empirical models. We then show that in a standard model of nominal rigidity estimated using impulse response matching, forward guidance shocks produce an elasticity of output with respect to expected interest rates similar to our empirical estimates. Our results suggest that standard macroeconomic models do not overstate the observed sensitivity of output to expected interest rates.

JEL Classification: E32, E52

Keywords: Forward Guidance, Monetary Policy Shocks, Zero Lower Bound, Impulse Response Matching

*We thank John Keating and Victor Valcarcel for helpful comments, Logan Hotz for excellent research assistance, Liz Willoughby for helpful expositional suggestions, and 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.

[†]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

Communication about future interest rates, known as forward guidance, has become a key feature of U.S. monetary policy over the past 25 years. The use of forward guidance especially intensifies when policy rates are constrained by the zero lower bound. When further reductions in the federal funds rate are infeasible, the Federal Open Market Committee (FOMC) typically provides explicit statements regarding the future path of interest rates. Standard macroeconomic models provide theoretical support for the use of forward guidance to overcome the zero lower bound as the forward-looking nature of households and firms implies that the entire path of interest rates determines current economic activity. By shaping expectations about future interest rates, these models suggest that policymakers retain significant influence over the economy even when the current policy rate is constrained.

After years of acceptance as a policy tool, a rapidly expanding literature challenges the efficacy of forward guidance implied by standard macroeconomic models. Recent contributions by Del Negro, Giannoni and Patterson (2015), Kiley (2016), and many others argue that output is too sensitive to future interest rates models with nominal rigidities. In response to this apparent excessive sensitivity, Del Negro, Giannoni and Patterson (2015), McKay, Nakamura and Steinsson (2016), Angeletos and Lian (2018), Hagedorn et al. (2019), and Michaillat and Saez (2019) suggest various tweaks to the standard forward-looking model which temper the economy's response to forward guidance. Through different mechanisms, these alternative models all give rise to an Euler equation with discounting, thereby reducing the elasticity of output with respect to changes in future interest rates.

Despite this rapidly expanding theoretical literature, little empirical evidence has been put forth regarding the object of central interest: the elasticity of output with respect to future interest rates. Without empirical estimates of this elasticity, we argue that researchers cannot properly evaluate whether output is too sensitive to changes in future interest rates in a standard forward-looking model versus various proposed alternatives.

In this paper, we present empirical estimates of the elasticity of output with respect to future interest rates from several vector autoregression (VAR) models. We first analyze the macroeconomic effects of forward guidance shocks during the zero lower bound period. Specifically, we use high-frequency changes in interest rate futures contracts around FOMC announcements to measure changes in central bank forward guidance. Embedding these shocks into a VAR, we find that, on average, an unexpected one basis point decline in 1-

year ahead interest rates increases output by about four basis points at its peak response. Thus, we estimate an elasticity of the peak response of output with respect to 1-year ahead expected policy rates around -4. We then compare this estimate with the elasticity implied by the VAR models of conventional monetary policy shocks developed by Romer and Romer (2004), Christiano, Eichenbaum and Evans (2005), and Gertler and Karadi (2015). Despite the use of alternative identification strategies and different sample periods, we find surprising agreement of the elasticity of output with respect to changes in future interest rates across VAR models of forward guidance shocks and conventional monetary policy shocks. The various models suggest that an unexpected one basis point decline in 1-year ahead future interest rates leads to a 3-5 basis point increase in output at its peak response. We also find similar agreement across these VAR models if we alternatively measure the elasticity of output with respect to the average movement in the path of interest rates over 2 years.

The narrow range of elasticity estimates provides a robust stylized fact which we can use to evaluate theoretical models. Then, we ask whether a model with nominal rigidities can match our estimates of the elasticity of output with respect to future interest rates. Using impulse response matching, we estimate the model's key parameters using our VAR estimates of the effects of forward guidance shocks during the December 2008 - December 2015 zero lower bound period as data. Our estimation routine leads to standard parameter values that govern the model's key frictions related to price rigidity, habits in consumption, investment adjustment costs, and variable capital utilization. Importantly, the estimated model produces an elasticity of output with respect to future interest rates of about -3.5, well within the range of our VAR estimates.

We find little evidence that a relatively standard macroeconomic model of price rigidity displays an excessive sensitivity of output to expected nominal interest rates 1- to 3-years ahead. However, we do *not* intend to suggest that these models are without flaws. For example, our estimated model predicts that the quantitative effects of a one-period innovation to the real interest rate, holding all other real rates fixed, increase as the real rate shock moves farther into the future. However, over the typical 1- to 3-year time horizon during which policymakers have historically provided rate guidance, our results suggest that a standard forward-looking macroeconomic model remains a reasonably good approximation to the actual economy following a forward guidance shock. In light of this finding, our results suggest that theoretical work addressing the forward guidance puzzle may best be targeted at eliminating potentially undesirable longer-term dynamics while leaving the model's short-run quantitative predictions in tact.

2 Elasticity of Output with Respect to Future Rates

In this section, we provide empirical estimates of the elasticity of output with respect to future interest rates from a range of VAR models. We first examine the implied elasticity from an empirical model estimated during the December 2008 - December 2015 zero lower bound period using a high-frequency-based measure of forward guidance shocks. In our view, this estimation strategy provides the purest assessment of how output responds to announced changes in future interest rates. Since we orthogonalize changes in the path of rates from changes in the target federal funds rate, our forward guidance measure is not contaminated by shocks to the current policy rate. Moreover, since we focus on the zero lower bound period, our baseline empirical estimates further align with the forward guidance thought experiment in theoretical models of changing future interest rates while "holding-fixed" the current policy rate.

Despite the conceptual appeal of our baseline estimates, Barakchian and Crowe (2013), Ramey (2016), and others in the monetary policy VAR literature show that quantitative results from any particular model rarely survive across different identification strategies and sample periods. Therefore, after presenting our preferred estimate of the elasticity as inferred from forward guidance shocks during the zero lower bound period, we estimate the elasticity using the benchmark VAR models of conventional monetary policy shocks developed in Romer and Romer (2004), Christiano, Eichenbaum and Evans (2005), and Gertler and Karadi (2015). Across these approaches, we find surprisingly similar estimates of the elasticity of output to the changes in future interest rates.

2.1 A VAR Model of Forward Guidance Shocks

We use a two-step procedure to examine the macroeconomic effects of forward guidance shocks during the zero lower bound period. First, we identify forward guidance shocks associated with FOMC meetings using high-frequency changes in a combination of federal funds and eurodollar futures contracts.[1] For each regularly-scheduled FOMC meeting from 1994–2015, we compute the daily change in the current month and 3-month ahead federal funds futures rates and the 2-8 quarter ahead eurodollar futures rates. 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

[1] The payoffs of these contracts depend on an underlying future short-term interest rate. The underlying interest rate is the effective federal funds rate for federal funds futures and the 3-month London Interbank Offered Rate (LIBOR) on dollar denominated deposits for eurodollar futures.

policy announcement. Following Gurkaynak, Sack and Swanson (2005), we extract a target and path factor which together summarize almost all of the variation in these futures rates around FOMC announcements. We focus our analysis on the path factor, which captures unexpected changes to the future path of policy rates unrelated to changes in the current policy rate, which we interpret as changes in the FOMC's forward guidance.[2] We scale the path factor such that it moves the 4-quarter ahead eurodollar futures rate one-for-one around FOMC meetings.

In the second step, we embed the cumulative sum of the path factor in a structural VAR to trace out the effects of a forward guidance shock.[3] This approach is similar to the external instruments approach to VAR identification. However, Ramey (2016) argues that the nature of forward guidance shocks, which are inherently "news" shocks or shocks to expectations, warrants caution in their use as an external instrument. Therefore, we take her recommendation and include the shocks to futures rates directly in the VAR, essentially using them as an internal instrument. This method is also akin to the approach outlined in Stock and Watson (2018) who note that ordering the shock of interest first in a Cholesky factorization of the VAR innovations yields consistent estimates of the impulse responses without needing to assume full invertibility of the structural VAR.

We estimate our baseline empirical model at a monthly frequency using several indicators of real economic activity, a measure of aggregate prices, and 1-year ahead expectations of future short-term interest rates. Specifically, we include a monthly measure of real GDP, a proxy for real equipment investment, capacity utilization, the GDP deflator, the path factor, and Blue Chip consensus forecasts for the 3-month Treasury bill rate. Including the 1-year ahead short-rate expectations allows us to read off the elasticity of the peak response of output with respect to future interest rates directly from the impulse responses of the VAR, which is our key object of interest for this paper. 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 core capital goods shipments, a series the Bureau of Economic Analysis uses to calculate the official quarterly investment data, deflated by the PPI price index for capital equipment. GDP, the GDP de-

[2]In the Appendix, we provide additional details on the construction of our target and path factors.

[3]This approach follows Romer and Romer (2004) and Barakchian and Crowe (2013) who discuss that including the cumulative sum of unexpected interest rate changes in a VAR is 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.

flator, and investment enter the VAR in the form of 100 times the natural log of the variable.

Following much of the previous VAR literature studying the effects of conventional monetary policy shocks, we assume that macroeconomic conditions adjust slowly to changes in expected policy rates but survey expectations may respond immediately. To this end, we order our forward guidance shock measure after real activity and the price level but before survey forecasts of interest rates using a recursive identification scheme. At a monthly frequency, this assumption that a monetary policy announcement today does not affect real activity or prices within the period seems plausible. However, we find similar results if we instead order the path factor first.

We estimate and conduct statistical inference on the VAR from a Bayesian perspective. Our primary interest is to estimate the impulse response of output following a forward guidance shock during the zero lower bound period. However, Gurkaynak, Sack and Swanson (2005) illustrate 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 to form our priors for the VAR parameters during the zero lower bound period, which we refer to as our empirical Bayes prior.[4] Using standard information selection criteria, we include 3 lags in the VAR.

Figure 1 plots the estimated impulse responses to an identified forward guidance shock along with their 90% probability intervals. A one standard deviation forward guidance shock lowers survey forecasts for interest rates 1-year ahead by about four basis points on impact and these forecasts continue to decline for roughly 1 year after the shock. This result suggests that interest rate expectations several years ahead decline on impact following a forward guidance announcement. The persistent decline in the expected path of rates induces a gradual expansion in output accompanied by a modest increase in inflationary pressures. By 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 next four years. At its peak response, output increases by about 0.15 percent. Investment and capacity utilization also exhibit hump-shaped responses, but peak a bit sooner than does output. Prices rise slowly over the horizon of the impulse response and level out after 3 or 4 years.

[4] 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 ordinary least squares (OLS) quantities. Then, we use this posterior distribution to define our prior for the VAR parameters at the zero lower bound period. Our implementation follows Koop and Korobilis (2010).

Overall, these results suggest that a forward guidance shock which lowers the expected path of rates leads to a modest but statistically significant increase in economic activity and prices.

2.2 Empirical Elasticity of Output Implied By Forward Guidance Shocks

Using the impulse responses following an identified forward guidance shock from the previous section, we can estimate our key object of interest: The elasticity of output with respect to changes in the path of expected interest rates. Specifically, we measure the elasticity of the peak response of output with respect to 1-year ahead expected interest rates. We choose this measure as it is easy to interpret and we believe it captures the horizon over which policy-makers historically provide guidance. However, during the December 2008 - December 2015 zero lower bound episode, rate guidance also extended over longer horizons. Therefore, we also calculate the implied elasticity of output using a longer 2-year ahead horizon in Section 2.4.

Our VAR model implies that a one standard deviation forward guidance shock reduces 1-year ahead expected policy rates by four basis points, which leads output to rise by about 15 basis points at its post-shock peak. These responses imply an elasticity of the peak response of output with respect to 1-year ahead expected policy rates of about -3.8 (0.15 ÷ -0.04). The top panel of Figure 2 illustrates these responses of output and expected rates graphically, in which the diamonds reflect the point estimates we use to compute this elasticity of interest. The top row of Table 1 contains the point estimates and 90% probability intervals for the peak output response, the response of 1-year ahead interest rates, and the resulting elasticity (the ratio of these two responses). Our VAR model of forward guidance shocks estimates that for every basis point decline in the 1-year ahead rates around an FOMC announcement, output increases by 3.8 percent at its peak response. Though the 90% error band is wide and skewed to the left, it excludes zero, suggesting that this estimated elasticity is statistically different from zero

Various perturbations of our baseline model along a number of dimensions lead to similar estimates of the elasticity of output with respect to 1-year ahead expected interest rates. Figure 2 reports the impulse responses for output and 1-year ahead expected interest rates for our baseline VAR model and variants of the model that: (i) relax the assumption that output doesn't respond within the month to a forward guidance announcement, (ii) include 12 lags, and (iii) are estimated using data solely from the zero lower bound period. The first robustness check orders our high-frequency forward guidance shock series first in the VAR model, which allows all variables to respond immediately to a forward guidance shock. The second robustness check increases the number of lags in the VAR to 12, a common "rule-of-thumb" lag length selection to ensure the VAR dynamics are richly specified. Finally, the third robustness check does not use any data prior to the onset of the zero lower bound period and instead centers the VAR parameters at the OLS estimates over the zero lower bound period.[5]

We observe similar dynamics of output and interest rate expectations across all four forward guidance model specifications. In each case, output increases in a hump-shaped manner and we observe a similar peak response, denoted by the red diamonds. The response of 1-year ahead rates is also similar across all of these specifications. Expected rates decline by roughly four basis points on impact and then fall further over the first year before returning to zero. Table 1 reports elasticity estimates for versions of the VAR model. The quantitative similarity of the output and expected rate responses leads to estimates of the elasticity of the peak response of output with respect to 1-year ahead expected interest rates that are near the point estimate of the baseline model (-3.8). Moreover, the 90% probability intervals of all three models exclude zero but include the point estimate from the baseline forward guidance VAR model. Therefore, the estimated output elasticity following a forward guidance shock appears stable with respect to the timing assumptions, the number of lags we include in the VAR, and the use of data only from the zero lower bound period.[6]

2.3 Empirical Elasticity of Output Implied By Conventional Monetary Policy Shocks

While questions about the quantitative effects of forward guidance motivate us to estimate the elasticity of output with respect to expected interest rates, a large previous literature on the effects of conventional monetary policy shocks can also inform estimates of this elasticity. For example, Romer and Romer (2004) describe the impulse responses to a conventional monetary policy shock: "An impulse response function for output or prices to a monetary policy innovation reflects both the effect of the initial innovation *and the effect of the predictable subsequent moves in the policy measure.*" (emphasis added)

[5]Given the limited sample, we include one lag in the VAR when we only use data from the zero lower bound. Lag-length selection criteria support this decision.

[6]Our estimated elasticity also appears to align with the elasticity implied by the estimates in D'Amico and King (2017). They report that a forward guidance-induced 10 basis point decline in Treasury bill rates over the next year increases output by about 0.6 percent, leading to an implied elasticity of roughly -6.0.

In other words, VAR models of the effects of conventional monetary policy shocks provide useful variation in the future path of interest rates which we can use to inform estimates of the output elasticity with respect to expected rates. Therefore, results from the prior literature on the effects of conventional policy shocks provide a useful cross-check on our estimates based on forward guidance shocks. Our results in this section below suggest that our preferred estimate of the elasticity as implied by forward guidance shocks is well within the range of the estimates of this elasticity implied by the previous literature. We conclude from this analysis that there is nothing peculiar in our VAR model that is driving our estimated elasticity. Instead, we find broad empirical support extending beyond our evidence from the zero lower bound period that a one basis point decline in 1-year ahead expected interest rates leads to a roughly four basis point increase in output at its peak response.

We compare our baseline estimate of the elasticity of output with respect to 1-year ahead expected rates to the same elasticity implied by the influential and highly-cited works of Romer and Romer (2004), Christiano, Eichenbaum and Evans (2005), and Gertler and Karadi (2015). Importantly, the VAR models estimated in these papers use a range of identification strategies including narrative, recursive, and high-frequency external instrument approaches. In addition, these studies span multiple sample periods including 1969-1996, 1964-1995, and 1983-2012 (respectively), which encompass several monetary regimes marked by the Great Inflation, post-Bretton Woods, the Great Moderation, and the beginning of the inflation targeting era.[7] Therefore, any similarity across our VAR model of forward guidance and these other works is not driven by a common identification strategy or sample period.

In the context of these models of conventional monetary policy shocks, we estimate the elasticity of output using the 1-year ahead movement in expected short-rates as predicted by the VAR following an identified policy shock. This approach is in line with the macro-finance

[7] We use the baseline empirical model and sample period in Christiano, Eichenbaum and Evans (2005). For Romer and Romer (2004), we estimate the implied elasticity using the VAR model from Section II.D of their paper. For the Gertler and Karadi (2015) results, we chose to work with their robustness exercise which uses the 1983-2012 (see their footnote 7 on page 54). We chose this sample for two reasons. First, over their original sample (1979-2012), we find much less persistence in the response of the funds rate following a monetary policy shock which leads to difficulties in calculating the implied elasticity with respect to future rates. However, we find that these transitory rate dynamics are apparently driven by the Monetarist Experiment period of 1979-1982. Second, Gertler and Karadi (2015) are able to estimate the version of their model with Blue Chip forecasts only over the 1982-2012 period, a specification that is a particularly interesting point of comparison for our analysis given that we also use these Blue Chip forecasts in our VAR model of forward guidance shocks.

literature that combines macroeconomic aggregates with interest rates to elicit interest-rate expectations, as in Ang and Piazzesi (2003), and has been used more recently in the monetary VAR literature by Gertler and Karadi (2015). In essence, using the VAR to forecast the response of 1-year ahead interest rates allows us to read this elasticity off of the impulse responses for output and the short-term policy rate. In the monthly-frequency VAR models of Romer and Romer (2004) and Gertler and Karadi (2015), we measure 1-year ahead policy rate expectations using the impulse response of the short-term policy rate 12 periods after the impact of the monetary policy shock. Similarly, we use the response of the federal funds rate four periods after the impact of the monetary policy shock in the quarterly VAR model of Christiano, Eichenbaum and Evans (2005) as our measure of 1-year ahead policy rate expectations. Figure 3 show the responses of output and interest rates in each VAR model and Table 1 contains the estimated point estimates and their associated confidence intervals. The diamonds again reflect the points we use to compute the elasticities of interest.

The estimated elasticity of output with respect to changes in future interest rates is quantitatively similar across our VAR models of forward guidance shocks as well as these three VAR models of conventional policy shocks. The last column of Table 1 suggests that each one basis point decline in 1-year ahead future interest rates leads to a 3-5 basis point increase in output at its peak response, implying that the estimated elasticity of -3.8 from our VAR model of forward guidance is well within the range of elasticities implied by other models. Figure 7 illustrates this agreement across models graphically by plotting the point estimate of the elasticity as well as the probability interval surrounding that point estimate for each VAR model.

The model of Romer and Romer (2004) generates an elasticity, both in terms of its point estimate and error band, closest to our baseline VAR model of forward guidance shocks. The Romer and Romer (2004) model is also the only model out of the three conventional policy shock specifications that produces precise estimates of this elasticity. The VAR models of Christiano, Eichenbaum and Evans (2005) and Gertler and Karadi (2015) produce wide probability intervals that contain zero. The wide probability intervals are driven by the uncertainty surrounding the path of 1-year ahead expected rates. In particular, as Figure 3 illustrates, in the Christiano, Eichenbaum and Evans (2005) and Gertler and Karadi (2015) models, short-term policy rates revert back toward zero quickly whereas in the model of Romer and Romer (2004), we observe more propagation in the policy rate following a shock.

Despite the uncertainty surrounding the implied path of short-term policy rates, the point estimates imply a similar elasticity across all the VAR models. This finding is surprising given the apparent disconnect that Coibion (2012) discusses across empirical models of the effects of conventional monetary policy shocks. However, we observe significant uncertainty around our estimates of the elasticity of output with respect to 1-year ahead interest rates for the Christiano, Eichenbaum and Evans (2005) and Gertler and Karadi (2015) models. Therefore, we now turn to an alternative elasticity estimate which can be more precisely estimated across various VAR models to verify the robustness of our findings.

2.4 Robustness to An Alternative Elasticity Calculation

Our results in the previous section show that our model of forward guidance shocks and a range of models of conventional policy shocks all generate broadly similar output elasticities in response to monetary policy shocks. However, one may be concerned that, due to conceptual differences between forward guidance and conventional policy shocks, our previous analysis is not an "apples-to-apples" comparison. Since we measure forward guidance shocks using a Gurkaynak, Sack and Swanson (2005)-style path factor at the zero lower bound, our measure captures changes in the future path of rates holding current policy rates fixed. By contrast, Figure 3 shows that in the models of Romer and Romer (2004), Christiano, Eichenbaum and Evans (2005), and Gertler and Karadi (2015), conventional policy shocks not only affect the expected path of rates but also affect the current policy rate at impact. Therefore, as Romer and Romer (2004) note, in VAR models of conventional monetary policy shocks, the impulse response functions for output capture the combined effects of both the initial innovation and the later forecastable policy moves.

To put our forward guidance shock model and the conventional monetary policy shock models on more equal footing, we now calculate the elasticity of output with respect to the change in interest rates over several years as opposed to using just a single point in the future. Specifically, rather than measuring changes in expected rates at just the 1-year ahead horizon, we now integrate the area under the policy rate response for 2 years when computing the elasticity for each model. This alternative elasticity calculation captures changes in the entire path of rates, not just changes in interest rates at a given horizon, which may better measure the total amount of monetary accommodation from a given shock.[8] For instance, as Woodford (2003) emphasizes, forward-looking theoretical models imply that the entire path of rates matters for determining current economic activity.[9]

[8] In unreported results, we extend the horizon to 3 years and find similar agreement regarding the elasticity of output across VAR models.

[9] By this logic, one could argue for integrating under the entire response of interest rates. However, in

This alternative measure of the elasticity of output with respect to future interest rates also addresses another thorny issue: the tendency for the response of policy rates to revert towards zero about 1 year after a conventional policy shock. For instance, in the models of Christiano, Eichenbaum and Evans (2005) and Gertler and Karadi (2015), the impulse response for the federal funds rate contain zero after 1 year which results in an imprecisely estimated elasticity of output with respect to 1-year ahead interest rates. Therefore, it would be difficult to either accept or reject any theoretical model on the basis of the estimated elasticity of output with respect to 1-year ahead interest rates from these models. However, in both the Christiano, Eichenbaum and Evans (2005) and Gertler and Karadi (2015) models, we can estimate the average change in the federal funds rate over 2 years with more precision.

While this alternative average elasticity is slightly harder to interpret than our baseline 1-year ahead elasticity, Figure 5 illustrates the calculations behind this alternative elasticity. For each model, we compute the elasticity using the peak response of output divided by the average change in expected interest rates 2 years following the shock. For our model of forward guidance shocks, we assume that policy rates remain fixed for the first year following the shock, which is consistent with the ex ante durations of the zero lower bound in Swanson and Williams (2014) over the December 2008 – July 2011 period. Thus, one can think of our forward guidance shock as similar to a conventional policy shock except that current policy rates cannot move at impact and only the path of rates changes.

We find more precise estimates and even stronger agreement in the output elasticity across forward guidance and conventional monetary policy VAR models using this alternative elasticity calculation. Table 2 reports the estimated output elasticity with respect to the change in expected rates over the next 2 years across VAR models. Our baseline VAR model of forward guidance shocks shares a very similar elasticity to Romer and Romer (2004) and Gertler and Karadi (2015). Unlike the 1-year ahead elasticity estimates, the 90% probability intervals around the 2-year average elasticity estimates do not contain zero for all of the conventional monetary policy shock models. Figure 8 illustrates this result graphically and shows the estimated elasticity of output with respect to the change in expected rates over the next 2 years, along with the associated error bands for each VAR model. As Figure 8 underscores, while the implied elasticity from Christiano, Eichenbaum and Evans (2005) is a bit larger than the other three VAR models, owing somewhat mechanically to the over-

practice, this elasticity measure can behave oddly (even switching sign) due to the tendency for the impulse responses to oscillate around zero. Ramey (2016) reaches the same conclusion regarding a similar measure.

shoot of the funds rate above zero after about 6 quarters, the probability interval around the Christiano, Eichenbaum and Evans (2005) estimate excludes zero and contains the point estimates from all of the other models, signaling no evidence of statistically significant differences in the output elasticity across models.

Moreover, the surprising agreement across models for both elasticity calculations suggests that the cumulative path of interest rates matters more than the exact timing of the policy accommodation, which is consistent with standard forward-looking theoretical models. For example, at least over the 2 year horizon, the results in this section suggest that each basis point of accommodation produces a similar response of output regardless of whether the current policy rate is constrained or not. These results qualitatively support the prediction from standard forward-looking models that policy makers retain influence over the economy, even at the zero lower bound, due to their ability to shape expectations about future policy rates. In the next section, we use these elasticity calculations as a robust stylized fact and ask whether a standard forward-looking model implies an excessive sensitivity of output with respect to changes in future policy rates.

3 A Theoretical Model of Nominal Price Rigidity

This section outlines a relatively standard dynamic stochastic general equilibrium model, which we use to analyze the theoretical responses to forward guidance shocks. The model shares features with the models of Ireland (2003, 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 align with the timing assumptions in our baseline empirical VAR analysis, we assume that consumption, investment, and firm pricing decisions in the model are made before these shocks are realized.[10]

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

[10]The Appendix provides details regarding all of the model's equilibrium conditions.

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 a zero net-supply 1-period nominal B_t bond, which pays one dollar and is purchased at a discounted price $1/R_t$, where R_t is the 1-period gross nominal interest rate. The household divides its income between consumption C_t and the amount of the bonds B_{t+1} and B^R_{t+1} to carry into next period.

The representative household maximizes lifetime utility by choosing $C_{t+s}, N_{t+s}, B_{t+s+1}$, and B^R_{t+s+1}, for all $s = 0, 1, 2, \dots$ 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,

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

λ_t denotes the Lagrange multiplier on the household budget constraint. In equilibrium, consumption habits are formed external to the household and are 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, so we can examine the effects of a forward guidance shock at the zero lower bound. The stochastic process for these fluctuations is as follows:

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

where ε^a_t is an independent and standard normal random variable.

3.2 Final Goods Producers

The representative final goods producer uses Y_{it} units of each intermediate good produced by the intermediate goods-producing firm $i \in [0, 1]$. The intermediate output is transformed into final output Y_t using the following constant returns to scale technology:

$$\left[\int_0^1 Y_{it}^{\frac{\theta-1}{\theta}} di \right]^{\frac{\theta}{\theta-1}} \geq Y_t,$$

where θ is the elasticity of substitution across intermediate goods. Each intermediate good Y_{it} sells at nominal price P_{it} and the final good sells at nominal price P_t. The finished goods producer chooses Y_t and Y_{it} for all $i \in [0, 1]$ to maximize the following firm profits:

$$P_t Y_t - \int_0^1 P_{it} Y_{it} di$$

subject to the constant returns to scale production function. Finished goods-producer optimization results in the following first-order condition:

$$Y_{it} = \left[\frac{P_{it}}{P_t} \right]^{-\theta} Y_t.$$

The market for final goods is perfectly competitive, and thus the final goods-producing firm earns zero profits in equilibrium. Using the zero-profit condition, the first-order condition for profit maximization, and the firm objective function, the aggregate price index P_t can be written as follows:

$$P_t = \left[\int_0^1 P_{it}^{1-\theta} di \right]^{\frac{1}{1-\theta}}.$$

3.3 Intermediate Goods Producers

Each intermediate goods-producing firm i rents labor N_{it} from the representative 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. We introduce a production subsidy $\Psi = \theta/(\theta - 1)$ to ensure that the steady state of the model is efficient. Firms rebate any profits to the household in lump sum each period.

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,

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

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:

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

subject to the following demand curve,

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

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 and capital as well as the same 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:

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

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 between the log of current and 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 (1993)-type policy rule with interest rate smoothing and 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.

3.6 Generating Model-Implied Forecasts of Future Interest Rates

In evaluating our theoretical model, we want to compute the same elasticity as in our empirical exercises in Section 2.2. To this end, we want to choose the appropriate values for our forward guidance shock process such that we generate the same movements in 1-year ahead interest rate expectations that we observed in our empirical VAR model of forward guidance shocks. In our VAR model, we measure the effect of forward guidance shocks on interest rate expectations using 1-year ahead Blue Chip consensus forecasts for the 3-month Treasury bill rate. Therefore, we generate a model counterpart to this survey measure.

In our piecewise-linear solution method, which we discuss below, the expectations hypothesis describes the evolution of interest rates solely as a function of expectations of future short-term interest rates. Therefore, we define the annualized 3-month Treasury bill rate as follows:

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

where r_{t+n} is the policy rate of the central bank in month $t+n$. Each month, Blue Chip forecasters are asked to provide a forecast for what the 3-month Treasury bill rate will average for the 3 months, 1 year from now. For example, in January 2008, Blue Chip forecasters submitted their forecast for what the 3-month Treasury bill will average across January, February, and March of 2009. We therefore define the model-implied Blue Chip forecast of the average 3-month Treasury bill rate for the 3 months, 1-year ahead, $\mathrm{E}_t^{BC}\left\{y_{t+12:t+15}^3\right\}$, as:

$$\mathrm{E}_t^{BC}(y_{t+12:t+15}^3) = \mathrm{E}_t\left\{ \frac{1}{3}\Big(y_{t+12}^3 + y_{t+13}^3 + y_{t+14}^3\Big) \right\}. \tag{6}$$

This model counterpart to the Blue Chip interest rate forecasts allows us to determine the appropriate-sized forward guidance shock to simulate in the model. To be consistent with the timing assumptions in our structural VAR, we assume that forecasts of interest rates 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 solution method allows us to model the occasionally-binding zero lower bound and solve for the model-implied interest rate expectations. The algorithm takes only a few seconds to solve the model, which permits us to estimate several key model parameters. The solution method constructs a piecewise linear approximation to the original nonlinear model. We have also solved a fully nonlinear, but simplified, version of our model with the policy function iteration method of Coleman (1990). We find that the Guerrieri and Iacoviello (2015) toolkit provides a good approximation dynamics of the full nonlinear economy after a forward guidance shock.

3.8 Estimation Strategy

Our primary interest is evaluating whether this relatively standard model implies an excess sensitivity of output relative to our empirical impulse responses following a forward guidance shock from Section 2.1. To measure the model-implied elasticity of output with respect to expected future policy rates, 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. 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. 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.

Our estimation strategy chooses model parameters such that the model's impulse responses come as close as possible to the empirical responses of output, investment, capacity utilization, the price level, and 1-year ahead interest rate expectations from Figure 1. This estimation strategy directly targets our elasticities of interest from the top rows of Figures 2 and 5. However, to ensure that our model does not try produce counterfactual dynamics while trying to fit these elasticities, we jointly aim to match impulse responses of output, investment, capacity utilization, the price level, and 1-year ahead interest rate expectations for 48 months, implying our model is over identified.

To implement this estimation strategy, we follow the spirit of 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 our primary empirical evidence 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 as well as the size of the adverse demand shock that initially sends the economy to the zero lower bound such that the model generates the same movement in the 4-quarter ahead of forecast of short-term interest 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 or aggregate demand shock to simulate in the model. Thus, even if one focuses on the elasticity of output with respect to future rates as we do (which should be somewhat invariant to the shock size), it is still imperative to discipline the size of the underlying forward guidance and aggregate demand shocks.

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 parameters using steady-state relationships or results from previous studies. Since the model shares features with the models of Ireland (2003, 2011), we calibrate some of our parameters to match his values or estimates. To match our VAR evidence, we calibrate the model at 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 2 percent and a 2 percent annualized inflation target. Table 3 summarizes these calibrated parameters.

We estimate the second set of model parameters which consists of the household habit parameter b, the probability that a firm can not re-optimize 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 5 variables in our empirical VAR stacked into a single vector with 240 (5×48) rows and let the diagonal matrix V^{-1} denote a measure of the precision of the empirical impulse responses.[11] Then, let $\psi(\gamma)$ denote the theoretical model's corresponding counterpart to $\hat{\psi}$. Following Christiano, Trabandt and Walentin (2010), we

[11]The diagonal of V^{-1} contains 1 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.

can write the approximate likelihood function as follows:[12]

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

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

$$f(\gamma \mid \hat{\psi}, V) \propto L(\hat{\psi} \mid \gamma, V) p(\gamma), \tag{7}$$

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

$$\max_{\gamma} f(\gamma \mid \hat{\psi}, V). \tag{8}$$

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 1 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.

[12] Christiano, Eichenbaum and Trabandt (2016) provide three reasons why this is only an approximate likelihood: (i) Standard asymptotic theory implies that under the assumption that the DSGE model is the correct data generating process with the true parameters γ_0, $\hat{\psi}$ converges only asymptotically to $N(\psi(\gamma_0), V)$ as the sample size grows arbitrarily large, (ii) our proxy for V is guaranteed to be correct only as the sample size grows arbitrarily large, and (iii) $\psi(\gamma)$ is approximated with a piece-wise linear DSGE model. A fourth reason in our application is due to the fact that in a non-linear model the IRFs are not a full summary of the model like they are in a linear model.

4 Estimated Responses & the Elasticity of Output

We now analyze the effects of a forward guidance shock in our estimated model and evaluate whether the model implies an excessive sensitivity of output with respect to the path of rates. Figure 6 plots impulse responses to a forward guidance shock both in the data and in our estimated dynamic, stochastic, general-equilibrium (DSGE) model.

At impact, a forward guidance shock causes the model-implied forecasts for the 1-year ahead Treasury Bill rate to decline by about 3.5 basis points which is nearly identical to our empirical findings. Then, the model-implied path of rates continues to decline for several months before reaching a trough 4 months after the shock. About 6 months after the shock, 1-year ahead rate forecasts gradually returning to their pre-shock level.

The interactions between the endogenous policy response to the adverse demand shock that sends the economy to the zero lower bound and expectations for eventually exiting the zero lower bound help the model reproduce the U-shaped path of expected rates that we observe in the data. Following the adverse demand shock but prior to the forward guidance shock, the economy is expected to remain at the zero lower bound for a little over 1 year. Therefore, when policymakers unexpectedly announce a new lower path of future rates (the forward guidance shock), forecasters can only modestly revise down 1-year ahead forecasts for the 3-month Treasury bill rate as short-term rates were already expected to be low. However, Treasury bill rate forecasts 16 months ahead have more space to decline since, by that time, rates were expected to be further above zero prior to the forward guidance shock. These elements cause the 16-month ahead Treasury bill rate forecasts to fall more than 12-month ahead Treasury bill rate forecast and allow the model to deliver the U-shaped path of forecast rates we see in the data despite the underlying forward guidance shock following a first-order autoregressive process.

Importantly, the mechanisms discussed above allow the model-implied path to closely mirror the VAR-implied path of short-term rate forecasts. Since the expected rate dynamics are nearly identical, we can cleanly compare the VAR- and model-implied elasticities of output under similar interest rate paths. Following the forward guidance shock, output, investment, and capacity utilization in the model all rise 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 similar to the peak response we estimate in our VAR model, although the model-implied peak is slightly lower and occurs a bit earlier earlier than in the data. Therefore, if anything, one might argue that the model features too low of a sensitivity of output to future policy rates. However, neither the magnitude nor timing of the peak output response differs from the empirical model given the uncertainty bands that surround the VAR impulse responses. Therefore, given that all of the model's impulse responses fall within the VAR's probability intervals, we conclude that the quantitative predictions from our relatively standard model of monetary policy are generally in line with the empirical responses of a forward guidance shock.

Our estimated model generates an elasticity of output with respect to future interest rates that is quantitatively consistent with our empirical evidence. The last row of Table 1 precisely quantifies the peak output response as well as the decline in 1-year ahead expected short-term interest rates from our estimated model. At its peak, output increase by 0.12 percent following a 3.5 basis point decline in the 1-year ahead forecast of short-term rates. These outcomes imply an elasticity of -3.3, which is quite close to our baseline VAR model of forward guidance shocks which implies an elasticity of -3.8 with a 90% probability interval ranging from -10.19 to -1.25. Therefore, the elasticity of output with respect to expected rates in our DSGE model is not statistically different from our VAR-implied elasticity.

More generally, the DSGE model-implied elasticity is well within the range of elasticities implied by the VAR models of both forward guidance and conventional monetary policy shocks. Figure 7 illustrates this close match by showing the implied elasticity of output with respect to 1-year ahead interest rates in all of the VAR models as well as in the estimated DSGE model. The estimated elasticity in the DSGE model is not only well within the error-band range of elasticity estimates across VAR models but is also very near the point estimates across these empirical models. Figure 8 shows similar synchronicity between the various VAR models and the estimated DSGE model when we measure the elasticity using the ratio of the peak output response to the average decline in interest rates over the next 2 years. Overall, our results suggest little disconnect between the elasticity of output implied by a range of VAR models and a relatively standard forward-looking theoretical model. These findings suggest that a standard model of nominal rigidity without any form of discounting in the Euler equation can reproduce the empirical estimates of the elasticity of output with respect to expected future interest rates.

4.1 Parameter Estimates

We now discuss the parameter estimates which deliver the close fit between the model-implied and empirical effects of forward guidance shocks. The model requires a mix of nominal as well as real rigidities to match the VAR evidence. Table 4 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), Eichenbaum and Fisher (2007), and Del Negro, Giannoni and Schorfheide (2015). We find essentially no role for lagged indexation of prices with $\chi = 0.03$, which likely reflects a decline in the persistence of inflation over time.

In addition to a moderate degree of nominal rigidity, consumption habits, investment adjustment costs, and variable capacity utilization help the model reproduce the empirical evidence. Our estimate of consumption habits b is higher than estimates of Christiano, Eichenbaum and Evans (2005) and Smets and Wouters (2007), 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 firms make more incremental adjustments in their capital stock at a monthly frequency than they do at a quarterly frequency.

We estimate a significant degree of desired-rate smoothing in the central bank's policy rule. However, our estimate of $\phi_r = 0.95$ doesn't significantly differ from its prior mode which suggests that ϕ_r may not be well-identified by our impulse response matching procedure. 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. We also estimate a fairly persistent process for the forward guidance innovation, suggesting that a prolonged reduction in the intercept of the central bank's policy rule is needed to match the persistent decline in 1-year ahead interest rate forecasts that we see in the VAR evidence.

5 Discussion

Our results support the use of standard forward-looking macroeconomic models as laboratories for examining the effects of central bank forward guidance. However, these results do not suggest that these models perfectly describe reality nor can they critically evaluate every possible policy experiment. More precisely, our results suggest for the 1-3 year ahead horizon, a relatively standard forward-looking model seems to be a reasonable description of the macroeconomic outcomes following changes in forward guidance. This horizon represents an important benchmark as policymakers typically communicate about the path of policy over the next 1-3 years. For example, in their Summary of Economic Projections, FOMC participants provide their appropriate policy path at the end of the current year and the following 2 years.

5.1 Policy-Relevant Horizon

Recent work by McKay, Nakamura and Steinsson (2016) suggests that the effects of far in the future forward guidance are too powerful and that these effects grow with the horizon of the guidance in standard forward looking models. Using a textbook model of nominal rigidities that abstracts from capital and other real rigidities, they argue that these features of standard forward-looking models represent a counterfactual flaw. As a fix, they introduce idiosyncratic risk and borrowing constraints which introduces discounting into the consumption Euler equation at the aggregate level and tempers the economy's response to changes in far forward real rates. Using our estimated model, we now reproduce their key forward guidance experiment. Figure 9 shows the responses of our estimated model if we simulate a 100 basis point decline in the real interest rate for a single period in the future, holding the real interest rate fixed in all other periods.[13]

As in McKay, Nakamura and Steinsson (2016), our model predicts that the quantitative effects of a one-period real rate shock increase as the shock moves farther and into the future. However, in the context of the empirical evidence we present in this paper, these experiments suggest two key insights. First, for a 3-year (36 month) ahead real rate shock, the model-implied responses of output in Figure 9 are of similar quantitative magnitude to the empirical impulse responses of our baseline empirical model of forward guidance shocks

[13] For these experiments, we replace our policy rule in Equation 2 with a real rate targeting rule: $r_t - r = 1.000001 * \mathrm{E}_t \left\{ \pi_{t+1} - \pi \right\} + \varepsilon^{12}_{t-11} + \varepsilon^{36}_{t-35} + \varepsilon^{60}_{t-59}$, where each ε represents either a 12-, 36-, or 60-period ahead real rate shock. We set the coefficient on expected inflation infinitesimally above 1 in order to get a determinant equilibrium. Also, we abstract from the zero lower bound in these experiments.

in Figure 1. Thus, our estimated forward-looking model produces reasonable quantitative magnitudes even under their forward guidance experiment for policy-relevant horizons up to 3 years. More importantly, however, the model-implied path of short-term *nominal* policy rates in these experiments looks quite different than any of the nominal rate paths we observe in the data following forward guidance or conventional monetary policy shocks (see the right column of Figure 3). Since we do not have an empirical counterpart against which to judge the model's predictions for these longer-horizon real-rate shocks, we argue that one cannot properly judge whether these predictions represent a fundamental flaw in the model or simply an empirically less-interesting policy experiment.

5.2 Different Characterizations of the Forward Guidance Puzzle

Regardless of our specific conclusions, we believe a key disagreement among the literature is the exact definition of the "forward guidance puzzle" used across papers. For example, in the seminal "Forward Guidance Puzzle" paper of Del Negro, Giannoni and Patterson (2015), these authors examine the effects on output and inflation following a roughly 1-year exogenous extension of the zero lower bound episode. This experiment is likely aimed to capture the adoption of the "mid-2015" language at the September 2012 FOMC meeting, which extended the previous rate guidance of "late 2014." McKay, Nakamura and Steinsson (2016) instead examine the effects of a one-period shock to real rates at a given point in the future, holding real rates fixed in all other periods. In contrast, Kiley (2016) expresses a "policy paradox" in a standard model with nominal rigidities by examining the model-implied elasticity of output with respect to changes in nominal interest rates. All of these variations of the puzzle require different sets of testable predictions in order to properly evaluate a standard macroeconomic model.

We prefer the Kiley (2016) formulation of the puzzle as elasticities are less sensitive to the size of the forward guidance shock and focus on nominal rather than real short-term rates, which are fully under the control of the central bank. We believe the focus on nominal short-rate thought experiments aligns with the types of quantitative analysis most useful to policymakers. However, we argue that, even if one wishes to study alternative formulations of the forward guidance puzzle, researchers need to use empirical evidence of the effects of changes in central bank forward guidance in order to properly evaluate the predictions of forward-looking macroeconomic models.

6 Conclusion

This paper studies the elasticity of output with respect to changes in the path of interest rates in both the data and a theoretical model. We find surprising agreement regarding this elasticity across a range of VAR specifications, which suggests that this elasticity represents a reasonably robust stylized fact against which to judge standard macroeconomic models. After estimating a standard DSGE model using impulse response matching, we find that the model-implied elasticity of output is quite consistent with our empirical estimates without the need to introduce discounting into the household's Euler equation. This evidence calls into question the need to develop new models that feature a reduced sensitivity of output to forward guidance.

References

Ang, Andrew, and Monika Piazzesi. 2003. "A No-Arbitrage Vector Autoregression of Term Structure Dynamics with Macroeconomic and Latent Variables." *Journal of Monetary Economics*, 50(4): 745–787.

Angeletos, George-Marios, and Chen Lian. 2018. "Forward Guidance without Common Knowledge." *American Economic Review*, 108(9): 2477–2512.

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

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

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, Martin S Eichenbaum, and Mathias Trabandt. 2016. "Unemployment and Business Cycles." *Econometrica*, 84(4): 1523–1569.

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

Coibion, Olivier. 2012. "Are the Effects of Monetary Policy Shocks Big or Small?" *American Economic Journal: Macroeconomics*, 4(2): 1–32.

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

Coleman, Wilbur John. 1990. "Solving the Stochastic Growth Model by Policy Function Iteration." *Journal of Business & Economic Statistics*, 8(1): 27–29.

D'Amico, Stefania, and Thomas B. King. 2017. "What Does Anticipated Monetary Policy Do?" Federal Reserve Bank of Chicago Working Paper 2015-10.

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.

Eichenbaum, Martin, and Jonas D.M. Fisher. 2007. "Estimating the Frequency of Price Re-Optimization in Calvo-style models." *Journal of Monetary Economics*, 54(7): 2032 – 2047.

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.

Hagedorn, Marcus, Jinfeng Luo, Iourii Manovskii, and Kurt Mitman. 2019. "Forward Guidance." *Journal of Monetary Economics*, 102: 1–23.

Ireland, Peter N. 2003. "Endogenous Money or Sticky Prices." *Journal of Monetary Economics*, 50: 1623–1648.

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

Kiley, Michael T. 2016. "Policy Paradoxes in the New Keynesian Model." *Review of Economic Dynamics*, 21: 1–15.

Koop, Gary, and Dimitris Korobilis. 2010. *Bayesian Multivariate Time Series Methods for Empirical Macroeconomics.* Now Publishers Inc.

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

Michaillat, Pascal, and Emmanuel Saez. 2019. "Resolving New Keynesian Anomalies with Wealth in the Utility Function." *Review of Economics and Statistics.* Forthcoming.

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

Ramey, Valerie A. 2016. "Macroeconomic Shocks and Their Propagation." Working Paper.

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.

Smets, Frank, and Rafael Wouters. 2007. "Shocks and Frictions in US Business Cycles: A Bayesian DSGE Approach." *American Economic Review*, 97(3): 586–606.

Stock, James H, and Mark W Watson. 2018. "Identification and Estimation of Dynamic Causal Effects in Macroeconomics Using External Instruments." *The Economic Journal*, 128(610): 917–948.

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.

Taylor, John B. 1993. "Discretion Versus Policy Rules in Practice." *Carnegie-Rochester Conference Series on Public Policy*, 39: 195–214.

Woodford, Michael. 2003. *Interest and Prices.* Princeton University Press.

Table 1: Estimated Output Elasticity With Respect to 1-Year Ahead Expected Rates

VAR Model	**Peak Output Response (Percent)**	**1-Year Ahead Expected Rate Response (APR)**	**Elasticity (Percent/APR)**
Forward Guidance Model: Baseline	0.16*	-0.04*	-3.82*
	(0.03, 0.33)	(-0.06, -0.02)	(-10.19, -1.25)
Forward Guidance Model: Path First	0.15*	-0.04*	-3.45*
	(0.02, 0.33)	(-0.06, -0.03)	(-9.29, -0.96)
Forward Guidance Model: 12 Lags	0.16*	-0.02*	-6.66*
	(0.04, 0.30)	(-0.04, -0.01)	(-19.46, -2.89)
Forward Guidance Model: ZLB Only	0.18*	-0.04*	-3.96*
	(0.02, 0.42)	(-0.07, -0.02)	(-12.45, -1.23)
Romer & Romer (2004)	0.54*	-0.17*	-3.23*
	(0.20,0.96)	(-0.25,-0.10)	(-6.44, -1.72)
Christiano, Eichenbaum, & Evans (2005)	0.49*	-0.10	-4.63
	(0.31,0.70)	(-0.33,0.10)	(-27.36, 20.57)
Gertler & Karadi (2015)	0.36*	-0.09	-4.07
	(0.02,0.51)	(-0.14,0.00)	(-22.82, 9.68)
DSGE Model			
Estimated DSGE Model	0.12	-0.04	-3.30

Notes: For the VAR models, 90% probability intervals are in parenthesis and * denotes estimates for which this interval excludes zero. Details on each VAR specification as well as the DSGE model are provided in the main text.

Table 2: Estimated Output Elasticity With Respect to the Change in Expected Rates Over the Next 2 Years

VAR Model	Peak Output Response (Percent)	2-Year Ahead Average Rate Response (APR)	Elasticity (Percent/APR)
Forward Guidance Model: Baseline	0.16*	-0.04*	-4.35*
	(0.03, 0.33)	(-0.06, -0.02)	(-12.27, -1.30)
Forward Guidance Model: Path First	0.15*	-0.04*	-3.80*
	(0.02, 0.33)	(-0.06, -0.02)	(-10.53, -1.00)
Forward Guidance Model: 12 Lags	0.16*	-0.03*	-5.31*
	(0.04, 0.30)	(-0.05, -0.01)	(-17.38, -2.11)
Forward Guidance Model: ZLB Only	0.18*	-0.02*	-8.49*
	(0.02, 0.42)	(-0.04, -0.01)	(-29.07, -2.27)
Romer & Romer (2004)	0.54*	-0.17*	-3.15*
	(0.20, 0.96)	(-0.24, -0.13)	(-5.68, -1.70)
Christiano, Eichenbaum, & Evans (2005)	0.49*	-0.15*	-3.17*
	(0.31, 0.70)	(-0.32, -0.01)	(-15.86, -1.03)
Gertler & Karadi (2015)	0.36*	-0.08*	-4.52*
	(0.02, 0.51)	(-0.12, -0.01)	(-23.20, -0.30)
DSGE Model			
Estimated DSGE Model	0.12	-0.03	-3.55

Notes: For the VAR models, 90% probability intervals are in parenthesis and * denotes estimates for which this interval excludes zero. Details on each VAR specification as well as the DSGE model are provided in the main text.

Table 3: Calibrated Model Parameters

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 4: Estimated Model Parameters

		Prior			Posterior	
Parameter	Description	Distribution	Mode	Std. Dev.	Mode	Std. Dev.
b	Habit Persistence	Beta	0.50	0.25	0.9459	0.0028
ω	Calvo Probability	Beta	0.93	0.01	0.9534	0.0015
χ	Degree of Lagged Indexation	Beta	0.50	0.25	0.0291	0.0187
ϕ_r	Policy Rate Smoothing	Beta	0.95	0.25	0.9489	0.0018
κ	Investment Adjustment	Gamma	2.48	60.0	39.1228	5.5102
σ_δ	Capacity Utilization Curvature	Gamma	0.01	60.0	0.0002	0.0002
ρ_ν	Policy Shock Persistence	Beta	0.50	0.25	0.8964	0.0060
$1200 \times \sigma_\nu$	Std. Dev. of Policy Shock	Gamma	25.0	1200	0.0416	0.0027

Figure 1: VAR Impulse Responses to a Forward Guidance Shock

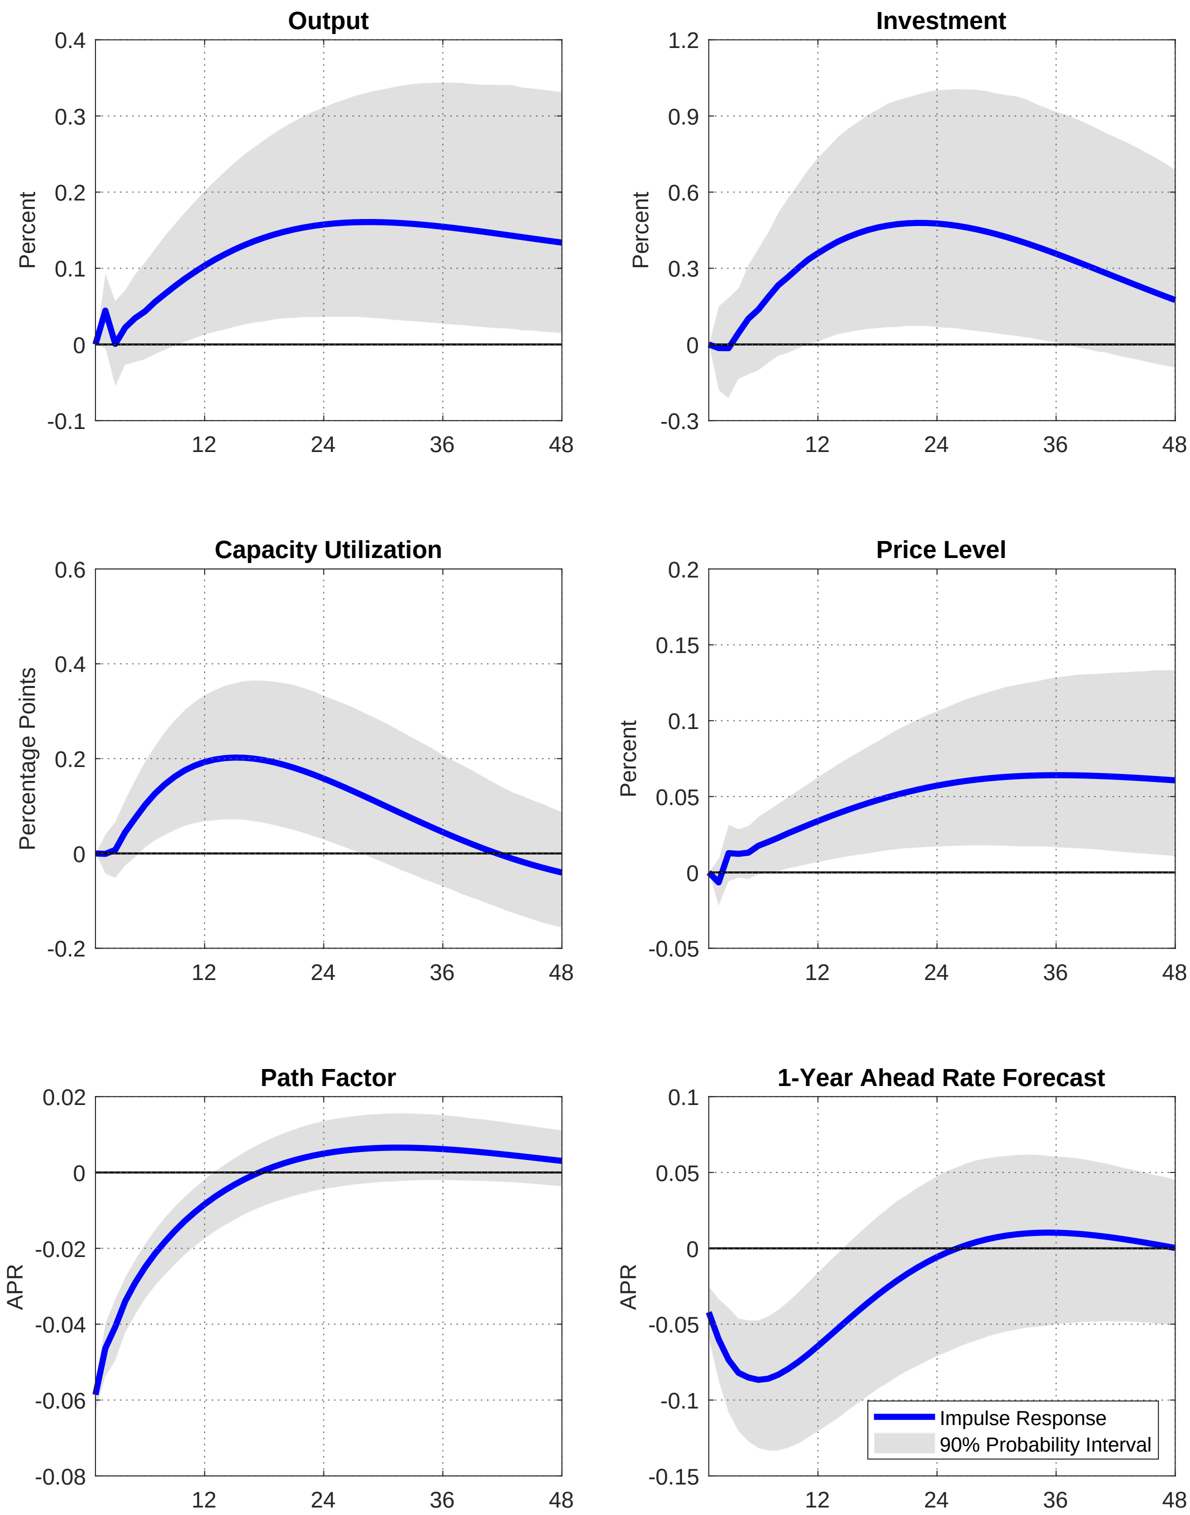

Note: The solid blue lines denote the empirical point estimate to a 1 standard deviation shock and the shaded areas denote the 90% probability interval.

Figure 2: VAR Impulse Responses to a Forward Guidance Shock: Robustness

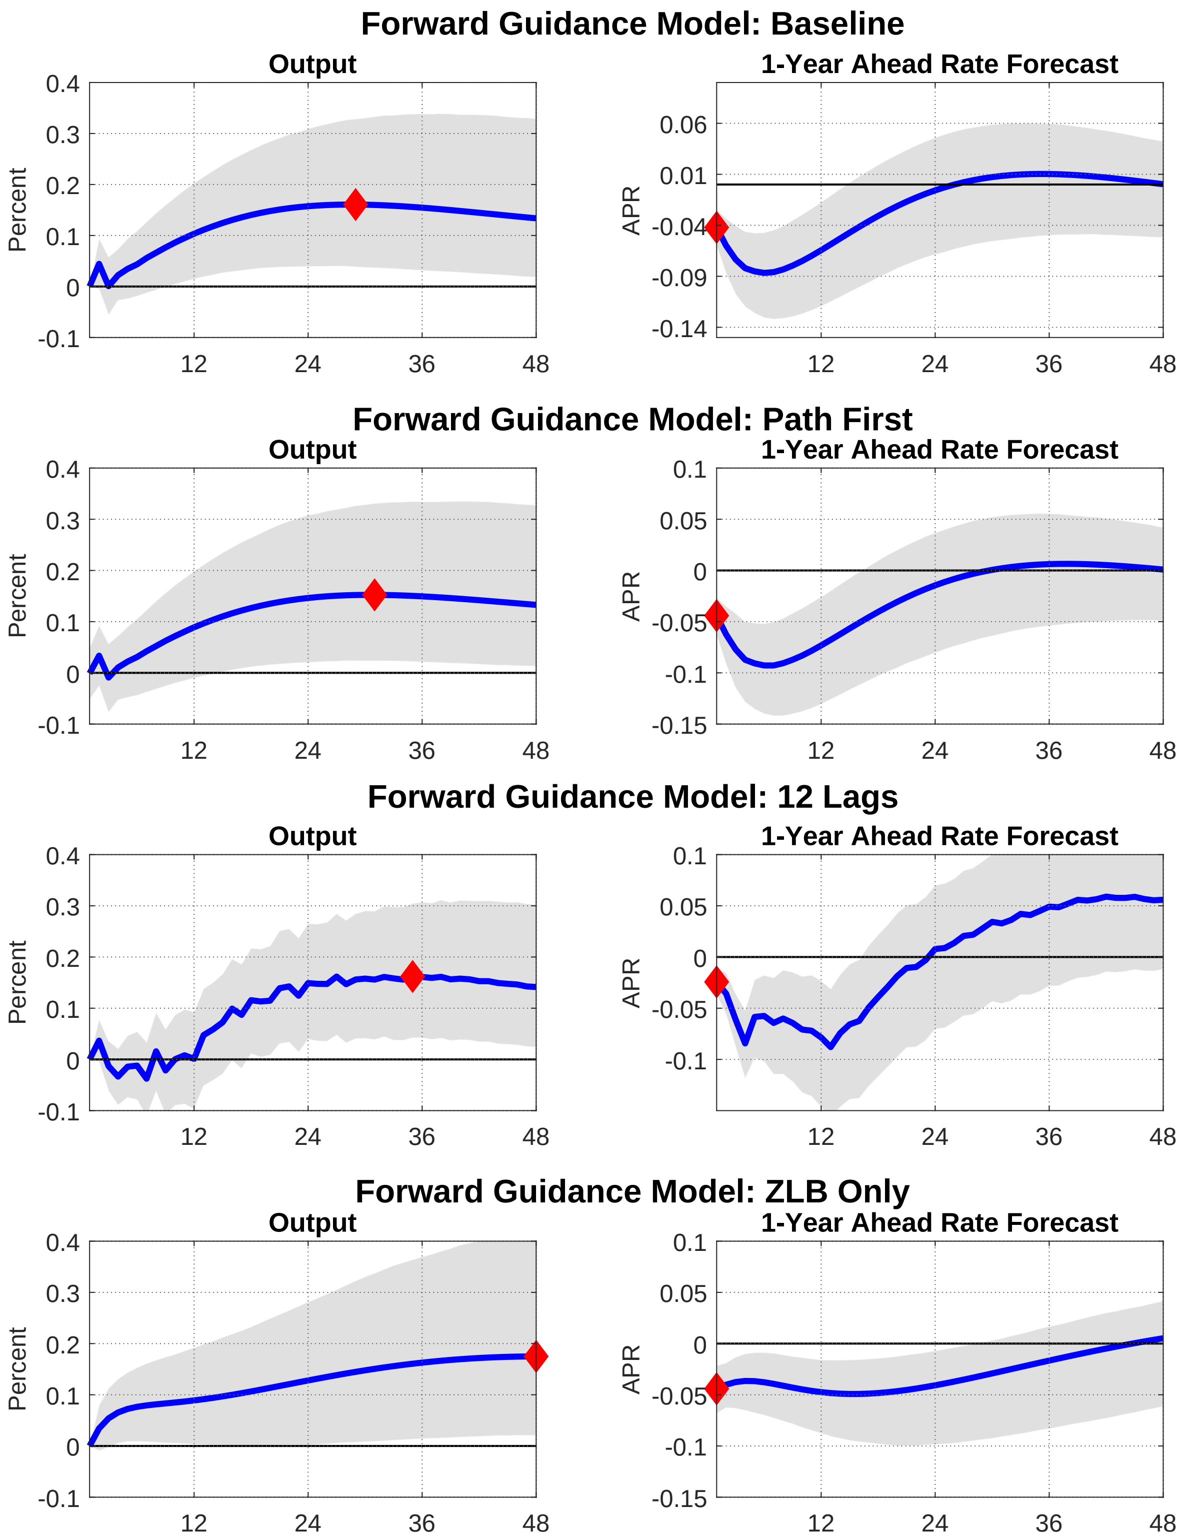

Note: The solid blue lines denote the empirical point estimate to a 1 standard deviation forward guidance shock and the shaded areas denote the 90% probability interval. For the output responses, the red diamonds denote the peak response and for the rate forecast the red diamonds denote the decline in 1-year ahead expected interest rates. Details for each VAR specification are provided in the main text.

Figure 3: VAR Impulse Responses to Forward Guidance & Conventional Policy Shocks

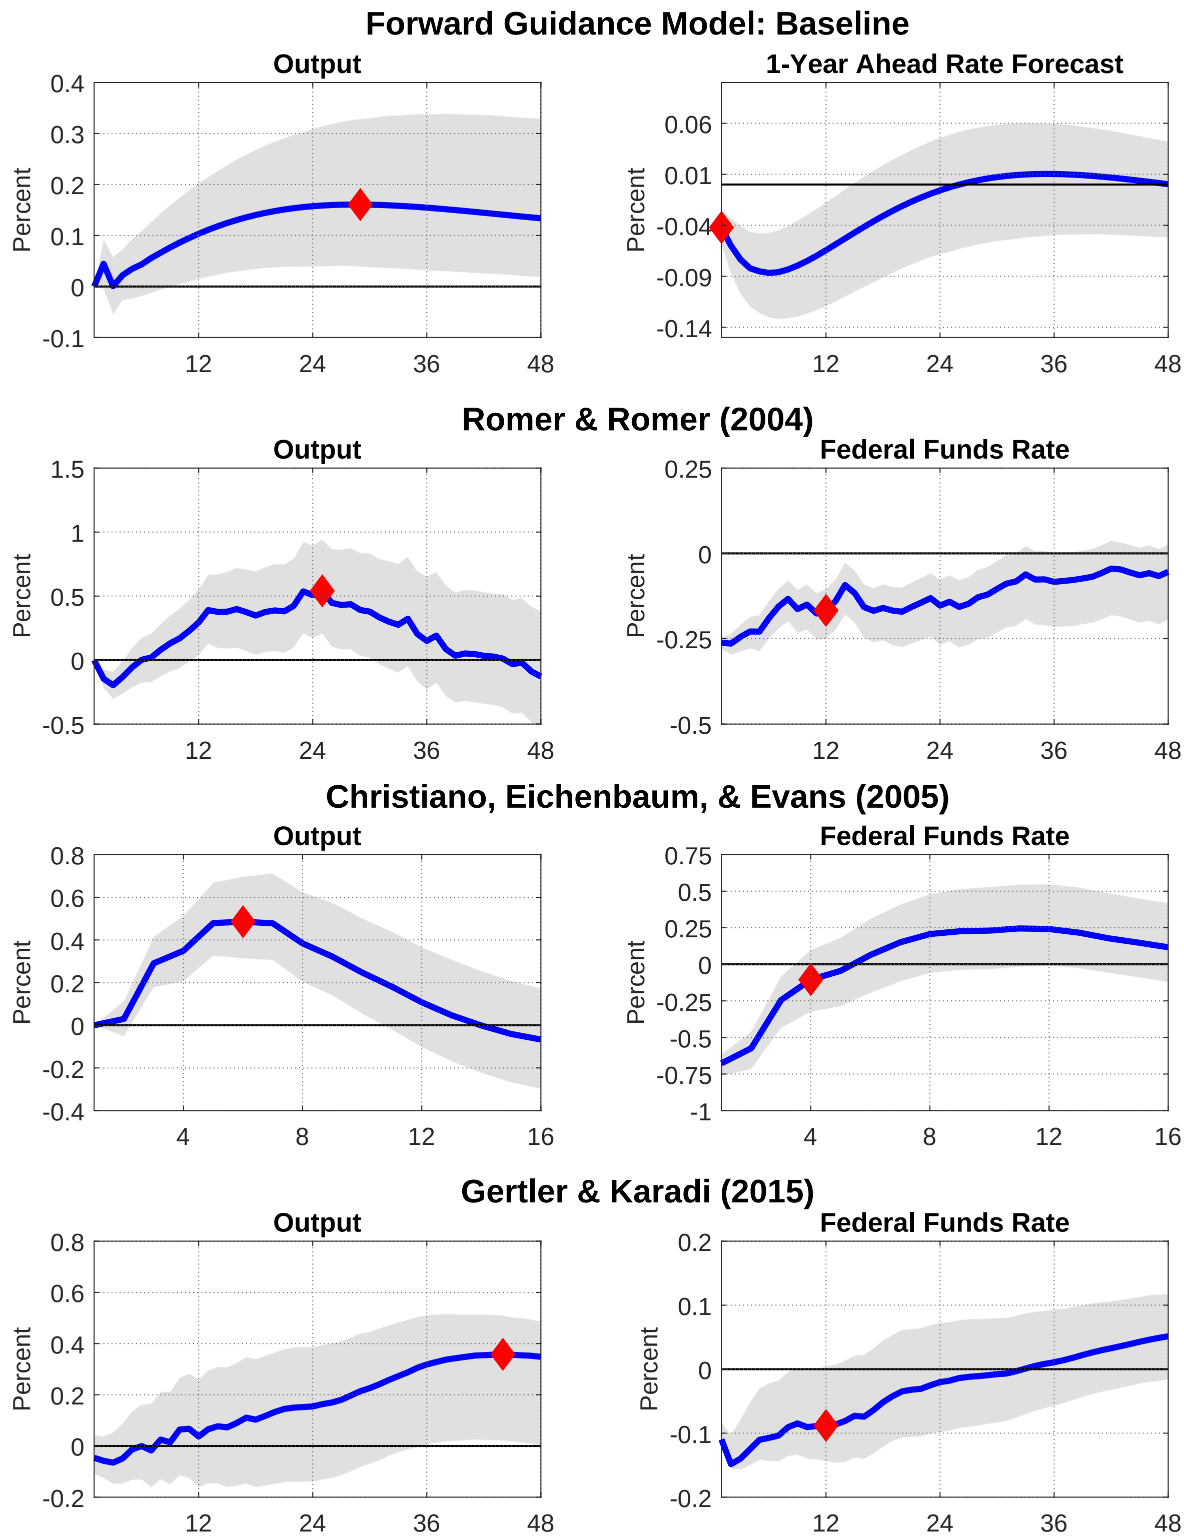

Note: The solid blue lines denote the empirical point estimate to a 1 standard deviation forward guidance or conventional monetary policy shock and the shaded areas denote the 90% probability interval. For the output responses, the red diamonds denote the peak response and for the rate forecast the red diamonds denote the decline in 1-year ahead expected interest rates. Details for each VAR specification are provided in the main text.

Figure 4: VAR Impulse Responses to a Forward Guidance Shock: Alternative Elasticity Calculation

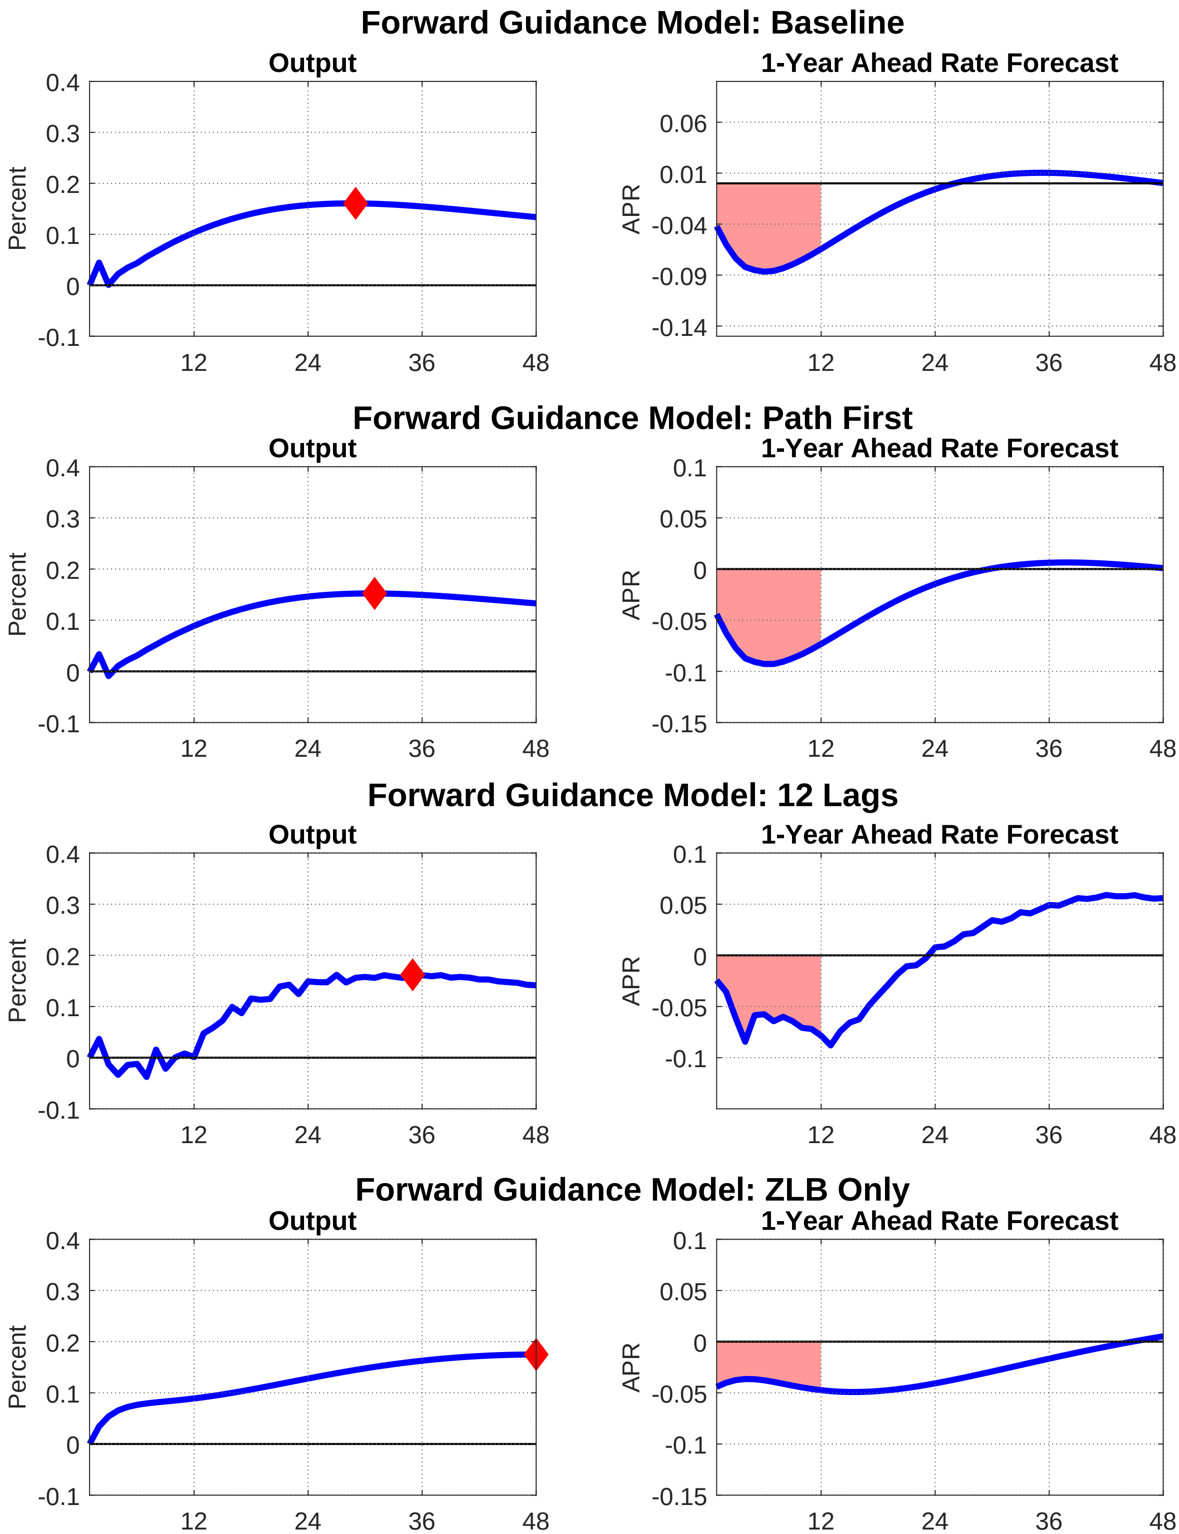

Note: The solid blue lines denote the empirical point estimate to a 1 standard deviation forward guidance shock. For the output responses, the red diamonds denote the peak response and for the rate forecast the red shaded area denotes the cumulative decline in expected rates over 2 years. Details for each VAR specification are provided in the main text.

Figure 5: VAR Impulse Responses to Forward Guidance & Conventional Policy Shocks: Robustness to Alternative Elasticity Calculation

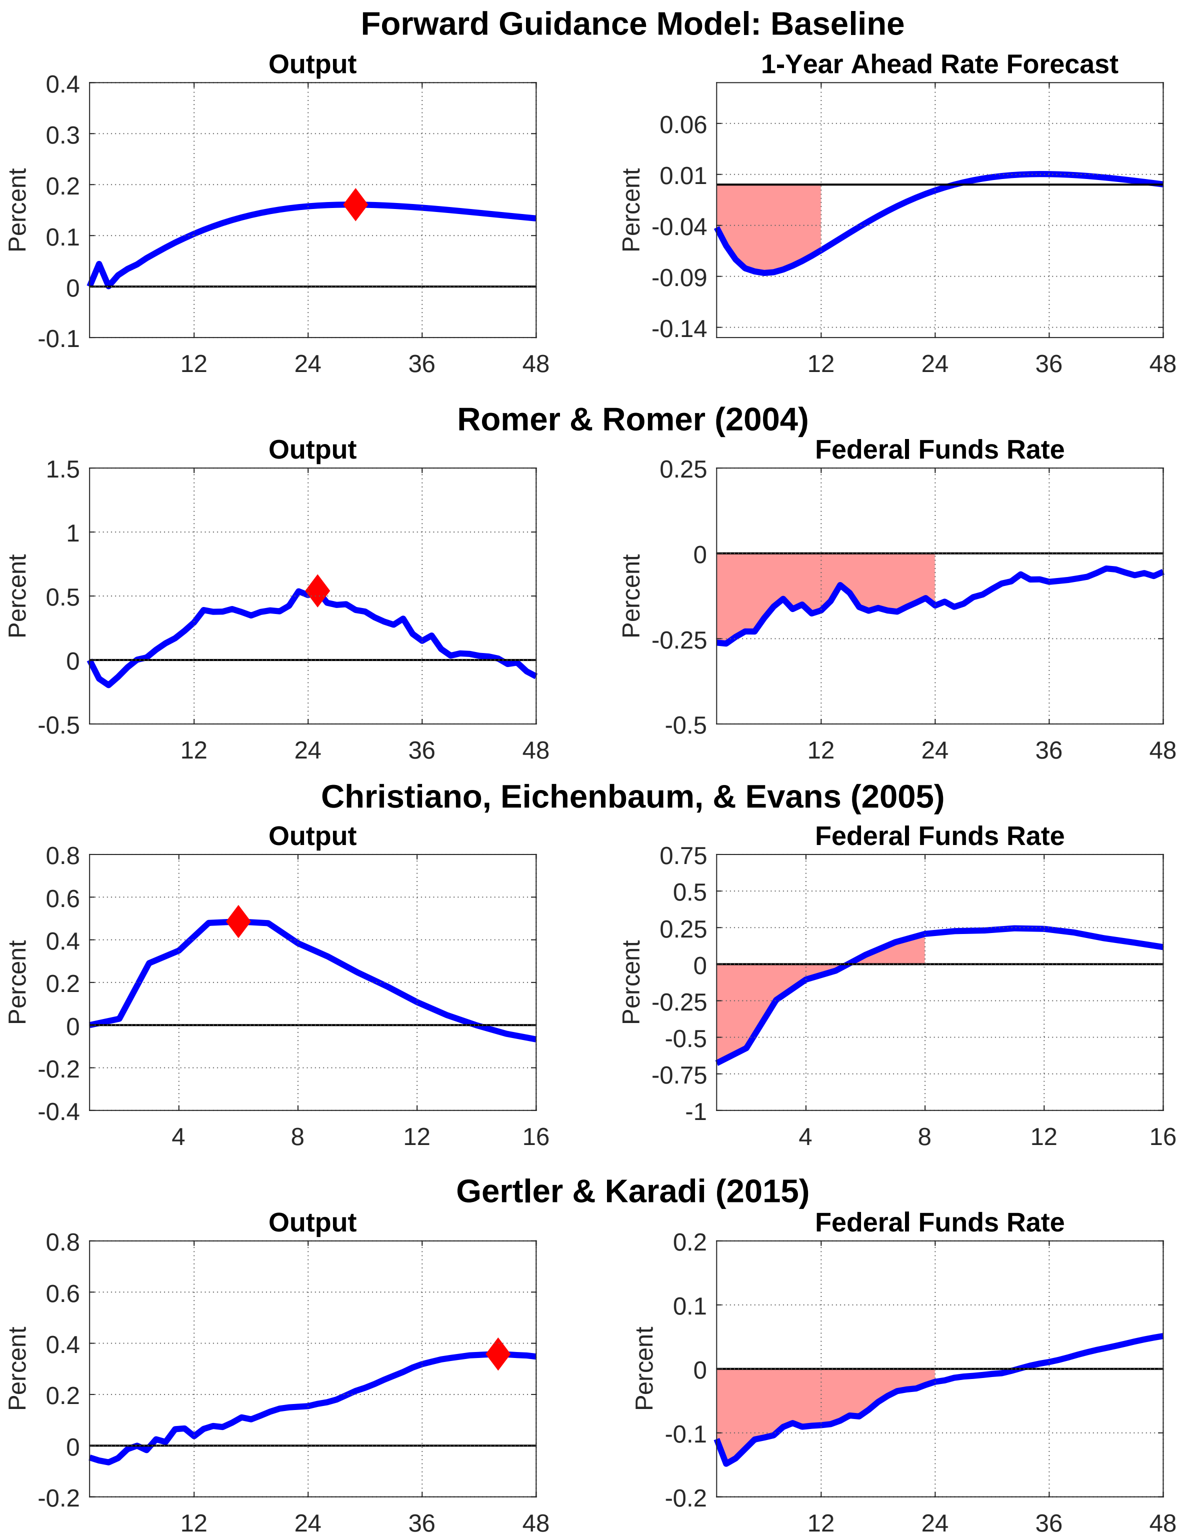

Note: The solid blue lines denote the empirical point estimate to a 1 standard deviation forward guidance or conventional monetary policy shock. For the output responses, the red diamonds denote the peak response and for the rate forecast the red shaded area denotes the cumulative decline in expected rates over 2 years. Details for each VAR specification are provided in the main text.

Figure 6: Implied Impulse Responses to Forward Guidance Shock: VAR and DSGE Model

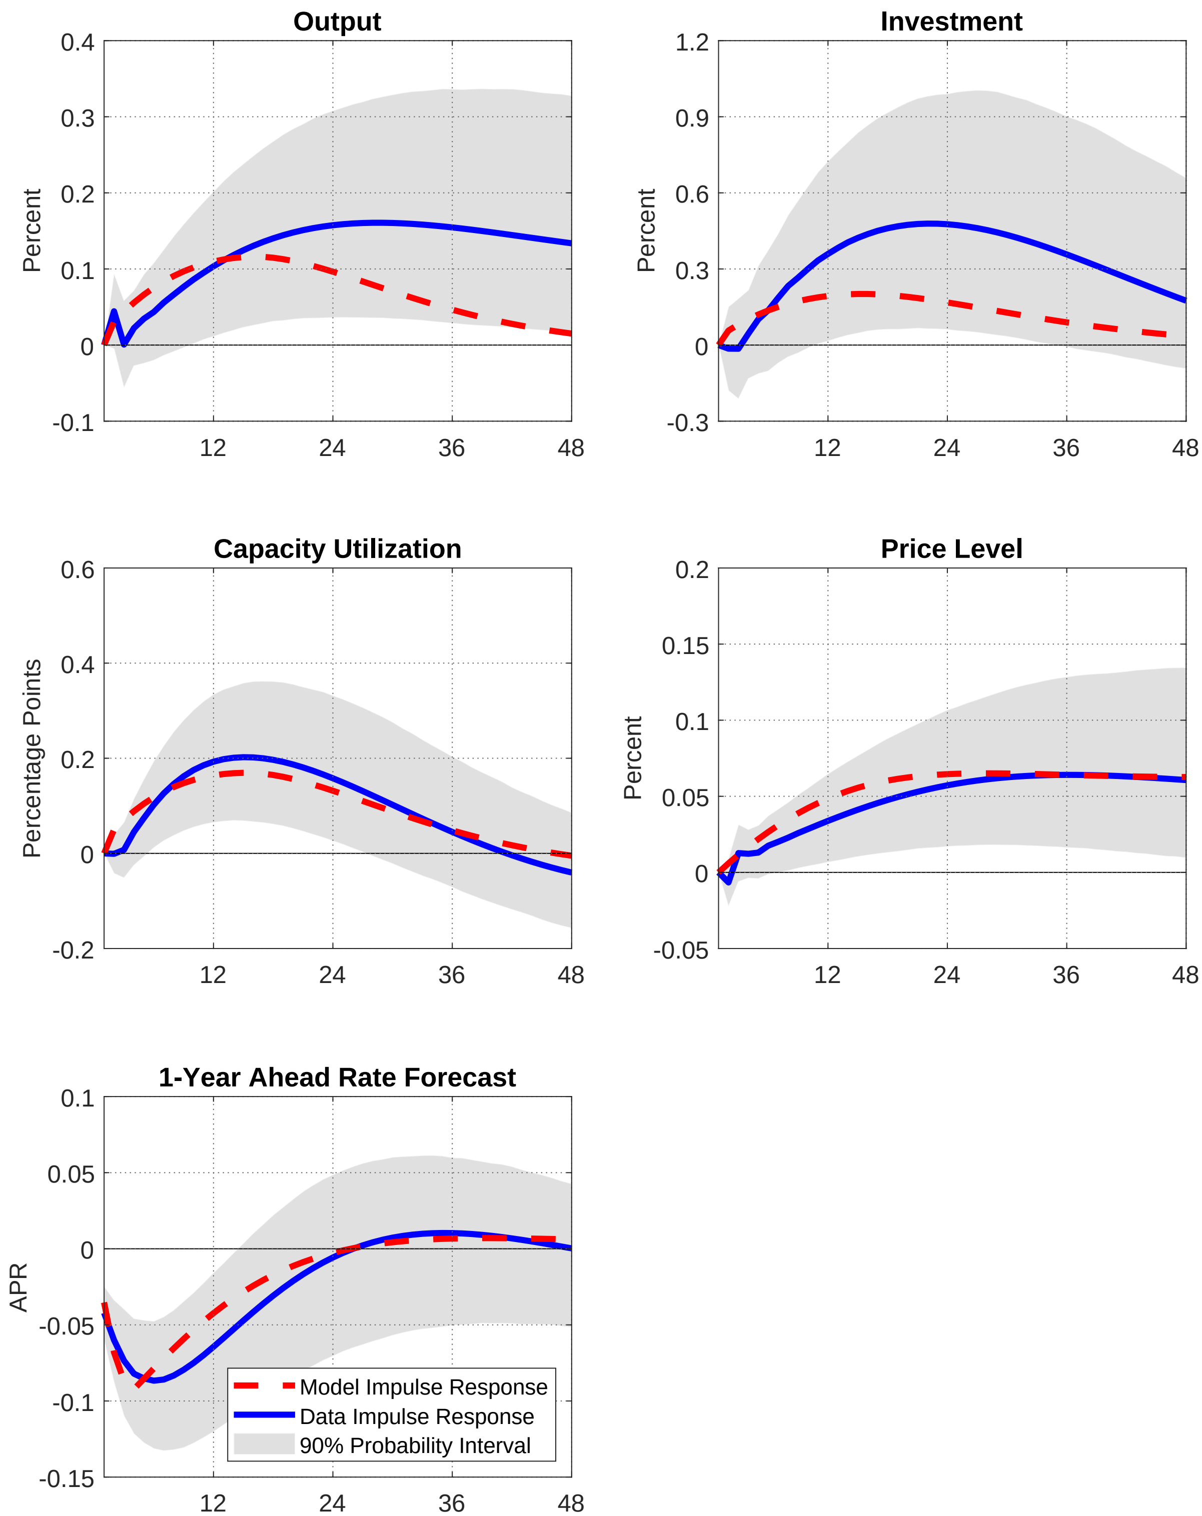

Note: The solid blue lines denote the point estimate to a 1 standard deviation forward guidance shock and the shaded areas denote the 90% probability interval of the posterior distribution from our baseline VAR model of forward guidance shocks, as reported in Figure 1. The red dashed lines denote the impulse responses from the estimated DSGE model.

Figure 7: Estimated Output Elasticity With Respect to 1-Year Ahead Expected Rates

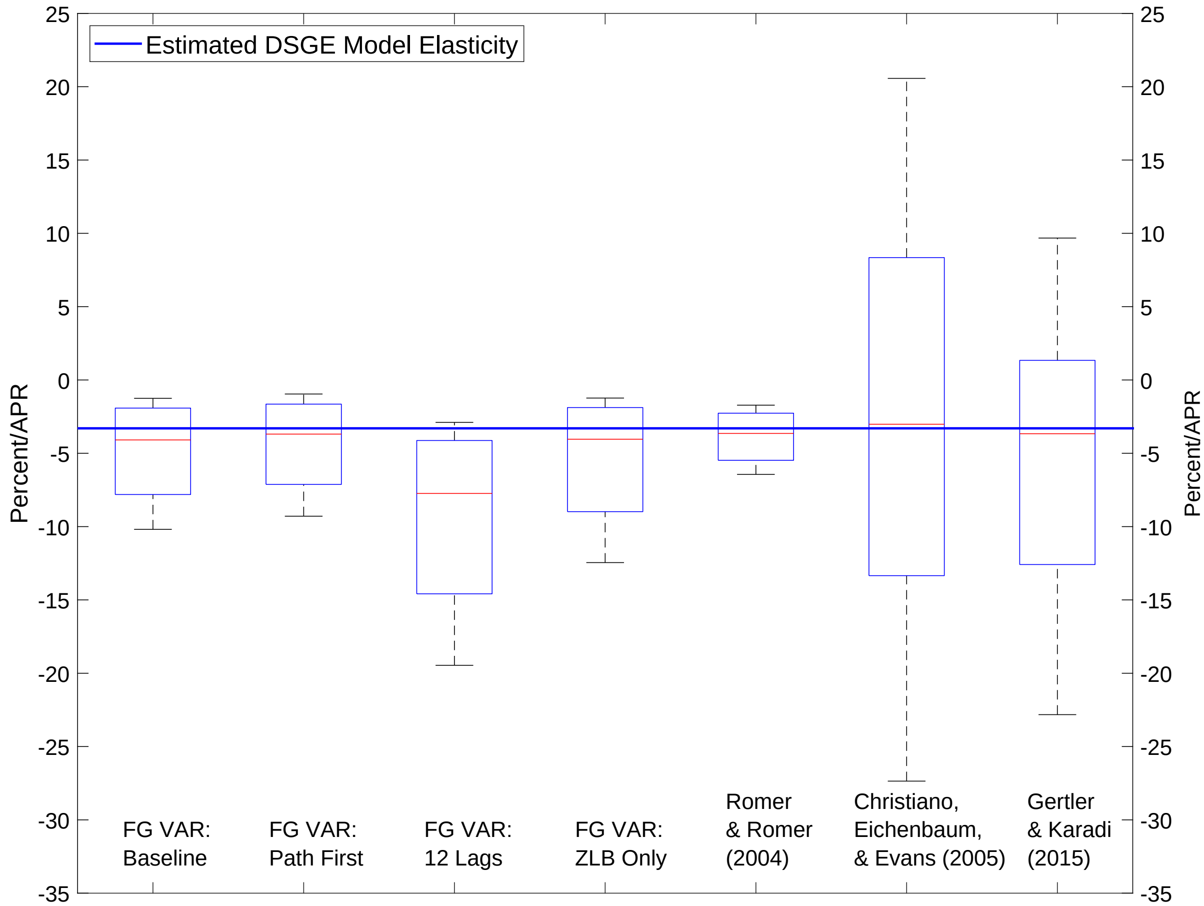

Note: This figure reports the estimated output elasticity with respect to 1-year ahead expected rates across VAR models as well as in the estimated DSGE model. The lines in the boxes denote the median estimate, the boxes denote the 68% probability intervals, and the whiskers denote the 90% probability intervals. The solid blue line denotes the elasticity in the estimated DSGE model.

Figure 8: Estimated Output Elasticity With Respect to the Change in Expected Rates Over the Next 2 Years

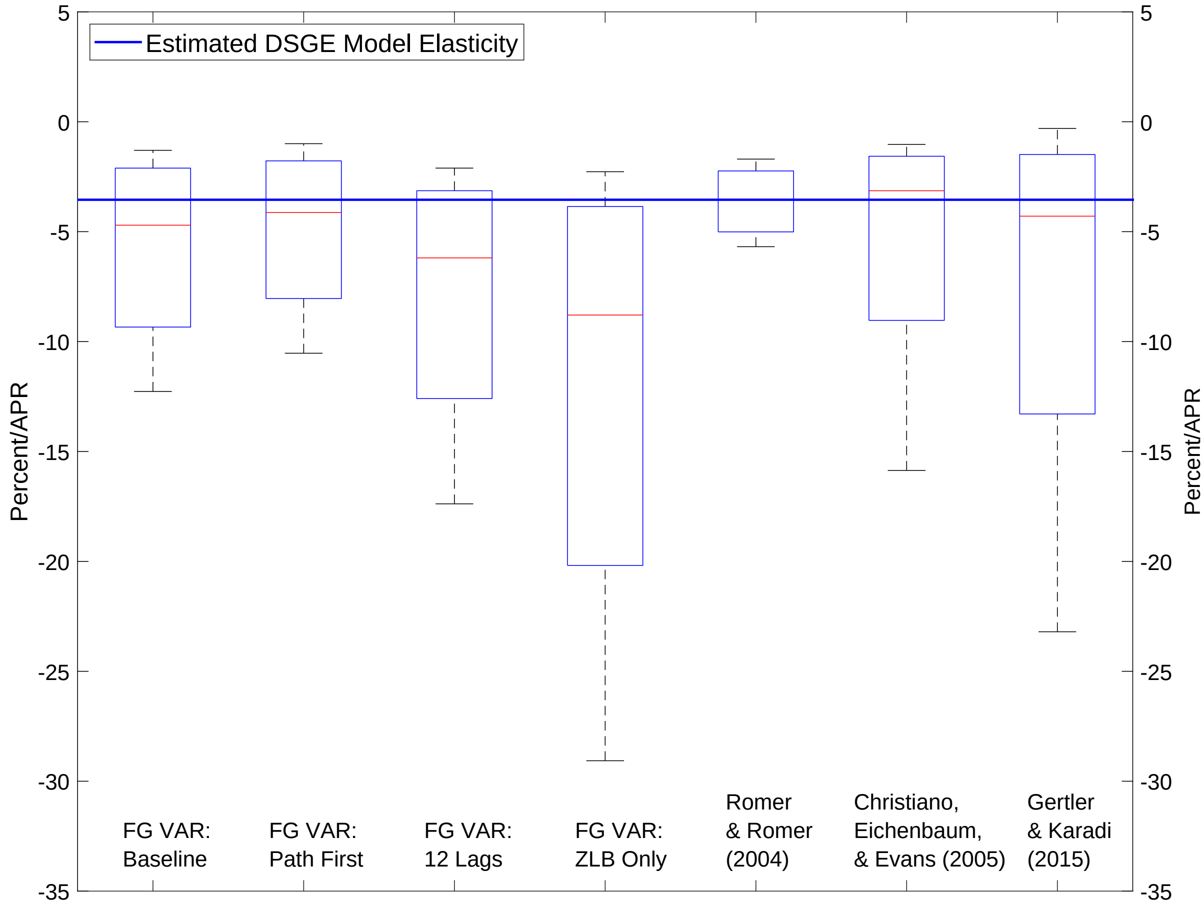

Note: This figure reports the estimated output elasticity with respect to the change in expected rates over the next 2 years across VAR models as well as in the estimated DSGE model. The lines in the boxes denote the median estimate, the boxes denote the 68% probability intervals, and the whiskers denote the 90% probability intervals. The solid blue line denotes the elasticity in the estimated DSGE model.

Figure 9: Responses to 1-Year, 3-Year, & 5-Year Ahead Real Rate Shocks

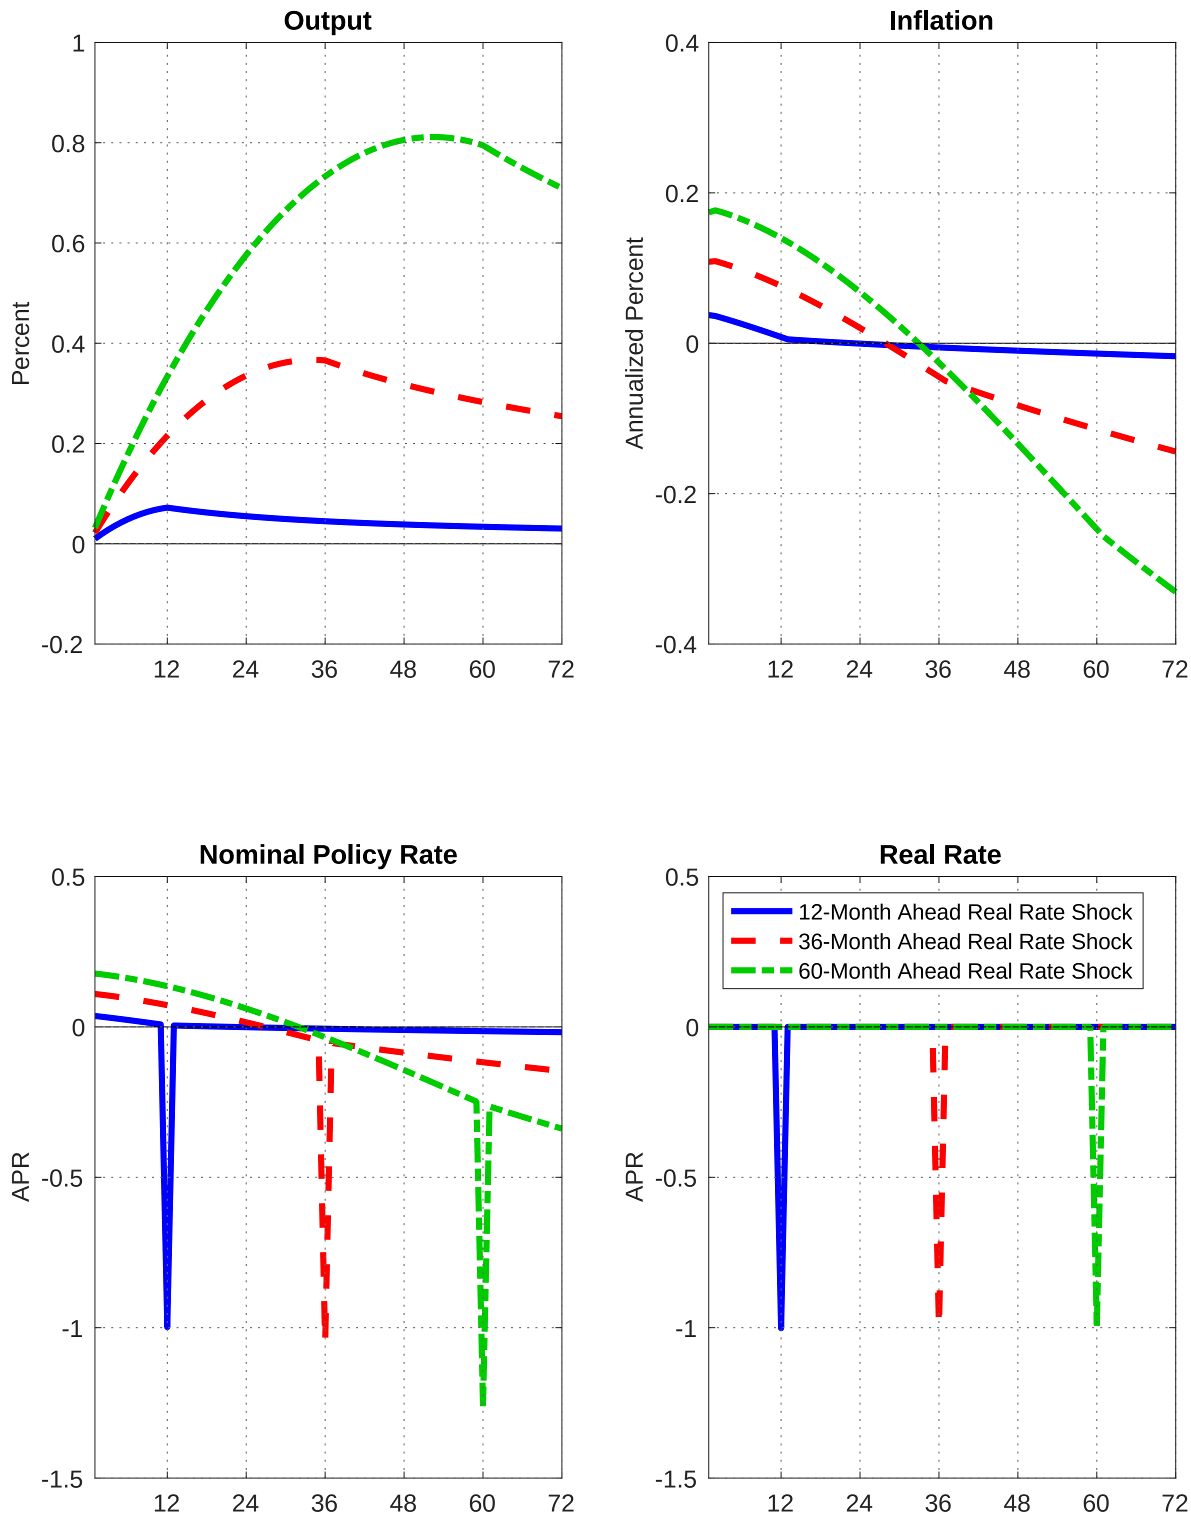

Note: This figure plots the responses of our estimated monthly-frequency DSGE model to 1-period innovations to the real rate at various horizons in the future, holding all other real interest rates fixed. See Section 5.1 for more details.