

A New Approach to Integrating Expectations into VAR Models Online Appendix*

Taeyoung Doh [†] A. Lee Smith[‡]

February 2022

*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, One Memorial Drive, Kansas City, MO 64198. Phone: (816) 881-2780, Email: taeyoung.doh@kc.frb.org.

[‡]Federal Reserve Bank of Kansas City, One Memorial Drive, Kansas City, MO 64198. Phone: (816) 881-2294, Email: andrew.smith@kc.frb.org.

A The Derivation of the Maximum Entropy Prior

In this section we detail the derivation of the maximum entropy distribution for our forecast consistent prior. Let π_0 be the Lebesgue measure on \mathbb{R} . Under the weak consistency, the prior $\pi(g)$ must satisfy the following two moment restrictions.

$$\int g\pi(g)dg = 0, \quad \int g^2\pi(g)dg = (\lambda W)^{-1}. \quad (1)$$

Maximizing the entropy of $\pi(g)$ under moment restrictions with respect to the reference measures π_0 is equivalent to minimizing the Kullbeck-Leibler distance between π and π_0 as follows:

$$\pi^*(g) = \operatorname{argmin}_{\pi(g)} \int \pi(g) \ln\left(\frac{\pi(g)}{\pi_0(g)}\right) dg - \mu_1 \left(\int g\pi(g)dg\right) - \mu_2 \left(\int g^2\pi(g) - (\lambda W)^{-1} dg\right). \quad (2)$$

The first-order condition for this is $\pi^*(g) = \pi_0(g)e^{1+\mu_1g+\mu_2g^2}$. Hence, $\pi^* \propto e^{\mu_1g+\mu_2g^2}$ and π^* is the normal distribution whose mean and variance are equal to $-\frac{\mu_1}{2\mu_2}$ and $-\frac{1}{2\mu_2}$. Since π^* must satisfy the above two moment restrictions, $\mu_1 = 0$ and $\mu_2 = -\frac{\lambda W}{2}$.

$$\pi^*(g) \propto e^{-\frac{\lambda W g^2}{2}} \rightarrow g \sim \mathcal{N}(0, (\lambda W)^{-1}). \quad (3)$$

B Full Impulse Responses for Alternative Specifications of the Forward Guidance SVAR Model

In Section 3.5 of the main text we present alternative specifications of our forward guidance VAR model. However, to conserve space, in Figure 6 of the main text we only present impulse responses for select variables for each alternative specification. In Figures B.1, B.2, and B.3, we present the full set of impulse responses for each alternative specification. For the ‘‘Full Sample’’ specification, we calibrate λ according to the marginal likelihood criterion which selects $\lambda = 8.03 \times 10^8$. For the ‘‘High-Frequency Calibration of Lambda’’ specification, we select λ to maximize the correlation between the SVAR forward guidance shocks and the high-frequency forward guidance shocks from Swanson (2021). We translate his daily forward guidance shocks to a monthly frequency using the time-aggregation procedure outlined in Gertler and Karadi (2015). This procedure selects $\lambda = 6.24 \times 10^8$, a value very near our baseline marginal likelihood calibration of $\lambda = 6.36 \times 10^8$. See the main text for further details of each alternative SVAR specification.

C Monte Carlo Simulations: Interpreting the Forecast Consistent Prior

This section provides further details on our Monte Carlo simulations, shown in Section 3.6 of the main text. We present detailed descriptions of the DSGE model and the calibration. We also present results from several variant of our baseline model.

C.1 Monte Carlo Experiment Design

Before describing the DSGE model in more detail, we provide an overview of how we conduct these Monte Carlo simulations.

1. Specify a structural New-Keynesian DSGE model in which the central bank’s interest-rate rule is subject to τ -period ahead news/forward guidance shocks, denoted by $\epsilon_{\tau,t-\tau}^{fg}$. These shocks are announced in period t but don’t directly affect the policy rate until period $t + \tau$. We also include in the model a preference shock, a supply shock, and a noise shock in the measurement of survey forecasts (the number of shocks equals the number of VAR variables). The DSGE model is fully specified below.
2. Simulate a time series of 50,000 observations of the vector $y_t = [x_t, \pi_t, i_t, \mathbb{E}_t^S i_{t+\tau}]'$, where x_t is output, π_t is inflation, i_t is the nominal policy rate, and $\mathbb{E}_t^S i_{t+\tau}$ is the τ -period (survey) forecast of the nominal policy rate.
3. Estimate a reduced form “population” VAR using the time series $\{y_{T-l}\}_{l=0}^{50,000}$ with the number of lags selected by the AIC.
4. Identify forward guidance shocks from this reduced form population VAR model, denoted by $\epsilon_t^{fg}(\lambda)$, by imposing the sign restrictions and varying degrees of tightness on the forecast consistent prior (the same identifying restrictions imposed in the empirical forward guidance application in the manuscript):

Sign Restrictions: $i_{t+\tau}^S$ increases and π_t falls for the first 6 periods following a forward guidance shock.

Forecast Consistent Prior: VAR expectations and survey expectations are (possibly) subject to some degree of forecast consistency, governed by the value of λ , with larger values of λ leading to a tighter prior over forecast consistency and $\lambda = 0$ reverting to the pure sign restrictions approach.

5. Calculate the correlation between the the SVAR-identified forward guidance shocks and each of the four underlying shocks from the DSGE model for alternative degrees of forecast consistency (alternative values of λ). Then scale each pairwise correlation by the norm of the four correlations.

C.2 Baseline DSGE Model: Noisy Survey Forecasts

Our baseline DSGE model used as the DGP for our Monte Carlo analysis is, at its core, a three equation New-Keynesian model. However, we model “survey” forecasts as being subject to information rigidities of the form documented in Coibion and Gorodnichenko (2012). For completeness, the model is characterized by the following equations:

$$x_t = \mathbb{E}_t x_{t+1} - (i_t - \mathbb{E}_t \pi_{t+1} - (a_t - \mathbb{E}_t(a_{t+1}))) \quad (4)$$

$$\pi_t = \kappa x_t + \beta \mathbb{E}_t \pi_{t+1} + \mu_t \quad (5)$$

$$i_t = \phi_i i_{t-1} + (1 - \phi_i) (\phi_x x_t + \phi_{\Delta x} (x_t - x_{t-1}) + \phi_\pi \pi_t) + \sigma_{fg} \epsilon_{\tau, t-\tau}^{fg} + \sum_{j=0}^{\tau-1} \sigma_j \epsilon_{j, t-j}^j \quad (6)$$

We set $\tau = 1$ for our baseline DGP and explore higher values of τ for robustness in this appendix. We calibrate the structural model parameters using the post-1990 estimates from Del Negro et al. (2020), a sample period which aligns well with our empirical application of identifying forward guidance shocks in the manuscript, and set $\kappa = 0.002$ and $\phi_i = 0.84$, $\phi_x = 0.22$, $\phi_{\Delta x} = 0.18$, and $\phi_\pi = 1.42$. The shock processes mirror those specified in models estimated by Ireland (2011) and therefore we rely on his estimates for $\rho_a = 0.98$, $\sigma_a = 0.08$, $\rho_\mu = 0$, and $\sigma_\mu = 0.002$. Ireland (2004a,b) estimates that σ_{mp} is near 0.0025 but has no forward guidance shocks in his models. Given the findings in Gurkaynak, Sack and Swanson (2005) that, since 1994, the dominant source of monetary policy innovations arise from “path shocks,” or surprises in the expected future path of policy rates, rather than “target shocks,” or surprise changes in the current federal funds rate, we set $\sigma_{fg} = 0.0025$, leaving the variance of monetary policy shocks essentially unchanged from Ireland (2004a,b). Finally, we set $\sigma_0 = 0.25 * \sigma_{fg}$ and which is sufficient to ensure that the current policy rate remains unchanged in response to a forward guidance announcement.

We define the τ -period ahead forecast for the nominal policy rate by $q_t = \mathbb{E}_t^S i_{t+\tau}$:

$$q_t = (1 - \rho) \mathbb{E}_t(i_{t+\tau}) + \rho q_{t-1} + \sigma_n \varepsilon_t^n, \quad (7)$$

where $\mathbb{E}_t(i_{t+\tau})$ is the rationale expectations forecast for the τ -period ahead nominal interest rate, $0 \leq \rho \leq 1$, and $\sigma_n \geq 0$. Therefore, unless $\rho = 0$ and $\sigma_n = 0$, survey forecasts are a noisy and imperfect measure of the rational expectations 12-period ahead forecast for interest rates. We set $\rho = 0.86$ so that, as documented in Coibion and Gorodnichenko (2012), survey forecasts only gradually incorporate new information. And we set $\sigma_n = 4 * \sigma_{fg}$ so the “noise” in the survey forecasts is sufficiently large relative the “signal” from the shock of interest. All of the innovations ϵ^a , ϵ_t^μ , ϵ_t^n , and $\epsilon_{\tau,t}^{fg}$ are i.i.d., mean zero, and unit variance shocks.

C.3 Expanded Description of Monte Carlo Analysis

Due to the specification of the survey forecast in equation (7), the survey forecast systematically differs from the VAR-implied forecast (and DSGE rational expectations). To be explicit, unconditional forecast consistency between VAR and survey forecasts does not hold in this model. To illustrate this point, Figure C.1 reports the (unconditional) root-mean squared error (RMSE) between the VAR and survey forecast of interest rates 1-period ahead. Even with a large number of lags in the VAR, a discrepancy persists between the survey and VAR interest rate forecasts. This feature of the model is critical as it enables the forecast consistent prior to shape structural shock identification even with a large number of lags in the VAR. For example, if in contrast $\mathbb{E}_t^S i_{t+\tau} = \mathbb{E}_t(i_{t+\tau})$, then, with a sufficient number of lags in the VAR, we find that survey and VAR-implied forecasts unconditionally align, rendering the forecast consistency restrictions following a forward guidance shock inconsequential.

In order to assess the ability of the forecast consistent prior to shape the identification of forward guidance shocks, we report the pairwise correlations between the SVAR-identified forward guidance shock and each of the four DSGE model shocks. We then scale these pairwise correlations by the norm of the four correlations so that the shock weights sum to unity in an L^2 sense. While not the same as the weights in Wolf (2020), these correlation-based weights are closely related to the regression coefficients from regressing the SVAR-identified forward guidance shock on the four (standardized) DSGE model shocks and therefore can be compared to give a sense of the relative relationship between our SVAR-identified forward guidance shock and each of the DSGE model shocks. Figure C.2 shows that as λ increases —and therefore the tightness of the forecast consistency restrictions increases —the SVAR-identified forward guidance shock places more weight on the true forward guidance shock. Moreover, as λ increases, the weight that the SVAR-identified shocks places on the DSGE noise and supply shocks is diminished. Intuitively, forecast consistency restrictions can deliver better alignment between the SVAR-identified and true forward guidance shocks

because forward guidance shocks in the DSGE model imply a greater degree of forecast consistency than does the linear combination of noise and other shocks that could masquerade as forward guidance shocks based on sign-restrictions alone. For example, we can construct a linear combination of the four DSGE shocks which places a relatively low weight on the true forward guidance shocks but nevertheless satisfies the sign restrictions.¹ Refer to this as the “masquerading” forward guidance shock. This masquerading shock produces a cumulative forecast discrepancy of about 19 basis points whereas the cumulative forecast discrepancy following the true forward guidance shock is just 3 basis points.²

The confusion of the sign-restricted SVAR model can be further illustrated by studying the DSGE and SVAR impulse responses to a forward guidance shock. Figure C.3 of this appendix shows the true impulse responses to a forward guidance shock from the DSGE model along with impulse responses estimated by the SVAR models with both sign restrictions and sign restrictions combined with forecast consistency restrictions. The responses of survey forecasts and inflation appears to largely reflect the weights this model places on noise and supply shocks to satisfy the sign restrictions. The resulting response of realized policy rates is revealing that the SVAR-implied expected path of rates is largely disconnected from the survey expectations. The result is an underestimation of the true output effects from forward guidance. In contrast, combining sign restrictions with forecast consistency restrictions results in SVAR impulse responses which closely mirror the true response of output. This close match is achieved by better aligning the realized paths of rates with survey expectations, thereby down-weighting linear combinations of noise and supply shocks which satisfy the sign restrictions. Therefore, imposing forecast consistency in the SVAR helps to distinguish forward guidance shocks from linear combinations of other shocks that could masquerade as a forward guidance shock based on sign-restrictions alone (as in Wolf, 2020).

C.4 Point Identification With the Forecast Consistent Prior

As we note in the main text, SVAR models identified by sign-restrictions on impulse responses only identify the model parameters up to a set, as discussed by Moon and Schorfheide (2012) and Uhlig (2017). This characteristic of sign restrictions has exposed this literature to multiple criticisms, especially related to conducting inference on this posterior set and interpreting

¹In particular, we use the results with $\lambda = 0$ from Figure C.2 to form a linear combination of the DSGE forward guidance, demand, supply, and noise shocks with weights (0.32, -0.50, -0.50, 0.63). This shock represents the “median” identified SVAR forward guidance shock using solely sign restrictions.

²To sharpen this intuition, below we consider the extreme case that, in the DSGE model, forward guidance shocks imply *perfect* forecasts consistency whereas other shocks do not.

median responses (See, among others, for a discussion Fry and Pagan, 2011; Baumeister and Hamilton, 2015; Kilian and Lütkepohl, 2017; Baumeister and Hamilton, 2018).

However, when using the forecast consistent prior along with sign restrictions for shock identification, achieving point identification is quite feasible. Intuitively, this is true because the forecast consistent prior induces a non-arbitrary distribution over the posterior set of SVAR models. Moreover, in practice, the “best draw” (the draw of the rotation matrices that comes the closest to satisfying forecast consistency) is a reasonable starting point for selecting one particular structural VAR model among the set of models which satisfy the sign restrictions. However, we can move beyond the “best draw” by minimizing the forecast discrepancy over the set of rotation matrices which satisfy the sign restrictions.

Table C.1 extends our Monte Carlo analysis by comparing three alternative approaches to identifying forward guidance shocks in a structural VAR in terms of their implied degree of (conditional) forecast consistency, as well as their relationship with the structural shocks from the DSGE model. The first column shows that a pure sign-restrictions approach results in a low degree of forecast consistency and a weak (inverse) relationship between the SVAR-implied and true forward guidance shocks. Instead, the SVAR-implied forward guidance shock is highly correlated with the DSGE noise shocks which, when combined with the model’s supply and demand shocks, can satisfy the sign restrictions implied by a forward guidance shock despite failing to correlate with the true forward guidance shock. The second column shows that augmenting sign restrictions with a tightly imposed forecast consistent prior (a large value of λ), by construction, reduces the discrepancy between VAR-implied and survey forecasts of interest rates. Moreover, the forecast consistent prior delivers SVAR-identified forward guidance shocks which are much less correlated with the noise shocks and, instead, share a large positive correlation with the true forward guidance shocks. Finally, the third row shows results from a point-identification strategy that chooses a single rotation matrix to minimize the degree of forecast discrepancy among the set of orthonormal matrices that satisfy the sign restrictions. By construction, numerically minimizing the forecast discrepancy results in forward guidance shocks with a greater degree of forecast consistency. The point-identified model also delivers SVAR-identified forward guidance shocks which are better linked with the true forward guidance shocks. However, relative to imposing the forecast consistent prior, the improvements are somewhat modest along both dimensions (forecast discrepancy and shock correlation). Nevertheless, achieving point identification is one way in which our approach addresses a common criticism of sign-restricted SVAR models.

C.5 Robustness of Monte Carlo Results to Alternative DGPs

We now examine the robustness of our Monte Carlo results to three alternative DGP specifications. The first two robustness checks examine the durability of our findings to alternative forward guidance shock strictures. First, we relax the assumption that current policy rates remain unchanged following a forward guidance announcement by assuming that $\sum_{j=0}^{\tau-1} \sigma_j \epsilon_{j,t-j}^j = 0$ so that, mechanically following the policy rule, the central bank partially offsets its own guidance. The second robustness check we consider sets $\tau = 4$ which, in this quarterly DSGE model, aligns better with the one-year ahead forward guidance shocks we study in our empirical application. In addition to setting $\tau = 4$, we continue to assume that $\sum_{j=0}^{\tau-1} \sigma_j \epsilon_{j,t-j}^j = 0$. Panels A and B in Table C.2 of this appendix shows the Monte Carlo results from these alternative DGPs.

In both settings, the forecast consistency restrictions meaningfully improve the ability of the SVAR to recover the true forward guidance results. The sign-restricted SVAR continues to confound forward guidance shocks with a combination of noise and supply shocks. In contrast, the combination of sign and forecasts consistency restrictions down-weights noise and supply shocks given the forecast discrepancies they generate and, instead, places greater weight on the true forward guidance shocks. However, the magnitude of improvement is somewhat diminished in both settings relative to our baseline Monte Carlo analysis as the forecast consistent SVAR somewhat struggles to distinguish forward guidance shocks from demand shocks. For some intuition, the impulse responses in Figure C.4 illustrate that the promise of future policy accommodation, either 1 or 4 periods ahead, leads to increases in inflation and output, inciting an immediate, mechanical tightening of policy in the present period through period $t + \tau - 1$. This mechanism is absent in our baseline DGP by the calibration of $\sum_{j=0}^{\tau-1} \sigma_j \epsilon_{j,t-j}^j$. However, when $\sum_{j=0}^{\tau-1} \sigma_j \epsilon_{j,t-j}^j = 0$, interest rates initially move in the opposite direction of output and inflation following a forward guidance shock, leading the SVAR to initially believe that a demand shock has occurred. As described in Wolf (2020, pg.27-28), once the forward guidance is realized in $(t + \tau)$, SVAR understands that a forward guidance shock had occurred. This is true for our forecast-consistent SVAR model too. In Table C.2, we also report the correlation-based the SVAR forward guidance shock places on the true forward guidance shock that was announced τ periods ago. This weight is significantly increased by forecast consistency restrictions. Reassuringly, Figure C.4 in this response suggests that this timing challenge does not prevent the forecast-consistent SVAR from largely recovering the true peak effects from forward guidance. Instead, it merely inhibits the estimated timing of the peak output response.

Finally, for the third robustness check, we adjust the information rigidity that accompanies forward guidance shocks. The baseline DSGE model used as the DGP for our Monte Carlo analysis assumes that forward guidance shocks, like other shocks, generate some degree of forecast disagreement. However, one could argue that forward guidance may be unique from other shocks in that it can better align interest rate forecasts. One counter argument however might be that this alternative model is inconsistent with the evidence in Coibion and Gorodnichenko (2012) who show that professional survey forecast errors display inertia in response to essentially all structural shocks. For this reason, we present results solely from the baseline model in the paper and present results from this alternative model in this appendix.

We execute the simulations from this alternative DGP using Dynare as follows. Within a single `.mod` file we create two identical model blocks. The parameters in the first model block are set to the values in our baseline model except we assume in this first model block that $\sigma_{fg} = 0$ such that there are no forward guidance shocks (however noise shocks remain). Then, in the second model block, we set $\sigma_{fg} = 0.0025$, its value in our baseline model, but we set $\sigma_a = \sigma_u = \sigma_n = 0$ such that there are only forward guidance shocks in the second model block. Moreover, in the second model block, we set $\rho = 0$ so that there are no information rigidities for the forward guidance shock and, therefore, conditional on a forward guidance shock, survey forecasts perfectly align with the rationale expectations forecast for interest rates. The econometrician observes the sum of output, inflation, nominal interest rates, and survey forecasts across the two blocks.³

Table C.2 shows results when using this alternative DGP. Comparing the two columns, the pure sign restrictions approach versus the combination of sign restrictions with the forecast consistent prior, Table C.2 reveals that we find a similar pattern of improved shock identification from the use of the forecast consistent prior in this alternative model as we found in our baseline model. The forecast consistent prior SVAR now delivers a forecast discrepancy of less than 1 basis point and, as one might expect given the DGP, very low weights on the noise and supply shocks.⁴

³Since all the structural shocks are orthogonal, the aggregate variables from the DSGE model are always equal to the sum of the variable conditional on the realizations of each individual structural shock. This logic would fail if the structural shocks were correlated.

⁴The fact that the weight on the true forward guidance shock is not exactly one in this alternative model could reflect the possibility that linear combinations of appropriately signed noise, supply, and demand shocks can deliver offsetting forecast discrepancies, allowing these masquerading shocks to still evade forecast consistency restrictions. Nevertheless, the pattern of improved identification suggests that forecast consistency restrictions greatly diminish the masquerading shocks problem in both DGP settings.

D Details of TVP-VAR Model

In Section 4 of the main text we present an application of our forecast consistent prior in a time-varying parameter vector autoregression (TVP-VAR) model. In this section, we provide further details of this model.

D.1 TVP-VAR Data

The specification of our TVP-VAR follows Clark and Davig (2011) closely by including long-term and near-term survey forecasts of inflation alongside realized inflation, a measure of real economic activity, and a measure of the policy rate. We use both 1 year and 10-year ahead forecasts for Consumer Price Index (CPI) inflation from the Survey of Professional Forecasters (SPF) as well as realized CPI inflation.⁵ We include the Chicago Fed National Activity Index to broadly measure real activity and the effective federal funds rate to account for the stance of monetary policy. Our formal estimation sample is 1982-2015 which includes the zero lower bound period. Therefore, to better account for the full spectrum of the FOMC’s policy actions from 2009-2015 we splice the federal funds rate together with Wu and Xia (2016) shadow federal funds rate.

We collect these five series together into a vector $y_t = [\pi_t^{S,L}, \pi_t^{S,S}, \pi_t, x_t, r_t]$ where $\pi_t^{S,L}$ denotes the long-term survey forecast of CPI inflation, $\pi_t^{S,S}$ denotes the 1 year ahead survey forecast of CPI inflation, π_t denotes the realized CPI inflation rate, measured as the quarter over quarter percent change at an annual rate, x_t denotes the CFNAI, and r_t denotes the short-term policy rate. Figure D.1 shows all the variables used in estimation. While the real economic activity measure does not show any trending behavior, all the other nominal variables exhibit downward trends since the early 1980s. The TVP-VAR model can accommodate time-varying trends in variables with the random-walk drifts of VAR coefficients.

⁵The 10-year ahead forecasts for CPI inflation from SPF are available beginning in 1991. Prior to 1991, we use long-run inflation forecasts obtained from the public release of the Federal Reserve Board of Governors’s FRB/SU econometric model which is constructed using alternative surveys and econometric estimates. We use realized inflation and inflation nowcasts to construct our inflation expectations measures to prevent overlap between long-term survey forecasts, near-term survey forecasts, and realized inflation.

D.2 TVP-VAR Model Details

We consider the following TVP-VAR(4) model with stochastic volatility for the five variables

$$y_t = [\pi_t^{S,L}, \pi_t^{S,S}, \pi_t, x_t, r_t]:$$

$$\begin{aligned}
y_t &= A_{D,t} + \sum_{j=1}^4 A_{j,t} y_{t-j} + u_t, \quad u_t \sim \mathcal{N}(0, B^{-1} \Sigma_{u,t} B^{-1'}), \\
\Sigma_{u,t} &= \begin{pmatrix} \sigma_{1,t}^2 & 0 & \cdots & 0 \\ 0 & \sigma_{2,t}^2 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \cdots & 0 & \sigma_{5,t}^2 \end{pmatrix}, \quad B = \begin{pmatrix} 1 & 0 & \cdots & 0 \\ B_{21} & 1 & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ B_{51} & \cdots & B_{54} & 1 \end{pmatrix}, \\
\tilde{y}_t &= \tilde{A}_{0,t} + \tilde{A}_{1,t} \tilde{y}_{t-1} + \tilde{u}_t, \\
\tilde{y}_t &= \begin{bmatrix} y_t \\ y_{t-1} \\ y_{t-2} \\ y_{t-3} \end{bmatrix}, \quad \mathbf{A}_{D,t} = \begin{bmatrix} A_{D,t} \\ 0_{(15 \times 5)} \end{bmatrix}, \quad \mathbf{A}_t = \begin{pmatrix} A_{1,t} & A_{2,t} & A_{3,t} & A_{4,t} \\ I_5 & 0_{(5 \times 5)} & 0_{(5 \times 5)} & 0_{(5 \times 5)} \\ 0_{(5 \times 5)} & I_5 & 0_{(5 \times 5)} & 0_{(5 \times 5)} \\ 0_{(5 \times 5)} & 0_{(5 \times 5)} & I_5 & 0_{(5 \times 5)} \end{pmatrix}, \quad \mathbf{u}_t = \begin{bmatrix} u_t \\ 0_{(15 \times 5)} \end{bmatrix}, \\
A_t &= [A_{D,t}, A_{1,t}, A_{2,t}, A_{3,t}, A_{4,t}]', \\
\text{vec}(A_t) &= \text{vec}(A_{t-1}) + \epsilon_t, \quad \epsilon_t \sim \mathcal{N}(0, Q), \\
\ln(\sigma_{i,t}^2) &= \ln(\sigma_{i,t-1}^2) + e_{i,t}, \quad e_{i,t} \sim \mathcal{N}(0, \sigma_{e,i}^2).
\end{aligned} \tag{8}$$

Since the forecast horizon of the 10-year forecast from the SPF changes only at the first quarter of each year, the number of quarterly forward inflation observations contained in the 10-year forecast varies depending on the quarter of year. Let $l(t)$ index the starting quarter for the long-horizon forward forecast where $l(t) = 5$ in the first quarter, $l(t) = 6$ in the second quarter, $l(t) = 7$ in the third quarter, and $l(t) = 8$ in the fourth quarter.⁶ Let π_t denote realized CPI quarterly annualized inflation and let $\pi_t^{e,L}$ and $\pi_t^{e,S}$ denote the long-term forward and short-term weighted averages of expected inflation at different horizons under

⁶In principle, we could treat four quarterly observations of the 10-year forecast of any given year as one quarterly observation for four different measures of long-term forecasts to keep the same forecast period for each measure of long-term forecasts. However, doing so would require us to estimate a mixed-frequency TVP-VAR which would substantially increase the dimension of parameters and latent variables. We leave this for future research.

the expectation operator E^e :

$$\begin{aligned}\pi_t^{e,L} &= \frac{\sum_{j=l(t)}^{40} E_t^e(\pi_{t+j})}{40 - l(t)} \\ \pi_t^{e,S} &= \frac{\sum_{j=1}^4 E_t^e(\pi_{t+j})}{4}\end{aligned}\tag{9}$$

where E_t^e is the survey expectation when $e = S$ and the VAR-based expectation when $e = VAR$.

Forecast consistency at both forecast horizons requires the following 42 cross-equation restrictions:

$$\begin{aligned}\pi_t^{S,L} - \pi_t^{VAR,L} &= e'_{\pi,S,L} \tilde{y}_t - e'_\pi \frac{[\sum_{h=5}^{40} \sum_{j=0}^{h-1} \mathbf{A}_t^j \mathbf{A}_{D,t} + \sum_{h=5}^{40} \mathbf{A}_t^h \tilde{y}_t]}{36}, \\ g(A_t)_L &= [-e'_\pi \frac{[\sum_{h=5}^{40} \sum_{j=0}^{h-1} \mathbf{A}_t^j \mathbf{A}_{D,t}]}{36}, e'_{\pi,L} - e'_\pi \frac{[\sum_{h=5}^{40} \mathbf{A}_t^h]}{36}]', \\ \pi_t^{S,S} - \pi_t^{VAR,S} &= e'_{\pi,S,S} \tilde{y}_t - e'_\pi \frac{[\sum_{h=1}^4 \sum_{j=0}^{h-1} \mathbf{A}_t^j \mathbf{A}_{D,t} + \sum_{h=1}^4 \mathbf{A}_t^h \tilde{y}_t]}{4}, \\ g(A_t)_S &= [-e'_\pi \frac{[\sum_{h=1}^4 \sum_{j=0}^{h-1} \mathbf{A}_t^j \mathbf{A}_{D,t}]}{4}, e'_{\pi,S} - e'_\pi \frac{[\sum_{h=1}^4 \mathbf{A}_t^h]}{4}]', \\ g(A_t) &= [g(A_t)_L, g(A_t)_S]'\end{aligned}\tag{10}$$

where e_i is a selection vector whose i th element is 1 while all the other elements are zeros. In the above calculations of the VAR-implied forecasts, we assume no future parameter drift:

$$E_t^{VAR}(\mathbf{A}_{D,t+h} | \mathcal{F}_t) = \mathbf{A}_{D,t} \qquad E_t^{VAR}(\prod_{k=1}^h \mathbf{A}_{t+k} | \mathcal{F}_t) = \mathbf{A}_t^h.$$

This ‘‘anticipated utility’’ approximation works well for mean forecasts. Since we impose consistency requirements only on point forecasts from surveys, using the approximation should be even less problematic.

D.3 Priors and Posterior Simulation

To determine priors for initial states ($A_{D,0}, A_{1,0}, A_{2,0}, A_{3,0}, A_{4,0}, \Sigma_{u,0}$) and parameters ($Q, B, \sigma_\varepsilon^2$), we use a training sample from 1970:Q2 to 1981:Q2. Initial values of time-varying coefficients (A_0) and the covariance matrix of innovations to time-varying coefficients (Q) are calibrated from the Ordinary Least Squares (OLS) estimates of a time-invariant VAR(4) using the training sample data. The prior mean of A_0 are the OLS estimates of coefficients from the

VAR(4) for the training sample data. The prior mean of Q is set to be proportional to the covariance matrix of the OLS estimates. As is common in the TVP-VAR literature, we assume that the initial values of latent states and other parameters are a-priori independent.

To impose forecast consistency requirements, we need to calibrate hyperparameters such as λ and W that determine the maximum entropy prior for $g(A_t)$. Since the prior for $g(A_t)$ follows a normal distribution, W needs to match the inverse of the covariance matrix of $g(A_t)$. We simulate multiple draws of A_t from the prior distribution and compute $g(A_t)$ for each draw. From these multiple realizations of $g(A_t)$, we compute the covariance matrix and set W^{-1} equal to it. In order to calibrate the hyperparameter λ , as is common in a hierarchical model, we estimate the TVP-VAR(4) model with different values of λ and pick the value that gives the best fit for the sample according to the marginal likelihood.

We obtain the posterior output of latent states and parameters using Gibbs sampling. While the joint posterior distribution of $(A^T, \Sigma_u^T, Q, B, \sigma_e^2)$ is difficult to characterize analytically, the distribution of one component conditional on all the other components is either analytically tractable or easy to simulate.⁷ Given hyperparameters (λ, W) , the joint posterior distribution of latent states and parameters can be obtained as the product of conditional posterior distributions as follows:

$$\begin{aligned}
p_{\lambda, W}(A^T, \Sigma_u^T, B, Q, \sigma_e^2 | y^T) &\propto p(Q | A^T, \Sigma_u^T, B, \sigma_e^2, y^T) p(Q) \\
&\times p(\sigma_e^2 | A^T, \Sigma_u^T, B, Q, y^T) p(\sigma_e^2) \\
&\times p(\Sigma_u^T | A^T, B, Q, \sigma_e^2, y^T) p(\Sigma_{u,0}) \\
&\times p(B | A^T, \Sigma_u^T, Q, \sigma_e^2, y^T) p(B) \\
&\times p_{\lambda, W} g(A^T | \Sigma_u^T, B, Q, \sigma_e^2, y^T) p(A_0).
\end{aligned} \tag{11}$$

In the last step, we first simulate M draws of A^T from $p_{\lambda, W} g(A^T | \Sigma_u^T, B, Q, \sigma_e^2, y^T) p(A_0)$ and then resample them using importance-sampling weights given by

$$w(A^T(j)) = \frac{\exp^{-0.5g(A^T(j))'(\lambda W)g(A^T(j))}}{\sum_{k=1}^M \exp^{-0.5g(A^T(k))'(\lambda W)g(A^T(k))}}. \tag{12}$$

To choose the hyperparameter λ that controls the tightness of forecast consistency prior restrictions, we calculate the following marginal data density for different values of λ . We calculate the marginal likelihood using the harmonic mean of the likelihood implied by posterior draws. The details are explained in the next section.

⁷We use superscript to denote an array of observations up to the point at the superscript. For example, A^T represents an array of $[A_0, \dots, A_T]$.

E The Calculation of the Marginal Likelihood

This section details the calculation of the marginal likelihood in both the TVP-VAR, as well as the forward guidance application.

E.1 Marginal Likelihood Calculation in TVP-VAR

Since draws of VAR parameters from prior distributions are likely to have very low likelihood, calculating the marginal likelihood from prior draws is quite inefficient and practically infeasible. We calculate the marginal likelihood by using importance sampling for posterior draws. Let's notice that the posterior density of VAR parameters can be expressed as the ratio of the posterior density kernel to the marginal likelihood.

$$\begin{aligned}
 p(A^T, \Sigma_u^T, B, Q, \sigma_e^2 | y^T, \lambda) &= \frac{p(y^T | A^T, \Sigma_u^T, B, Q, \sigma_e^2) p(A^T, \Sigma_u^T, B, Q, \sigma_e^2 | \lambda)}{\int p(y^T | A^T, \Sigma_u^T, B, Q, \sigma_e^2) p(A^T, \Sigma_u^T, B, Q, \sigma_e^2 | \lambda) d(A^T, \Sigma_u^T, B, Q, \sigma_e^2)}, \\
 &= \frac{p(y^T | A^T, \Sigma_u^T, B, Q, \sigma_e^2) p(A^T, \Sigma_u^T, B, Q, \sigma_e^2 | \lambda)}{p_\lambda(y^T)}, \\
 p(A^T, \Sigma_u^T, B, Q, \sigma_e^2 | \lambda) &= p(A^T, \Sigma_u^T, B, Q, \sigma_e^2) h(A^T, \Sigma_u^T, B, Q, \sigma_e^2 | \lambda), \\
 h(A^T, \Sigma_u^T, B, Q, \sigma_e^2 | \lambda) &= \frac{p(g(A^T, \Sigma_u^T, B, Q, \sigma_e^2) | \lambda)}{\int p(g(A^T, \Sigma_u^T, B, Q, \sigma_e^2) | \lambda) d(A^T, \Sigma_u^T, B, Q, \sigma_e^2)}.
 \end{aligned} \tag{13}$$

Here, h stands for the density for the weak forecast consistency prior with respect to VAR parameters. Obtaining the exact value of the marginal likelihood is challenging in this case because the normalizing constant for the forecast consistency prior is difficult to compute, even by simulation. But what we are interested in is the relative evaluation of the marginal likelihood across multiple values of λ . Therefore, we only need $\frac{p_{\lambda_1}(y^T)}{p_{\lambda_0}(y^T)}$, which is feasible to approximate by simulation. To do this calculation, first let's take the harmonic mean of the likelihood for posterior draws of VAR parameters.

$$\begin{aligned}
 &\int \frac{1}{p(y^T | A^T, \Sigma_u^T, B, Q, \sigma_e^2)} p(A^T, \Sigma_u^T, B, Q, \sigma_e^2 | \lambda, y^T) d(A^T, \Sigma_u^T, B, Q, \sigma_e^2) \\
 &= \frac{\int p(A^T, \Sigma_u^T, B, Q, \sigma_e^2 | \lambda) d(A^T, \Sigma_u^T, B, Q, \sigma_e^2)}{p_\lambda(y^T)}.
 \end{aligned} \tag{14}$$

Since the tail of the inverse of the likelihood can be thick, we truncate the region of posterior draws that we calculate the harmonic mean of the likelihood to between the 16%

quantile and the 84% quantile of VAR parameters in terms of the likelihood. Let's denote this region by χ . Then, we can compute the marginal likelihood as follows.

$$\begin{aligned}
p_\lambda(y^T) &= \int_\chi p(A^T, \Sigma_u^T, B, Q, \sigma_e^2 | \lambda) d(A^T, \Sigma_u^T, B, Q, \sigma_e^2) \\
&\times \left[\int_\chi \frac{1}{p(y^T | A^T, \Sigma_u^T, B, Q, \sigma_e^2)} p(A^T, \Sigma_u^T, B, Q, \sigma_e^2 | \lambda, y^T) d(A^T, \Sigma_u^T, B, Q, \sigma_e^2) \right]^{-1}, \quad (15) \\
&= P_\lambda(\text{prior}) \times L(\text{likelihood}).
\end{aligned}$$

Therefore, the posterior odds ratio is reduced to the multiplication of the prior odds ratios and the ratio of the harmonic mean of the likelihood. Notice that χ potentially depends on λ .

$$\begin{aligned}
\frac{p_{\lambda_1}(y^T)}{p_{\lambda_0}(y^T)} &= \frac{\int_{\chi(\lambda_1)} p(A^T, \Sigma_u^T, B, Q, \sigma_e^2) p(g(A^T, \Sigma_u^T, B, Q, \sigma_e^2) | \lambda_1) d(A^T, \Sigma_u^T, B, Q, \sigma_e^2)}{\int_{\chi(\lambda_0)} p(A^T, \Sigma_u^T, B, Q, \sigma_e^2) p(g(A^T, \Sigma_u^T, B, Q, \sigma_e^2) | \lambda_0) d(A^T, \Sigma_u^T, B, Q, \sigma_e^2)} \\
&\times \frac{\left[\int_{\chi(\lambda_1)} \frac{1}{p(y^T | A^T, \Sigma_u^T, B, Q, \sigma_e^2)} p(A^T, \Sigma_u^T, B, Q, \sigma_e^2 | \lambda_1, y^T) d(A^T, \Sigma_u^T, B, Q, \sigma_e^2) \right]^{-1}}{\left[\int_{\chi(\lambda_0)} \frac{1}{p(y^T | A^T, \Sigma_u^T, B, Q, \sigma_e^2)} p(A^T, \Sigma_u^T, B, Q, \sigma_e^2 | \lambda_0, y^T) d(A^T, \Sigma_u^T, B, Q, \sigma_e^2) \right]^{-1}}. \quad (16)
\end{aligned}$$

Unlike $h(A^T, \Sigma_u^T, B, Q, \sigma_e^2 | \lambda)$, we can easily compute $p(g(A^T, \Sigma_u^T, B, Q, \sigma_e^2) | \lambda)$. We can replace the integral by using simulated posterior draws. When we do not impose the forecast consistency prior (in other words, $\lambda = 0$), we set $h(A^T, \Sigma_u^T, B, Q, \sigma_e^2 | \lambda) = 1$.⁸

The log marginal likelihood in Figure E.1 for $\lambda \in [0, 1.5]$ suggests that marginal likelihood is maximized at $\lambda = 1.42$. We truncate the value of λ at 1.5 because the effective sample size becomes too small (less than 2 percent of posterior draws) above that value. The rapid increase of the marginal likelihood above a positive threshold of λ in Figure E.1 indicates that imposing a modest degree of forecast consistency restrictions improves the time series fit of the TVP-VAR model.

E.2 Marginal Likelihood Calculation in Forward Guidance Application

We can apply a similar idea to our forward guidance application. Here, we have two sets of reduced-form parameters: VAR coefficients (α) and the Cholesky factor of the covariance matrix (Σ) and the forecast consistency prior imposes restrictions on the impulse response up

⁸In practice, we used $p(g(A^T, \Sigma_u^T, B, Q, \sigma_e^2) | \lambda)^{\frac{1}{n}}$, which is a geometric average of the prior at each point of time. Considering the joint prior for each trajectory of time-varying parameters led to sample depletion in which only a few draws have non-negligible importance sampling weights.

to horizon H . On top of sign restrictions, the forecast consistency prior induces restrictions on (α, Σ, Q) where Q is an orthonormal matrix to identify structural shocks. As in the above example, we use the harmonic mean of the likelihood from posterior draws to calculate the posterior odds ratio as follows:

$$\begin{aligned} \frac{p_{\lambda_1}(y^T)}{p_{\lambda_0}(y^T)} &= \frac{\int_{\mathcal{X}(\lambda_1)} p_{\text{sign restriction}}(\alpha, \Sigma, Q) p(g(B, C, Q|H)|\lambda_1) d(\alpha, \Sigma, Q)}{\int_{\mathcal{X}(\lambda_0)} p_{\text{sign restriction}}(\alpha, \Sigma, Q) p(g(B, C, Q|H)|\lambda_0) d(\alpha, \Sigma, Q)} \\ &\times \frac{[\int_{\mathcal{X}(\lambda_1)} \frac{1}{p(y^T|\alpha, \Sigma)} p(\alpha, \Sigma, Q|\lambda_1, y^T, H) d(\alpha, \Sigma, Q)]^{-1}}{[\int_{\mathcal{X}(\lambda_0)} \frac{1}{p(y^T|\alpha, \Sigma)} p(\alpha, \Sigma, Q|\lambda_0, y^T, H) d(\alpha, \Sigma, Q)]^{-1}}. \end{aligned} \quad (17)$$

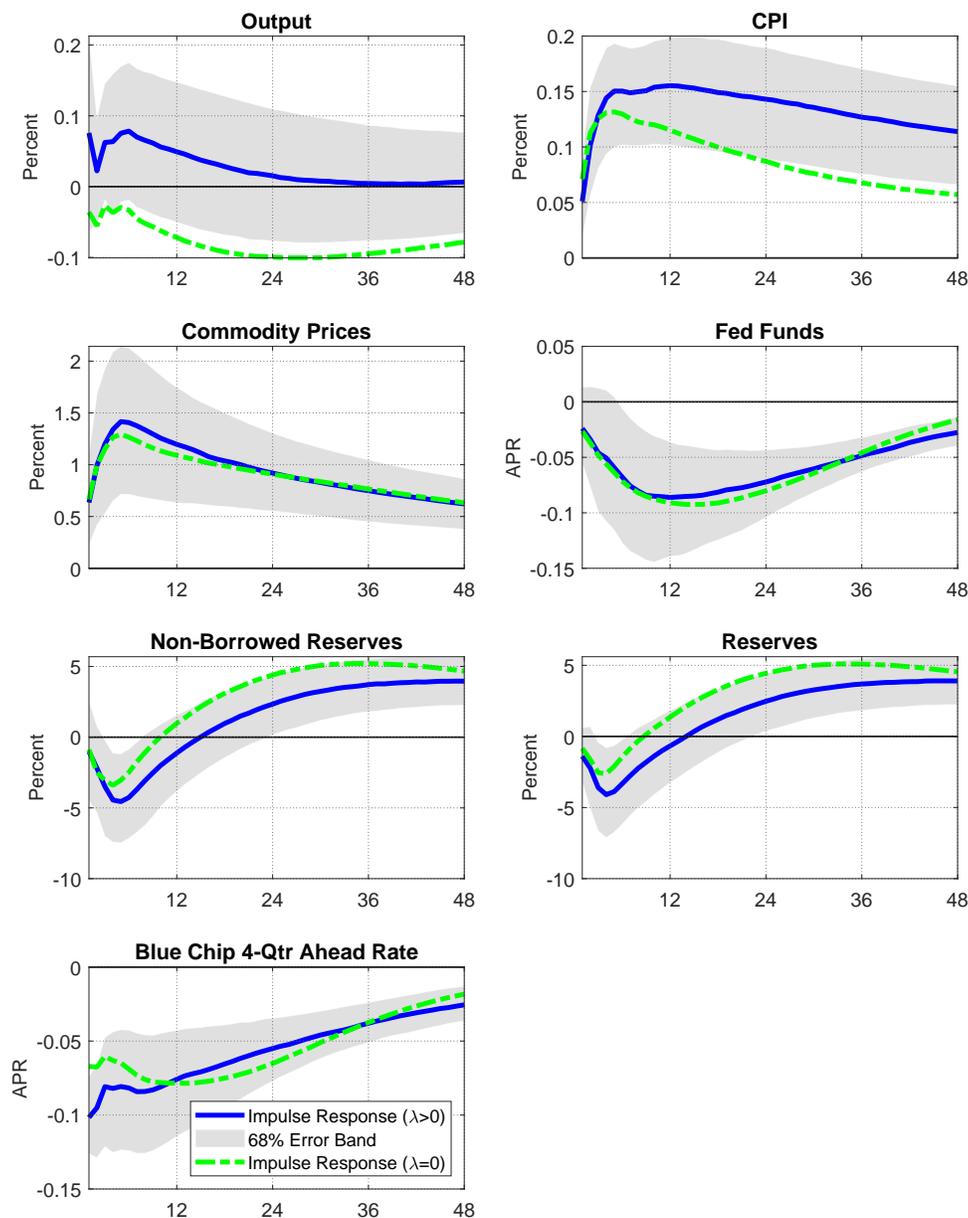
As the top chart in Figure E.2 shows, the marginal likelihood displays significant curvature across moderate values of λ and then becomes flat as λ further increases. While the number of unique posterior draws decrease in resampling as we increase λ , sample depletion does not appear to drive the flatness of the marginal likelihood because the number of resampled draws also stabilizes.

References

- Baumeister, Christiane, and James D Hamilton.** 2015. “Sign restrictions, structural vector autoregressions, and useful prior information.” *Econometrica*, 83(5): 1963–1999.
- Baumeister, Christiane, and James D Hamilton.** 2018. “Inference in structural vector autoregressions when the identifying assumptions are not fully believed: Re-evaluating the role of monetary policy in economic fluctuations.” *Journal of Monetary Economics*, 100: 48–65.
- Clark, Todd E, and Troy Davig.** 2011. “Decomposing the Declining Volatility of Long-term Inflation Expectations.” *Journal of Economic Dynamics and Control*, 35(7): 981–999.
- Coibion, Olivier, and Yuriy Gorodnichenko.** 2012. “What Can Survey Forecasts Tell Us about Information Rigidities?” *Journal of Political Economy*, 120(1): 116–159.
- Del Negro, Marco, Michele Lenza, Giorgio E Primiceri, and Andrea Tambalotti.** 2020. “What’s Up with the Phillips Curve?” *Brookings Papers on Economic Activity*.
- Fry, Renée, and Adrian Pagan.** 2011. “Sign Restrictions in Structural Vector Autoregressions: A Critical Review.” *Journal of Economic Literature*, 49(4): 938–60.
- Gertler, Mark, and Peter Karadi.** 2015. “Monetary Policy Surprises, Credit Costs, and Economic Activity.” *American Economic Journal: Macroeconomics*, 7(1): 44–76.
- Gurkaynak, Refet S, Brian Sack, and Eric T Swanson.** 2005. “Do Actions Speak Louder Than Words? The Response of Asset Prices to Monetary Policy Actions and Statements.” *International Journal of Central Banking*.
- Ireland, Peter N.** 2004a. “Money’s Role in the Monetary Business Cycle.” *Journal of Money, Credit, and Banking*, 36(6): 969–983.
- Ireland, Peter N.** 2004b. “Technology shocks in the New Keynesian Model.” *Review of Economics and Statistics*, 86(4): 923–936.
- Ireland, Peter N.** 2011. “A New Keynesian Perspective on the Great Recession.” *Journal of Money, Credit and Banking*, 43(1): 31–54.
- Kilian, Lutz, and Helmut Lütkepohl.** 2017. *Structural vector autoregressive analysis*. Cambridge University Press.

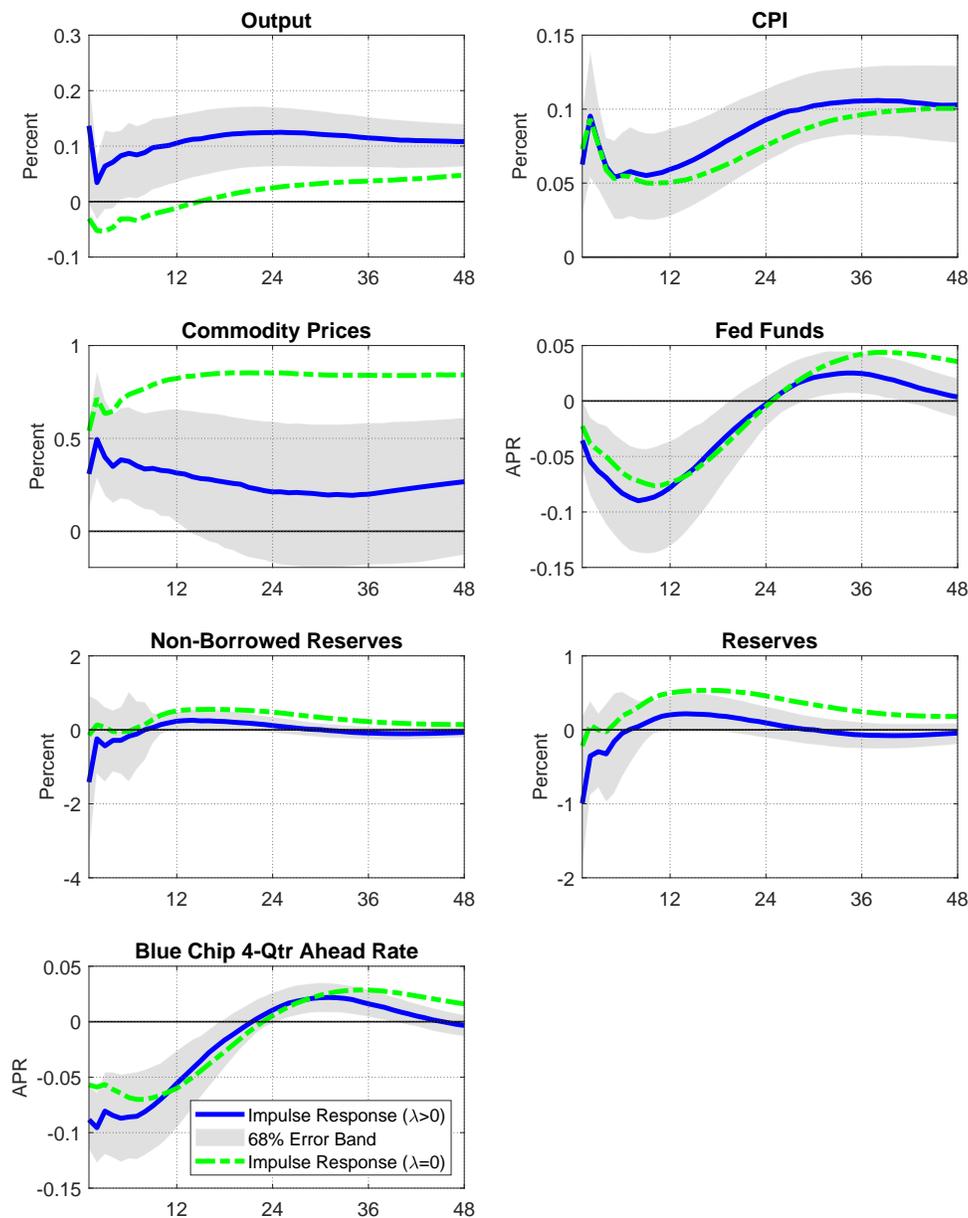
- Moon, Hyungsik Roger, and Frank Schorfheide.** 2012. “Bayesian and frequentist inference in partially identified models.” *Econometrica*, 80(2): 755–782.
- Swanson, Eric T.** 2021. “Measuring the effects of Federal Reserve forward guidance and asset purchases on financial markets.” *Journal of Monetary Economics*, 118: 32–53.
- Uhlig, Harald.** 2017. “Shocks, Sign Restrictions, and Identification.” Vol. 2, 95, Cambridge University Press.
- Wolf, Christian K.** 2020. “SVAR (Mis)identification and the Real Effects of Monetary Policy Shocks.” *American Economic Journal: Macroeconomics*, 12(4): 1–32.
- Wu, Jing Cynthia, and Fan Dora Xia.** 2016. “Measuring the macroeconomic impact of monetary policy at the zero lower bound.” *Journal of Money, Credit and Banking*, 48(2-3): 253–291.

Figure B.1: FORWARD GUIDANCE SHOCK ALTERNATIVE SVAR SPECIFICATION: FULL SAMPLE ESTIMATES



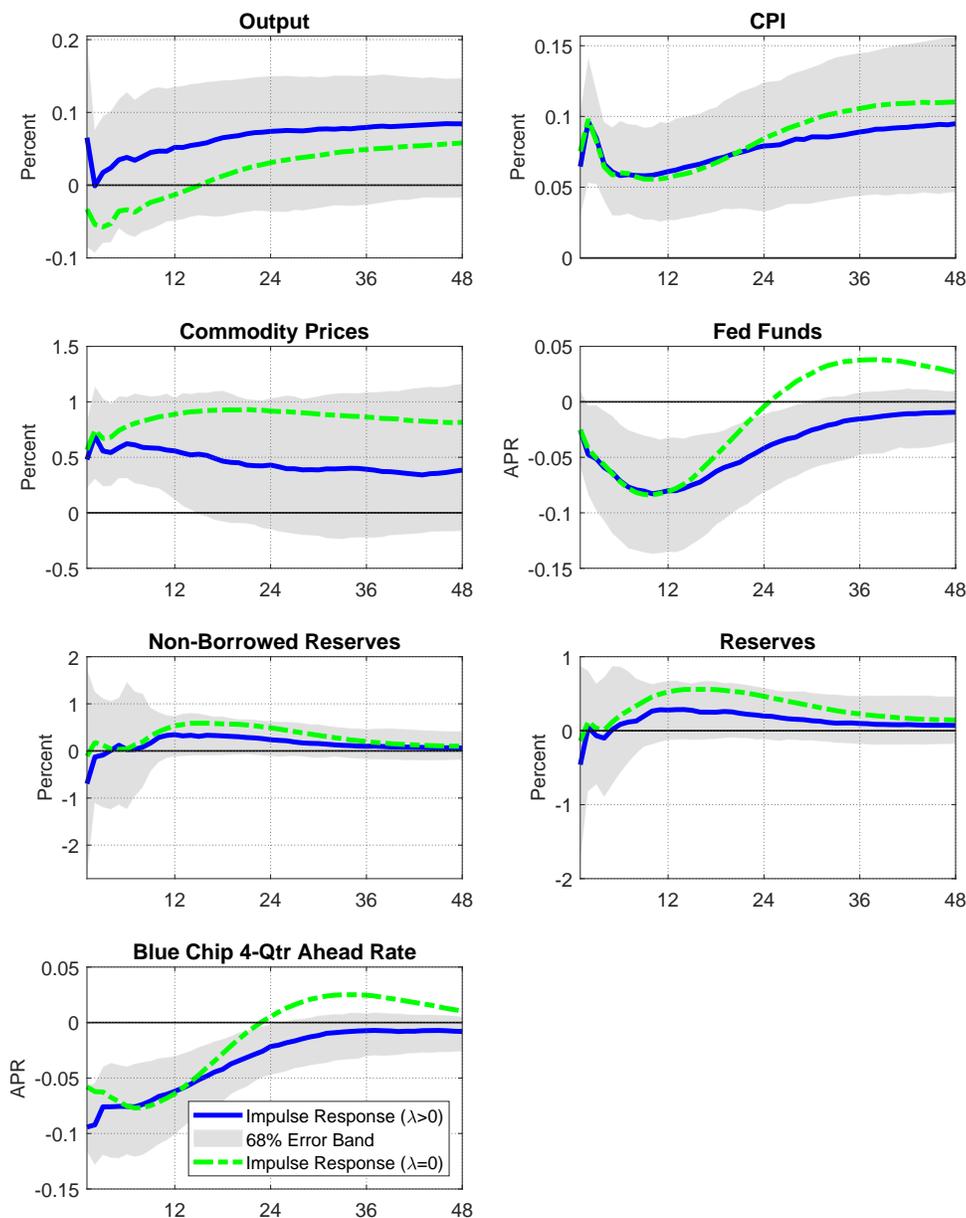
Notes: This figure shows the impulse responses to an identified forward guidance shock using sign restrictions as well as our forecast consistent prior. The solid blue line is the median response and the shaded region is the 68% error band. The green-dashed line shows the median impulse response to an identified forward guidance shock using only sign restrictions. The estimation sample period is 1994-2015.

Figure B.2: FORWARD GUIDANCE SHOCK ALTERNATIVE SVAR SPECIFICATION: HIGH-FREQUENCY CALIBRATION OF LAMBDA



Notes: This figure shows the impulse responses to an identified forward guidance shock using sign restrictions as well as our forecast-consistent prior now calibrated to maximize the correlation between our structural VAR forward guidance shocks and high-frequency financial market measure of forward guidance shocks as constructed by Swanson (2021). The solid blue line is the median response and the shaded region is the 68% error band. The green-dashed line shows the median impulse response to an identified forward guidance shock using only sign restrictions. The estimation sample period is 1994-2007.

Figure B.3: FORWARD GUIDANCE SHOCK ALTERNATIVE SVAR SPECIFICATION: JOINT MCMC DRAWS



Notes: This figure shows the impulse responses to an identified forward guidance shock using sign restrictions as well as our forecast-consistent prior over the full set of VAR parameters. The solid blue line is the median response and the shaded region is the 68% error band. The green-dashed line shows the median impulse response to an identified forward guidance shock using only sign restrictions. The estimation sample period is 1994-2007.

Table C.1: Comparing Alternative Approaches to Identifying Forward Guidance Shocks
 Baseline DGP: NK DSGE Model with Noisy Survey Forecasts

	Pure Sign Restrictions	Forecast Consistency with Sign Restrictions	Minimum Forecast Discrepancy with Sign Restrictions
Identification	Set	Set	Point
Forecast Discrepancy (RMSE, annualized bps) [†]	10.37 [‡]	5.85 [‡]	5.52
	Weight of SVAR-identified Forward Guidance Shock on DSGE Shocks		
Forward Guidance Shock	0.32 [‡]	0.82 [‡]	0.84
Demand Shock	-0.50 [‡]	-0.48 [‡]	-0.48
Supply Shock	-0.50 [‡]	0.00 [‡]	0.04
Noise Shock	0.63 [‡]	0.31 [‡]	0.23

[†] *Conditional* discrepancy following an SVAR-identified 100 basis point forward guidance shock (RMSE, annualized bps).

[‡] Denotes median across posterior set. In the pure sign restrictions case, we therefore caution that the median correlations are purely an artifact of the Haar prior (Baumeister and Hamilton, 2015).

Notes: Each VAR is a 5-lag VAR estimated on one sample consisting of 50,000 DSGE model-generated observations with the number of selected lags based on the AIC.

Table C.2: Robustness of Monte Carlo Results to Alternative DGPs

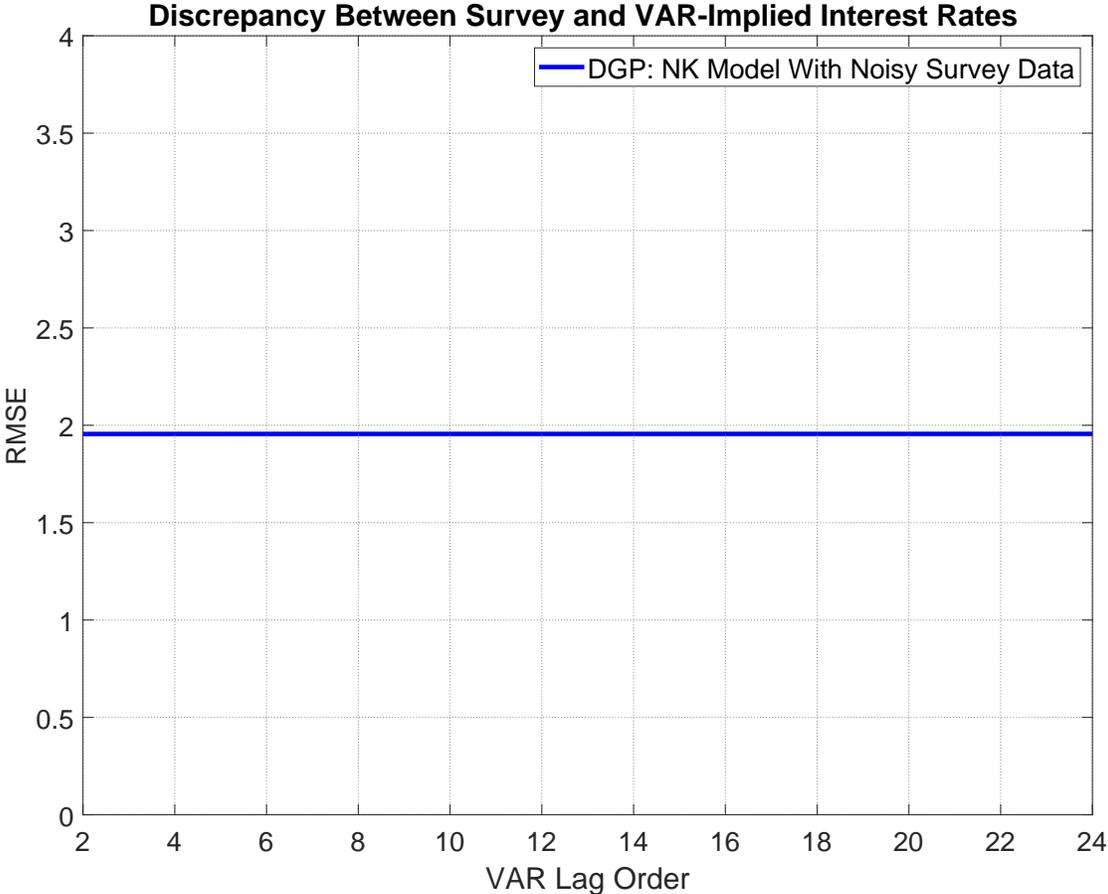
A: 1-period Ahead Forward Guidance Allowing For Policy Reaction to Forward Guidance		
	Pure Sign Restrictions	Forecast Consistency with Sign Restrictions
Forecast Discrepancy [†]	6.91 [‡]	3.59 [‡]
	Weight of SVAR-identified Forward Guidance Shock on DSGE Shocks	
Forward Guidance Shock	0.31 [‡]	0.49 [‡]
Demand Shock	-0.48 [‡]	-0.76 [‡]
Supply Shock	-0.52 [‡]	-0.37 [‡]
Noise Shock	0.64 [‡]	0.21 [‡]
Forward Guidance Shock (t- τ)	-0.20 [‡]	0.81 [‡]
B: 4-period Ahead Forward Guidance Allowing For Policy Reaction to Forward Guidance		
	Pure Sign Restrictions	Forecast Consistency with Sign Restrictions
Forecast Discrepancy [†]	6.87 [‡]	3.60 [‡]
	Weight of SVAR-identified Forward Guidance Shock on DSGE Shocks	
Forward Guidance Shock	0.23 [‡]	0.42 [‡]
Demand Shock	-0.50 [‡]	-0.90 [‡]
Supply Shock	-0.53 [‡]	0.02 [‡]
Noise Shock	0.64 [‡]	0.15 [‡]
Forward Guidance Shock (t- τ)	-0.23 [‡]	0.84 [‡]
C: Noisy Survey Forecasts But Perfect Forecast Consistency for Forward Guidance Shocks		
	Pure Sign Restrictions	Forecast Consistency with Sign Restrictions
Forecast Discrepancy [†]	9.82 [‡]	0.28 [‡]
	Weight of SVAR-identified Forward Guidance Shock on DSGE Shocks	
Forward Guidance Shock	0.38 [‡]	0.81 [‡]
Demand Shock	-0.47 [‡]	-0.58 [‡]
Supply Shock	-0.52 [‡]	0.03 [‡]
Noise Shock	0.61 [‡]	0.09 [‡]

[†] *Conditional* discrepancy following an SVAR-identified 100 basis point forward guidance shock (RMSE, annualized bps).

[‡] Denotes median across posterior set. In the pure sign restrictions case, we therefore caution that the median correlations are purely an artifact of the Haar prior (Baumeister and Hamilton, 2015).

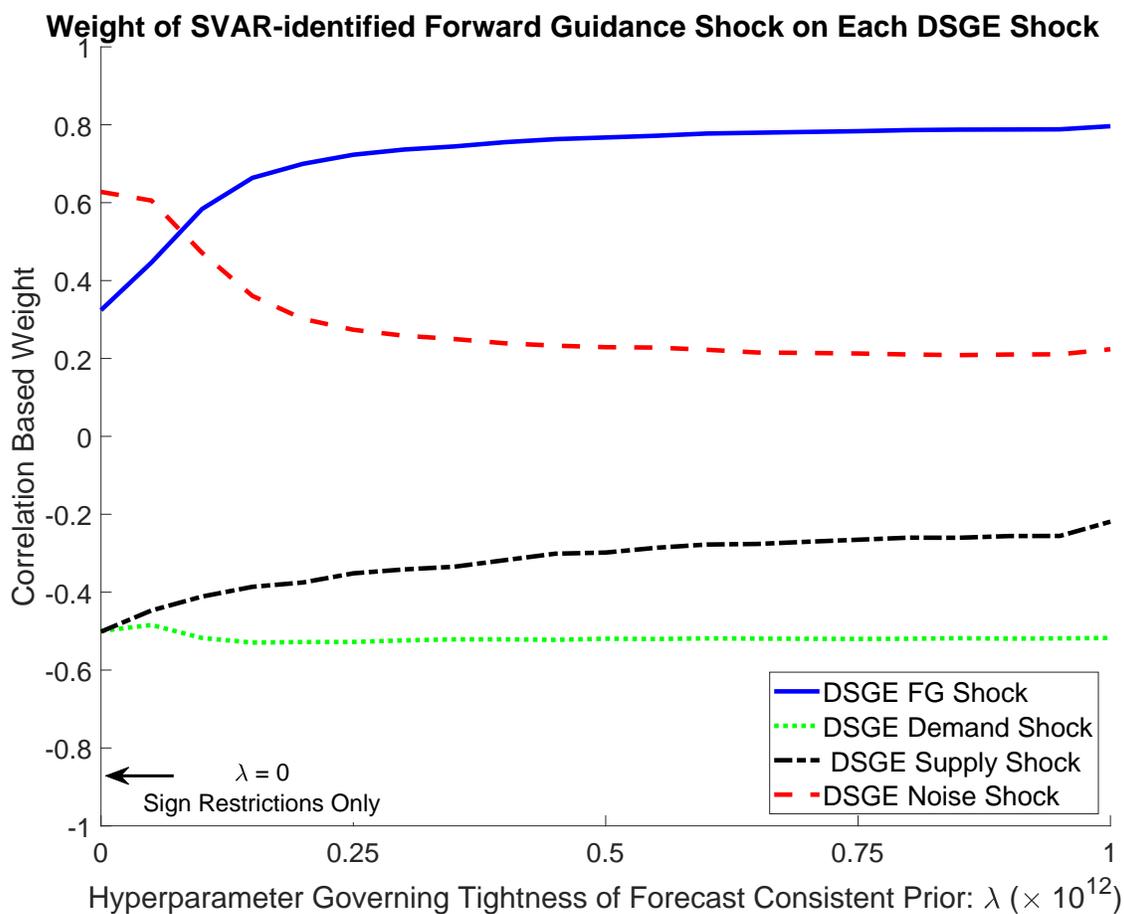
Notes: Each VAR is estimated on one sample consisting of 50,000 DSGE model-generated observations with the number of selected lags based on the AIC.

Figure C.1: MONTE CARLO SIMULATION RESULTS: DEGREE OF UNCONDITIONAL FORECAST CONSISTENCY



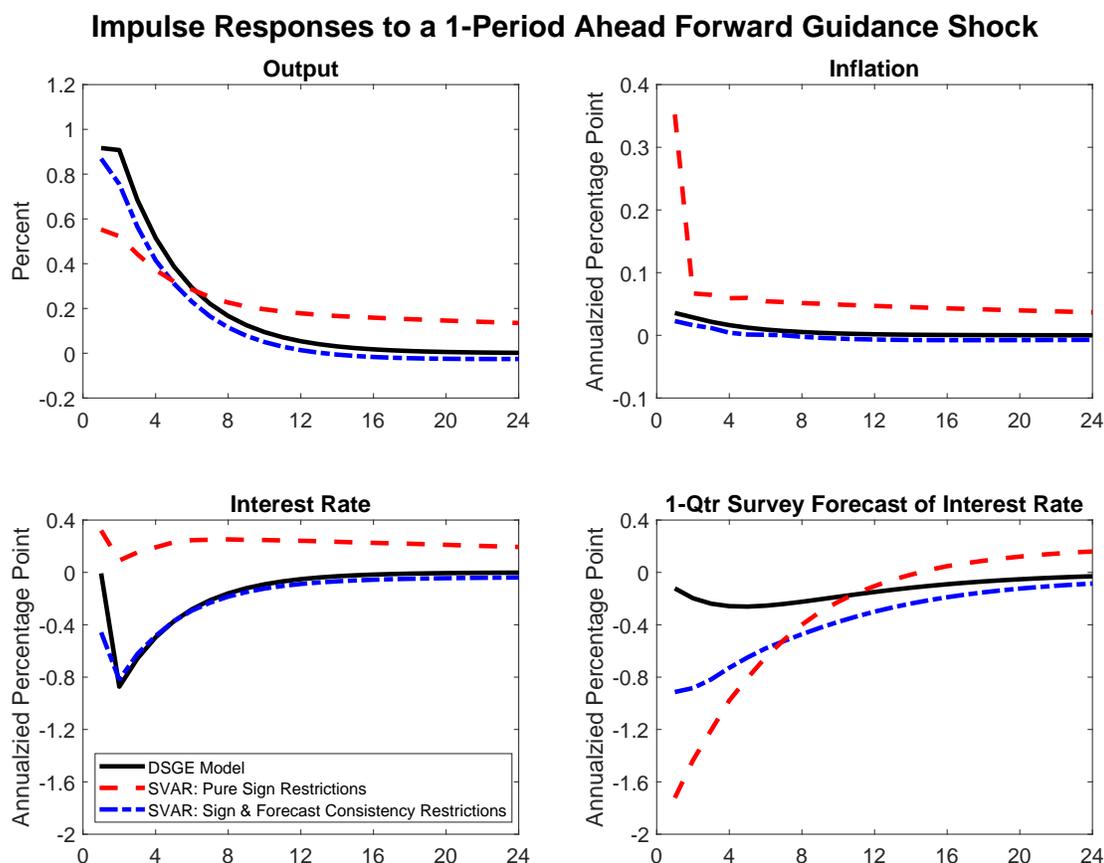
Notes: This figure shows the root-mean squared error between the (noisy) survey and VAR-implied 1-period ahead interest rate forecasts across various lag-lengths of the VAR. Each VAR lag-specification is estimated on one sample consisting of 50,000 DSGE model-generated observations.

Figure C.2: MONTE CARLO SIMULATION RESULTS: RECOVERING FORWARD GUIDANCE SHOCKS WITH SIGN RESTRICTIONS AND THE FORECAST CONSISTENT PRIOR



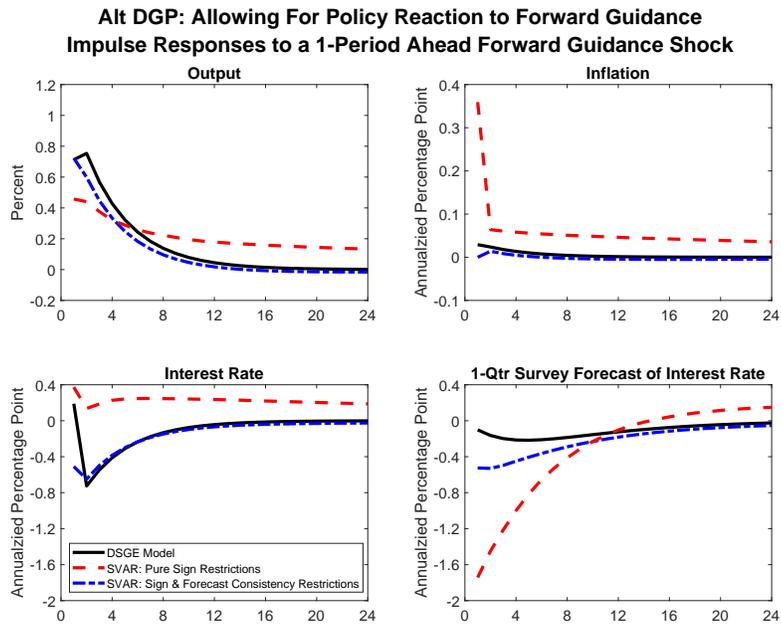
Notes: This figure shows the correlation-based weights that forward guidance shocks identified from SVAR models place on various structural shocks from a New-Keynesian DSGE model with noisy survey forecasts which serves as the data-generating process. These weights are shown for alternative values of λ , which governs the tightness of the forecast consistent prior. When $\lambda = 0$, only sign restrictions are used to identify forward guidance shocks. The VAR is a 5-lag VAR estimated on one sample consisting of 50,000 DSGE model-generated observations with the number of selected lags based on the AIC.

Figure C.3: MONTE CARLO SIMULATION RESULTS FROM A NK DSGE MODEL WITH NOISY SURVEY FORECASTS: ESTIMATING FORWARD GUIDANCE SHOCK IMPULSE RESPONSES WITH SIGN RESTRICTIONS AND THE FORECAST CONSISTENT PRIOR

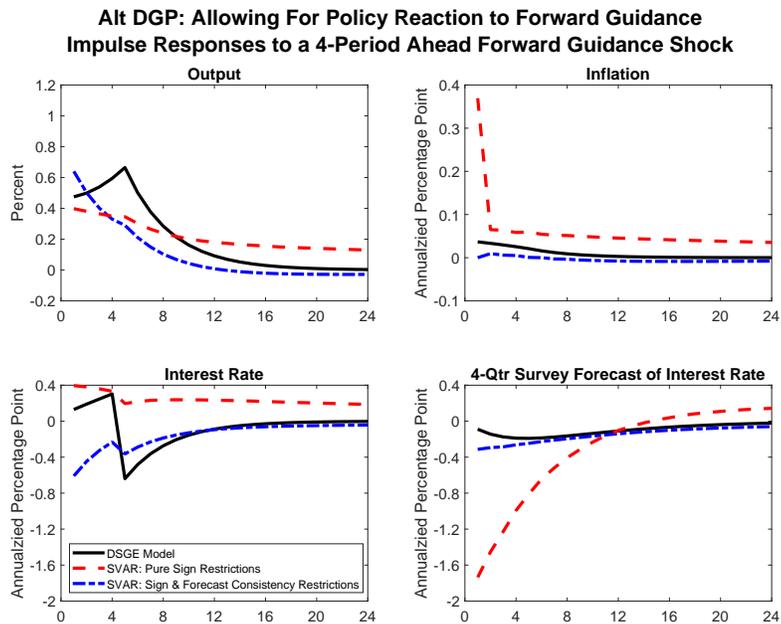


Notes: This figure shows the estimated impulse responses to a forward guidance shocks from the DSGE model which serves as the DGP, the SVAR model estimated on simulated data and identified using solely sign restrictions, and the SVAR model estimated on simulated data and identified using sign restrictions and forecast consistency restrictions. Each SVAR is a 5-lag VAR estimated on one sample consisting of 50,000 DSGE model-generated observations with the number of selected lags based on the AIC. The DSGE model is the noisy survey forecasts model wherein survey forecasts exhibit information rigidities along the lines empirically documented by Coibion and Gorodnichenko (2012).

Figure C.4: ROBUSTNESS OF MONTE CARLO RESULTS TO ALTERNATIVE DGPs: IMPULSE RESPONSES



(a) ALLOWING FOR ENDOGENOUS RESPONSE FROM THE POLICY RULE TO FORWARD GUIDANCE



(b) 4-PERIOD AHEAD FORWARD GUIDANCE SHOCKS WHILE ALLOWING FOR ENDOGENOUS RESPONSE FROM THE POLICY RULE TO FORWARD GUIDANCE

Notes: This figure shows the estimated impulse responses to a forward guidance shocks from alternative DSGE models which serve as the DGP, the SVAR model estimated on simulated data and identified using solely sign restrictions, and the SVAR model estimated on simulated data and identified using sign restrictions and forecast consistency restrictions. Each SVAR is estimated on one sample consisting of 50,000 DSGE model-generated observations with the number of selected lags based on the AIC.

Figure D.1: SAMPLE DATA 1970:Q2-2017:Q4

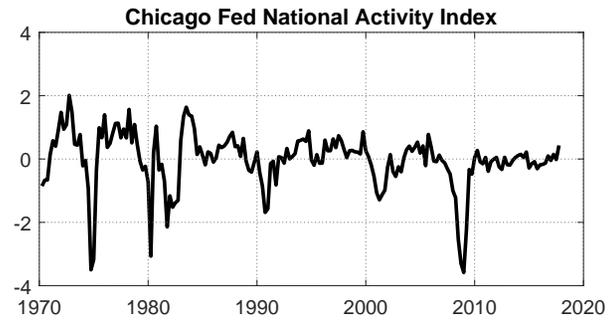
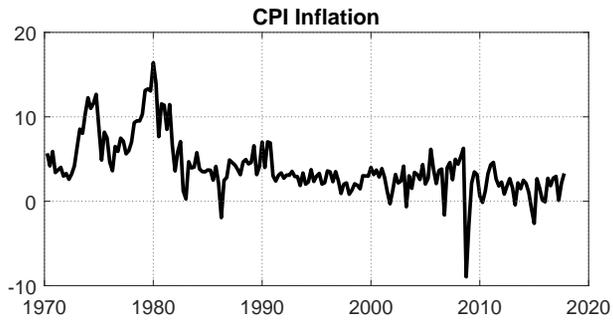
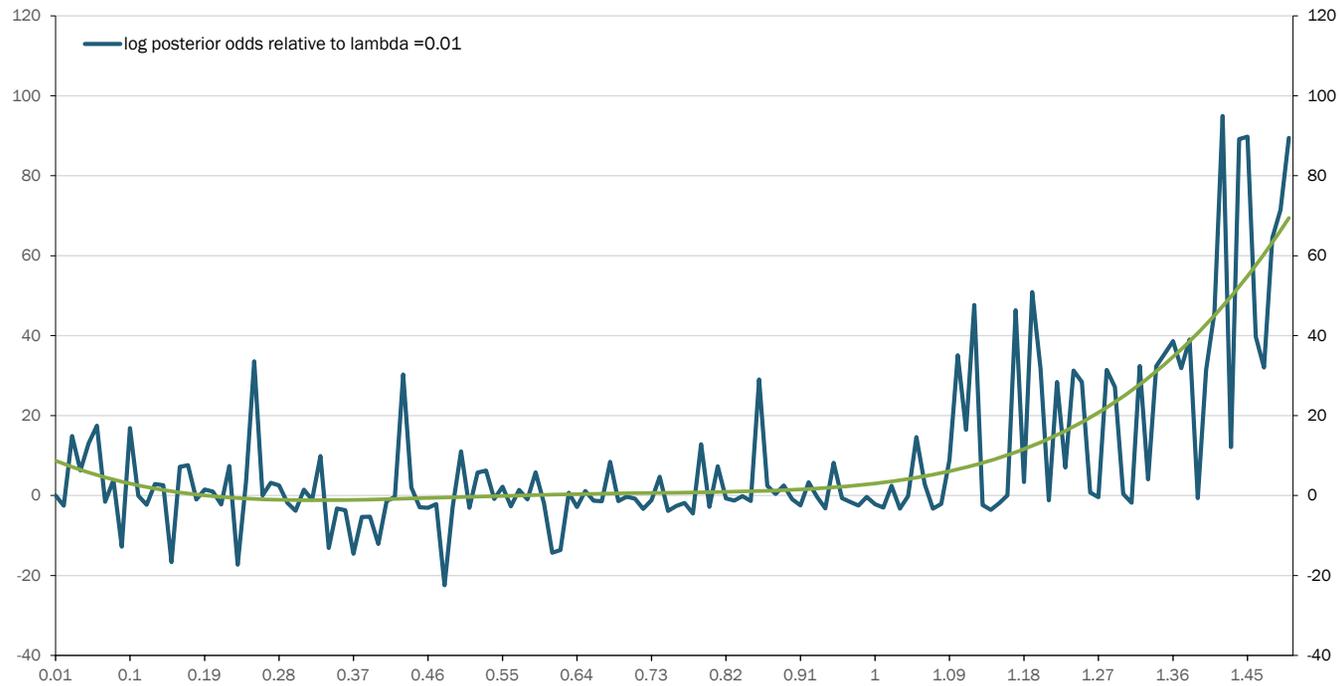
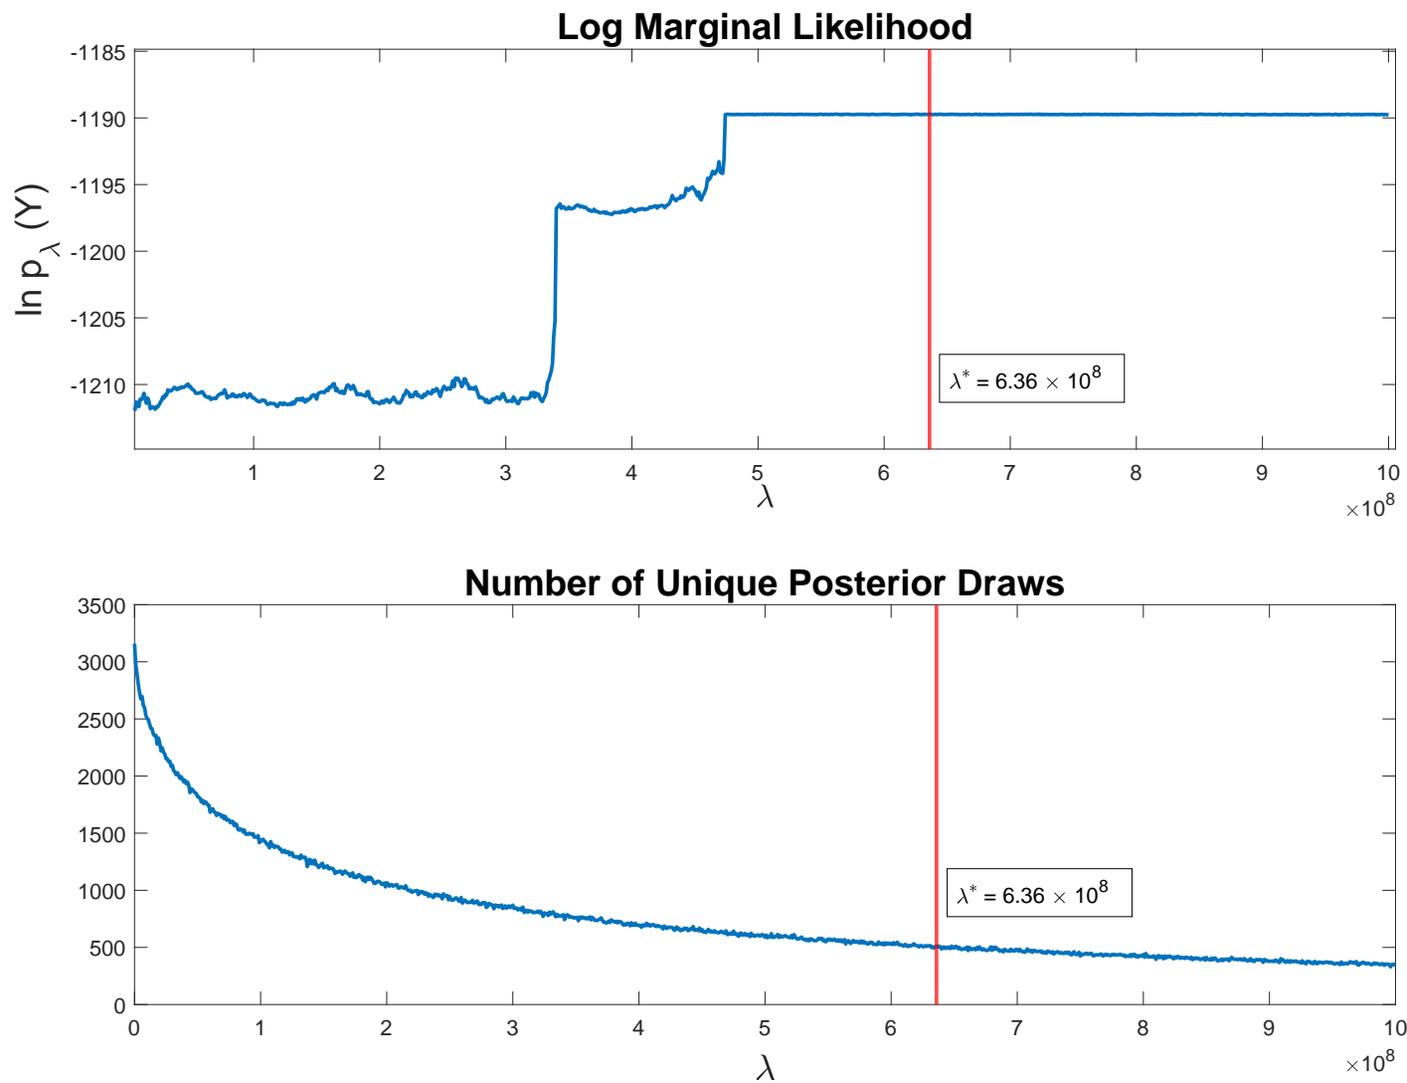


Figure E.1: LOG POSTERIOR ODDS RATIO: TVP-VAR APPLICATION



Notes: The blue line represents the log posterior odds ratio as a function of λ while the green line describes the smooth trend line using a 4-th degree polynomial of λ . The log posterior odds ratio is calculated by: $\ln\left(\frac{p_{\lambda}(y^T)}{p_{\lambda=0.01}(y^T)}\right)$.

Figure E.2: MARGINAL LIKELIHOOD AND THE NUMBER OF UNIQUE RESAMPLED DRAWS: FORWARD GUIDANCE APPLICATION



Notes: The top chart shows the log marginal likelihood as a function of λ . The bottom chart shows the number of unique posterior draws when resampled with alternative values of λ . Without resampling, the number of total posterior draws is 5000.